Vision, Modeling, and Visualization (2014)
J. Bender, A. Kuijper, T. von Landesberger, H. Theisel and P. Urban (Eds.)

Automated Detection of Anatomical Regions in Magnetic Resonance Images

Márton J. Tóth[1], Tamás Blaskovics[2], László Ruskó[2], Gaspar Delso[3], and Balázs Csébfalvi[1]

[1]Budapest University of Technology and Economics, Budapest, Hungary
[2]GE Healthcare Division, Szeged, Hungary
[3]Klinik für Nuklearmedizin Universitätsspital, Zürich, Switzerland

Abstract

Recognition of body parts in three-dimensional medical images is an important task in many clinical applications. It can facilitate image segmentation, registration methods and it can be the first step of an automatic image-processing workflow. In this paper, we propose an automated method to classify the axial slices of three-dimensional magnetic resonance image series according to the body part they belong to. We apply the Zernike transform to obtain feature vectors representing the structural information of the axial slices. Using machine learning tools statistical correlation is found between the extracted feature vectors and the position of the slices within the human body. The initial classification is filtered by a dynamic programming based error correction method that takes the correct sequence of anatomy regions into consideration to eliminate the false recognitions. Using our approach, different body regions can be recognized at high precision rate.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Object recognition

1. Introduction

Magnetic Resonance Images (MRI) are now important sources of information in many cases of clinical decision making. Thanks to the development of data acquisition methods and computer technology, the clinical application of medical imaging and image processing shows an increasing rate. This trend includes the increasing amount of data as well. Therefore, the physicians have to spend more and more time with the analysis of these images. Thus the amount of time spent on analyzing an image is getting more and more limited.

This makes it necessary to develop Computer Aided Diagnostic (CAD) systems that can automatically provide as much information as possible. This can significantly aid the work of the physicians who will have more time in this way to make the diagnosis and to treat the patients.

The images usually do not contain any information about the represented body region [GWKT02], as the corresponding DICOM header tag should be filled in manually by the operator of the imaging device. However, it would be really useful to know the body regions seen in the image as the medical image processing algorithms are usually organ or body region specific. By using the proposed method, it is easier to initialize segmentation and registration algorithms, to optimize the image-processing workflow or the user interface of clinical software applications (e.g. by offering only the tools that are useful in a certain anatomical region). The proposed method can be used as one of the building blocks of CAD systems as it can provide some basic information for an automatic or semi-automatic diagnostic workflow.

2. Related Work

Although classifying slices within a human body could be useful in various medical image-processing workflows, this area of research has not received much attention yet. Even though there are existing methods to determine the body region from a topogram [Bür08], the general way to navigate through a body scan is to localize landmark points as references and use them as a relative coordinate system. For instance, Haas et al. [HCS*08] introduced an approach to create a navigation table using eight landmark points. Other

DOI: 10.2312/vmv.20141279

similar methods are also known. Seifer et al. [SBZ*09] proposed a method to detect invariant slices and landmark points in full body scans by using HAAR image features and Probabilistic Boosting Trees. Their method is able to detect nineteen salient and robust landmarks within a full body scan but it cannot be used to localize and to classify single slices as the method detects only the landmark points and further processing would be needed to perform the classification itself.

Other techniques were developed to estimate the location of axial-slices in the human body, but these methods try to give the position of the axial slices and do not provide the labeling itself. For Computer Tomography (CT) scans Feulner et al. [FZS*09] developed a method which can determine the relative position of the examined slice series. Their method requires approximately 50 slices and the average error of the position estimation lies between 44mm for small query volumes and 16.6mm for large query volumes. According to the authors, their method does not perform well when localizing single slices only. Emrich et al. [EGK*10] have managed to determine the position of a single slice within the body with an overall error of 2.83cm, but their algorithm works only on CT images.

For CT images Nakamura et al. [NLIS08] proposed a method that can directly label the axial slices of an image series. Their results are highly promising, but it cannot be directly applied to MRI images. Their feature vector construction is based on the intensity values of the CT scans that are much more reliable than the values of the MRI scans. The CT intensities are standardized by the Hounsfield scale and they do not fluctuate so unpredictably as the intensities in MRI images.

In summary, the existing methods to determine the body regions are only for CT modality and they are not applicable to MRI images. They usually use landmark points for localization and they do not classify the individual slices. In this paper, we propose a method that can process MRI images, classifies the slices and extends the reliability of the classification by an error-correction method.

3. Materials

The proposed algorithm was trained and tested on a set of 49 MRI images. The images were acquired for research purposes by the University Hospital Zürich. The studies were acquired with the Discovery MR750w scanner using the same (LAVA-Flex sequence, T1 weighted, FA/TR = 5°/3.7ms, acquisition matrix = 256x128, 75% phase FOV, scan time 17s, TE1/TE2 = 1.15ms/2.3ms) MRI protocol. The majority of the images covered all anatomy regions including head, neck, chest, abdomen, pelvis, and some parts of the legs. The test images involved male and female patients of different age (adults), level of obesity, and the position of the arms were varying among examinations. The resolution of each image slice was 512 x 512 pixels. The slice number of the exams ranged between 26 and 864 (average 447). The pixel size varied between 0.39mm and 1.37mm (average 0.91mm), and the slice thickness was between 0.47mm and 6.80mm (average 5.79mm). All test cases have been annotated manually. The following region labels were defined: LEG, PELVIS, ABDOMEN, CHEST, HEAD and EMPTY. The selection of the labels was based on the basic structure of the human body. The LEG region is defined from the beginning of the image series to the point where the two legs join. There starts the PELVIS that ends at the top of the pelvic bone. The ABDOMEN area ends at the beginning of the lungs. The CHEST lasts till the neck and above it is the HEAD. EMPTY areas are usually above the HEAD but sometimes they could appear under the LEG as well.

4. Our Method

A successful solution of the recognition task is a multi-step process. First, an appropriate image representation has to be identified that can provide meaningful feature vectors, as these vectors are used as the input for a machine learning tool that solves the primary classification task. After the initial classification is done, an error correction method is applied. As the primary slice classification runs only on individual slices, we lose the connection between the slices and an error correction method is needed to filter the outliers and provide a reliable classification along the whole image series.

The proposed algorithm performs five steps to solve the classification problem. Each step processes the output of the previous stages and provides input for the next level. The first four steps run on individual slices and the last step uses the whole image series to perform the correction. The steps are the following:

1. Slice normalization.
2. Feature extraction (Zernike transform).
3. Principal Component Analysis (PCA) for dimension reduction.
4. Random forest classification.
5. Continuity check and error correction.

Figure 1 demonstrates the proposed workflow of the training and the classification.

4.1. Slice Normalization

MRI images have a wide range of variety in image properties. The imaging devices have several acquisition modes, which can affect the appearance of the resulting images. The choice of pulse sequence and its associated parameters can significantly modify the results. As the selected protocol has a direct impact on the intensity values in the image and can change the visibility of the structures within the body, our approach deals only with T1 weighted images (which are common in clinical practice). The selection of the appropriate field of view (FOV) value is also a very important aspect.

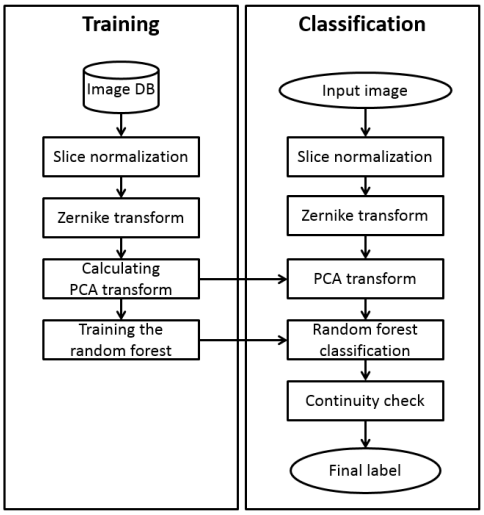

Figure 1: The structure of proposed method.

Different FOV settings provide images with different zoom factors. The FOV values are stored in the DICOM header by the device itself, so it is reliable and can be used to normalize the images for a standardized FOV. We had two types of series. One type was recorded with 70cm FOV and the other one was recorded with 48cm FOV. The 48cm FOV was selected as the standard FOV value. In this case the meaningful area of the image is larger as the image of the patient fills the available area to a larger extent.

Other image artifacts can be caused by the positioning of the examined patient. As it will be described later in this paper, the Zernike transform can be formed to be rotation invariant so it is not necessary to take any further steps to determine the precise orientation of the patient, as a rotation of few degrees is tolerated by the transform. However the translation of the patient in the transverse plane can cause major changes in the transform, so it has to be eliminated from the image. This can be done by calculating the $(p+q)$ [Hu62] order regular moment of the image, which can be formalized in the following way:

$$m_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x^p y^q f(x,y)dxdy, \qquad (1)$$

where m_{pq} is the moment and $f(x,y)$ is the function that defines the image. To obtain translation invariance, the center of mass of the image is moved into the center of the image. This can be done by a transform from $f(x,y)$ to $f(x+(x^c-x^s,y+(y^c-y^s)))$ where $x^s=\frac{m_{10}}{m_{00}}$, $y^s=\frac{m_{01}}{m_{00}}$, while x^c and y^c indicate the coordinates of the center of the image.

Figure 2 demonstrates the steps of the slice normalization. In the first image the original input can be seen, followed by the result of the translation normalization, and finally the FOV normalization is shown.

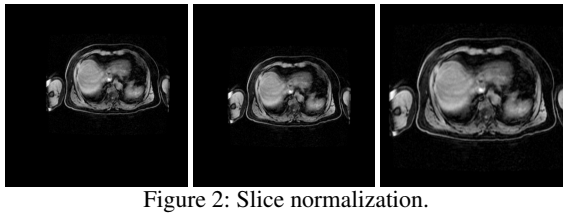
Figure 2: Slice normalization.

4.2. Zernike transform

For a human observer the slice classification task is obvious. One does not need to examine the image in details, only the large scale structures of the images contain enough information to perform the anatomy classification. Therefore a representation was chosen that can capture the topological information itself. Topological descriptors are not so widely affected by the typical artifacts of the MRI images as long as the topology is recognizable and they are not so sensitive [BSA91] to the fluctuation of the image intensities. This Zernike structure description is useful for the representation of the slices of three-dimensional medical images, as their FOV is circular and the axial cross sections of the human body can be represented by elliptic shapes as well.

The Zernike transform [Zer34] is one of the state of the art global shape description methods. It was originally introduced for image description by Teague et al. [Tea80]. Since its publication, it has been used in several applications [KH90] [QRAH07] [SFPS11] and has a wide literature. Its properties have been studied extensively and they have been summarized and compared with other moments by Belkasim et al. [BSA91].

According to the results of Belkasim, the Zernike moments have several favorable properties over other geometrical moments. They can represent the image with no redundancy or overlapping information between the moments. They can be formalized in a rotation invariant way with polar coordinates. Rotating the image does not change the magnitude of the moments. This could be an important property if the patients were not laid fully on their back. Another main advantage of the Zernike moments is that they can be constructed to an arbitrarily high order. This feature, with the orthogonality, is very important to the reconstruction of the original image from the moments. The orthogonality enables to separate the information content of each moment and the reconstruction can be achieved by a simple addition of these individual contributions. The selection of the required maximum order can be done by the evaluation of the image representation ability of the features. The reconstructed image can be compared to the original image and the applied maximum order is sufficient if the reconstructed image is close to the original image.

The Zernike moments can be expressed as a set of complex polynomials $V_{mn}(x,y)$ over the interior of the unit disk

Order of the transform	5	10	15	25	35
Number of moments	24	72	144	364	684

Table 1: Number of Zernike moments.

i.e., $x^2+y^2 \leq 1$. In polar coordinates the form of the polynomials is:

$$V_{mn}(x,y) = V_{mn}(r,\theta) = R_{mn}(r)exp(jm\theta). \quad (2)$$

where n is a non-negative integer, m is integer, $n-|m|$ is even and $|m| \leq n$; $r = \sqrt{x^2+y^2}$ is the length of the vector from the origin to pixel (x,y); $\theta = arctan(y/x)$ is the angle between the vector to the point (x,y) and the x axis. R_{mn} is the Radial polynomial defined as:

$$R_{mn}(r) = \sum_{k=0}^{(n-|m|)/2} \frac{(-1)^k (n-k)!}{k!((n+m)/2-k)!((n-m)/2-k)!} r^{n-2k}. \quad (3)$$

Then the two-dimensional Zernike moment of order n and m for a function $f(x,y)$ can be defined as:

$$Z_{mn} = \frac{n+1}{\Pi} \iint_{unitdisk} f(r,\theta) V^*_{nm}(r,\theta) dr d\theta, \quad (4)$$

where $V^*_{mn}(x,y) = V_{n,-m}(x,y)$. The applied maximum order of the transform controls the amount of the captured details. Using higher orders make the reconstruction more accurate as the transformed values contain more information, but the number of moments grows rapidly. The number of the moments is shown in Table 1 according to the maximal order of the Zernike transform.

To visually evaluate the captured information content by the transform, a back transform is required. In this way we can specify the applied maximal order of the transform.

Suppose that one knows all the moments Z_{mn} of the function $f(r,\theta)$ up to a given order N_{max}. To reconstruct a function $\hat{f}(r,\theta)$ whose moments exactly match those of $f(r,\theta)$ up to the order N_{max} one should use the following equation:

$$\hat{f}(r,\theta) = \sum_{n=0}^{N_{max}} \sum_{m} Z_{nm} V_{mn}(r,\theta), \quad (5)$$

where the same constraints are applied as in (2). Note that as N_{max} approaches infinity, $\hat{f}(r,\theta)$ approaches $f(r,\theta)$. The results of the reconstruction are shown in Figure 3.

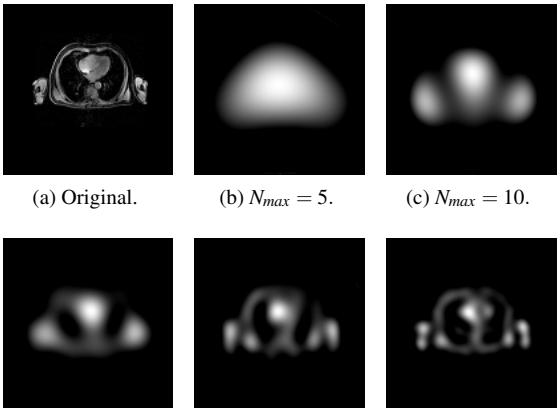

(a) Original. (b) $N_{max} = 5$. (c) $N_{max} = 10$.

(d) $N_{max} = 15$. (e) $N_{max} = 25$. (f) $N_{max} = 35$.

Figure 3: Reconstruction from Zernike moments.

4.3. Principal Component Analysis

Although, the Zernike-transform captures all the necessary details needed for slice classification, it does not mean that all of its attributes are necessary. It is worthwhile to find a way to filter the unnecessary data from the observations to produce an optimal feature vector. This is an important step. The original feature vector contains the information we need, but the high number of dimensions makes the machine learning algorithm complex, and makes the number of parameters too high. Therefore a dimension reduction shall be applied. PCA [Jol05] is a mathematical procedure that uses an orthogonal transform to map a series of observations of correlated variables into an uncorrelated set of variables. These variables are called principal components. The PCA is used for dimension reduction, while the variance of the samples is preserved as much as possible. The procedure is defined in a way that the first principal component has the highest variance and each succeeding component has the largest possible variance under the constraint that it has to be orthogonal to the previous components.

4.4. Random Forest

A random forest [Bre01] is composed of multiple independently trained decision trees. Decision trees are popular classification tools, but they are known to suffer from overfitting. It means that their generalization ability is poor. However, it has been shown [Bre01] that a composition of many randomly trained decision trees, as known as random forest, has much better generalization ability and it can maintain the advantages of the conventional decision trees. Other classifiers (Neural Network, Support Vector Machine) were also tested during the development, but their results were not significantly better than the results of the random forest. For development purposes we used the WEKA [HFH*09] toolkit, but in the final system we had to use our own implementation and therefore the simplicity of the learning tool was an important aspect. For these reasons we chose the random forest as our primary classification tool.

4.5. Continuity Check and Error Correction

The primary classification done by the random forest, works with single slices only, thus we lose valuable information. The correlation of the relative positions of the slices play a key role in the correct recognition process. As the sequence

of the regions and their sizes within the human body are not random, the label, assigned to a specific slice, is highly correlated with the labels of the nearby slices. If we take this information into consideration, we can achieve a more precise, outlier-free and reliable classification.

To achieve an optimal labeling, we propose a graph based shortest path finding method, which is based on Dijkstra's algorithm [Dij59]. The random forest provides a confidence value for each slice for each body region. We can define the order of the body regions as LEG, PELVIS, ABDOMEN, CHEST, HEAD and EMPTY. Let $R = \{r_1, r_2, \ldots, r_6\}$ denote the set of anatomy labels respectively and $S = \{s_1, s_2, \ldots, s_M\}$ the set of the slice indices, where M is the number of slices. The order of the slice sequence is from the leg to the head. Given these structures, we can define a graph in the following way. Each node represents a classification of a slice. The set of the nodes is denoted by $N = \{(r_i, s_j) : r_i \in R, s_j \in S\}$. The directed edges can be defined as $E = \{((r_k, s_{j-1}), (r_i, s_j)) : (r_i, s_j) \in N \wedge (r_k, s_{j-1}) \in N \wedge (k = 1 \vee k = i-1)\}$. This structure defines a graph, where each node represents a possible classification and the edges connect only the allowed transitions. In this way, the classification of the whole image series is forced to follow the order of the body regions. The structure of the defined graph is illustrated in Figure 4.

To perform a shortest path finding, we have to define the costs of moving along the edges. This can be done in several ways. We defined the cost function as $C((r_k, s_{j-1}), (r_i, s_j)) = 1 - \text{confidence}((r_i, s_j))$, where $\text{confidence}((r_i, s_j))$ refers to the output of the random forest and it defines the confidence value of that the slice s_j is in the body region r_i. This function describes how much the decision of the path finding algorithm differs from the results of the random forest. To obtain a labeling from this graph structure, an optimal path has to be found from one of the nodes belonging to the first slice to one node of the last slice. The solution of this task can be completed by dynamic programming. We can define a two-dimensional distance matrix D, with M rows and 6 columns. Each cell $D[i, j]$ corresponds to node (r_j, s_i) in the graph. The first row and column has to be defined manually, as this step initializes the algorithm. This can be done as follows:

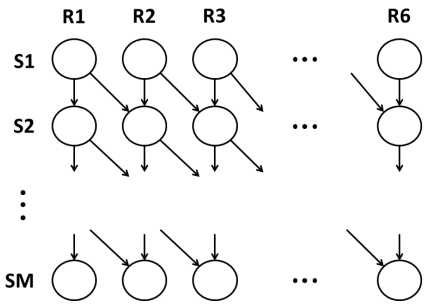

Figure 4: Correction-Graph structure.

$$D[1, j] = 1 - confidence((r_j, s_1)), 1 < j \leq 6, \quad (6)$$

$$D[i, 1] = D[i-1, 1] + 1 - confidence((r_1, s_i)), 2 < i \leq M. \quad (7)$$

Equation 6 represents that the first image of the series can be positioned anywhere in the body and we can rely on the results of the machine learning tool only. However, this uncertainty will be corrected later in the process, as the shortest path finding algorithm will trace to the optimal choice. Equation 7 defines the fill in role for the first column of the table. This is necessary to get a simple representation of the general rule for filling the matrix. From this point we can fill the matrix cell by cell, starting it from the second column of the second row. The general form to compute element $D[i, j]$ of the distance matrix can be formalized as follows:

$$D[i, j] = min(c_1 + D[i-1, j], c_2 + D[i-1, j-1]), \quad (8)$$

where $c_1 = C((r_j, s_{i-1}), (r_i, s_j))$ and $c_2 = C((r_{j-1}, s_{i-1}), (r_i, s_j))$. To find the optimal path in the graph, its end point has to be identified first. This can be done by finding the region of the last slice with the lowest total cost. This will be the final label of the last slice. Starting from that, one can find the optimal path by tracking back from the end point to a label of the first slice. The optimal path defines our final classification. This process is illustrated in Figure 5.

	R1	R2	R3	R4	R5	R6
S1	1	0	1	1	1	1
S2	2	0,1	0,3	2	2	2
S3	3	0,1	0,6	0,8	3	3
S4	4	0,2	0,8	1,6	3,1	3,8
S5	5	1,5	0,2	2	3,2	4,1
S6	6	2,6	0,3	2,1	3,6	4,5
...	8,3	8,5	8,2	8,1	1,1	8,3
...	9,4	9,4	8,4	8,3	1,2	8,5
...	10,1	11,2	9,5	8,7	9	1,2
...	12,5	13,5	10,1	9,8	9,7	1,3
SM-1	13,8	13,7	12,6	10,7	10,2	1,3
SM	16,7	15,5	14,5	13,7	12,3	1,4

Figure 5: Shortest path finding.

5. Training

There are several properties of the proposed method that have to be set. The first parameter is the order of the Zernike transform. As it can be seen in Figure 3, the different orders produce different reconstructed images and the feature

vector contains different amount of information. As our inspection showed, the $25th$ order transform collects all the required details that are necessary for an acceptable initial classification. This choice can be also confirmed by a visual comparison.

To avoid the problem of dealing with high dimensional feature vectors, a PCA transform was applied. The size of the transformed feature vector was set to 25. This was motivated by the orthogonal property of the Zernike transform. Its basis functions are orthogonal to each other, therefore the PCA transform is likely to find 25 significant components. This assumption was verified during the training and testing process.

The training of the random forest model required a few considerations. We split the image series database into two groups. Each image series, as a whole, was selected for either training or for testing. This was necessary to avoid the learning of the image specific noise effects, that could influence the testing process. We used 18 image series for training and the data of 31 patients for testing. The larger scale of previously unseen data ensured the robustness of the algorithm, while the applied ten-fold cross validation during the learning process showed the generalization skill of the method. The training process has basically two parameters relevant to the random forest. The first one is the number of the trees and the second one is the number of the variables to test at each split. According to Breiman [Bre01] the latter parameter should be set to $log_2(N+1)$, where N is the length of the feature vector. The number of the trees was chosen to be one hundred. The arrangement used for training and validation is illustrated in Figure 6.

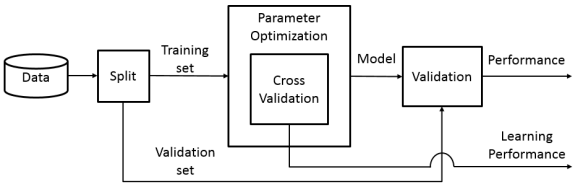

Figure 6: The training process.

6. Results

To measure the overall accuracy of the random forest classification on unseen data, a validation process was applied. Table 2 shows the results of the classification made by the random forest during the validation process. This table is not corrected with the continuity check. As it can be seen, the class precision values vary from 82.59% to 98.50%. The most reliable rate was achieved at the head region and the worst was at the pelvis. The diagonal cells of the table show the number of the correctly labeled slices. The other cells mean the confusion of the labeling. Most of the errors were made between the neighboring regions and only a small percentage of the errors were made because of the complete miss of the labeling. The overall accuracy was 92.95% that makes the primary classification a good basis for further processing.

However, most of the errors were made because of the permutation of the neighboring regions, Table 2 does not show us the problem of discontinuity. This kind of error occurs on the borders of the regions, as the slices are getting more similar to the slices of the neighboring region. Thus the classification tool cannot provide us a reliable decision. The labeling begins to alternate between the two regions and we do not receive a continuous labeling. This problem is illustrated in Figure 7. The image presents the direct result of the random forest classification without the continuity check. The alternating gray strips illustrate the automatically assigned different labels and the bright white lines indicate the borders of the regions denoted by the manual labeling. This continuity problem made it necessary to add the graph

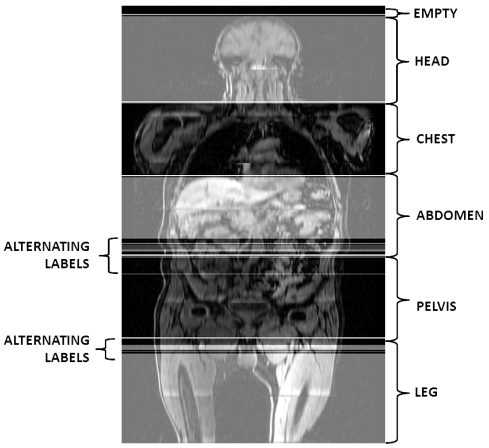

Figure 7: Alternating labels.

based post processing phase to the system. With the proposed method, the labeling of the nearby slices can be taken into consideration when assigning a label to one slice. This method successfully eliminated the problem of the alternating labels and solved the outlier artifacts as well. The results after the application of the correction method are illustrated in Figure 8 and Table 3.

As it can be seen in Table 3, the class precision values of different body parts are between 85.79% and 99.03%. The overall accuracy is increased to 94.20%. The lowest rate is achieved with the pelvis. Thanks to the continuity check, there are no outliers in the labeling and the problem of the alternating labels is also solved. In each test case, the order of the labels is correct and only small shifts can be seen, usually around the borders between the ADBOMEN and the PELVIS regions. However, these small errors between the neighboring regions can be tolerated in most of the clinical applications. Further results can be seen in Figure 9.

	true LEG	true PELVIS	true ABDOMEN	true CHEST	true HEAD	true EMPTY	class precision
pred. LEG	2533	43	4	12	0	0	97.72%
pred. PELVIS	157	2395	334	13	0	1	82.59%
pred. ABDOMEN	1	211	2367	33	0	2	90.55%
pred. CHEST	7	12	53	2372	55	0	94.92%
pred. HEAD	0	0	6	21	3077	20	98.50%
pred. EMPTY	1	0	0	0	16	461	96.44%
class recall	93.85%	90.00%	85.64%	96.78%	97.74%	96.25%	

Table 2: The confusion matrix results of the testing process without correction.

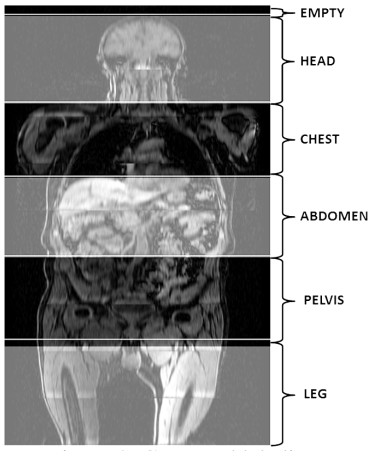

Figure 8: Corrected labeling.

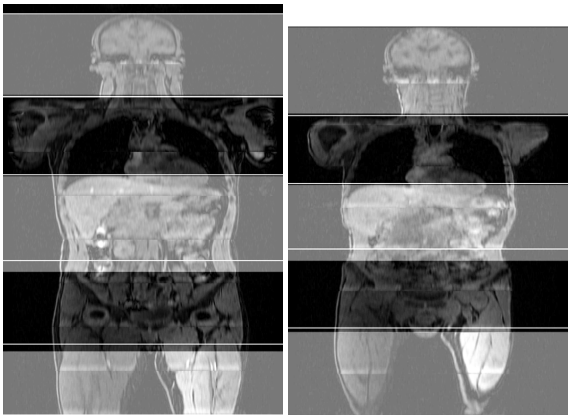

(a) (b)

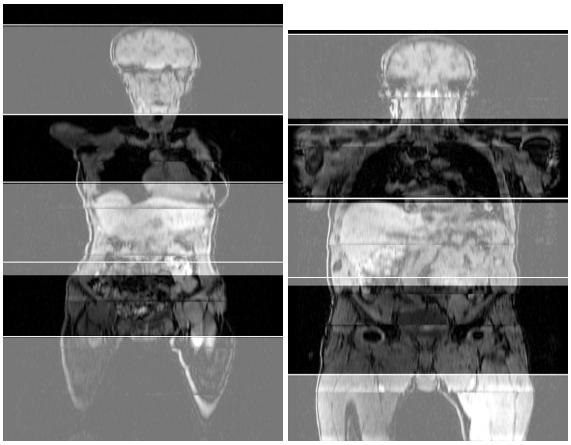

(c) (d)

Figure 9: Examples from the final results.

7. Conclusion

We have demonstrated that our method can solve the labeling task with an overall accuracy of 94.20% for whole body MRI image series. The proposed image normalization and feature vector extraction method can provide an image descriptor that is robust against imaging artifacts and the shape variance of the patients. This can serve as a good basis for the classifications. The results of the random forest tool are corrected with a dynamic programming based continuity check. With this postprocessing, the achieved labeling is continuous and does not contain any outliers. The final results have been tested on large number of MRI series and they have showed that the method is suitable for clinical application.

8. Future work

Our future work will aim the adaptation of the method to CT modality and we have plans to extend our method to partial body scans as well.

9. Acknowledgment

This work was supported by projects TÁMOP-4.2.2.B-10\1-2010-0009 and OTKA 101527.

	true LEG	true PELVIS	true ABDOMEN	true CHEST	true HEAD	true EMPTY	class precision
pred. LEG	2548	43	0	0	0	0	98.34%
pred. PELVIS	151	2426	251	0	0	0	85.79%
pred. ABDOMEN	0	192	2466	35	0	0	91.57%
pred. CHEST	0	0	47	2395	58	0	95.8%
pred. HEAD	0	0	0	21	3077	9	99.03%
pred. EMPTY	0	0	0	0	17	475	96.54%
class recall	94.41%	91.17%	89.22%	97.72%	97.62%	98.14%	

Table 3: Results on the test set with correction.

References

[Bre01] BREIMAN L.: Random forests. *Machine learning 45*, 1 (2001), 5–32. 4, 6

[BSA91] BELKASIM S. O., SHRIDHAR M., AHMADI M.: Pattern recognition with moment invariants: a comparative study and new results. *Pattern recognition 24*, 12 (1991), 1117–1138. 3

[Bür08] BÜRGER C.: Automatic localisation of body regions in CT topograms. *VDM Verlag Dr. Müller, Saarbrücken* (2008). 1

[Dij59] DIJKSTRA E. W.: A note on two problems in connexion with graphs. *Numerische mathematik 1*, 1 (1959), 269–271. 5

[EGK*10] EMRICH T., GRAF F., KRIEGEL H.-P., SCHUBERT M., THOMA M., CAVALLARO A.: CT slice localization via instance-based regression. In *SPIE Medical Imaging* (2010), International Society for Optics and Photonics, pp. 762320–762320. 2

[FZS*09] FEULNER J., ZHOU S. K., SEIFERT S., CAVALLARO A., HORNEGGER J., COMANICIU D.: Estimating the body portion of CT volumes by matching histograms of visual words. In *SPIE Medical Imaging* (2009), International Society for Optics and Photonics, pp. 72591V–72591V. 2

[GWKT02] GÜLD M., WEIN B., KEYSERS D., THIES C.: A distributed architecture for content-based image retrieval in medical applications. URL: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.16.2324. 1

[HCS*08] HAAS B., CORADI T., SCHOLZ M., KUNZ P., HUBER M., OPPITZ U., ANDRÉ L., LENGKEEK V., HUYSKENS D., VAN ESCH A., ET AL.: Automatic segmentation of thoracic and pelvic CT images for radiotherapy planning using implicit anatomic knowledge and organ-specific segmentation strategies. *Physics in medicine and biology 53*, 6 (2008), 1751. 1

[HFH*09] HALL M., FRANK E., HOLMES G., PFAHRINGER B., REUTEMANN P., WITTEN I. H.: The WEKA data mining software: An update. *SIGKDD Explorations 11* (2009). URL: http://www.cs.waikato.ac.nz/ml/weka/. 4

[Hu62] HU M.-K.: Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory 8*, 2 (1962), 179–187. 3

[Jol05] JOLLIFFE I.: *Principal component analysis*. Wiley Online Library, 2005. 4

[KH90] KHOTANZAD A., HONG Y. H.: Invariant image recognition by Zernike moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*, 5 (1990), 489–497. 3

[NLIS08] NAKAMURA K., LI Y., ITO W., SHIMURA K.: A machine learning approach for body part recognition based on CT images. *Proc. SPIE 6914, Medical Imaging 2008: Image Processing 6914* (Mar. 2008), 69141U–69141U–9. URL: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1329303, doi:10.1117/12.768480. 2

[QRAH07] QADER H. A., RAMLI A. R., AL-HADDAD S.: Fingerprint recognition using Zernike Moments. *Int. Arab J. Inf. Technol. 4*, 4 (2007), 372–376. 3

[SBZ*09] SEIFERT S., BARBU A., ZHOU S. K., LIU D., FEULNER J., HUBER M., SUEHLING M., CAVALLARO A., COMANICIU D.: Hierarchical parsing and semantic navigation of full body CT data. In *SPIE medical imaging* (2009), International Society for Optics and Photonics, pp. 725902–725902. 2

[SFPS11] SOUMELIDIS A., FAZEKAS Z., PAP M., SCHIPP F.: Generic Zernike-based surface representation of measured corneal surface data. In *Medical Measurements and Applications Proceedings (MeMeA), 2011 IEEE International Workshop on* (2011), IEEE, pp. 148–153. 3

[Tea80] TEAGUE M. R.: Image analysis via the general theory of moments*. *J. Opt. Soc. Am. 70*, 8 (Aug 1980), 920–930. 3

[Zer34] ZERNIKE V. F.: Beugungstheorie des schneidenverfahrens und seiner verbesserten form, der phasenkontrastmethode. *Physica 1*, 7 (1934), 689–704. 3