

Dense Rotation Invariant Brain Pyramids for Automated Human Brain Parcellation

Henrik Skibbe^{1,3}, Marco Reisert²

¹Department of Computer Science, University of Freiburg, Germany

²Dept. of Diagnostic Radiology, Medical Physics, University Medical Center, Freiburg

³Center for Biological Signalling Studies (BIOSS), University of Freiburg
skibbe@informatik.uni-freiburg.de, marco.reisert@uniklinik-freiburg.de

Abstract: The automatic parcellation of the human brain based on MR imaging is in several areas of high interest. In particular, identifying corresponding brain areas between different subjects is an indispensable prerequisite for any group analysis. But also, simple segmentations into different tissue types is an important preprocessing step. We present a generic framework for describing and automatically parcellating high angular resolution diffusion-weighted magnetic-resonance images (HARDI) of the human brain. Based on an initial training step our approach is capable to segment the images into coarse parcellations or detailed fine grain regions of interest. In contrast to existing model-free methods [SSK⁺09] we are not only using the raw measurements at each position, but we are also including neighboring measurements in a rotation invariant way.

1 Introduction

The study of Magnetic Resonance (MR) imaging modalities is of great interest in fundamental neuroscience and medicine. MR imaging offers a wide range of different contrasts, providing scientists insights into anatomical and functional properties of the human brain. One may distinguish between anatomical contrasts (T1-weighted, T2-weighted) and derived contrasts that use combinations of different measurement parameters to infer underlying tissue properties. Examples are Diffusion Weighted Imaging (DTI), High Angular Resolution Diffusion Imaging (HARDI [TWBW99]), functional MRI (fMRI) and Arterial Spin Labeling. The main difference to anatomical images is the rather bad image quality in terms of resolution and signal to noise ratio.

For several applications it is necessary to transfer highly resolved anatomical information to, for example, low resolution DTI images. Usually co-registration is employed. One scenario are group studies, where the determination of corresponding anatomical regions is necessary for inter subject comparisons. High quality anatomical images are used to find inter-subject correspondences, which are transferred to the low resolution contrasts by co-registration. Another application scenario is seed region generation for fiber tracking analysis. On a manually labeled atlas seeds are selected, which are again transferred to the DTI or HARDI contrast by co-registration with a template.

In this work we want to point out an alternative to the usual approach of co-registration and segmentation. The idea is to parcellate the brain into a large number of different unique regions by a machine learning approach. That is, every 'brain region' gets a unique label. Suppose that it is possible to compute the same type of labeling for several different contrasts, then co-registration becomes superfluous. By the labeling we already know 'where' we are in the brain. To achieve this goal the task is formulated as a machine learning problem. For each voxel, a discriminative feature vector is computed which encodes the voxel and its 'surrounding'. Based on the feature vector a supervised classifier assigns an unique label to the voxel. Thus, the brain parcellation task is formulated as a large voxel by voxel classification problem. To enable the classifier to do his job the features have to be expressive. It is not enough that a feature encodes just the current voxel and its nearby surrounding, then, there will be several ambiguous areas. The feature must represent nearly the whole brain, otherwise it is quite difficult to infer the absolute position in the brain, which the label is actually standing for. That is, we will need a kind of scale pyramid: large scale features encoding the position coarsely, and fine scale features providing accuracy for the determination of the region boundaries. Another demand is rotation invariance. Indeed, the head is usually in more or less the same pose, but changes of the head pose must not alter the classification. Thus, the features must not depend on the rotational pose of the head.

This paper presents one early instantiation of the proposed idea on HARDI images. We focus on the feature design and use random forests [Bre01] for classification. To fulfill the above formulated feature demands we utilize spherical tensor algebra (STA) [Ros95, RB09], which allows to compute dense, multi-scale feature images from HARDI data in an efficient and rotation invariant way. The use of STA is quite nearby, because it is a common way to represent HARDI images by Spherical Harmonics (SH). The orientation distribution in each voxel is encoded in the SH basis providing a memory efficient and smooth way to handle the data. STA provides two basic operations to handle such kind of SH representations with an important property: the operations do not alter the rotation behavior of the SH representation, that is, they allow to compute rotation invariant features in a simple and systematic way. The first class of operations are finite difference operators that connect SH representations of different degrees by differentiations, called spherical derivatives. The second class of operations are products that connect two different SH representations to form a new field with a different degree, called spherical products. Based on these primitives the feature extraction procedure can be concluded as follows: 1) the HARDI images in SH representation are smoothed successively by Gaussians to sense different scales. 2) Onto the smoothed images spherical derivatives are applied to represent the 'neighborhood' of each voxel. The size of the neighborhood is implicitly defined by the Gaussian applied beforehand. 3) Spherical products are used to obtain a set of invariant features.

One problem we have not mentioned yet is the question, how to generate the training data which is needed by the classifier during the learning stage. It is quite difficult or even impossible to obtain 'real' ground truth information. In this work we use the usual co-registration approach to generate the training data: a T1-weighted anatomical image is

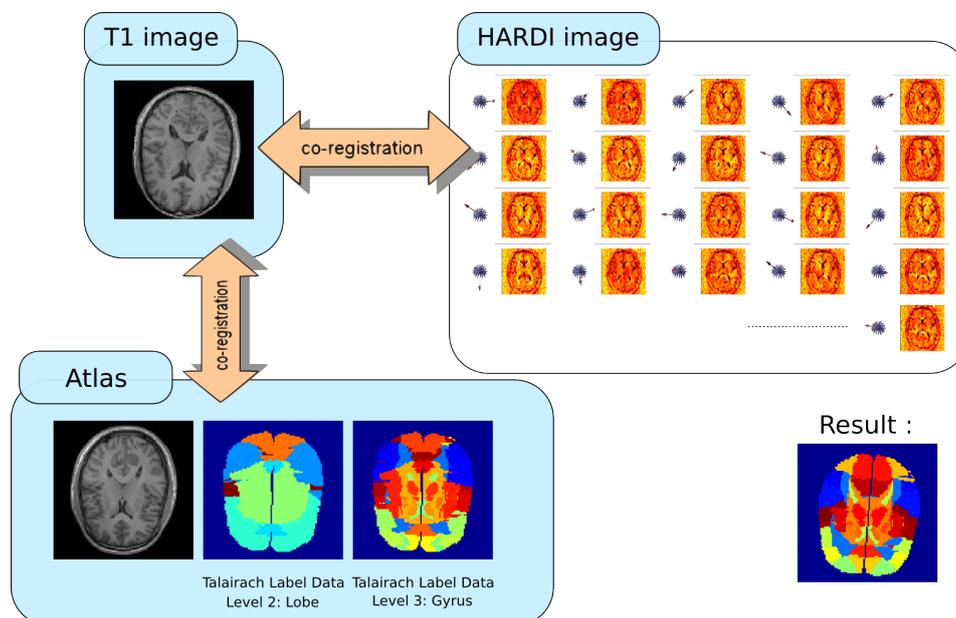


Figure 1: The standard co-registration approach is used to generate the training data: a T1-weighted anatomical image is co-registered with a template which also provides atlas information. Based on co-registrations it is possible to transform the labeling information from the atlas into the native space of the HARDI data.

co-registered with a template which provides atlas information (WFU pickatlas¹ [Lan97]). Simultaneously the T1-weighted image is also co-registered with the B0-weighted image from the HARDI data. The co-registrations were performed with the SPM5 toolkit². Based on the co-registrations it is possible to transform the labeling information from the atlas into the native space of the HARDI data (see Figure 1 for visualization).

The paper is organized as follows: We first introduce the rotation invariant brain pyramids in section 2, densely representing the local appearance of structures in the human brain based on the HARDI signal. In section 3 we show how to use the rotation invariant representation of the brain pyramids to realize the automated brain parcellation. In section 4 we evaluated our method on seven in vivo HARDI datasets leading to very promising results. A conclusion is given in section 5.

¹WFU PickAtlas, A software that provides a method for generating ROI masks based on the Talairach Daemon database. <http://fmri.wfubmc.edu/software/PickAtlas>

²SPM (Statistical Parametric Mapping version 5) software package for the analysis of brain imaging data sequences. <http://www.fil.ion.ucl.ac.uk/spm/>

2 Rotation Invariant Brain Pyramids

2.1 Mathematical Notation

We write vectors $\mathbf{v} \in \mathbb{C}^n$ in bold letters. We denote the complex conjugate of \mathbf{v} by $\bar{\mathbf{v}}$ and the transpose of \mathbf{v} by \mathbf{v}^T . Unit-length vectors $\mathbf{n} = (x, y, z)^T \in \mathbb{R}^3$, $\|\mathbf{n}\| = 1$ are considered w.l.o.g as points on the unit-sphere S_2 which we denote by $\mathbf{n} \in S_2$. Hence we equivalently can represent \mathbf{n} in spherical coordinates (θ, ϕ) , where $\theta = \arccos(z)$ and $\phi = \text{atan2}(y, x)$. The convolution is denoted by $*$.

2.2 Rotation Invariant Brain Pyramids

We consider HARDI-images as functions $\mathbf{f} : \mathbb{R}^3 \times S_2 \rightarrow \mathbb{R}$, where S_2 denotes the unit-sphere in \mathbb{R}^3 . This means at each voxel-position $\mathbf{x} \in \mathbb{R}^3$ we have an angular dependent measurement $\mathbf{f}(\mathbf{x}, \mathbf{n})$ represented as a function on the unit-sphere. This function represents the diffusion weighted MR-signal with respect to different diffusion directions \mathbf{n} . One important fact is that since diffusion is symmetric the signal shows an even symmetry. It follows that $\mathbf{f}(\mathbf{x}, \mathbf{n}) = \mathbf{f}(\mathbf{x}, -\mathbf{n})$. Asymmetric components within the measured signals are not providing any reliable information about the true matter and can be considered as noise.

For a further analysis of the signal we propose a decomposition of the signal into its basic frequency components. This offers several advantages. First, the angular resolution of the raw HARDI images is usually quite high thus the images are consuming a large amount of memory. In contrast, for a representation of the HARDI signal in terms of its basic frequency components only a small number of coefficients is sufficient because the most reliable information is represented by lower frequency components. Higher frequency components are mostly representing noise. This means a low dimensional subspace in the frequency domain is sufficient to represent all characteristics of the signal offering a memory efficient representation. The orthogonal basis functions representing the different frequency components for functions on the sphere are the Spherical Harmonic (SH) basis functions $Y_m^\ell : S_2 \rightarrow \mathbb{C}$ (see e.g. [Ros95] for further informations). With each index ℓ a certain frequency is represented by a set of functions with indices m ranging from $m = \{-\ell \dots \ell\}$. We arrange them in a vector-valued function $\mathbf{Y}^\ell : S_2 \rightarrow \mathbb{C}^{2\ell+1}$. An even index ℓ indicates a symmetric pattern, if odd it indicates an asymmetric pattern (similar to the cos and sin functions in 1D). As mentioned before the signals are symmetric, hence we only need to consider SH associated with an even index to completely represent the signal in the SH-domain. Consequently we can represent \mathbf{f} in terms of \mathbf{Y}^ℓ by

$$\mathbf{f}(\mathbf{x}, \mathbf{n}) = \sum_{\substack{\ell=0 \\ \ell \text{ even}}}^{\infty} \mathbf{a}^\ell(\mathbf{x})^T \mathbf{Y}^\ell(\mathbf{n}) \quad , \quad (1)$$

where $\mathbf{a}^\ell(\mathbf{x}) \in \mathbb{C}^{2\ell+1}$ are the vector valued expansion coefficients representing the HARDI signal of \mathbf{f} at image position $\mathbf{x} \in \mathbb{R}^3$ in the SH-domain. We compute the coefficients \mathbf{a}^ℓ

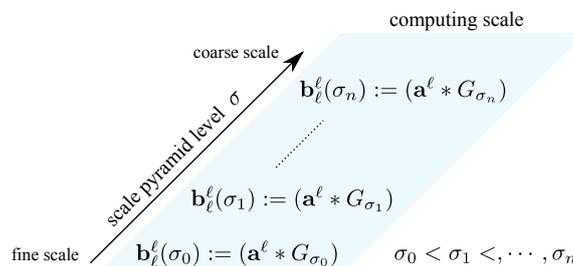


Figure 2: We propose a scale space pyramid of the single HARDI signal components \mathbf{a}^ℓ representing the characteristics of the data in a fine (the pyramid base) to coarse order.

by orthogonal projection of the discrete HARDI signal voxel by voxel onto the subspace spanned by \mathbf{Y}^ℓ (details can be found e.g. in [SRB11]). A decomposition of \mathbf{f} allows further analysis of the HARDI signal in an analytical way. \mathbf{a}^0 for instance, represents the mean of the signal. $\mathbf{a}^2(\mathbf{x})$ gives us information about the diffusion direction at the current voxel position. $\mathbf{a}^4(\mathbf{x})$ and further higher order coefficients are necessary to represent more complex local diffusion properties like crossings. Furthermore, by representing the HARDI signal in terms of SH we become independent from a discrete angular representation which is typical for the raw HARDI data allowing us to easily cope with datasets having different angular resolutions.

2.2.1 The Brain Scale Pyramid

For a local analysis of the HARDI images \mathbf{f} we propose to compute a scale space pyramid of the single signal components \mathbf{a}^ℓ of \mathbf{f} representing the appearance of the data in a fine (the pyramid base) to coarse order. We compute the scale space pyramid by successively smoothing the images \mathbf{a}^ℓ with an isotropic Gaussian $G_\sigma(\mathbf{r}) := e^{-\mathbf{r}^T \mathbf{r} / (2\sigma^2)}$, $\sigma \in \mathbb{R}_{>0}$. We call the resulting scale space representatives $\mathbf{b}_\ell^\ell(\sigma) : \mathbb{R}^3 \rightarrow \mathbb{C}^{2\ell+1}$, where

$$\mathbf{b}_\ell^\ell(\sigma) := (\mathbf{a}^\ell * G_\sigma) \quad . \quad (2)$$

The computation of the scale space pyramid is illustrated in figure 2.

2.2.2 Tensor Derivatives of the Brain Pyramid

The coefficients $\mathbf{a}^\ell(\mathbf{x})$ are so-called *spherical tensors* and the coefficient images $\mathbf{a}^\ell : \mathbb{R}^3 \rightarrow \mathbb{C}^{2\ell+1}$ are so-called *spherical tensor fields* allowing the usage of spherical tensor algebra (For further details we recommend [RB09]). Consequently, the scale space representatives $\mathbf{b}_\ell^\ell(\sigma)$ are spherical tensor fields, too. We use *spherical tensor down-derivatives* ∇_1 and *spherical tensor up-derivatives* ∇^1 to derive further tensor fields locally describing the scale space representatives $\mathbf{b}_\ell^\ell(\sigma)$ itself in an analytical way. Spherical tensor derivatives are derivative operators acting on spherical tensor fields. Similar to the ordinary Cartesian differential operators the resulting derivatives are representing local

characteristics of the initial spherical tensor field. After applying these operators we gain a huge number of tensor fields, each representing a certain characteristic of the local brain structure and its surrounding.

We first use spherical down-derivatives to successively create fields representing characteristics of its predecessor. Meanwhile the dimension of the fields decreases until we end up with scalar valued fields \mathbf{b}_0^ℓ :

$$\underbrace{\underbrace{\underbrace{\underbrace{\nabla_1(\nabla_1(\dots \nabla_1(\underbrace{\mathbf{a}^\ell * G_\sigma}_{\mathbf{b}_\ell^\ell(\sigma) \in \mathbb{C}^{2\ell+1}}))}_{\mathbf{b}_{(\ell-1)}^\ell(\sigma) \in \mathbb{C}^{2(\ell-1)+1}})}_{\mathbf{b}_1^\ell(\sigma) \in \mathbb{C}^3}}_{\mathbf{b}_0^\ell(\sigma) \in \mathbb{C}^1}} \quad . \quad (3)$$

In a second step we utilize the spherical up-derivatives to successively create higher order tensor fields encoding further characteristics and constellations of local neighboring HARDI signals within the human brain image. We limit this expansion by an upper boundary $L \in \mathbb{N}$:

$$\underbrace{\underbrace{\underbrace{\underbrace{\nabla^1(\nabla^1(\dots \nabla^1(\underbrace{\mathbf{a}^\ell * G_\sigma}_{\mathbf{b}_\ell^\ell(\sigma) \in \mathbb{C}^{2\ell+1}}))}_{\mathbf{b}_{(\ell+1)}^\ell(\sigma) \in \mathbb{C}^{2(\ell+1)+1}})}_{\mathbf{b}_{(\ell+L-1)}^\ell(\sigma) \in \mathbb{C}^{2(\ell+L-1)+1}}}}_{\mathbf{b}_{(\ell+L)}^\ell(\sigma) \in \mathbb{C}^{2(\ell+L)+1}} \quad . \quad (4)$$

As a result we gain for each tensor order ℓ and scale σ a set of spherical tensor valued feature images

$$\left\{ \underbrace{\mathbf{b}_0^\ell(\sigma), \dots, \mathbf{b}_{\ell-1}^\ell(\sigma)}_{\text{down-derivatives of } \mathbf{b}_\ell^\ell(\sigma)}, \underbrace{\mathbf{b}_\ell^\ell(\sigma)}_{\text{scale space representative}}, \underbrace{\mathbf{b}_{\ell+1}^\ell(\sigma), \dots, \mathbf{b}_{\ell+L}^\ell(\sigma)}_{\text{up-derivatives of } \mathbf{b}_\ell^\ell(\sigma)} \right\} \quad . \quad (5)$$

2.2.3 Rotation Invariant Feature Representation

Changes of the head pose should not alter our parcellation. Since we aim at utilizing the tensor valued features derived from the brain scale space pyramid a rotation invariant representation of these local brain descriptions is required. We achieve rotation invariance by utilizing so-called spherical tensor products [RB09] to couple the tensor valued feature images (eq. (5)) to form scalar valued feature images. Moreover, following [SRS⁺10] a certain set of invariants are corresponding to the power-spectra of the tensor valued feature images:

$$\left\{ \|\mathbf{b}_0^\ell(\sigma)\|^2, \dots, \|\mathbf{b}_{\ell-1}^\ell(\sigma)\|^2, \|\mathbf{b}_\ell^\ell(\sigma)\|^2, \|\mathbf{b}_{\ell+1}^\ell(\sigma)\|^2, \dots, \|\mathbf{b}_{\ell+L}^\ell(\sigma)\|^2 \right\} \quad , \quad (6)$$

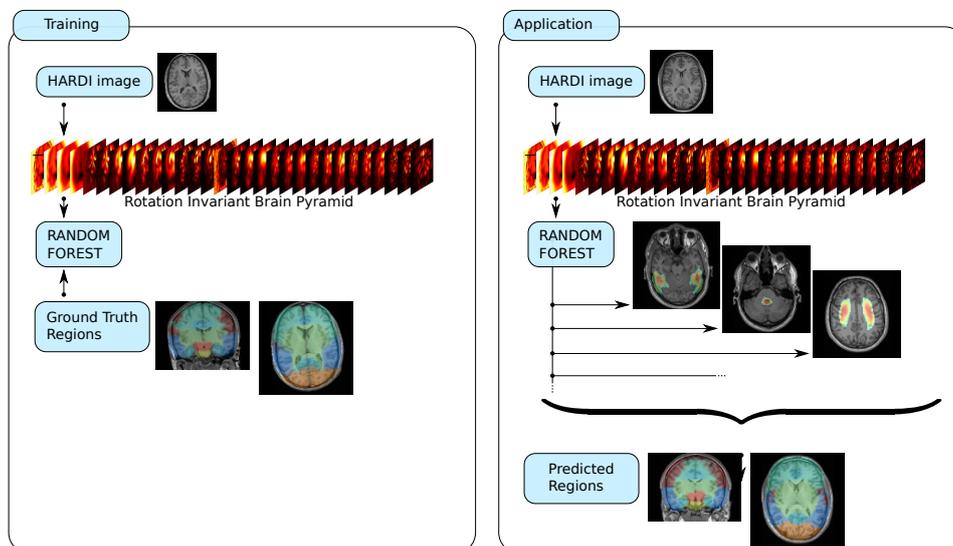


Figure 3: In a training step we provide training images (HARDI) together with a ground truth labeling and the rotation invariant brain pyramids. A random forest learns the connection between feature appearance and labels that is most consistent with the provided HARDI data (left image). Once the random forest has been initialized, we can utilize it to predict the region labels in unclassified HARDI data (right image).

which we use in this work. It is worth noting that the power-spectra of the expansion coefficients $\|\mathbf{a}^\ell\|^2$ are rotation invariant, too [KFR03, SRS⁺10].

3 Automated Brain Parcellation

For an automated parcellation of the brain we utilize random forest classifiers [Bre01] to learn the connection between region labels and local appearance. The random forest classifier consists of a huge number of decision trees, each predicting a class label. A decision is made by majority.

In a training step (see figure 3) we provide the rotation invariant brain pyramids from several training images (HARDI) (the feature images according to eq. (6)) together with a ground truth labeling (in our case the co-registered region labels). Using a large number of training images is of particular importance when coping with uncertainties in the ground truth. The random forest learns a connection between feature appearance and labels that is most consistent with the provided HARDI data.

Once the random forest has been initialized, we can utilize it to predict the region labels in unclassified HARDI data (figure 3). For this we first compute the rotation invariant brain pyramids and then ask the random forest for predicting the region labels. We gain some kind of probability map for each region label according to the number of decision trees in

the random forest voting for this class. A prediction for a region label is made voxel by voxel by choosing the label with the most votes.

Our approach gives us both a heat map for each region label representing the uncertainty of our prediction and a parcellation of the human brain most consistent with respect to the appearance of the underlying HARDI data.

4 Experiments

For our experiments seven in vivo diffusion measurements were acquired on a Siemens 3T TIM Trio using an SE EPI sequence with a TE of 95 ms and a TR of 8.5 s and an effective b-value of 1000 mm. One voxel corresponds to $2mm^3$. The number of diffusion directions differs between the datasets (see table 1) thus the images are strongly vary in quality. For evaluating our approach we used a coarse and a fine parcellation of the human brain based on the Talairach Daemon database [Lan97], namely a parcellation with respect to the Talairach Labels of Level 2 (Lobes, 11 regions + background) and Level 3 (Gyrus 55 regions + background). The ground truth was generated by co-registration as illustrated in figure 1. Then we performed a leave-one-out cross validation according to the following procedure:

We used 1 image out of the seven HARDI images for testing and the remaining six for training the random forest classifier as illustrated in figure 3. We got much better results when initially smoothing the raw HARDI data ones with a small Gaussian ($\sigma = 2$). We further initially used a gray matter / white matter mask to get rid of the skull and the noisy background before computing the features using a similar model-free learning based approach [SRB11]. This is not an essential step but noticeably increases the performance. We then expanded the HARDI signal in terms of spherical harmonics up to order $\ell \leq 4$. Then we computed the rotation invariant brain pyramid features for seven different scales, namely $\sigma = \{1, 2, 4, 6, 8, 10, 12\}$. The upper limit for the tensor derivative order was set to $L = 8$. We additionally included the power-spectra of the expansion coefficients a^ℓ in the feature vector. After computing all feature images we took the square-root out of the single power-spectra and normalized the feature vectors voxel by voxel with respect to the l_2 -norm. The parameters (number of trees in the forest and number of variables to split at each node in the decision trees) are determined experimentally on the training set by minimizing the OOB error rate [Bre01]. We set the number of variables to split at each node (the parameter **mtry**) to two times the square root of the feature dimension. The total number of trees within the forest was set to 1000. The whole experiments where repeated for all seven datasets.

We evaluated our approach by comparing the co-registered region labels with our predicted labels. For this we generated precision-recall graphs (PR-graph) using the number of trees voting for a certain class for thresholding (by dividing the votes by the total number of 1000 trees we gain a probability for a region). The results for the Talairach Level 2 labeling are shown in figure 5. The Talairach Level 3 labeling consist of 54 different regions thus showing the corresponding PR-graph would not be helpful. Hence we decided to show the

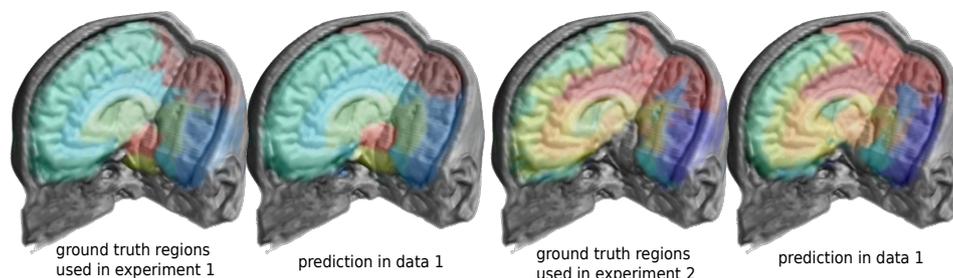


Figure 4: The ground truth regions that we used to train and evaluate our algorithm shown together with our algorithm's regions prediction. We can clearly see that our predictions are much more consistent with the data.

Table 1: Size and the angular resolution of the used datasets.

	size	directions
data 1	104 × 104 × 76	81
data 2	104 × 104 × 76	31
data 3	104 × 104 × 76	81
data 4	112 × 112 × 51	93
data 5	124 × 124 × 61	61
data 6	104 × 104 × 81	61
data 7	81 × 65 × 60	61

mean, best and worst PR curves within the PR-graph in figure 7. In figure 6 we show the predictions for several regions in data1, data2 and data7. Figure 4 exemplarily shows final predictions for data1.

Considering the PR-graphs we can see that we have a strong similarity for most of the co-registered labels and the predicted region labels. However, they are not identical and sometimes they differ strongly. This is in particular true for small regions (in particular when considering the Talairach Level 3 labeling). For instance, considering the worst PR-curve in the PR-graph in figure 7 corresponding to data7 we must assume that our algorithm false completely. However, considering the votes of the random forests (figure 8) we can observe that the region predicted in data7 is very similar to the region predicted in data2. This is of particular interest because the two datasets strongly differ in visual appearance and quality (data7 61 directions vs. data2 31 directions, see table 1). We observed that especially for small regions the co-registration often false and leads to slightly misplaced labels. In contrast, our algorithm includes the information of all training data ensuring that the connection between label and local appearance is consistent.

Further probability maps corresponding to the Level 3 regions for data1 and data2 can be found in figure 9 and 10.

5 Conclusion

A model-free parcellation of the brain has several advantages. In contrast to a registration based approach, where the data is registered to a reference brain in an initial step before the

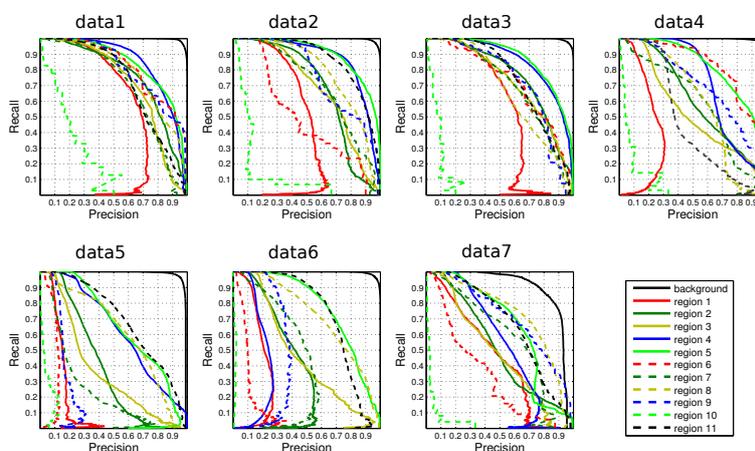


Figure 5: Results corresponding to the Talairach Daemon database regions of Level 2 (Lobes, 11 regions + background). The PR-curves are representing the similarity of the co-registered region masks and the predictions produced by our algorithms.

brain regions can be determined and categorized, our approach segments the brain based on the structural appearance of the data independently of the brains global orientation and position. Furthermore, our experiments have shown that due to the statistical learning approach based on random forests, our algorithm automatically learns the natural variation in human brains and thus can extrapolate these information to the individual structures in unclassified images. As a result our algorithm becomes robust against errors in the training data leading to better segmentation results.

Acknowledgment

This study was partly supported by the Excellence Initiative of the German Federal and State Governments (EXC 294). Marco Reisert is indebted to the “Baden-Württemberg Stiftung” for the support of this research project by the “Eliteprogramme for Postdocs”. We thank Susanne Schnell, Constantin Anastosopoulos and Irina Mader for providing the data.

References

[Bre01] Leo Breiman. Random Forests. *Mach. Learn.*, 45:5–32, October 2001.

[KFR03] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *SGP '03: Proc. of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164. Eurographics Association, 2003.

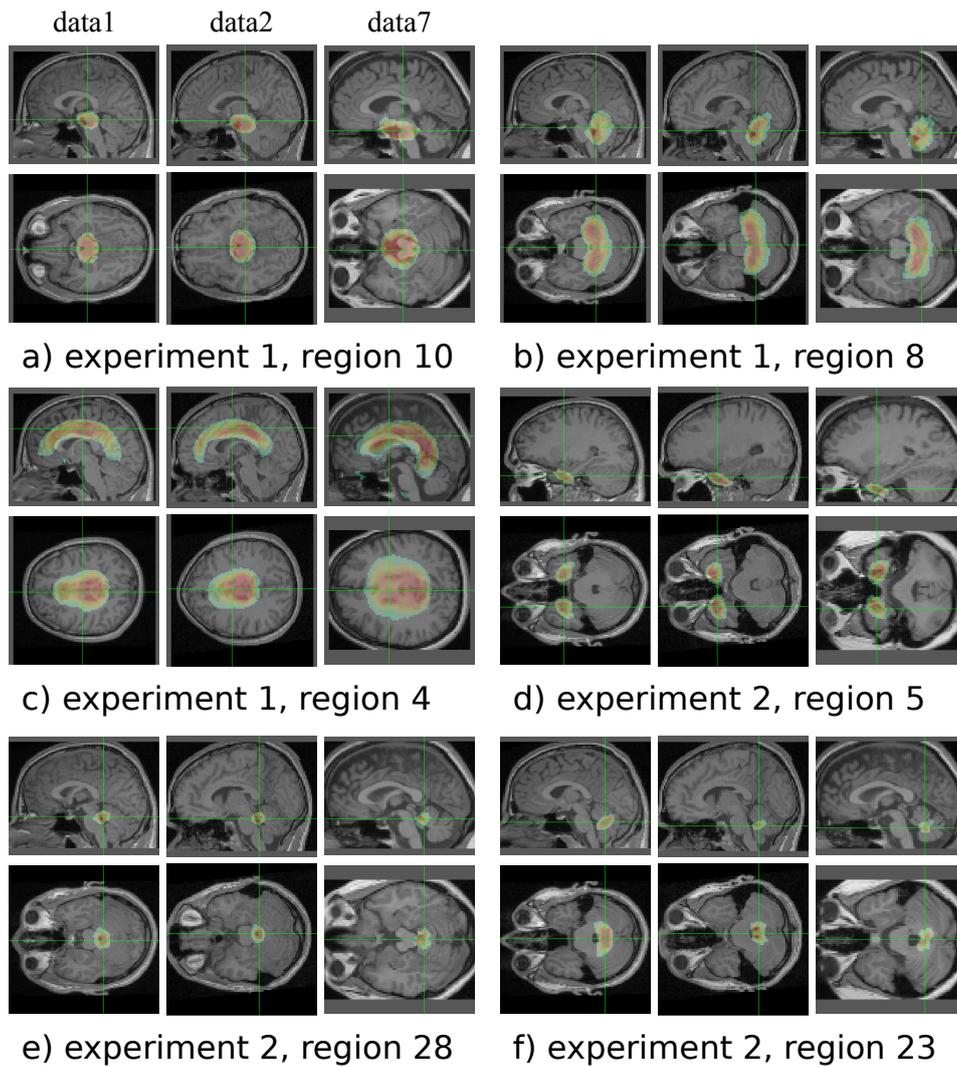


Figure 6: Heat maps representing the probability for a certain region. We exemplarily show detections for the datasets data1, data2 and data7.

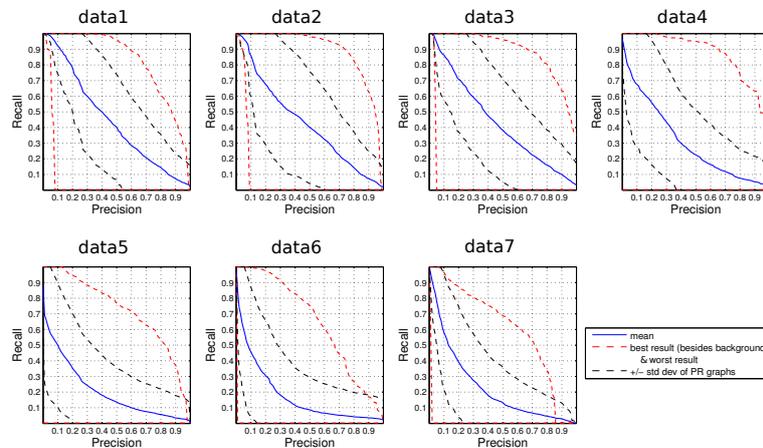


Figure 7: The Talairach Level 3 labeling consist of 54 different regions. We show the mean, best and worst PR curves within the PR-graph. We can observe that for some regions the results are very bad (see e.g. the curve representing the worst performance for data7). However, considering the corresponding output of the random forest (figure 8) we can observe that the prediction is quite good but the ground truth label was incorrectly co-registered. Moreover, considering figure 9 and figure 10 we see that most of the predicted regions are corresponding within different brains leading to the conclusion that our approach includes the information of all training data, automatically learns a data consistent relationship between region labels and appearance and thus becomes robust against errors in training data.

[Lan97] Summerlin J.L. Rainey L. Freitas C.S. Fox P.T. Lancaster, J.L. The Talairach Daemon, a database server for Talairach Atlas Labels. *Neuroimage*, 1997.

[RB09] M. Reisert and H. Burkhardt. Spherical Tensor Calculus for Local Adaptive Filtering. In S. Aja-Fernández, R. de Luis García, D. Tao, and X. Li, editors, *Tensors in Image Processing and Computer Vision*. Springer, 2009.

[Ros95] M. Rose. *Elementary Theory of Angular Momentum*. Dover Publications, 1995.

[SRB11] Henrik Skibbe, Marco Reisert, and Hans Burkhardt. Gaussian Neighborhood Descriptors for Brain Segmentation. In *Proc. of the 12th International Conference on Machine Vision Applications (MVA 2011)*, Nara, Japan, 2011.

[SRS⁺10] Henrik Skibbe, Marco Reisert, Thorsten Schmidt, Klaus Palme, Olaf Ronneberger, and Hans Burkhardt. 3D Object Detection using a Fast Voxel-Wise Local Spherical Fourier Tensor Transformation. In *Proc. of the DAGM*, pages 412–421, Darmstadt, Germany, 2010. LNCS, Springer.

[SSK⁺09] S. Schnell, D. Saur, B.W. Kreher, J. Hennig, H. Burkhardt, and V.G. Kiselev. Fully automated classification of HARDI in vivo data using a support vector machine. *Neuroimage*, 46:642–651, 2009.

[TWBW99] DS. Tuch, RM. Weisskoff, JW. Belliveau, and VJ. Wedeen. High angular resolution diffusion imaging of the human brain. In *Proceedings of the 7th Annual Meeting of the ISMRM*, Philadelphia, USA, 1999.

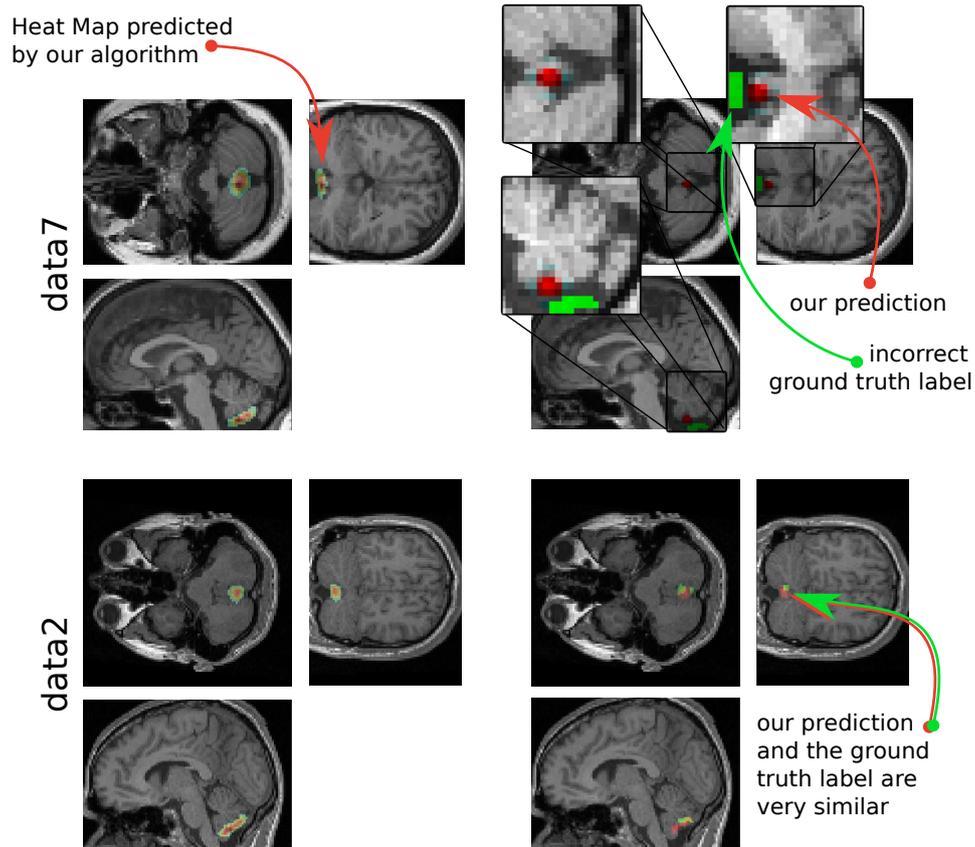


Figure 8: Our approach includes the information of all training data, automatically learns a data consistent relationship between region labels and appearance and thus becomes robust against errors in training data: Here we show the heat maps corresponding to the region in experiment 2 with the worst performance in data7. We can clearly see, that our algorithm detected the region correctly (see the upper left heat map and for comparison the heat map corresponding to data2 in the lower left images). However, the ground truth labeling has been registered incorrectly to the data which is illustrated in the upper right images (green corresponds to the ground truth region, red is our predicted region.)

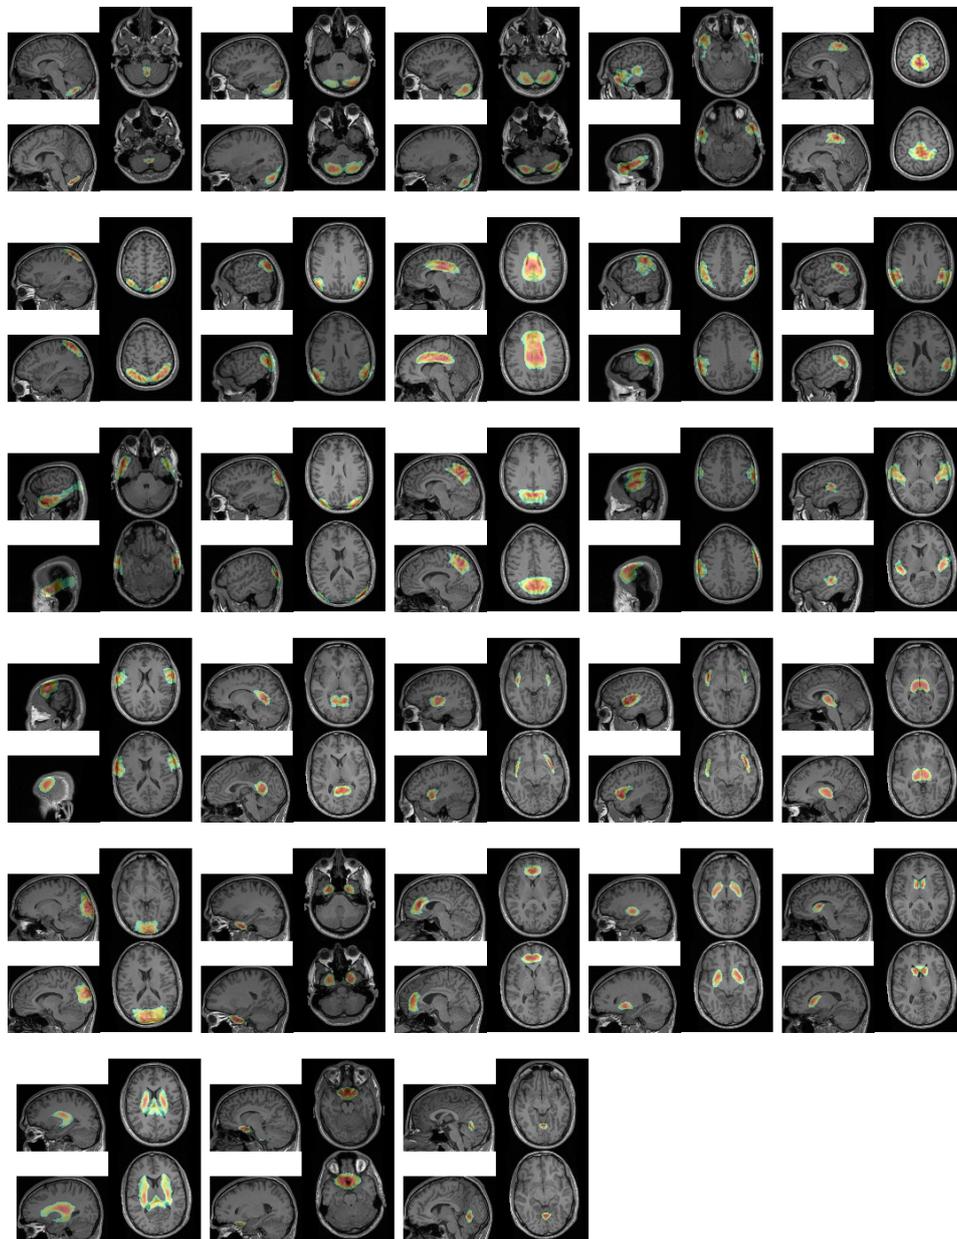


Figure 9: Heat maps representing the probability for all regions used in experiment 2 exemplarily shown for the datasets data1 and data2 (continued in figure 10).

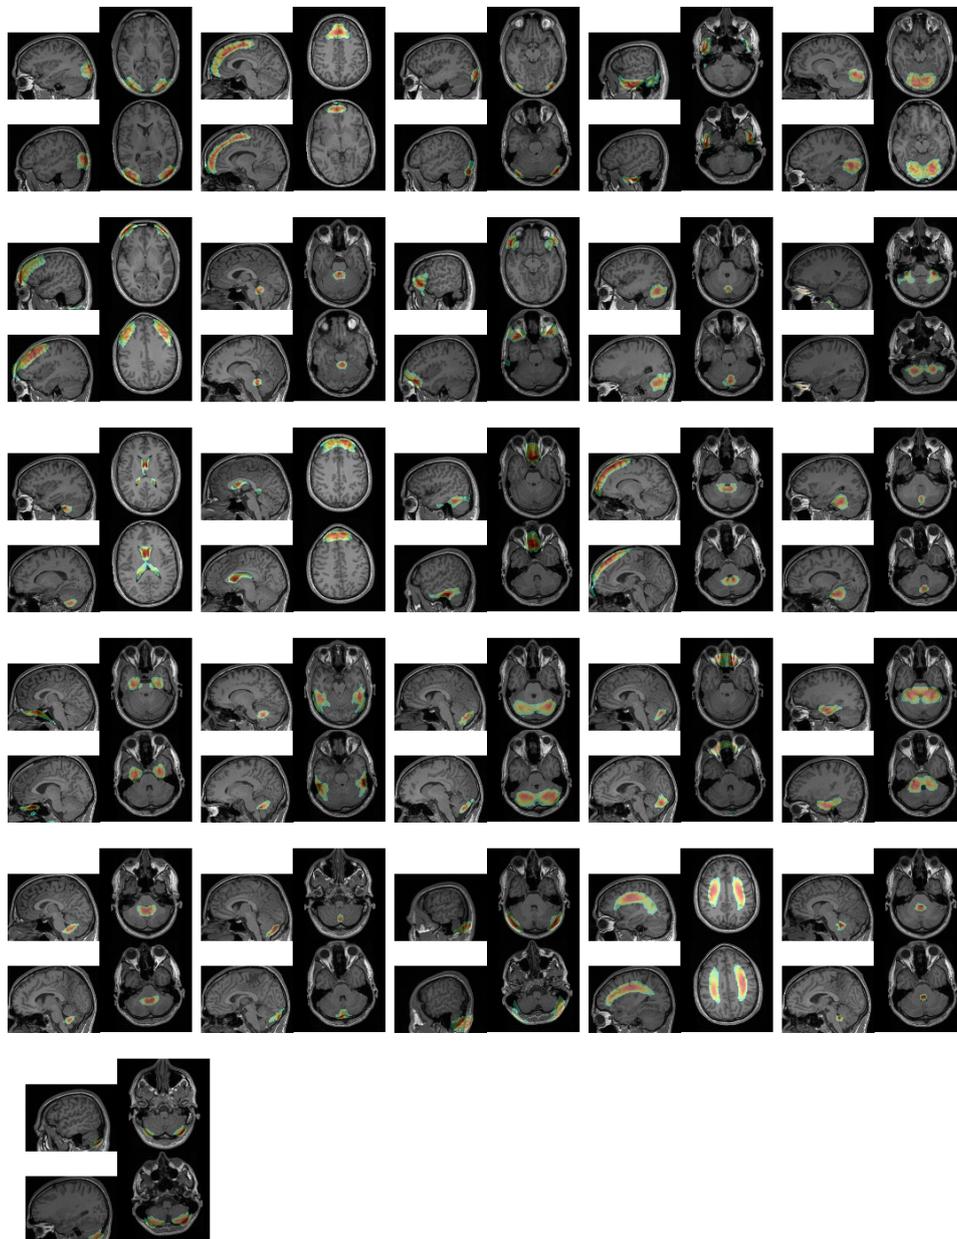


Figure 10: Heat maps representing the probability for all regions used in experiment 2 exemplarily shown for the datasets data1 and data2 (starting in figure 9).