Fast Multiple Organ Detection and Localization in Whole-Body MR Dixon Sequences

Olivier Pauly[1], Ben Glocker[2], Antonio Criminisi[2], Diana Mateus[1], Axel Martinez Möller[3], Stephan Nekolla[3], and Nassir Navab[1]

[1] Computer Aided Medical Procedures, Technische Universität München, Germany
{pauly,mateus,navab}@cs.tum.edu
[2] Microsoft Research Ltd, CB3 0FB, Cambridge, UK
{glocker,antcrim}@microsoft.com
[3] Nuklearmedizin, Klinikum Rechts der Isar, Technische Universität München, Germany
{a.martinez-moller,s.nekolla}@lrz.tu-muenchen.de

Abstract. Automatic localization of multiple anatomical structures in medical images provides important semantic information with potential benefits to diverse clinical applications. Aiming at organ-specific attenuation correction in PET/MR imaging, we propose an efficient approach for estimating location and size of multiple anatomical structures in MR scans. Our contribution is three-fold: (1) we apply supervised regression techniques to the problem of anatomy detection and localization in whole-body MR, (2) we adapt random ferns to produce multi-dimensional regression output and compare them with random regression forests, and (3) introduce the use of 3D LBP descriptors in multi-channel MR Dixon sequences. The localization accuracy achieved with both fern- and forest-based approaches is evaluated by direct comparison with state of the art atlas-based registration, on ground-truth data from 33 patients. Our results demonstrate improved anatomy localization accuracy with higher efficiency and robustness.

1 Introduction

Following the success of combined PET/CT, the possibility of combining PET with MRI has gained increased interest, as significant advantages are expected compared to PET/CT for many imaging tasks in neurology, oncology and cardiology [1]. However, before its introduction in the clinical practice, a technical challenge impacting the quality of PET/MR imaging needs to be solved: the attenuation correction of 511 keV photons according to the radiodensity of the tissues. While in PET/CT [2], radiodensity information provided by CT at X-ray energies can be converted into attenuation information, MR does not provide any information on the tissue density. Therefore, methods have been investigated to generate an attenuation correction map directly from MR. For brain imaging, atlas-based solutions using registration were evaluated in [3,4]. For whole-body imaging, different approaches based on the classification of tissues into 4 classes (background, lungs, fat, and soft tissue) have been investigated, for instance in [5]. While previous methods showed promising results for attenuation correction of whole body imaging with PET/MR, they propose only a coarse tissue classification, not accounting for organ-specific attenuation and for the attenuation introduced by bones. To further improve the quality of whole-body PET data reconstruction, we aim at generating organ-specific attenuation information directly from MR. Therefore, the position

of the organs which impact the attenuation of photons need to be known. This paper presents an approach for simultaneously localizing multiple organs in multi-channel whole-body MR. It builds upon state-of-the-art non-linear regression techniques. We adapt random ferns for regression and compare them to random regression forests. Experiments on 33 patient scans demonstrate better performance than atlas-based techniques in terms of accuracy, speed, and robustness.

2 Related Work

Classical object detection algorithms are based on sliding windows and classifiers whose role is to predict whether a voxel belongs to the object of interest or not. In [6], Viola and Jones introduced a fast detection approach based on a cascade of classifiers trained using Adaboost. Built as a succession of classifiers taking sequentially more and more features into account. This approach achieved impressive performance for real-time face detection. In medical applications, there has been an increasing interest in regression-based solutions for organ localization. Since the human body consists of a specific arrangement of organs and tissues, it can be expected that voxels, based on their contextual information, can predict the surrounding anatomy. For instance, if the neighborhood of a voxel shows an appearance which is typical of heart tissue, besides the position of the heart, this voxel can provide an estimate of position of the nearby lungs. In [7], Zhou *et al.* introduced an approach based on boosting ridge regression to detect and localize the left ventricle (LV) in cardiac ultrasound 2D images. There, the learned function predicts the relative position, scale and orientation of the LV based on Haar-like features computed on 2D images. Impressive results are demonstrated on echocardiogram sequences. To detect and localize the heart chambers in 3D cardiac CT, Zheng *et al.* proposed in [8] an approach called marginal space learning (MSL). To break down the complexity of learning directly in the full 3D similarity transformation space, the authors demonstrate that training a classifier on projections of the original space effectively reduces the search space. Using this idea, they build a cascade of classifiers based on probabilistic boosting tree (PBT) to predict first the position, then the position-orientation and finally the full 3D pose. In [9], the authors push this idea further to non-rigid marginal space learning using statistical shape models. Although these approaches have shown very good performance on CT scans, building such a cascade of classifiers is a computationally intensive learning procedure which requires large training sets. In this paper we avoid intensive training by building a single regressor predicting simultaneously the position of multiple organs. In [10], Criminisi *et al.* proposed a regression approach based on random forests for the localization of organs in 3D CT scans. The authors showed that their method achieves better performance than atlas registration, and this, while benefiting from fast training and testing. While in [10], the authors could rely on absolute radiodensity values provided by CT, here, we deal with MR images which provide only relative values and suffer from field inhomogeneities. To tackle this challenging problem, we adapt the regression forest framework by introducing 3D LBP descriptors. Additionally, we implement a random ferns approach [11] and compare it with forests. Both regression techniques are evaluated and compared to an atlas-based registration approach.

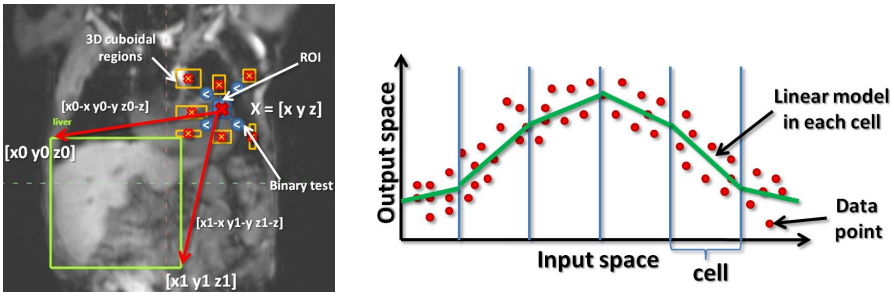

Fig. 1. (Left) Each voxel predicts its relative displacement to all organ bounding boxes represented here as a green box. Multi-scale textural information is extracted using LBP-like feature representation computed over 3D cuboidal regions. (Right) 1D regression example: data samples in red, blue lines represent the partition built over the input feature space. In each cell, simple linear models (in green) are fitted to the points. Their combination over the full space results in a complex non-linear predictor.

3 Proposed Method

This section describes details of our organ detection and localization approach. First, we cast this problem as a regression task. Second, we introduce our feature representation based on water and fat channels computed from MR Dixon sequences. Third, we present regression forests and explain how to adapt random ferns for regression. Finally, we show how to combine voxel predictions to localize all organs of interest in one shot.

3.1 Problem Statement

In our framework each voxel votes for the ***relative position*** of all organs of interest. The individual votes will produce very noisy predictions. But the probabilistically weighted combination of all votes will produce an accurate output (see Fig.1). Since the relative displacements we want to predict are continuous values, we use a ***regression*** paradigm. In the following, we introduce new features based on the fat and water MR Dixon channels, and then present a new non-linear regression approach based on random ferns. Next, we introduce the general problem of organ localization as a regression task.

Input Space: Formally, let us consider the water W and fat F channels computed from MR Dixon sequences defined by the two intensity functions $I^{(W)}, I^{(F)} : \Omega \rightarrow \mathbb{R}$, $\Omega \subset \mathbb{R}^3$ being the image domain. While $\mathbf{x} = [x, y, z]$ represents a voxel location in this domain Ω, $\Psi(I^{(W)}, I^{(F)}, \mathbf{x}) = \mathbf{X}$ denotes a function mapping the voxel location to a feature space according to both intensity functions $I^{(W)}$ and $I^{(F)}$. The role of feature representation $\mathbf{X}$ is to encode contextual information in the neighborhood of location $\mathbf{x}$ computed using $I^{(W)}$ and $I^{(F)}$. $\mathbf{X}$ is the input of our regression function.

Output Space: Let us now consider a set of K organs of interest contained in bounding boxes $\mathcal{O} = \{O_1, \cdots, O_k, \cdots, O_K\}$. Each bounding box O_k is represented by a vector $O_k = \left[x_k^0,\ y_k^0,\ z_k^0,\ x_k^1,\ y_k^1,\ z_k^1\right]$. The relative displacement $\mathbf{v}_k$ between voxel location $\mathbf{x}$ and bounding box O_k is parametrized as:

$$\mathbf{v}_k = \left[x_k^0 - x,\ y_k^0 - y,\ z_k^0 - z,\ x_k^1 - x,\ y_k^1 - y,\ z_k^1 - z\right] \tag{1}$$

We denote by $\mathbf{V} = [\mathbf{v}_1, \cdots, \mathbf{v}_k, \cdots, \mathbf{v}_K]$ the vector containing relative displacements between $\mathbf{x}$ and ***all*** organs of interest. $\mathbf{V}$ is the output of our regression function. On Fig.1, these relative displacements between the point $\mathbf{x}$ and the liver bounding box are represented by the red arrows. Here, we consider the following organs: head, left lung, right lung, heart and liver.

Regression: We assume a training set $\left(\mathbf{X}^{(n)}, \mathbf{V}^{(n)}\right)_{n=1}^{N}$ computed over a set of N patient MR volumes. We could think of modeling the posterior distribution $p(\mathbf{V}|\mathbf{X})$ linking the input and output spaces. However, in such high-dimensional feature spaces, modeling the posterior distribution directly is very difficult. To break down the complexity of this problem, we can first subdivide the input feature space by building a partition $\mathcal{P}$ over it. Indeed, by subdividing the feature space, we obtain cells containing data points which are easier to model even with simple mathematical models such linear or constant functions. As illustrated by the low-dimensional toy example on Fig. 1 (right), the combination of these models over the whole partition results then in a complex non-linear model. Formally, $\mathcal{P}$ is defined as an ensemble of T cells $\mathcal{P} = \{\mathcal{C}_t\}_{t=1}^{T}$. With $\mathcal{P}$ given, we propose to model the posterior in each cell $\mathcal{C}_t$ as follows:

$$p(\mathbf{V}|\mathbf{X} \in \mathcal{C}_t, \mathcal{P}) = \mathcal{N}_t(\mathbf{X}|\mu_t, \Sigma_t) \tag{2}$$

where $\mathcal{N}_t$ is a multivariate Gaussian distribution whose parameters are estimated during the training phase. In fact, this choice permits to model the full distribution as a piecewise Gaussian distribution. In contrast to fitting a Gaussian mixture model, partitioning is here performed in the input feature space and not in the output space. Moreover, in the case of trees, this partitionning is performed hierarchically. Based on this, we can model the probability distribution of $\mathbf{V}$ over the full feature space according to partition $\mathcal{P}$ as:

$$p(\mathbf{V}|\mathcal{P}) = \sum_{t=1}^{T} p(\mathbf{V}|\mathcal{C}_t, \mathcal{P}) p(\mathcal{C}_t) \tag{3}$$

Clearly, the quality of the posterior approximation depends on the partition $\mathcal{P}$. If its number of cells T is low, then the posterior approximation will be very rough. On the other hand, if T is high, each cell will include few training points. In this case, the partition $\mathcal{P}$ tends to overfit the training data and suffers from poor generalization. In [12], Breiman demonstrates that replacing a single partition $\mathcal{P}$ with an ensemble of independent random partitions $\{\mathcal{P}_z\}_{z=1}^{Z}$ leads to an ensemble regressor achieving better generalization. In this paper, we apply regression forests and adapt random ferns for regression to construct multiple independent partitions $\{\mathcal{P}_z\}_{z=1}^{Z}$. The posterior estimates from the different partitions of the ensemble are then combined using averaging. Finally, we can estimate in one shot the position of all organs of interest contained in variable $\hat{\mathbf{V}}$ using the mathematical expectation: $\hat{\mathbf{V}} = \int_{\mathbf{V}} \mathbf{V} p(\mathbf{V}) d\mathbf{V}$. Before going into the details of our regression techniques in section 3.3, let us first describe the feature representation we use in the problem of organ localization in MR Dixon sequences.

3.2 Feature Representation

As described in [13], MR Dixon imaging techniques are based on the one shot acquisition of a so-called "in phase" scan where water and fat signals are in-phase and an "opposite phase" scan where water and fat signals are $180°$ out-of-phase. Using these 2

scans from the same patient, water and fat signals can be separated to construct a water $I^{(W)}$ and a fat $I^{(F)}$ channel. Since these 2 channels are perfectly registered, we propose to take advantage from their complementary nature and design a feature representation based on both water and fat information. While in CT intensity information is directly related to the underlying tissue distribution, MR intensity information is not absolute and suffers from variability between different images. For this reason, we will not rely on intensities as in [10], but on textural information by employing Local Binary Patterns (LBP) [14]: we propose to extract textural context variations at different scales (see Fig. 1). Let us consider a 3D region $\mathcal{R}^s_{\mathbf{x}}$ at scale s centered on voxel location $\mathbf{x}$ and a set $\{\mathbf{N}^{s,q}_{\mathbf{x}}\}_{q=1}^{Q}$ of Q 3D asymmetric cuboidal regions having different sizes, orientations and offsets in the neighborhood of $\mathbf{x}$. Using this, we can extract two binary feature vectors $\mathbf{X}^{(W)}_s$ and $\mathbf{X}^{(F)}_s$ from the two channels where each entry is the result of the following binary test comparing average intensities within regions $\mathbf{N}^{s,q}_{\mathbf{x}}$ and $\mathcal{R}^s_{\mathbf{x}}$:

$$\mathbf{X}^{(i)}_s[q] = \frac{1}{|\mathbf{N}^{s,q}_{\mathbf{x}}|} \sum_{\mathbf{x}' \in \mathbf{N}^{s,q}_{\mathbf{x}}} I^{(i)}(\mathbf{x}') < \frac{1}{|\mathcal{R}^s_{\mathbf{x}}|} \sum_{\mathbf{x}' \in \mathcal{R}^s_{\mathbf{x}}} I^{(i)}(\mathbf{x}'), \tag{4}$$

and this, $\forall q \in \{1, \cdots, Q\}$ and $i \in \{W, F\}$. Repeating this operation at several scales results in two feature vectors $\mathbf{X}^{(W)}$ and $\mathbf{X}^{(F)}$ describing the multi-scale textural context for both channels in the neighborhood of voxel location $\mathbf{x}$. Since $\mathbf{X}^{(W)}$ and $\mathbf{X}^{(F)}$ are binary vectors, they can be further encoded to reduce their dimensionality. Finally, they are concatenated in one feature vector: $\mathbf{X} = \left[\mathbf{X}^{(W)},\ \mathbf{X}^{(F)}\right]$.

3.3 Ensemble Regression Approaches

This section explains how to use forests and ferns to efficiently partition the input data. While regression forests have been used for detecting organs in CT [10], there exists little work on ferns-based regression. In [15], Dollar *et al.* use a ferns-based regressor in a cascade fashion for pose detection of objects in 2D images. In contrast, we use a single ensemble regressor.

Forests and ferns: ***Random forests*** [12] are constructed as an ensemble of independent random trees. Each tree is a set of decision functions that split feature vectors at each node towards the left or the right branch. ***Random ferns***, later introduced by Özuysal *et al.* [11], are an ensemble of *constrained* trees. While a tree applies a different decision function at each node, a fern systematically applies the same decision function for each node of the same level. Results of these random tests are finally stored as binary values. Ferns benefit from a more compact and simple structure for an accuracy that is similar to those of random trees [11]. In fact, while $2^{(N-1)}$ operations are needed to grow a tree of 2^N leaves, only N operations are needed to grow a fern of 2^N leaves.

Let us now describe these partitioning approaches more formally. We denote by $\mathcal{F} = \left\{\mathbf{F}^{(z)}\right\}_{z=1}^{Z}$ an ensemble of trees or ferns. Each element of this ensemble induces an independent partition $\mathcal{P}^{(z)} = \left\{\mathcal{C}^{(z)}_1, \cdots, \mathcal{C}^{(z)}_T\right\}$ of the input feature space. Each tree or fern $\mathbf{F}^{(z)}$ is defined as a set of L nodes, each node being equipped with a linear function f_l and an associated threshold τ_l, $l \in \{1, \cdots, L\}$. The output of evaluating a pair $(f^{(z)}_l, \tau^{(z)}_l)$ on a visual feature vector $\mathbf{X}$ is binary, that is $f^{(z)}_l(\mathbf{X}, \tau^{(z)}_l) : \mathbf{X} \mapsto \{0, 1\}$. While in a tree, the output value denoted by $b^{(z)}_l$ decides whether $\mathbf{X}$ gets pushed towards

the left or the right branch, in a fern, $\mathbf{X}$ is evaluated at all nodes. The corresponding outputs are then stored in binary vector $\mathbf{b}^{(z)} = [b_1^{(z)}, \ldots, b_l^{(z)}, \ldots, b_L^{(z)}]^\top$. In the end, while $\mathbf{X}$ is pushed through the whole tree until it reaches a leaf (which is a cell $\mathcal{C}_t^{(z)}$ of the partition $\mathcal{P}^{(z)}$), in a fern, the full vector $\mathbf{b}^{(z)}$ encodes the cell index of the partition where the vector falls.

Training/Testing: During the training of a fern, the whole training data is used at each node. This is in contrast to trees where only a subset is considered at each node. If we consider a training set $\left(\mathbf{X}^{(n)}, \mathbf{V}^{(n)}\right)_{n=1}^{N}$ computed over a set of different patient scans, all feature vectors are pushed through the ferns ensemble and fall into the cells of the different partitions. Finally, the parameters of each Gaussian can be estimated for each cell $\mathcal{C}_t^{(z)}$ using the subset $\left\{\mathbf{V}^{(n)} | \mathbf{X}^{(n)} \in \mathcal{C}_t^{(z)}\right\}_{n=1}^{N}$ of training data that fell into $\mathcal{C}_t^{(z)}$. In the current paper, we do not use optimization in the construction of our ferns regressor, *i.e.* the linear functions and their associated thresholds are chosen randomly. While this permits to have a very fast training procedure, it provides independency from the training set. This can be an advantage for instance in the case of noisy data. Once the training has been performed, all node functions and thresholds are frozen. During the test phase, an unseen data point $\mathbf{X}$ is pushed through the whole ensemble until it reaches a cell in each partition. Then, each cell contributes to the final prediction using its stored Gaussian model as seen in section 3.1. Next, we describe how to combine the predictions to localize all organs of interest.

3.4 Anatomy localization

Let us consider the water $I^{(W)}$ and fat $I^{(F)}$ channels of an unseen patient. From both channels, a set of feature vectors $\left\{\mathbf{X}^{(n)}\right\}_{n=1}^{N}$ is extracted from voxel locations $\left\{\mathbf{x}^{(n)}\right\}_{n=1}^{N}$. By pushing this set of feature vectors through the regression ensemble, predictions $\left\{\hat{\mathbf{V}}^{(n)}\right\}_{n=1}^{N}$ are computed as described in section 3.1. They correspond to the relative displacements $\hat{\mathbf{V}}^{(n)} = [\hat{\mathbf{v}}_1, \cdots, \hat{\mathbf{v}}_k, \cdots, \hat{\mathbf{v}}_K]$ between each location $\mathbf{x}^{(n)} = \left[x^{(n)}, y^{(n)}, z^{(n)}\right]$ and all organ bounding boxes $\mathcal{O} = \{O_1, \cdots, O_k, \cdots, O_K\}$. The bounding box of organ O_k can be finally estimated as follows:

$$O_k = \sum_{n=1}^{N} w_n \left(\hat{\mathbf{v}}_k^{(n)} + \left[\mathbf{x}^{(n)}, \mathbf{x}^{(n)}\right]\right) \tag{5}$$

where each w_n weights the contribution of voxels according to the confidence of their predictions. Note that $\sum_{n=1}^{N} w_n = 1$. In this paper, we discard contributions having low confidence and perform averaging on the remaining predictions.

4 Experiments and Results

In this section, we compare our approaches based on regression forests and random ferns with the current state-of-the-art multi-atlas registration.

Data: Our dataset currently consists of scans from 33 patients who underwent a 3-Tesla whole-body MR Dixon sequence. All patients have cancer (mostly neck, lung, liver cancer) and show a high variability in their anatomy partially due to their disease. For the

detection and localization of organs, we use the water and fat channels. In each scan, we manually delineated the bounding boxes for following organs: head, left lung, right lung, liver and heart. The size of the volumes are $192 \times 124 \times 443$ and the pixel spacing is $2.6 \times 2.6 \times 2.6$ mm.

Regression approach: 100 runs of cross-validation experiments have been conducted where each experiment consists of a training phase on 20 patients chosen randomly and a test phase on the 13 remaining patients. For both forests and ferns, all parameters (number of trees/ferns and tree depth/number of nodes) have been tuned by performing grid-search within the same range for both techniques. Note that node optimization has been performed for random forests based on information gain (cf. [10]). For prediction, each fourth pixel is used and described using 3D LBPs computed over 26 cuboidal regions chosen at 3 different scales.

Multi-atlas registration: 100 runs of cross-validation experiments have been performed. Each experiment is defined as follows: a set of 20 patients are chosen randomly as multi-atlas database and 1 patient is randomly chosen as test case. All 20 patients from the database are registered to the test patient using affine registration. Then, using the ground truth position of the bounding boxes of the test patient (which is not available in reality), we evaluate the theoretical lower and upper bounds of the error by using the patients in the database who provide the lowest and highest localization error. The mean error is computed over the whole database.

Results: Results reported on Tab.1 shows that we achieve an accuracy which is better than the "best case" atlas accuracy, while providing an increased robustness. Taking a look at the localization error per organ, one can notice that the lowest error for our approach is achieved for the localization of the head, which is due to the fact that the head is surrounded by a lot of air which makes it easier to localize. While the heart shows second lowest error, lungs and liver were more difficult to localize. This is mainly due to the high inter-patient variability of the shape of these organs. The best results were obtained with 14 ferns/6 nodes for random ferns, and 6 trees/depth of 8 for regression forests. On a laptop with MATLAB 64 Core Duo 2.4 GHz, the training/testing time on 20/13 patients is 0.7/0.5 s for random ferns. Random Forests need 25/1 s. Concerning atlas registration, each single affine registration needs 12.5 s. To conclude, our approach provides a fast and robust solution for organ detection and localization and thus fulfills our requirements towards organ-specific attenuation map.

5 Conclusion

Our contribution is a supervised regression approach based on random ferns and random forests to detect and localize in one shot multiple organs in whole-body multi-channel MR images. Experiments conducted on a dataset of 33 patients show that our approach achieves an accuracy which is better than atlas-based methods, while providing higher robustness (lower uncertainty) and faster training/prediction times. Furthermore, this approach can be also useful to integrate semantic information *i.e.* incorporating organ labels in further applications such as registration, image navigation or image retrieval. In future work, we plan to investigate the online performance of the proposed approach to enable a fast updating of our organ localization system, and then we will move towards the construction of organ-specific attenuation correction maps.

MEAN LOCALIZATION ERRORS (mm)						
Organs	**Head**	**Left lung**	**Right lung**	**Liver**	**Heart**	**Overall**
Random ferns	9.82 ± 8.07	14.95 ± 11.35	16.12 ± 11.73	18.69 ± 13.77	15.17 ± 11.70	$\mathbf{14.95 \pm 11.33}$
Random forests	10.02 ± 8.15	14.78 ± 11.72	16.20 ± 12.14	18.99 ± 13.88	15.28 ± 11.89	$\mathbf{15.06 \pm 11.55}$
Atlas lower bound	18.00 ± 14.45	14.94 ± 11.54	15.02 ± 13.69	18.13 ± 16.26	13.31 ± 11.03	$\mathbf{15.88 \pm 13.40}$
Atlas upper bound	70.25 ± 34.23	60.78 ± 29.47	63.95 ± 30.13	70.59 ± 32.88	60.38 ± 28.90	$\mathbf{65.19 \pm 31.12}$
Atlas Mean	35.10 ± 13.17	30.41 ± 11.39	29.85 ± 12.62	31.74 ± 13.49	29.82 ± 12.23	$\mathbf{31.38 \pm 12.58}$

Table 1. Comparative results: Compared to atlas-based method, our approaches based on random ferns and forests achieve better accuracy and lower uncertainty.

References

1. Judenhofer, M.S., et al.: Simultaneous pet-mri: a new approach for functional and morphological imaging. Nature Medicine (2008) 459–465
2. Kinahan, P., Hasegawa, B., Beyer, T.: X-ray-based attenuation correction for positron emission tomography/computed tomography scanners. Semin. Nuc. Med. (2003)
3. Kops, E.R., Herzog, H.: Template-based attenuation correction of PET in hybrid MR-PET. Journal of Nuclear Medicine (2008)
4. Hofmann, M., Steinke, F., Scheel, V., Charpiat, G., Farquhar, J., Aschoff, P., Brady, M., Schölkopf, B., Pichler, B.J.: MRI-Based Attenuation Correction for PET/MRI: A Novel Approach Combining Pattern Recognition and Atlas Registration. Journal of Nuclear Medicine (2008)
5. Martinez-Möller, A., Souvatzoglou, M., Delso, G., Bundschuh, R.A., Chefd'hotel, C., Ziegler, S.I., Navab, N., Schwaiger, M., Nekolla, S.G.: Tissue Classification as a Potential Approach for Attenuation Correction in Whole-Body PET/MRI: Evaluation with PET/CT Data. Journal of Nuclear Medicine **50** (2009) 520–526
6. Viola, P., Jones, M.J.: Robust real-time face detection. International Journal on Computer Vision **57** (2004) 137–154
7. Zhou, S.K., Zhou, J., Comaniciu, D.: A boosting regression approach to medical anatomy detection. In: IEEE Conference on Computer Vision and Pattern Recognition. (2007)
8. 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 Transactions on Medical Imaging **27** (2008) 1668–1681
9. Zheng, Y., Bogdan, Comaniciu, D.: Marginal Space Learning for Efficient Detection of 2D/3D Anatomical Structures in Medical Images. In: Information Processing in Medical Imaging. (2009)
10. Criminisi, A., Shotton, J., Robertson, D., Konukoglu, E.: Regression Forests for Efficient Anatomy Detection and Localization in CT Studies. In: Jiang T., Navab N., Pluim J., Viergever M. (Eds.) MICCAI 2010 workshop in Medical Computer Vision. (2010)
11. Özuysal, M., Calonder, M., Lepetit, V., Fua, P.: Fast Keypoint Recognition using Random Ferns. IEEE Transactions on Pattern Analysis and Machine Intelligence **32** (2010) 448–461
12. Breiman, L.: Random forests. Machine learning **45** (2001) 5–32
13. Ma, J.: Dixon Techniques for Water and Fat Imaging. J Mag. Res. Im. (2008)
14. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on featured distributions. Pattern Recognition **29** (1996) 51–59
15. Dollar, P., Welinder, P., Perona, P.: Cascaded pose regression. In: IEEE Conference on Computer Vision and Pattern Recognition. (2010)