Bibliography index
Martin Styner, Guido Gerig, *Evaluation of 2D/3D bias correction with 1+1ES-optimization*, 1997, TR 179, Technical Report Image Science Lab, ETH Zürich, October 1997.

Evaluation of 2D/3D bias correction with 1+1ES-optimization

Technical Report BIWI-TR-179

Martin Styner, Guido Gerig
Communication Technology Laboratory, Image Science
ETH-Zentrum, CH-8092 Zurich, Switzerland
email: {styner,gerig}@vision.ee.ethz.ch

Abstract

This paper presents tests of a new approach for the correction of image inhomogeneities directly from the corrupted image data.The distortion of the image brightness values by a low-frequency bias field often occurs in MR imaging and impedes visual inspection and intensity-based segmentation. The inhomogeneity problem is even more pronounced in surface-coil images. The new correction method is based on a simplified model of the imaging process and the inhomogeneity field. The appropriateness of a polynomial model of the bias field in MR imaging is tested by analyzing the MR images of a phantom with known ground truth.

The new bias correction scheme has been applied to large series of 2D and 3D MRI image data sets as well as to different types of synthetic images, demonstrating the robustness and generic nature of the algorithm. This paper additionaly contains two sections about the apropriate choice of parameters for the correction.

Evaluation of 2D/3D bias correction with 1+1ES-optimization

Technical Report BIWI-TR-179

Martin Styner, Guido Gerig
Communication Technology Laboratory, Image Science
ETH-Zentrum, CH-8092 Zurich, Switzerland
email: {styner,gerig}@vision.ee.ethz.ch

Abstract

This paper presents tests of a new approach for the correction of image inhomogeneities directly from the corrupted image data.The distortion of the image brightness values by a low-frequency bias field often occurs in MR imaging and impedes visual inspection and intensity-based segmentation. The inhomogeneity problem is even more pronounced in surface-coil images. The new correction method is based on a simplified model of the imaging process and the inhomogeneity field. The appropriateness of a polynomial model of the bias field in MR imaging is tested by analyzing the MR images of a phantom with known ground truth.

The new bias correction scheme has been applied to large series of 2D and 3D MRI image data sets as well as to different types of synthetic images, demonstrating the robustness and generic nature of the algorithm. This paper additionaly contains two sections about the apropriate choice of parameters for the correction.

1 General remarks

The optimization with the 1+1ES-method has the advantage of a automatic adjustment of the step size. This adjustment is dependent on the value of the c_{grow}-factor. For a big c_{grow} the optimization is fast, but more likely to end in a local minimum compared to a little c_{grow}. But the choice of a little c_{grow} leads to rather slow optimization.

This report is about the examination of this adjustment factor c_{grow} on one hand, and on the other hand about the general handling of the parameters of the bias correction with 1+1ES-optimization.

Normally a distorting bias field is not added, but multiplied to the original image. Therefore the bias correction is computed normally with the "log"-flag on. When using the "log"-flag the bias correction first transforms each pixel of the original image and all parameters concerning pixel values by the natural logarithm-function. The bias correction is then computed as an additive bias field to the log-transformed image. The corrected image is retransformed by applying the exponential-function to each pixel.

All tests were performed with a $c_{shrink} = c_{grow}^{1/4}$, which proved to be a good value for c_{shrink}. A small evaluation of choosing c_{shrink} other than this, showed just small improved efficiency in the best cases. When c_{shrink} was chosen to far away of $c_{grow}^{1/4}$ efficiency even heavily decreased.

2 2D-image of a step edge

2.1 General Method

A two dimensional image is created with a step edge function in x-direction. Then this image is distorted by a bias field composed of the 2D Legendre polynoms of second degree. Gaussian and uniform distributed noise of different amplitude is added to the distorted image (no "log"-flag).

The step edge function is constructed of two classes located at $\mu_1 = 100$ and $\mu_2 = 140$. The parameter σ of the valley-function of the energy function is chosen to the value 6.8.

2.2 Detailed examination: Simple coefficients, no noise

2.2.1 Minima

Using the coefficients $[0, 0, 0, 0, 100, 0]$ for the 2D Legendre polynoms, we observed multiple local minima, the local minimum at $[9.0984, 0, 0, 19.844, 100, 15.8584]$ was chosen to be further examinated. The existence of a local minimum is the effect of choosing the value of the fifth coefficient very high compared to the distance between the two classes of the step edge function.

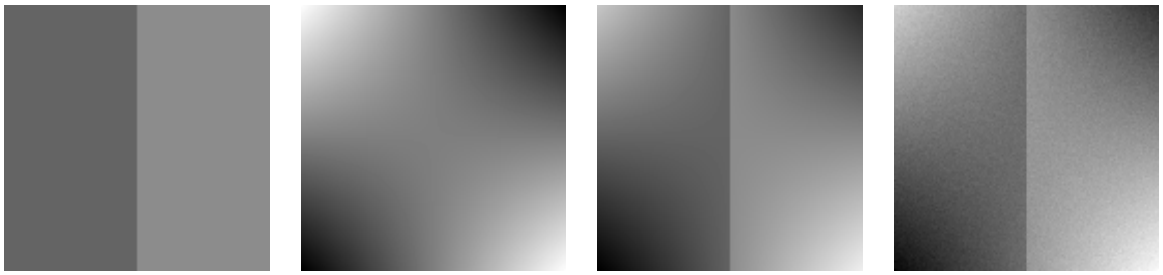

Figure 1: (from right to left) (*i*): Image of a step edge function along x-axis (gray values 100,140), (*ii*): bias field, (*iii*): image *i* added to bias field *ii*, (*iv*): medium noise added to the distorted image *iii* (Gaussian distribution, level 10)

As an evaluation, we did a 3D visualization of the isosurfaces of energies around the minima, when changing the first, the forth and the sixth coefficient of the Legendre polynoms (see figure 2.2.1). The evaluation showed that the area connecting local and global minimum has relatively high energy values. Also is the distance from the local minimum to the points of lower energy big.

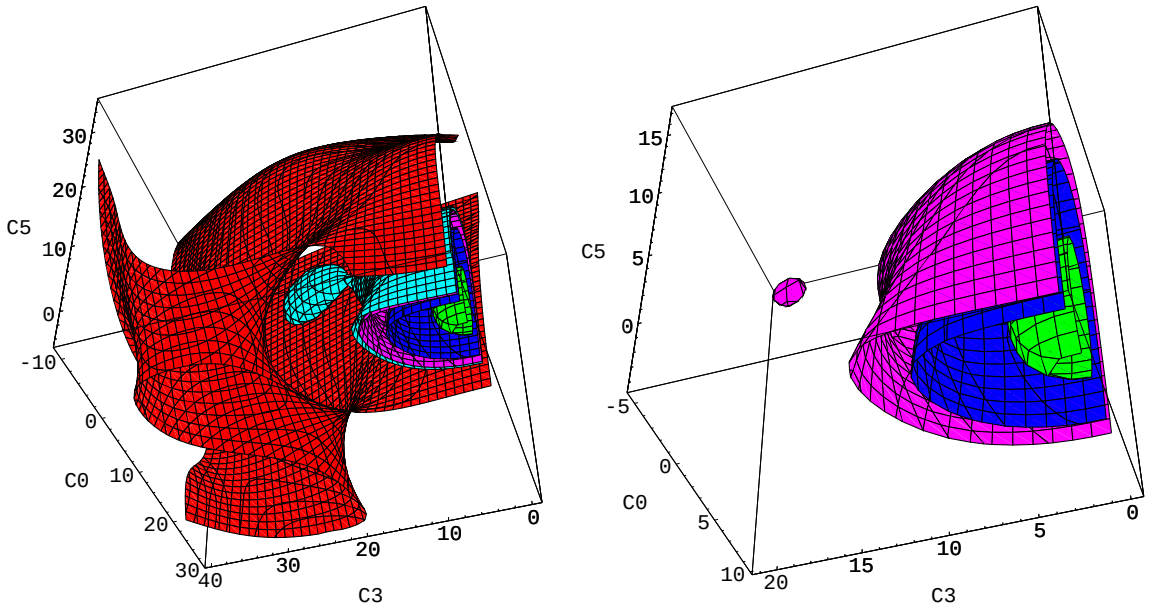

Figure 2: Plots of the energy function of the bias-distorted step edge image without any noise around the coefficients $[0, 0, 0, 0, 100, 0]$. axes: first coefficient ('0'), forth coefficient ('y'), and fifth coefficient ('y^2') of the 2D Legendre polynoms (degree 2). Isosurfaces are drawn in the left image at the energies 1000, 2000, 2800,3000, and 6000, in the right at the energies 1000, 2000, and 2870.Clearly can the local minimum be seen.

We calculated also the Hessian matrix at the local minimum and plotted the values of the

energy function along the eigenvectors of the Hessian matrix. The Hessian matrix at the local minimum is regular and the plots along the eigenvectors showed no other minimum although the global minimum is in the region spanned by the plotted range of the eigenvectors (see figure 3).

As comparison we also plotted the values of the energy function along the vector connecting local and global minimum (see figure 4). In the zoomed plot around the local minimum it can clearly be seen, that the energyfunction along the eigenvector of the smallest eigenvalue is less steep than along the direct line from local to global minimum.

When adding noise to the stepedge image, the plots around the local minimum do not change, that means the energyfunction stays smooth, just the position of the minimum is slightly shifted ! We can therefore conclude that the noise does not affect the smoothness of our energyfunction, but it affects the position of the global minimum.

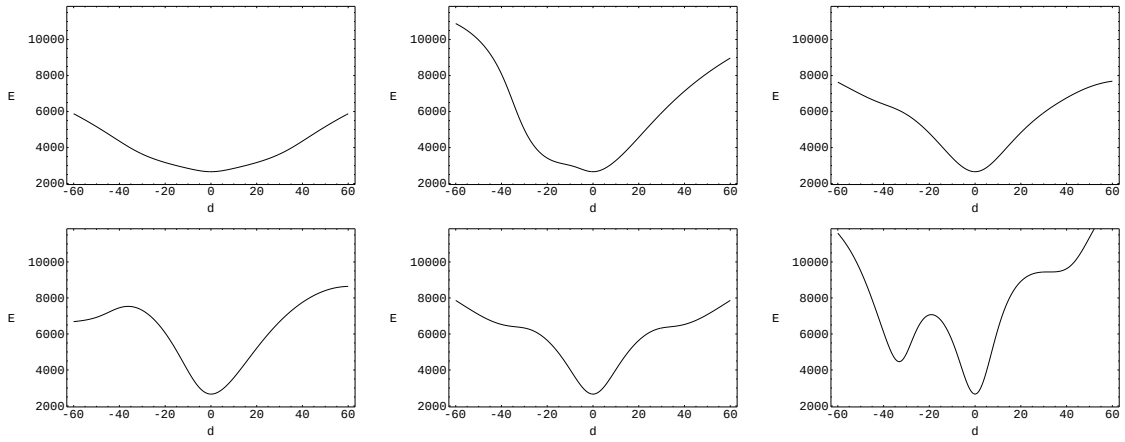

Figure 3: Plots of the energy function of a bias-distorted step edge image with Legendre polynoms coefficients $[0, 0, 0, 0, 100, 0]$. Plots were obtained along the eigenvectors of the Hessian matrix at the local minimum $[9.0984, 0, 0, 19.844, 100, 15.8584]$. Plots are ascending sorted by the eigenvalues of the corresponding eigenvector.

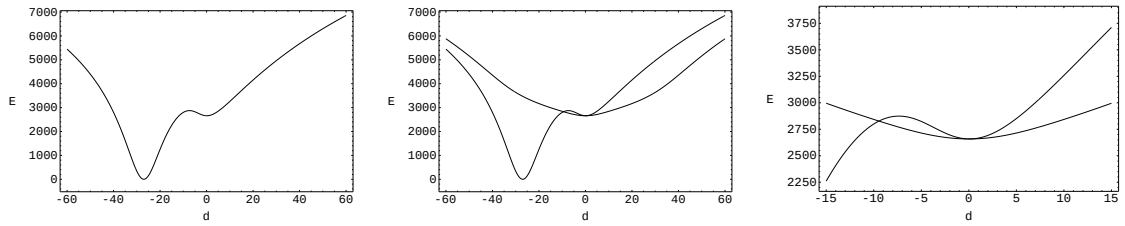

Figure 4: Plots of energy function of a bias-distorted step edge image with Legendre polynoms coefficients $[0, 0, 0, 0, 100, 0]$. Left (i): Plot along the line connecting local and global minimum; null position is at the local minimum $[9.0984, 0, 0, 19.844, 100, 15.8584]$. Middle (ii): Combined plot of i and energies along first eigenvector (smallest eigenvalue) of Hessians matrix. Right (iii): Zoomed plot of ii around the local minimum.

Other local minima were observed at around $[-16, 0, 0, 43, 100, 23]$, $[-8, 0, 0, 20, 100, -15]$, $[9, 0, 0, 20, 100, 16]$, and $[17, 0, 0, 43, 100, -22]$. All those minima possesed different energies, but none of them has a energy as low as the local minimum discussed before.

2.2.2 Optimization with 1+1ES

In table 1 it can be seen, that a good c_{grow}-value is at 1.05, where we have an acceptable computation time and a very good chance to find the global minimum. When c_{grow} is chosen too high the calculation takes even more time than with a smaller c_{grow} and only little chances exist to find the global minimum. The computation time was measured on a SUN-ultra 1, initial step size was chosen at 100, which is the maximal value of a coefficient of this bias field.

c_{grow}	Elapsed time	Optimum
2.0	$\approx$ 4 min	$<$ 50 %
1.5	$\approx$ 3 min	79 %
1.1	$\approx$ 4 min	81 %
2 x 1.1	$\approx$ 8 min	98 %
1.05	$\approx$ 7 min	99 %
1.01	$\approx$ 14 min	100 %

Table 1: Time/Optimum tradeoff: the first column shows the used c_{grow}-factor, the second the average elapsed time and the third column shows the percentage, where the optimization found the global minimum. The average elapsed time belong to an optimization on a SUN-ultra 1. In the second row the optimization was invoked twice, in the second attempt with the found minimum as starting point. All other optimization used the null vector as starting point.

2.3 Various coefficients, various noise

We added several bias fields to the step edge image and also uniform and Gaussian distributed noise of different levels. The bias correction then run on the noisy images without further preprocessing. The c_{grow} parameter was chosen at 1.05, the initial step size at 40.

The following different coefficients of the Legendre polynoms were used:

degree	coefficients
2	$[0.0, 0.0, 0.0, 0.0, 100.0, 0.0]$
2	$[10.0, 3.45, 6.6, 2.3, 12.4, 25.1]$
2	$[1.0, 2.0, 3.0, 4.0, 5.0, 6.0]$
2	$[16.734, 8.2, 0.4, -23.5, -12.3, 10.342]$
2 (log)	$[0.3, 1.4, 0.123, -1.2, -0.3, 0.65]$
3 (log)	$[0.2, -0.02, -0.34, -0.34, 0.23, 1.23, 0.23, 1.4, -1.053, 0.89]$

The different noise levels were chosen as a fraction of the difference d of the mean values of the two classes, which is in our step edge image 40. The following levels were therefore used: $d/8 = 5$, $d/4 = 10$, $3 \cdot d/8 = 15$, $5 \cdot d/8 = 25$, $7 \cdot d/8 = 35$, and $5 \cdot d/4 = 50$.

At the noise level $5 \cdot d/8$, the two classes start to overlap even without the bias field. The images having a Gaussian distributed noise of level up to $7 \cdot d/8$ and those having a uniform distributed noise level up to $5 \cdot d/8$ were all very well corrected with a probability of 0.99 (with the given parameters). Whereas the bias field of images with higher noise was not computed correctly. We observed in those images always a drift from left to right (orthogonal to the step edge) (see figure 5). This bias field is a global minimum of our energyfunction! The energy at the right solution is therefore lower than on the calculated wrong solution.

When doing an anisotropic filtering before the bias correction, then the problem of images with a too high noise level does not occur (see also figure 5) and the correct bias field is computed. The calculated bias field of the filtered image is then applied to the original image.

The calculation of the bias field took between 5 and 8 minutes, but generally it can be said that the more noise an image had the more computation time it took, although the increase is relatively small.

3 2D-image and 3D-image of a checkerboard

A checkerboard was generated, consisting of two classes ($\mu_1 = 100$, $\mu_2 = 140$). Different bias field and noise levels were added. Then the bias correction was applied to all those images. The bias fields were all computed correctly up to a high level of noise. Even at the noise level $5 \cdot d/4$, with d as the difference of the mean class values, the bias field was computed correctly (!).

4 2D-Image of a head obtained by a surface coil

This section is about the bias correction of the image of a head obtained by a surface coil. The sensibility of a surface coil decreases quadratic with distance, we choose therefore the bias field to be composed of Legendre polynoms to the second degree. In figure 7 is the original image shown. It can clearly be seen, that a computerized analysis of the image is extremely difficult without a correction, since the pixel values of the classes of tissues strongly overlap.

The results of the bias correction are satisfying, as in figure 8 can be seen.

In the corrected head image we observed on regions that should consist of homogeneous intensity a sawtooth function was added. This is caused by the addition of a continuous bias field to the discrete original head image. In figure 9 this effect and in figure 10 the origin of the sawtooth function is shown.

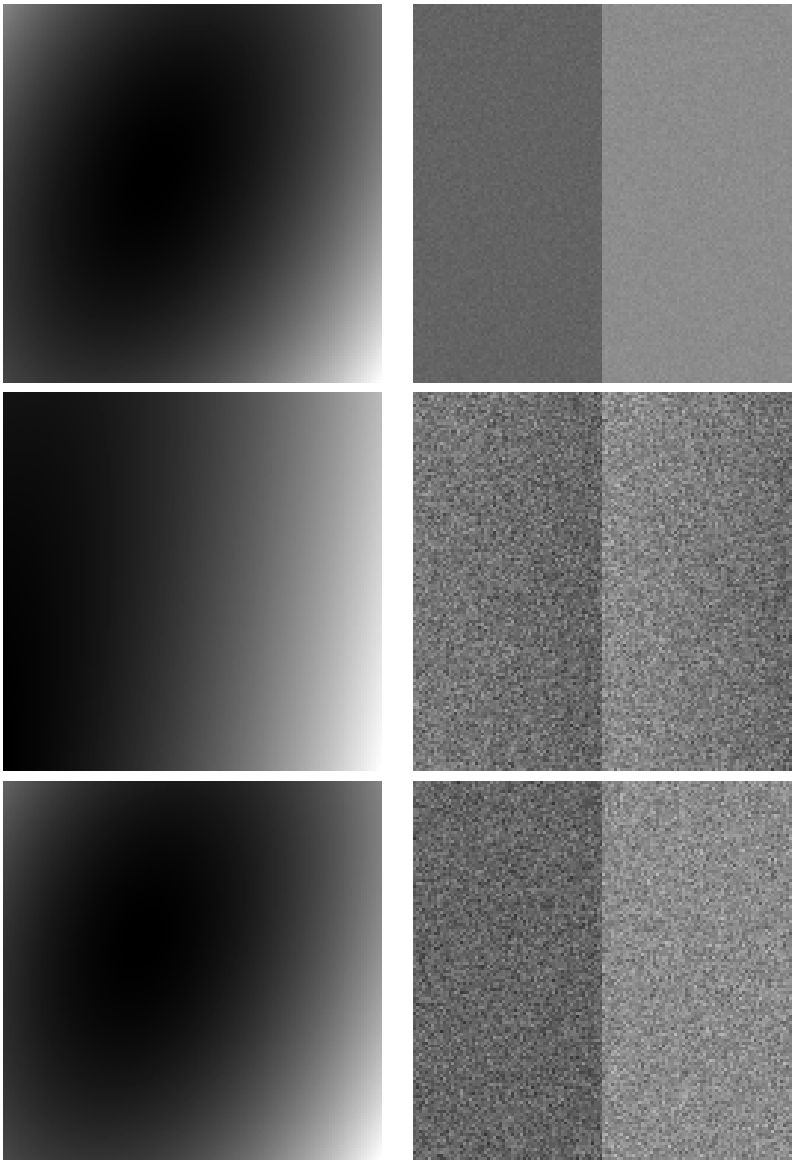

Figure 5: Correct and incorrect calculation of the bias field in the step edge image depending on the noise level:
top row: noise level $3 \cdot d/8 = 15$ with correct bias field,
middle row: noise level $5 \cdot d/4 = 50$ with incorrect bias field,
bottom row: noise level $5 \cdot d/4 = 50$ and anisotropic filtering with correct bias field.
Right column: calculated bias field, left column: corrected image.

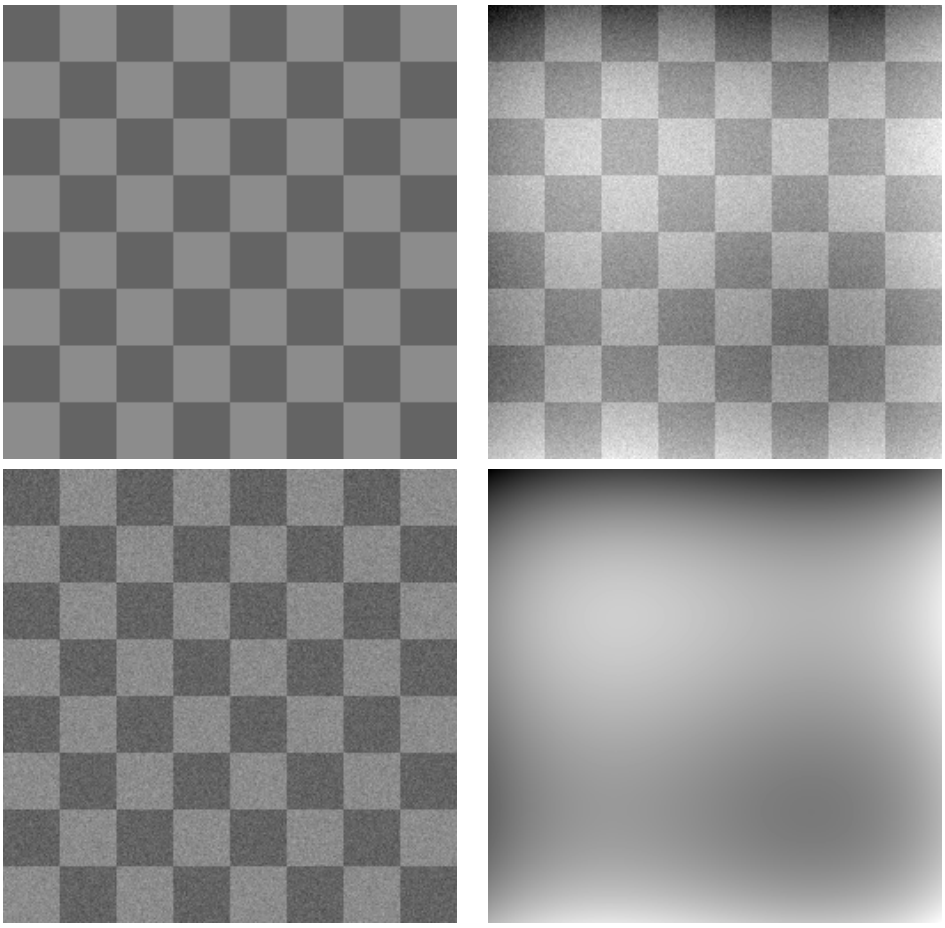

Figure 6: bias field correction of a checkerboard (Gaussian noise, level $5 \cdot d/8$) : t.l.: original image, t.r: distorted image, b.l.: corrected image, b.r: calculated bias field

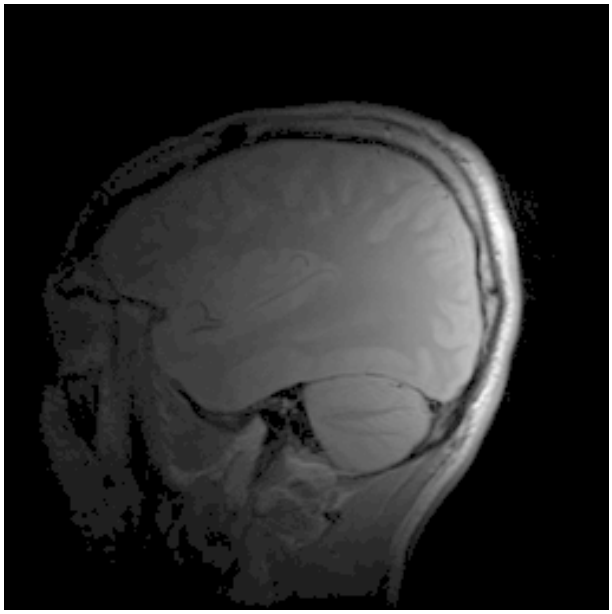

Figure 7: Original image of a head obtained by a surface coil.

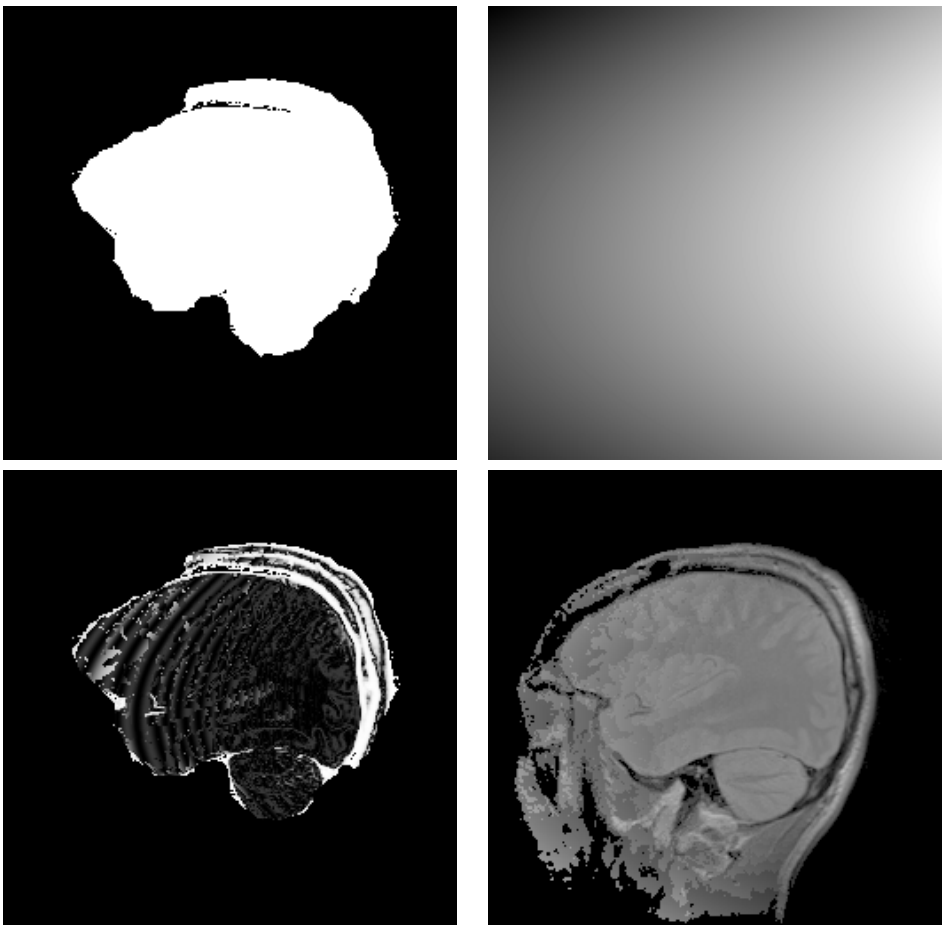

Figure 8: bias field correction of the head image: t.l.: mask, t.r: computed bias field, b.l.: energy image, b.r: corrected image

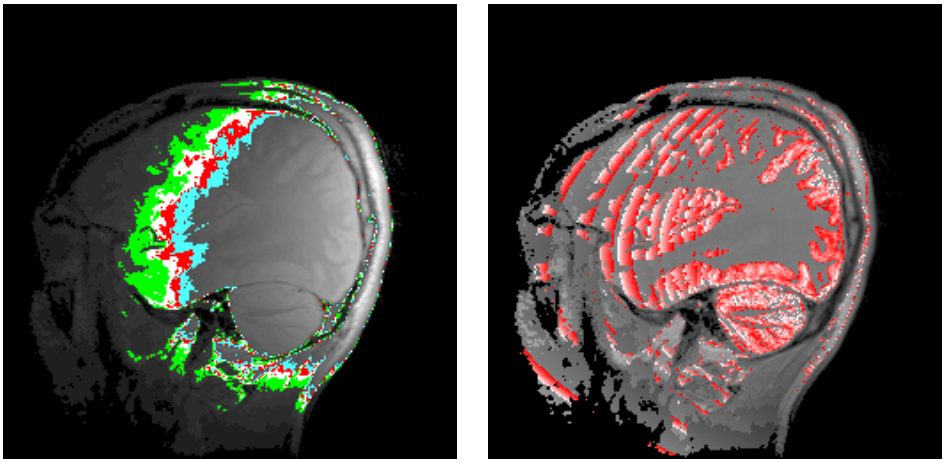

Figure 9: On the original head image (left) large regions with the same pixel value are observed due to the discreetness. Three such regions are highlighted by color. This discreetness causes a sawtooth function (in red) in the regions of the corrected image (right) that should be homogeneous.

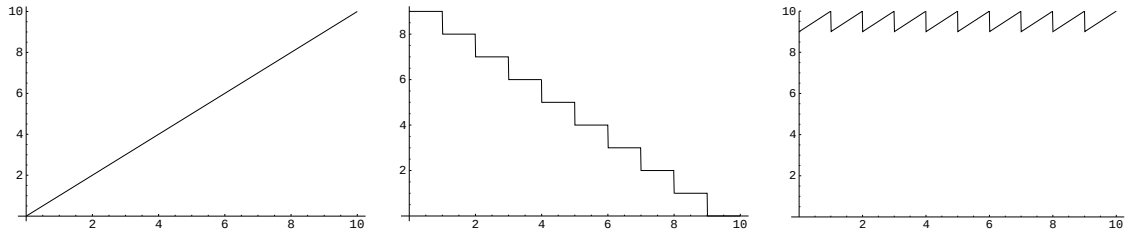

Figure 10: Origin of a sawtooth function (in 1D): (from left to right) (*i*) continuous bias field, (*ii*) discrete image, (*iii*) corrected image (addition of *i* and *ii*).

5 2D-Image of a breast obtained by mammography

A breast image (see figure 11) is obtained by mammography and similar to a MRI image obtained by a surface coil, we observe regions with heavy decreasing pixel values. This causes that the range of pixel values of the breast tissue is large and intersects with values of other tissues or blood vessels.

The results of the bias correction are satisfying, as in figure 12 can be seen. The calculated bias field consists of Legendre polynoms up the third degree. Two classes were used with $\mu_1 = 25, \sigma_1 = 5, \mu_2 = 120, \sigma_2 = 17$.

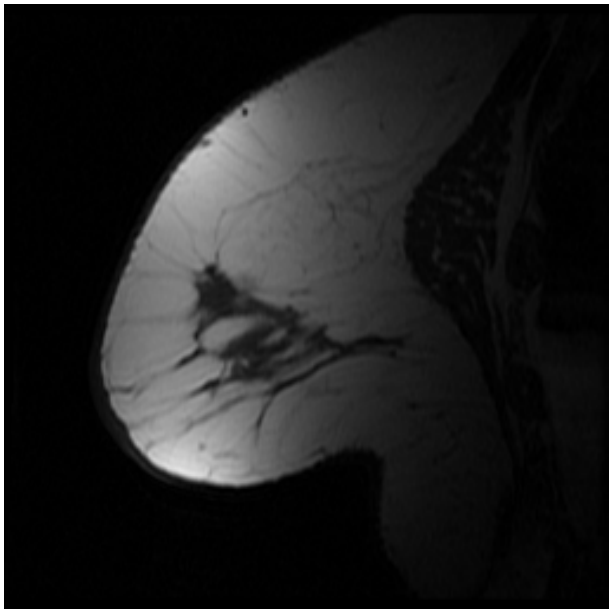

Figure 11: Original image of a breast obtained by mammography.

6 3D-image of onion like phantom

A onion like phantom was constructed that consists of layers/skins of different intensity. The skins of the onion are ellipsoids. In this example, the onion consists of 3 skins with the intensities $\mu_1 = 120, \mu_2 = 150, and \mu_3 = 180$ and a resolution of 64x64x64 voxels.

Since we're working in 3D, we use subsampling to calculate first approximately the bias field on the subsampled volume and second calculate the exact bias field on the whole volume.

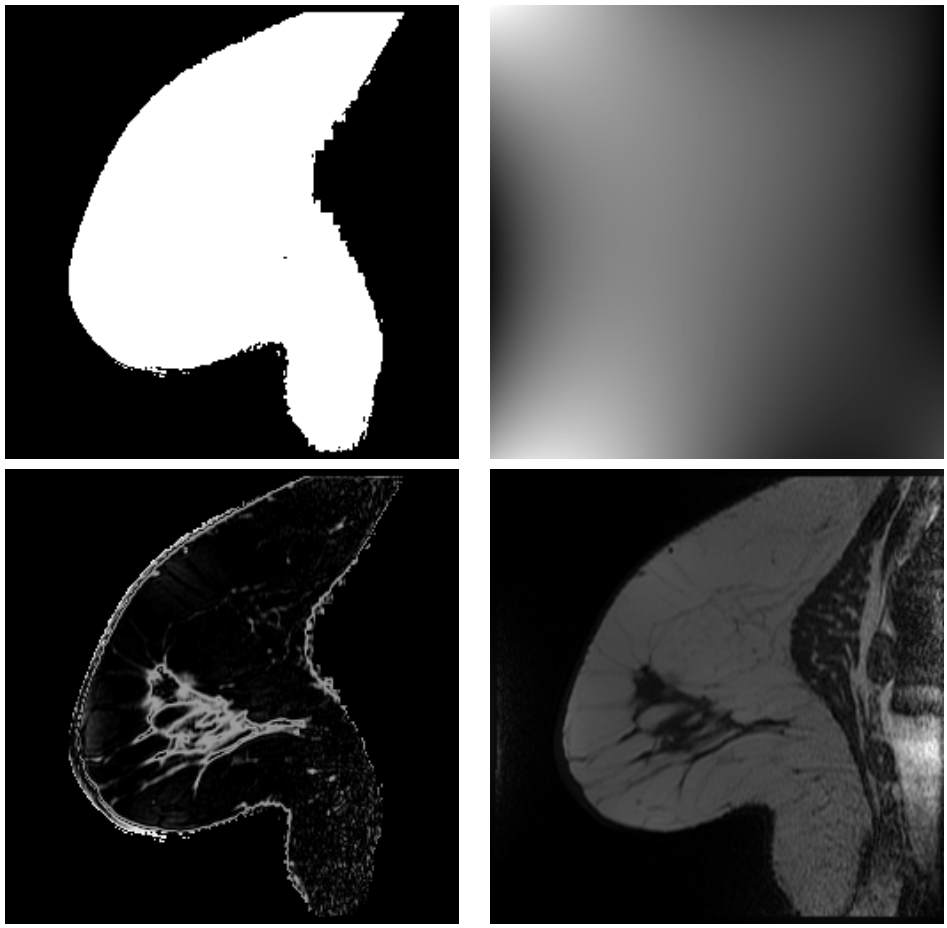

Figure 12: bias field correction of the breast image: t.l.: mask, t.r: computed bias field, b.l.: energy image, b.r: corrected image

Here we observed that when choosing the subsampling factor too high (more than factor 2), the optimization was sometimes trapped in a local minima. Without subsampling all applied bias fields and noises were computed correctly! But subsampling reduced the computation time by the factor 2. Mean computation time of this example without subsampling is between 1 and $1\frac{1}{2}$ hour on a SUN ultra 1 (Legendre polynoms up to second degree).

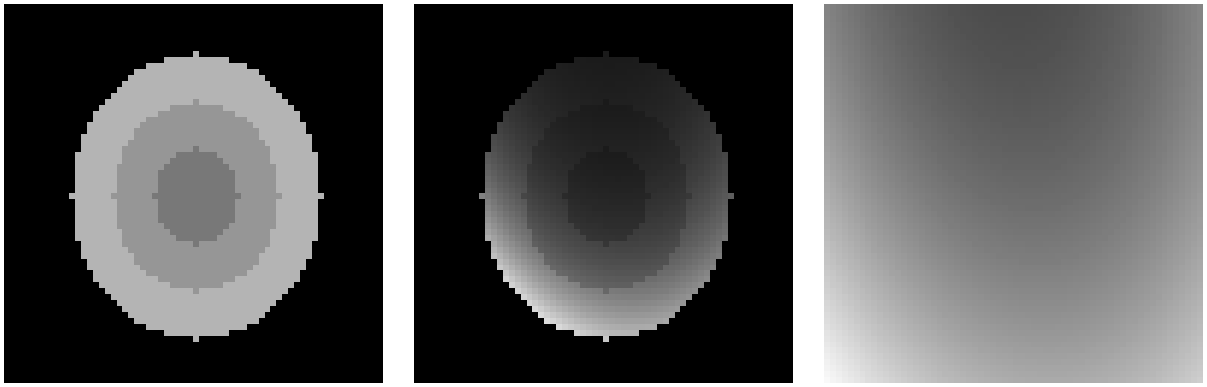

Figure 13: bias field correction of an onion like phantom (slice in the middle of volume) without any noise: from left to right: original image, distorted image, calculated bias field

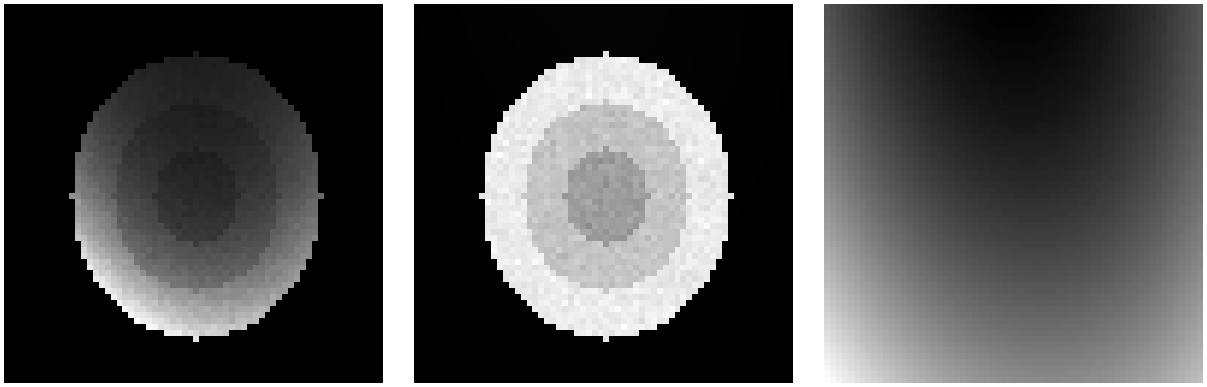

Figure 14: bias field correction of an onion like phantom (slice in the middle of volume) without noise of level $\frac{d}{2}$: from left to right: distorted image, corrected image, calculated bias field

7 3D-image of the head of a MS-patient

This section is about the correction of an MRI-volume of a head of a patient with Multiple Sclerosis (MS) (data Test01_MS_PD_001 from the biomorph project). We found out that the different slices of the volume each had a individual brightness. To correct this with our three dimensional bias correction the degree of the Legendre polynoms needed would be at least as much slices as we have. This would slow down the correction extremely to some days as calculation time. Also the calculated bias field wouldn't be correct, because we'd then no longer have a smooth bias field and the bias field would accomodate too much to anatomical structures. We therefore choose to use a multi-step approach:

1. The brightness of each slice is corrected using a bias correction with the Legendre polynom of the zeroth degree.

2. The corrected volume is smoothed by anisotropic filtering (small kappa and few iterations).

3. A three dimensional bias correction is calculated on the filtered volume (degree 2).

4. The calculated bias field is applied to the corrected, unfiltered volume.

The volume generated in step 2 by the anisotropic filtering process serves just to obtain a denoised volume for the three dimensional bias correction. For the first, two dimensional bias correction such a volume with much noise removed is not necessary, since we just use the Legendre polynom of zeroth degree, so that the bias field can not adapt to the noise. But when using Legendre polynoms up to the second or third degree, the calculated bias field could adapt to the noise, which is not desired. That's why we do such an anisotropic filtering. But our resulting volume can not be an the corrected, denoised volume because on one hand the anisotropic filtering filters not just the noise away, to some amount it also filters away anatomical structures or edges, which is of course not desired when using the volume for diagnosis purpose. On the other the filtering makes relatively big parts of the volume homogeneous, which leads together with a continuous bias field to a sawtooth function in those homogeneous areas (see also section 4 / figures 10, 9).

To simplify and speed up the correction we concentrated on two classes, the white matter and the gray matter. Everything else was masked out. Masks were generated by a coarse statistical classification (with PCI) based on a training set on one slice. The volume contains some slices with a very small amount of voxels belonging to white or gray matter, these slices were masked out fully. These slices all are positioned at the boundary of the volume (slices 0-7 at the bottom and slice 23 at the top).

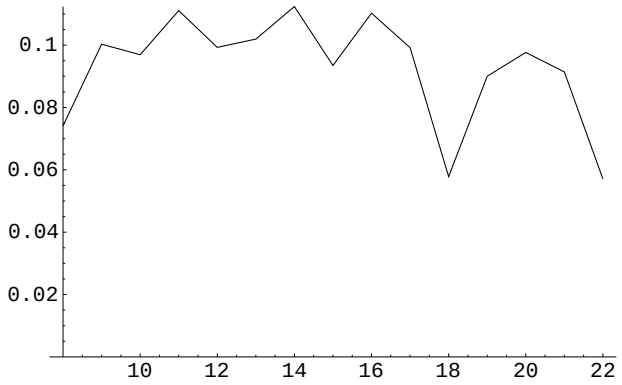

Figure 15: Coefficients of the calculated bias field with Legendre polynoms of zeroth degree. Each coefficient belongs to a slice (shown slices 8-22) of the volume of the MS-patient. The individual slice brightness can be seen clearly.

The calculated coefficients of the first bias correction correspond with the subjective impression of the brightness of the slices: brighter slices have a lower coefficient that darker slices (see figure 15). The difference between the lowest and the highest coefficient is 0.055. This difference is measured in the log-scale, it means a difference of about 15-20 gray levels in our volume (values 0-500).

The computation of the correction of the individual slice brightness was at half a minute per slice, which summoned to 12 minutes per volume. In contrast the three dimensional bias correction lasted around 6 hours (!) per volume with Legendre polynoms up to the third degree. When using Legendre polynoms up to the second degree the dimension of the parameter space decrease from 20 to 10, and the computation time goes down to 1 hour and 45 minutes.

The results of the corrections with Legendre polynoms up to the third degree did only slightly differ from those with Legendre polynoms up to the second degree. The difference of the two solutions was just in those slices of significance that were masked out fully. The result with the Legendre polynoms up to the second degree was smoother and and looked similar to its succeeding slices, that were not masked out fully. Whereas the result with the Legendre polynoms up to the third degree shaped the bias field in these slices not so smooth and diverged in the slices at the bottom of the volume (see figure 16). Also the geometry of the head coil used to scan the volume suggest to choose the Legendre polynoms up to the second degree and not up to the third degree.

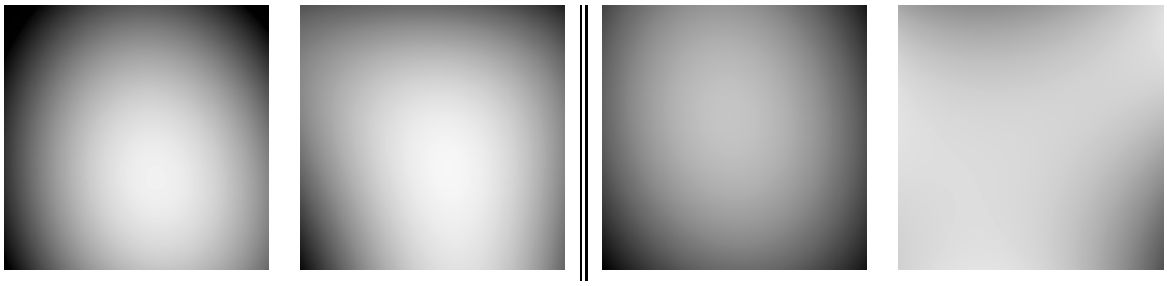

Figure 16: Bias fields of a slice not masked out (slice 16, left side) and of a slice masked out (slice 0, right side). On the left the bias fields of the correction with Legendre polynoms up to the second degree is shown, on the right those up to the third degree.

The 2D-correction, which corrected the individual slice brightness, could be a problem of exporting the data properly from the scanner-own format to a raw-data or GIPL format. The calculated bias fields from the three dimensional correction look like the shape of the head coil, which is circumcentral like a cylinder with decreasing intensity from the center to the boundary. When viewing the uncorrected and the corrected volume there is no big difference noticeable (see figure 18), because the detected bias field is too low.

Figure 17: Bias field (left), input mask (middle) and output mask (right) of slice 17. The output mask masks out all background voxels, the input mask all voxels that surely do not belong to white matter or gray matter.

The effect of the correction is good viewable in figure 19, where the values of the voxels are plotted along a line in the white matter area in the uncorrected and in the corrected volume. The direction of the line is chosen approximately orthogonal to the isosurface of the bias field. On one hand the shift by the first, two dimensional correction can be seen, on the other hand the bias along the plotted line.

To demonstrate the effect of the bias correction, we also did a simple thresholding on the volumes to segment the white matter. For both volumes (uncorrected and corrected) the same

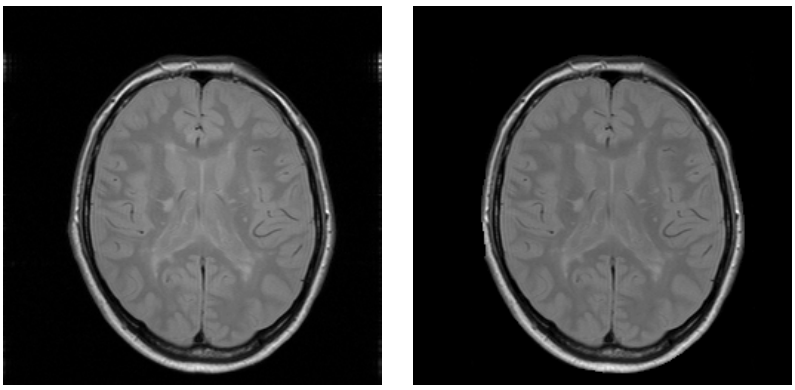

Figure 18: Slice 17 of uncorrected volume (left) and corrected volume (right). A slight shift of intensity can be observed between uncorrected and corrected image, result of the two dimensional bias correction of zeroth degree . But the three dimensional bias correction can hardly be seen by eye.

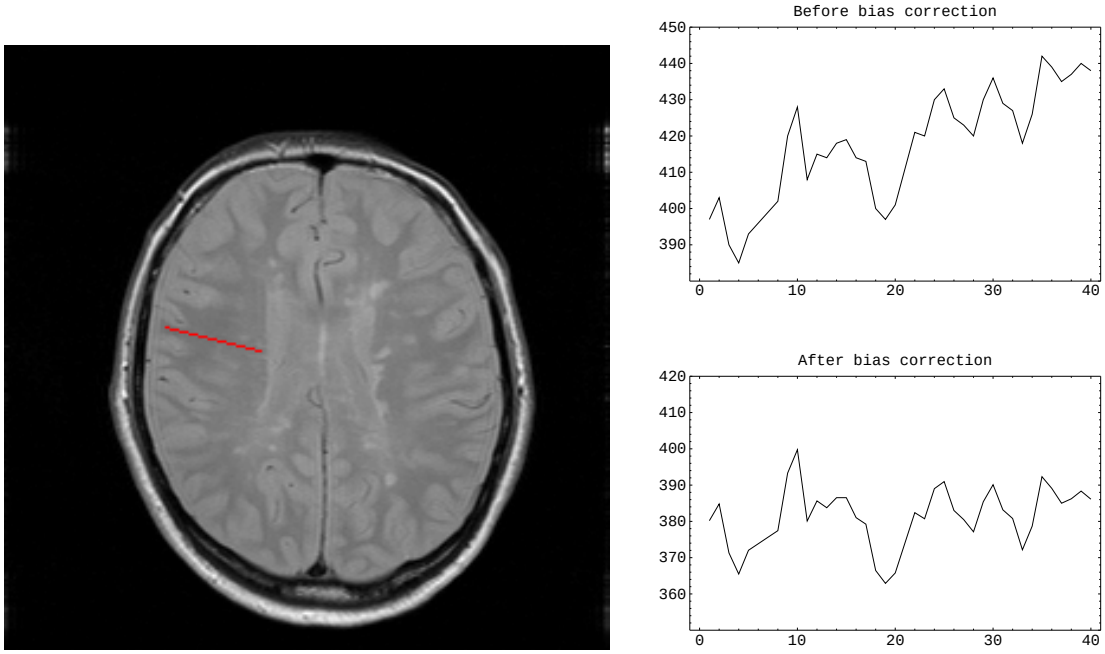

Figure 19: Plots of the voxel values along a radial line in slice 18. Left: Slice of the uncorrected volume with plot line. Right: plot of the values along the plot line, top: in the uncorrected volume, bottom: in the corrected volume.

width of range for the thresholding as applied. The threshold values were adapted manually to the best result (see figures 20 and 21). The segmentation of the uncorrected volume seemed to be more noisy and we were not able (even by increasing the width of range) to find a thresholding that delimited well the boundary of the white matter. Also showed the segmentation of the uncorrected volume empty parts (parts with no white matter) that should contain white matter. The segmentation of the corrected volume was much nicer: better boundaries, not so noisy appearance, no bigger missing parts.

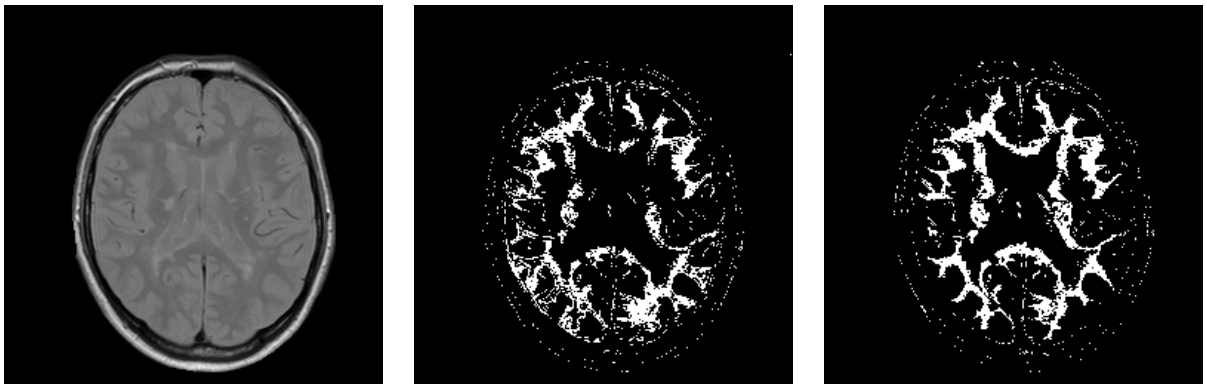

Figure 20: Segmentation of white matter by simple thresholding (manually adapted to the best solution) in slice 17 (left). In the middle the segmentation of the uncorrected volume and in the right the segmentation of the corrected volume is shown. Clearly the much nicer segmentation of the corrected volume can be seen.

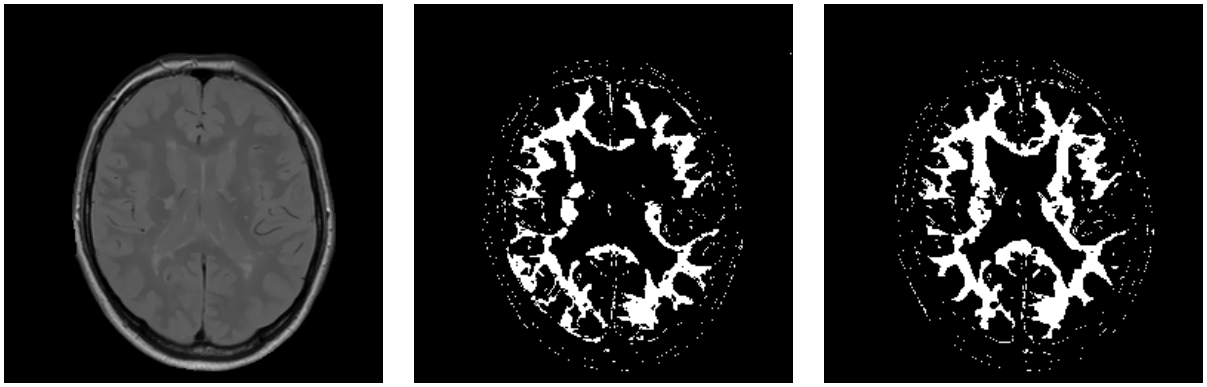

Figure 21: Segmentation of white matter by simple thresholding (manually adapted to the best solution) in slice 17 (left) after applying a three dimensional anisotropic filter to the volume. In the middle the segmentation of the uncorrected volume and in the right the segmentation of the corrected volume is shown. Again the much nicer segmentation of the corrected volume can clearly be seen.

Also a segmentation of the gray matter on slice 18 was performed the same way as the segmentation of the white matter. In the uncorrected image there are some heavy misclassified region in the center of the image, where the bias field reaches its maximal value. The cerebrospinal fluid for example is classified there as gray matter and the gray matter (nucleus caudatus) near the cerebrospinal fluid is not classified as gray matter (see figure 22).

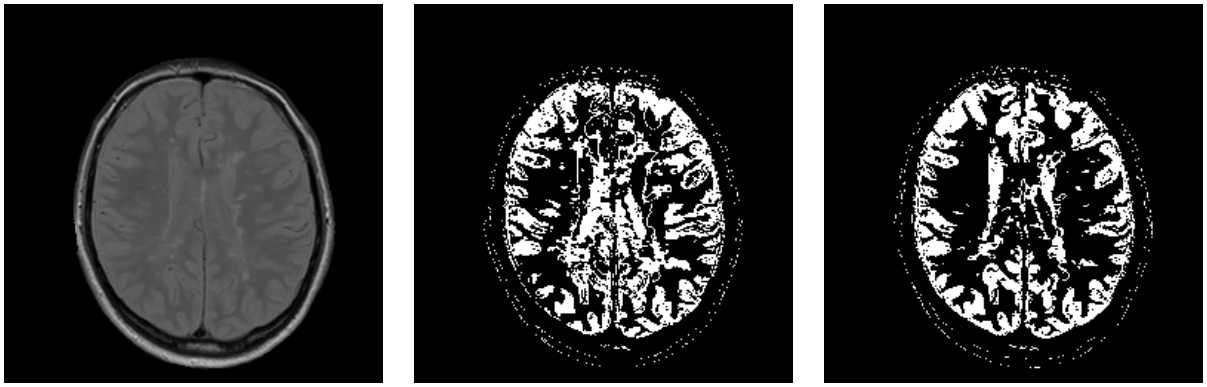

Figure 22: Segmentation of gray matter by simple thresholding (manually adapted to the best solution) in slice 18 (left) after applying a three dimensional anisotropic filter to the volume. In the middle the segmentation of the uncorrected volume and in the right the segmentation of the corrected volume is shown. The classification of the nucleus caudatus changes from incorrect to correct.

8 Anisotropic filtering as a preprocessing step

Sometimes images are very noisy and the position of the global minimum of the energyfunction is moved to a wrong position. The calculated bias field is then therefore wrong as well. We recommend therefore to insert a preprocessing step before the bias correction that reduces the noise in the image. Best results were observed with the anisotropic filter using a relatively low kappa (smaller than the half of the distance of two neighbouring classes) and a low number of iterations. In the simulated tests the bias fields were always computed correctly when using this filter as preprocessing step.

The calculated bias field of the filtered image is then applied to the original image. This preserves all details of the image and does not introduce a saw-tooth function in bigger homogenous areas (see subsection 4 and figures 10 and 9).

9 Choice of the parameters of the Legendre polynoms

9.1 Multiplicative/Additive bias fields

Usually the observed bias field is a multiplicative field. Therefore normally the bias correction is invoked with 'log'-flag on to calculate an additive bias field to the log-transformed image. This the calculation of the additive bias field of a log-transformed image equals the calculation of the multiplicative bias field.

9.2 Class mean values and deviation

This subject is a very important one, choosing these values not appropriately leads often to the calculation of a wrong bias field.

There are some thumbrules to obtain good values for the class parameters:

- Don't use too much classes. Too much classes on one hand drastically slow down the calculation of the bias field on the other hand the calculated field is more likely to be wrong. The amount of classes should be around 2 or 3 classes. Choose those classes that are widespread (in different parts of the image) and occupy a lot of area in the image. Mask the other classes out as good as you can.

- There are two ways to determine the class mean values:

  1. Assign a set of points (as much as possible) to every class, and compute the mean value over each set (with PCI for example).
  2. Choose a part of the image, where the different classes come together on a small area and pick out the pixel values of the different classes in this area. Use those pixel values as mean values.

- A good approximative value for the class standard deviation is $\sigma = d/6$, where d means the difference of class mean value between two neighbouring classes in the histogram. The values of the deviation should be controlled by watching the energy function. There each energy pot of a class should occupy the same amount of area. Also should the value of the energy function exceeds 0.5 between two classes to distinct well the two corresponding energy pots (see subsection 9.2.1).

9.2.1 Visualization of the energy function to control the class values

The actual version of the bias correction uses the formula :

$$E = \prod_{i=1}^{n} valley(\frac{x - \mu_i}{\sigma_i}) \tag{1}$$

$$valley(x) = 1 - \frac{1}{1+x^2/3}$$

You can visualize this function for example with Mathematica, Maple, or Matlab. The code for the visualization in Mathematica is printed below (output shown in figure 23):

```
valley0[d_] := 1-1/(1+d^2/3);
my = {70,140};
sigma = {7,14};
maxval = 255;

energy[x_] := Product[valley0[(x-my[[i]])/sigma[[i]]],{i,1,Length[my]}];
energyplotA = Plot[energy[x],{x,1,maxval+1},PlotRange->{0,1},
                   PlotLabel->"Additive bias"];

sigmalog = Log[1+sigma / (my + 1.0)];
mylog = N[Log[1+my]]
energy[x_] := Product[valley0[(x-mylog[[i]])/sigmalog[[i]]],{i,1,Length[my]}];
energyplotM = Plot[energy[x],{x,1,Log[maxval+1]},PlotRange->{0,1},
                   PlotLabel->"Multiplicative bias"];
```

9.3 Subsampling

To calculate an approximate solution of the bias field in a short time, subsampling is useful instrument. But when using subsampling it is more likely to end in a local minimum (Caution!). Therefore the subsampling factor should not be chosen more than 4, a good choice is a factor 2. The calculation with a subsampling factor x is speeded up by a factor of x^2 in 2D and x^3 in 3D compared to the calculation with no subsampling factor.

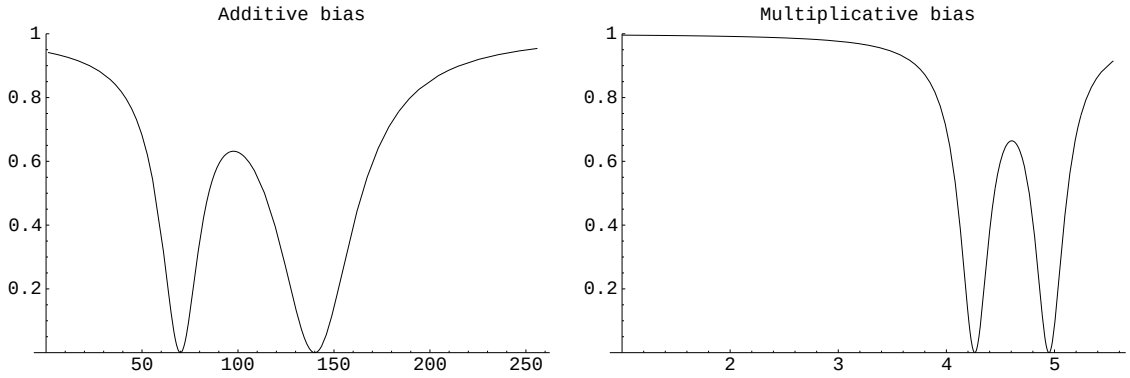

Figure 23: Energy function used by the bias correction. $\mu_1 = 70$, $\mu_2 = 140$, $\sigma_1 = 7$, $\sigma_1 = 14$. The values for the standard deviation are chosen for a multiplicative bias field. As it can be seen, they would not be appropriate for an additive bias field.

9.4 Masks

There are two different masks for different purposes in this bias correction. There is the input mask and the output mask. The input mask serves to mask out all voxels not belonging to the classes on which the bias correction runs. The output mask masks out the background when applying the calculated bias field to generate the corrected volume/image. This is necessary since there is no bias on the background, just noise.

Input masks are an effective instrument to improve on one hand the calculation time and on the other the quality of the calculated bias field. Well designed masks should input mask out all area not belonging to the defined classes, so the efficiency of a mask strongly depends on the choice of the classes. Input masks should cover as much parts of the image as possible, otherwise the calculated bias field may be incorrect in uncovered parts (if those are big enough).

9.5 Degree of polynoms

Choosing the degree of the Legendre polynoms too high results in high computation time and in inconstant coefficients for the bias field, although each of the calculated bias fields represent a good correction. If not chosen more than one degree too high, this should be no critical parameter.

Try first computing an approximate bias field (by subsampling) with a degree of relatively high order, usually this is either third degree or forth degree. Seldom a higher degree than forth is needed. Then look at the computed bias field and compare it with the table shown in figure 24. If the impression rises, that a bias field of a lower degree (or a combination of the bias fields of lower degree) look the same as the calculated bias field, then the bias correction should be run again with a lower degree.

10 How to choose the parameters for the 1+1ES-optimization

10.1 c_{grow} and c_{shrink}

Several tests revealed that the best values for c_{shrink} is chosen around $c_{grow}^{1/4}$. Since c_{shrink} has this by default, when invoking the bias correction, there is normally no need to type it in as a command line argument.

The tests done in this report with changing c_{grow} showed, that it should not be chosen to high. A value between 1.01 and 1.1 is a good choice. The higher the degree of the Legendre

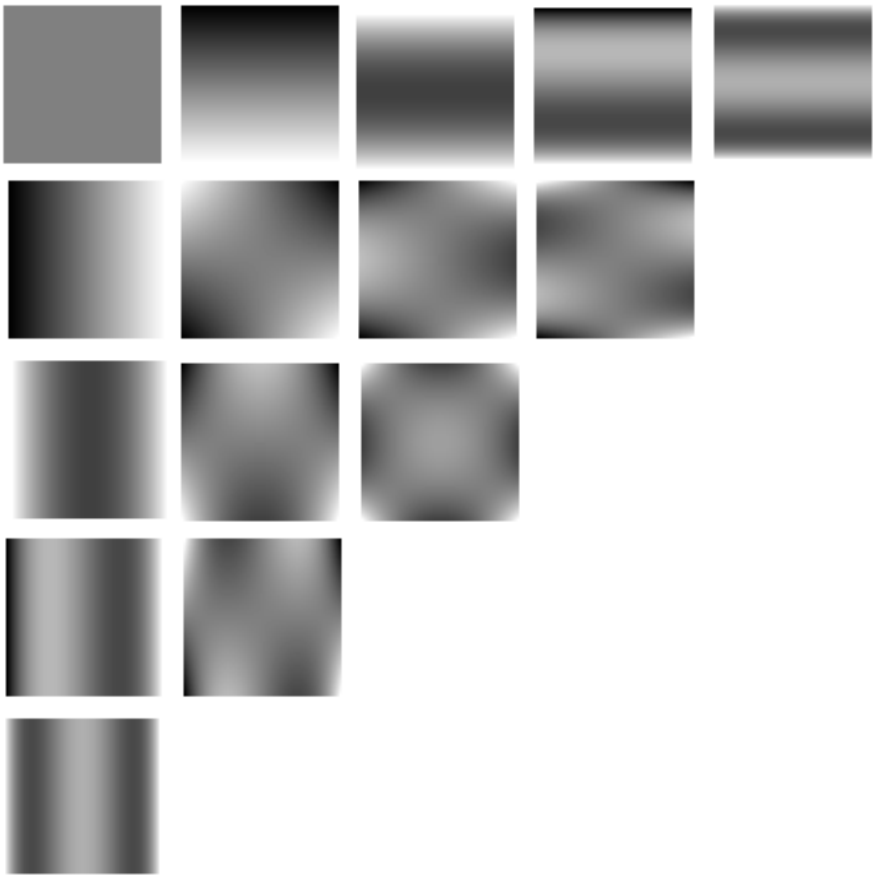

Figure 24: Table of Legendre polynoms for two dimensions (0.th to 4.th degree): polynoms are ordered the same way as coefficients of bias correction (from left to right increasing y-degree, from top to bottom increasing x-degree).

polynoms the more coefficients exists and therefore the higher dimensional is our search space. This leads to the obvious conclusion, which was also observed in our tests, that the higher the degree of our Legendre polynoms the lower c_{grow} should be.

10.2 Initial step size

A good choice for the initial step size is the maximal distance of two neighbouring classes in the histogram. When choosing a smaller step size, the calculation time is speed up, but is more likely to end in a local minimum.

10.3 Maximal number of iterations

The maximal number of iterations should not be chosen too high, since the calculated bias field is always an approximation and there is no need to have an approximative bias field calculated to a very high precision. When using subsampling a value between 2000 and 5000 should be enough. If the parameter space is high dimensional, that means when Legendre polynoms of high order or when a three dimensional bias field should be computed, then this value is too low. There you should choose a value between 10000 and 20000. Normally there is no need to choose the maximal number of iterations higher than 20000. Most tests in this paper used 10000 maximal iterations for two dimensional and 20000 maximal iterations for three dimensional fields.

11 Computation time after speeding up the bias correction

The bias correction application was speeded up by precalculating energies in tables instead of calculating them each time again. Also some slight corrections were made that increased the efficiency of the application. The amount of gain of speed is around a factor of two.

For example the calcuation of the bias field for the 3D-images of the head of a MS-patient was reduced from around 1:45 hours to around 45 minutes when using Legendre Polynoms up to the second degree.