### 4.5.3 Improving accuracy by means of 2–D to 3–D self-registration

The root mean squared error of the calibration, as reported in section 4.4.5, was roughly 3 mm, about 30 pixels. With this numbers, we expected that little could be seen in the composed image. In reality, the positioning error had perhaps been overestimated, and the composed image was fairly regular and coherent.

Still there was an obvious chance to improve the results: to register each video 2–D image to the compounded 3–D volume. This should be possible also because the Minibird already provides us with an initial estimation of the position of the image in the cuberille. Otherwise, the field of view would probably be too small compared to the size of the volume image to permit any registration. Now, given the initial position, only a small transformation, correcting the positioning errors of tracker, should be expected.

There are several similarity measures in literature for 2-D–3-D registration. Penney in [73] compares the performance for six different measures (cross correlation, gradient-based and mutual information-based) regarding robustness and accuracy to register a CT scan of a spine phantom to a fluoroscopy image of the phantom. His results were however not applicable to our case, because the modalities and features involved are completely different.

Similarly to any registration algorithm (page 7), there were several issues to decide:

- **the model of the transformation to optimise**

  We modelled the position error as a rigid transformation given the matrix:

  $$M^{ERR} = Tras(T_x, T_y, T_z) * Rot(\phi_x, \phi_y, \phi_z) \tag{4.11}$$

  And we modified equation (4.2) to include this adjustment:

  $$C_{\overline{\mathbf{x}}} = P_{\overline{\mathbf{x}}} \cdot M^{ERR} \cdot M_U^R \cdot M_R^T \cdot M_T^C \tag{4.12}$$

  For nomenclature purposes, we will refer as $\mathcal{B}$ the B-scan of ecography image output, and $\mathcal{C}$ the 2D image in the cuberille with the coordinates given by the previous equation.

  The new matrix $M^{ERR}$ measures the error in the position of the slice. We decided to include it in the product before $M_U^R$ because then the units of the transformation would be related to $\mathcal{B}$ , i.e., pixels, and not $mm$ as it would be the case had we included it after $M_T^C$ .

- **the features to be compared**

  The natural choice was to adapt the already tested algorithms, which worked for 3D–3D and 2D–2D, to work also with 2D–3D. This is, nevertheless, a very suitable choice because again the images to be compared may be very different. To illustrate this point, see figure 4.27: the gaps in the picture on the right are caused by multiple overlappings. An optimisation algorithm using the direct correlation or the mutual information of the two images is bound to fail for this reason. The creaseness image suppresses the effect of the gaps and makes them comparable.

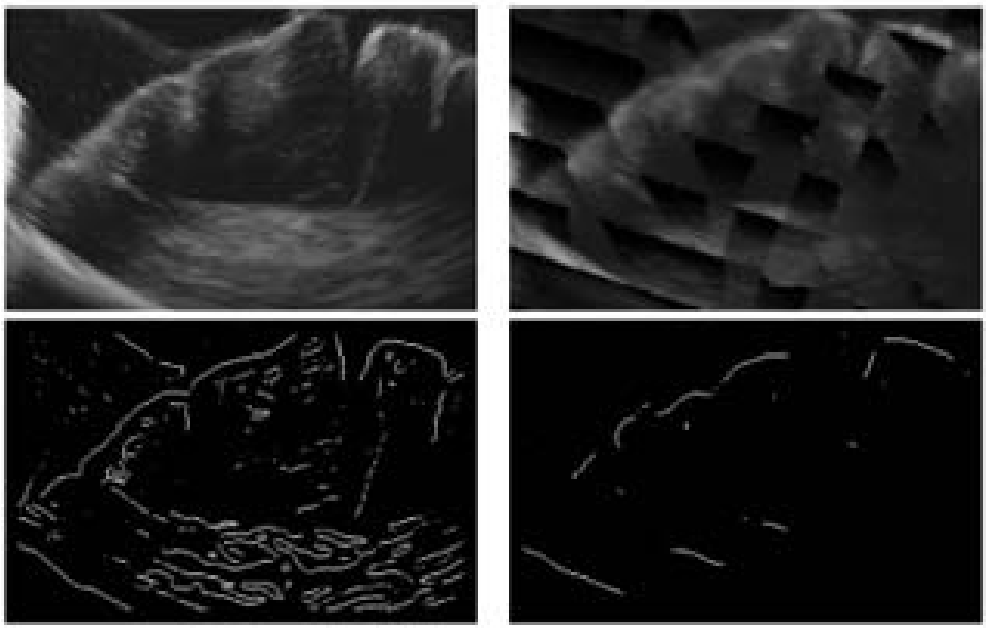

**Figure 4.27:** Top: The original image from the ecographer (left) may be very different from its corresponding slice in the volume image (right) due to multiple overlappings. Bottom: the creaseness image, however, depicts comparable features.

- **the scheme of iterative process**

  The next step was to modify our algorithms to run with a 2D–3D scheme. The iterative step could be very similar, the only additional step being the computation of $\mathcal{C}$ , i.e. the slice to be compared in the cubic image. The iteration would modify the values of $M^{ERR}$ , which in turn would change the contents of $\mathcal{C}$ , until the desired convergence has been achieved.

  But the initial step (exhaustive search in the Fourier domain, see section 2.4.2) had to be redesigned, as the dimensionality of the images was different. This initial step could not be suppressed because otherwise the iteration could get trapped in some local maximum. We took the approach to run the 2D–2D registration with the two initial images, the video frame and the corresponding cutout, and then use the result as the first estimation for the 2D–3D algorithm. This approach is schemed in figure 4.29.

We run successfully this algorithm for all the series we acquired. The main problem was the time it took for each frame: 190 seconds. While this algorithm was meant to be only an initial prototype, for the large number of frames to be registered would make experiments very lengthy. The main time consuming task was the extraction of the creaseness image for each iteration, and this step could be suppressed by pre-computing the creaseness of the whole volume image prior to the registration.

The creases extracted in this fashion had the inconvenient to be too wide (about 10 pixels) when seen at the pixel size of the video frame; ideally, they ought to be extracted after the cubic image had been reformatted to the B–scan small pixel size. However, when we tried this idea, the computer ran out of memory as the images were very large (about 150 Mb).

In order to solve this problem, we programmed a new type of image, based upon the template library of the C++ language (as implemented in gcc 2.96, Redhat 7.0), which would require only a bit per pixel. After all, this is all we need to store: whether a particular pixel is a crease or not. Also, we modified the creaseness extraction algorithm to run on small sub-images of the large initial image, and to store the result in one bit per pixel format.

Another improvement which made this idea possible was to compute the target coordinates $C_{\overline{\mathbf{x}}}$ only for non-void pixels $P_{\overline{\mathbf{x}}}$ in the video image, which were usually only 5% of the total. The result was rewarding: the new algorithm ran within 10 secons, about 20 times faster. See in figure 4.30 an scheme of the algorithm without and with these improvements.

The set of transformation values given by the registration algorithm follow some sort of pattern, this is, the direction is maintained along some frames, until another direction appears. See the 2D graph in figure 4.28. The same figure (bottom) also shows the transformation values as vectors, where each starts at the coordinate of the middle point of the B-frame in the cuberille, and ends at the same coordinate after applying the correction. Therefore, the picture gives a graphical idea of the length and directions of misregistrations.

Figure 4.33 shows the successful convergence for a few frames. Sometimes, the algorithm fails because the compared creases are too dissimilar. Other times, as seen in 4.33, last row, large artifacts appearing in $\mathcal{C}$ mislead the search. Actually, these are the proper results of the creaseness step, only that now the slice is extracted containing the whole surface, instead of a single line as previously. Since the optimisation searches the highest correlation value, the search is lead to areas with higher creaseness content. This effect occurs when the initial transformation is poorly estimated because of the lack of reliable landmarks, as it is the case for $\mathcal{B}$ depicting border areas of the brain.

We have manually classified each individual registration as valid or non valid, in order to measure the influence of the error described in the last paragraph. The criteria to label a registration as valid was somewhat restrictive:

| a registration is valid when the registered creases are visibly closer than before |
|---|

therefore rejecting the registration when the match is only visually similar.

We have visually compared the compounded image without and with position correction. To compound the cuberille with the registered values, we have taken the choice to use the original, unregistered transformation, for those frames labelled as not properly registered. This approach permits an easier visualisation of the effects of the registration.

Figure 4.34 shows a selection of relevant regions of interest for experiment C. Regions shown compare favourably for the self-registered version: the contrast is clearer, there are less artifacts and contours are more continuous. This is the case

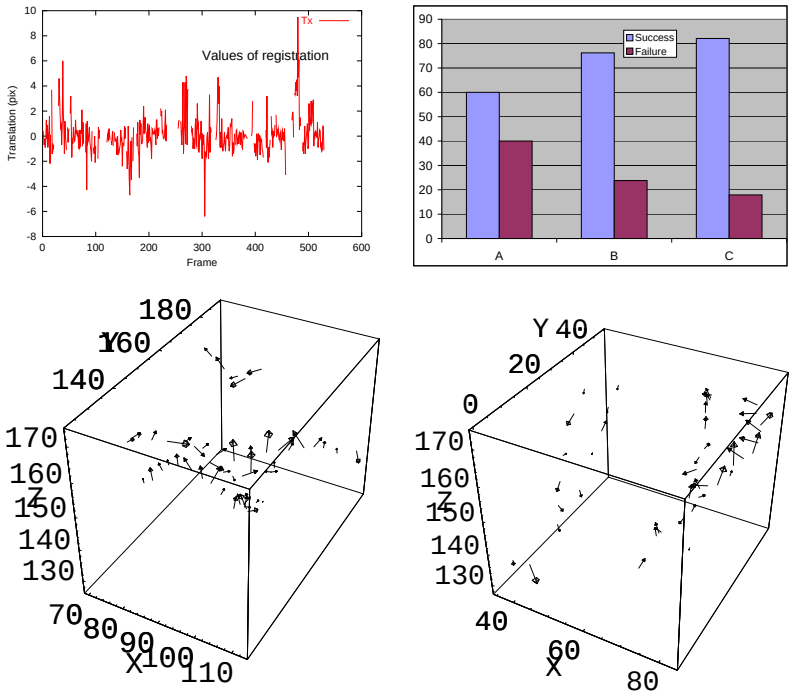


**Figure 4.28:** Top left: profile of the translations in $z$ for experiment C. Top right: percentage of successful and failure registrations for experiments A, B, C. Bottom: Corrections in the position for a number of frames in two areas of the volume. Arrows start at the position of the middle pixel of the frame and end where at the position given by the self-registration. Note that arrows in the same groups point to the same direction.

for most cases, but occasionally the contrary also happens, mainly for experiment A, where the small field of view restricts the chances of relevant landmarks to appear (recall samples in figure 4.5.2). In effect, if only a small round or flat surface can be seen, there is large number of possible matches, whereas the wider view provided by the other transducers is less likely to get trapped.

Note that the axial view of both images compares unfavourably to the self-registered version. This is caused by the misleading effect of the creases at the border, which was explained above. Despite our efforts to suppress wrong registrations, still the algorithm has been influenced at a number of cases classified as valid.

The last row of the figure 4.34 shows how this effect would influence the compounded image if we had not discarded wrong results. The view on the left is on an image compounded with all registered values, while that on the right is the same as the previous. Note the blurring and the gaps in the image on the left.

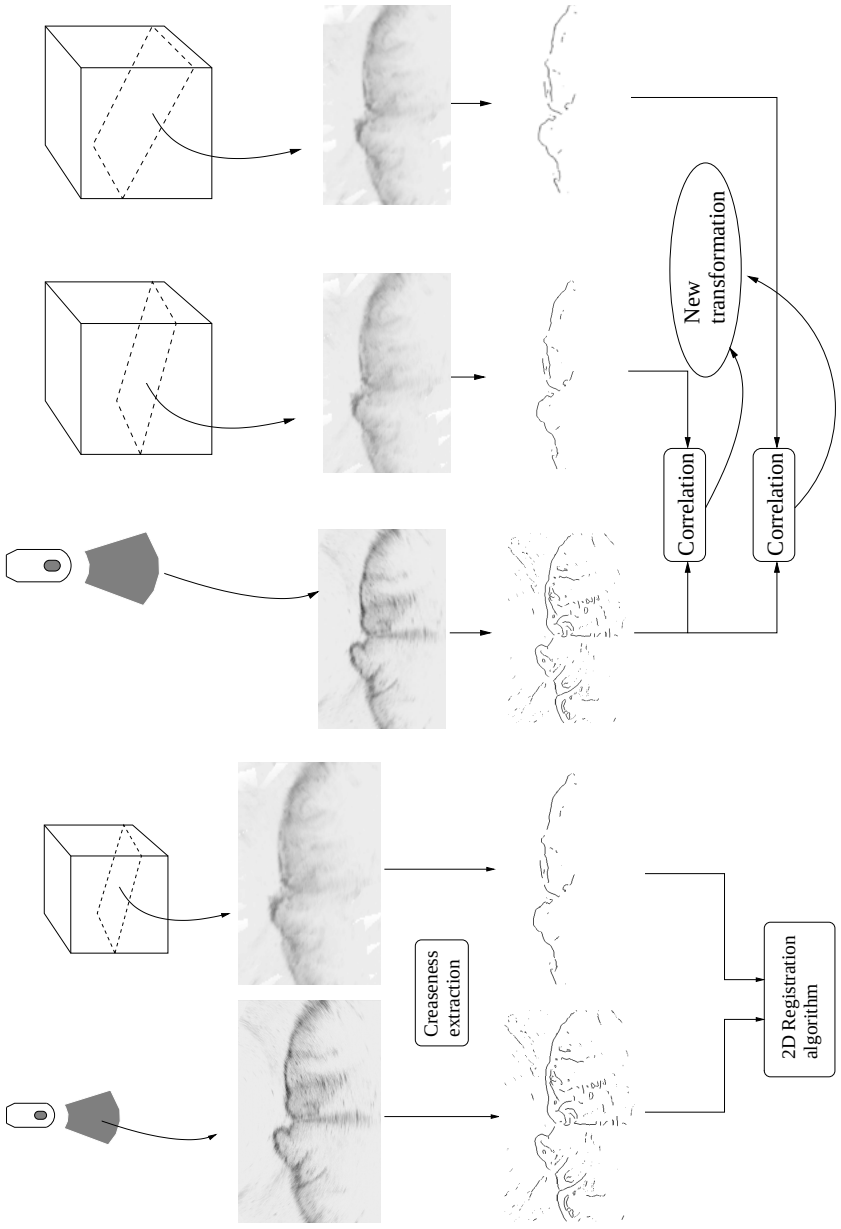


**Figure 4.29:** The algorithm for 2D–3D registration starts with a 2–D estimation (left). Then, the full 3D registration is iterated on the whole volume image (right). (images in figures have been inverted to enhance visualisation)

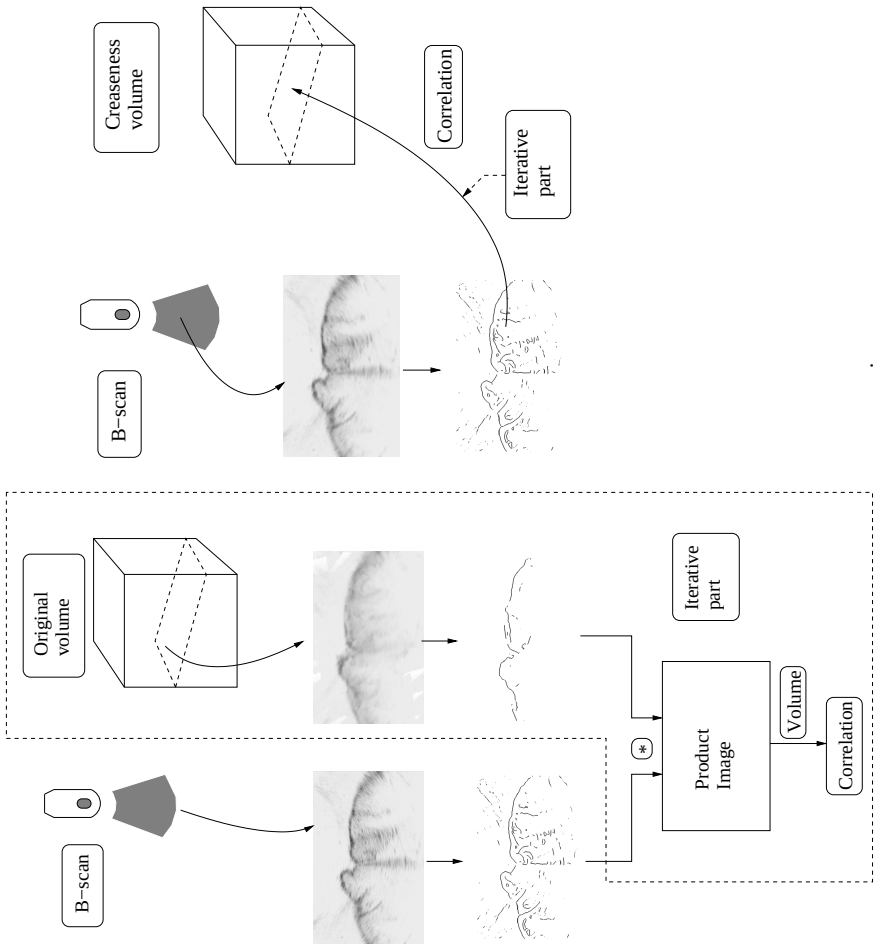


**Figure 4.30:** The iterative part of registration consumes a lot of time: locating the corresponding slice in the volume, computing the creaseness image, the product image and finally the volume (left image, dotted part). With the optimised version of the algorithm (right image), the iterative part is much shorter: computing the corresponding 3D position of the creases only.

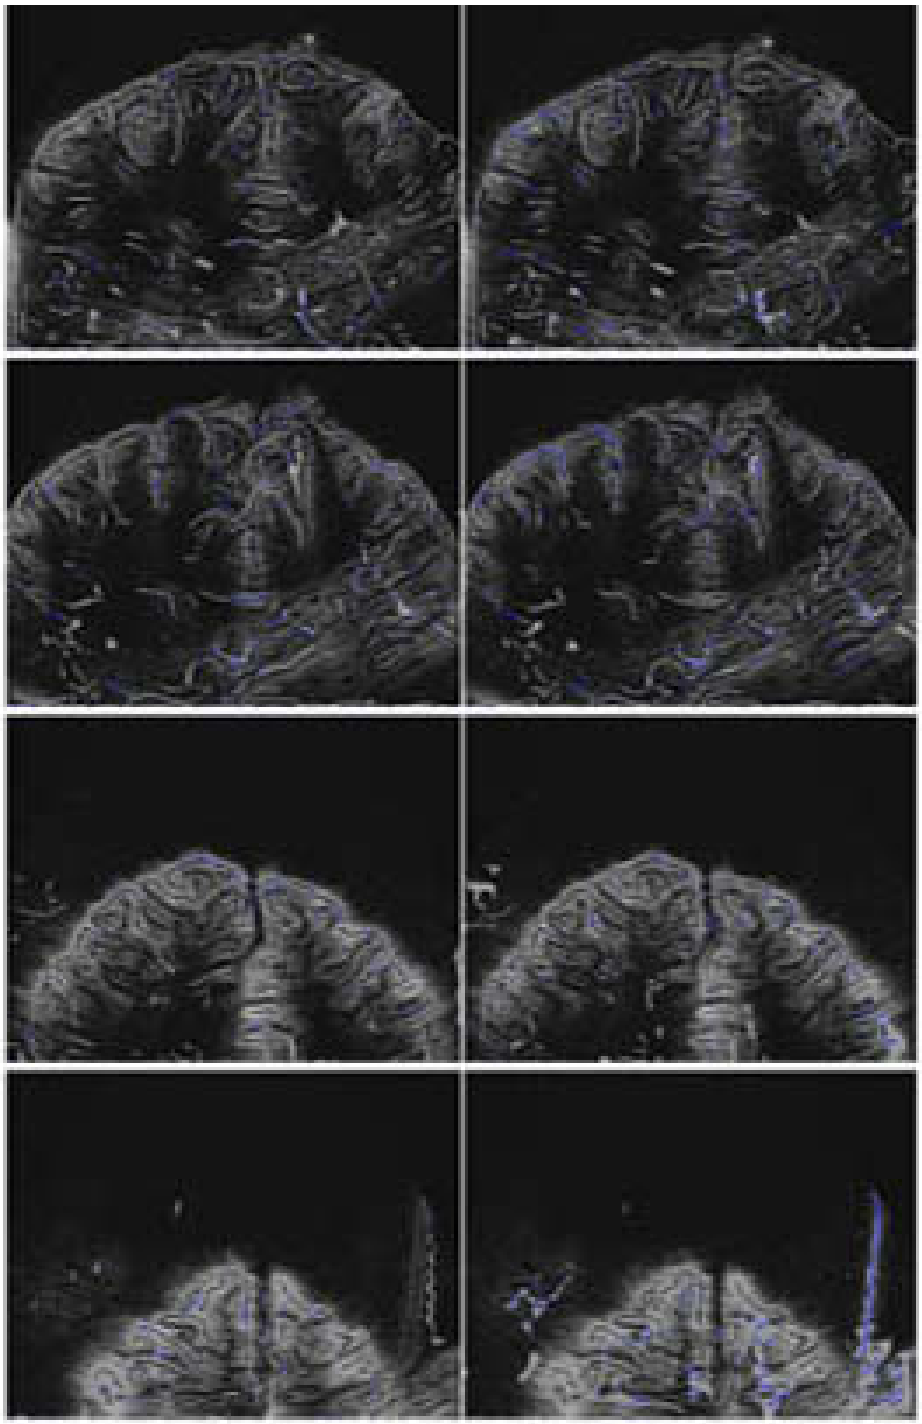

**Figure 4.31:** 2-D to ecography volume registration with the 3.5 MHz transducer (Experiment B). First two rows: successful registrations. Last row: wrong registrations, due to large crease surfaces. (see text for details)

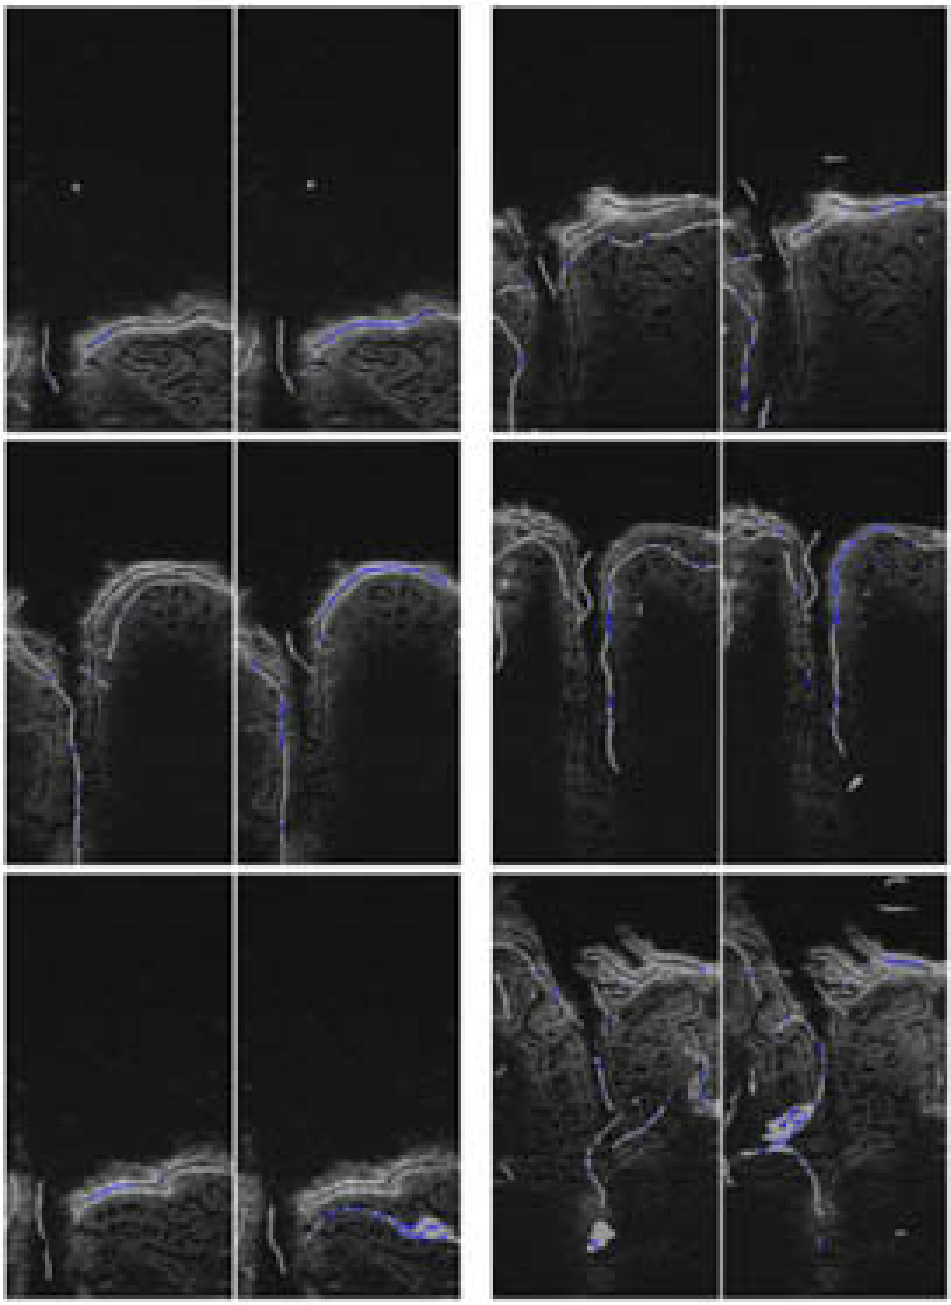

**Figure 4.32:** 2-D to ecography volume registration with the 10 MHz transducer (Experiment A). First two rows: successful registrations. Last row: wrong registrations, due to large crease surfaces. (see text for details)

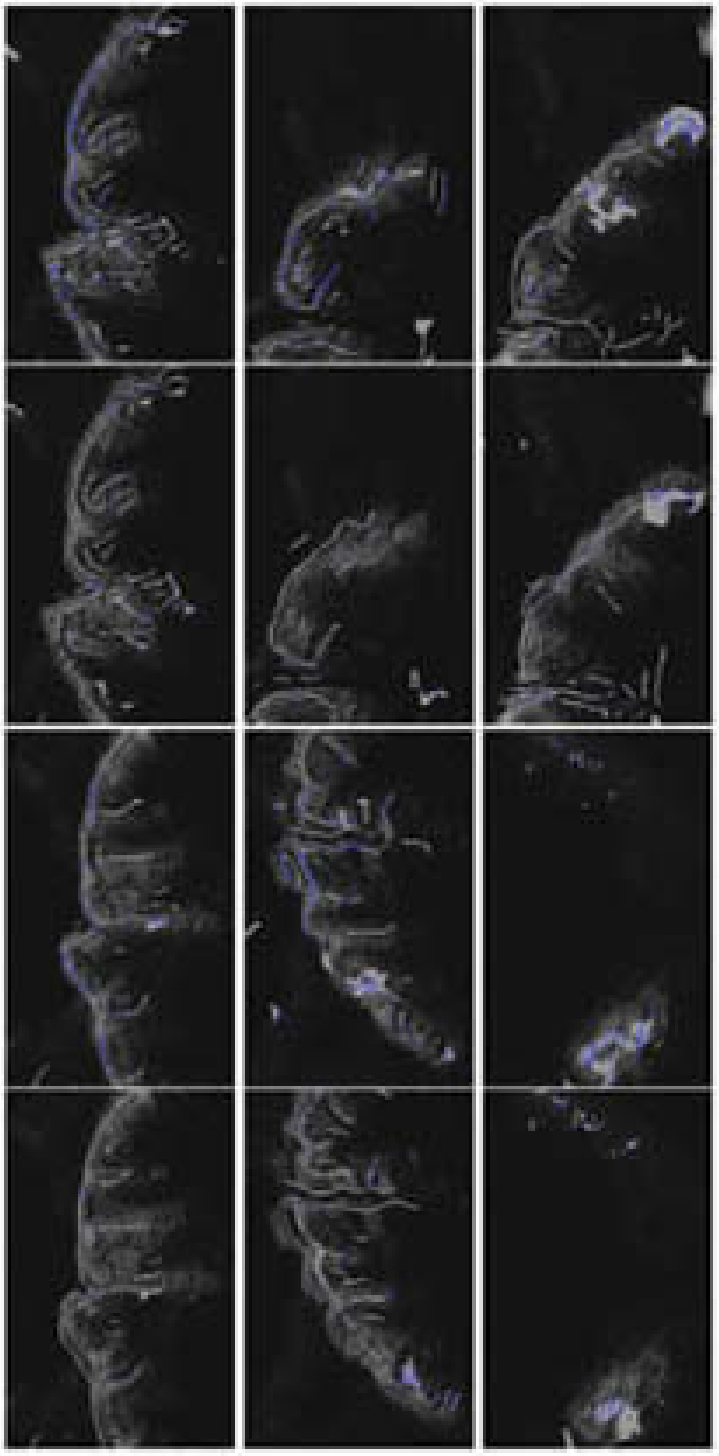

**Figure 4.33:** 2-D to ecography volume registration with the 6.5 MHz transducer (experiment C). First two rows: successful registrations. Last row: wrong registrations, due to large crease surfaces. (see text for details)

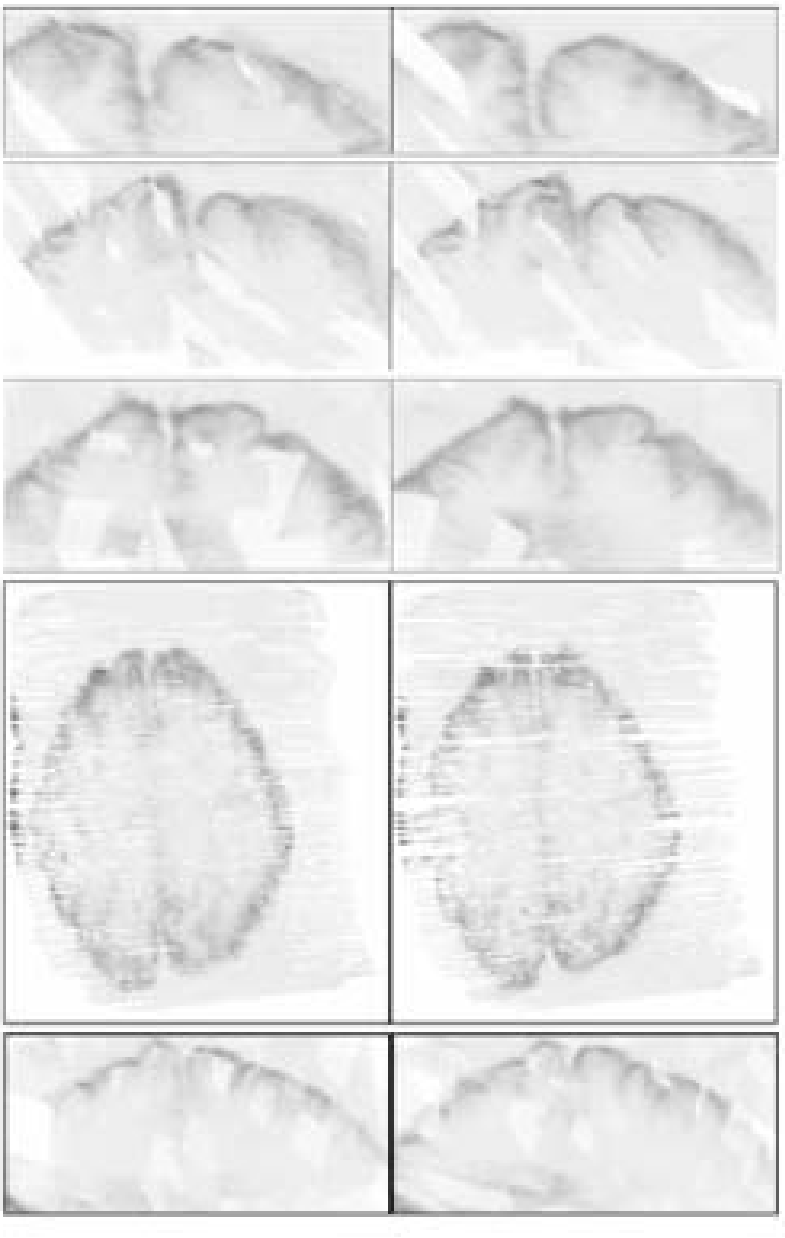

**Figure 4.34:** (Experiment C) Rois of the compounded volume with corrected position (right) compared to the original (left). The first three rows show the improvement for coronal slices: see the better contrast. Note, however, that gaps occur more frequently in the distal parts of the image as seen in the axial slice. This effect is also visible when comparing to the cuberille composed without classification (last row, left). (Graymap inverted)