Chapter 2

CT to MR volume registration

This chapter presents an algorithm for CT to MR volume image registration based on the alignment of landmarks, namely, the skull, extracted from both images by a differential operator. After an initial broad alignment, the algorithm builds a hierarchical structure for each image, and an exhaustive search is performed at the coarser level. Successful seeds are refined through the hierarchy until the highest value has been iterated. Our tests have included a comparison with another, well known method, as well as an evaluation, for a large database of images, reported independently by an external university.

2.1 Introduction

Perhaps the biggest attractive of feature-based methods is their similitude to the human process of image matching and recognition. When we are presented with two images and asked to align them, our procedure is, when possible, to search for common features. Conversely, a suitable approach for computer registration consists on comparing the alignment of equivalent features extracted from the images.

In a recent survey, Maintz [59] has classified an exhaustive list of publications according to 9 different criteria. Automatic intrinsic methods can be divided as volume-based or as surface-based. Volume-based methods statistically compare the similarity between the intensity values of the two images, while surface-based methods match surfaces previously extracted from the images.

West [114] evaluated the accuracy of 12 of these methods for 6 pairs of images (from CT and PET to 3 MR settings) from 7 patients (in sum, 42 pairs). As a group, for the data in the study, volume-based methods clearly showed better performance and have become the preferred approach.

We have developed a surface-based method based on an idea first presented by van den Elsen [102]. The general scheme is shown in figure 2.4. We take the skull —a structure visible both in CT and MR brain images— as the reference landmark. If we take as a landscape an axial section of a head image, the skull forms a ridge in the CT image, and a valley in the MR image. The same idea applies to 3–D, but is is more difficult to visualise.

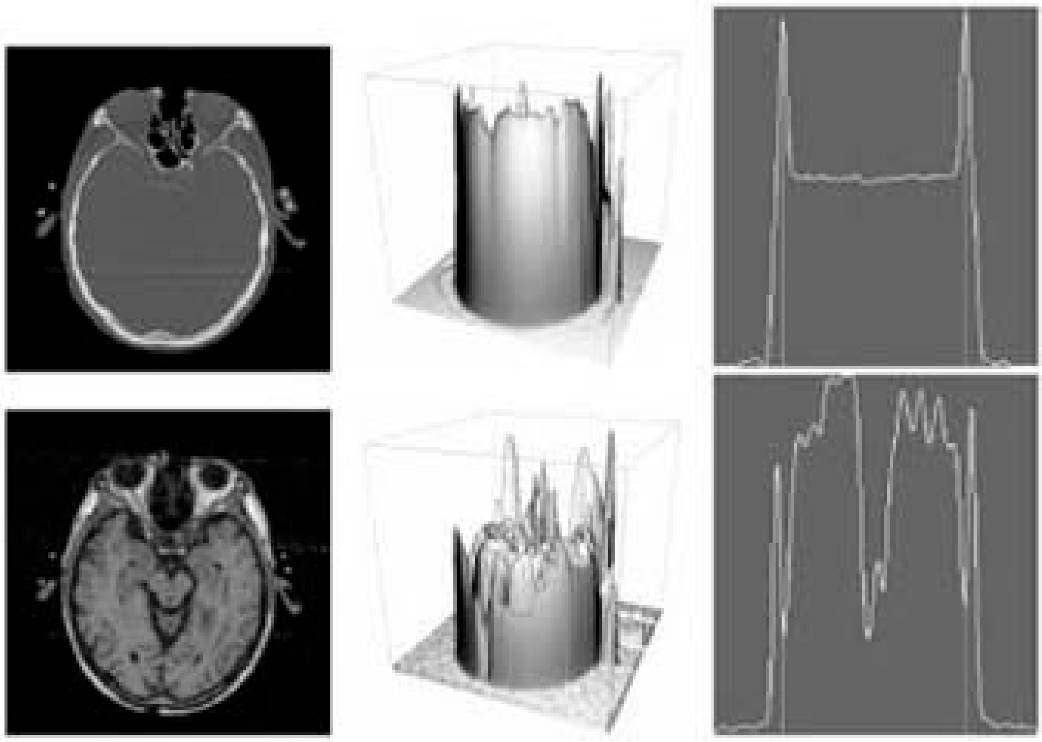

Figure 2.1: 2–D images (left) can be seen as a landscape (middle). In this case, they show that bone is a ridge in CT (top) and a valley in MR (bottom). The blue line on the left images marks the section shown on the right. Images have been previously registered to show the coincidence (red lines), and blurred with a Gaussian $\alpha = 2$ pixels to enhance visualisation.

Figure 2.1 shows an image of each modality depicted as a landscape. The bone clearly forms the described structures, which is made clearer by taking a section. Moreover, since the skull is an undeformable structure, it is ideal for matching purposes because only rigid transformations need to be considered. Only in some cases, where the skull has been drilled for operational purposes, the creaseness operator seems to be affected locally, while the rest remains unaltered.

A volume image can be processed in two fashions: as a stack of 2–D images, and therefore applying the 2–D operator to each individual slice, or as a full volume, thus employing also for each pixel the neighbouring pixels at other slices. This later fashion provides more regular response. See in figure 2.2 that the 3–D operator is more continuous, and also skips a number of false responses compared to the 2–D counterpart.

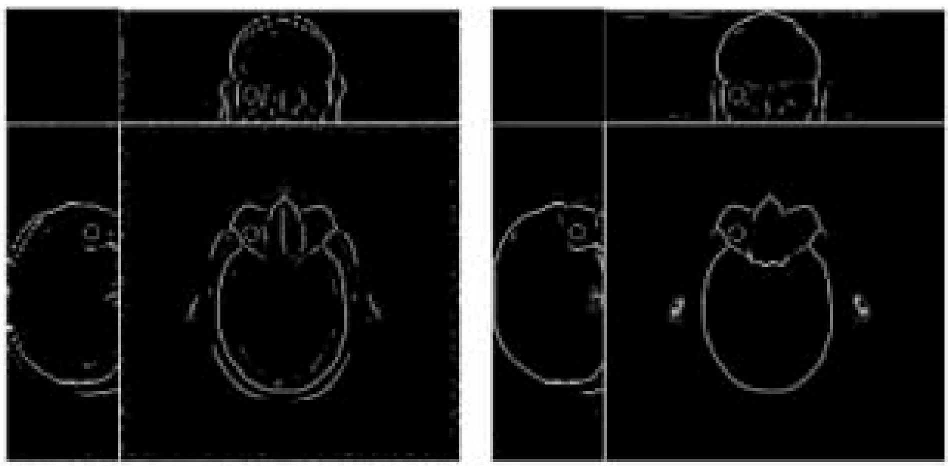

Figure 2.2: The segmentation of the bone is more reliable when employing the full volume (right) rather than as a stack of 2–D slices (left)

Figure 2.3: The thickness of the bone varies, making the segmentation more difficult. Left: slice of a CT image. Middle and right: profiles at selected zones.

The result of the creaseness procedure is a pair of images depicting for each point its creaseness. Because of the design of our creaseness operator, only the relevant features are usually selected. That means that CT and MR images have been transformed into two clouds of points, seen as well as surface patches, which must be aligned.

The alignment of two surfaces has received large attention from researchers because it is a common problem at many areas. However, this case must be considered specially difficult due to the nature of the images to be matched. The thickness of the bone is not constant throughout the skull, and therefore the scale of the features to be matched differs. As a result, some areas depict walls thicker than others, while other areas appear void (figure 2.3)

Another trying problem is the different field of view of each image. For medical purposes, often one image scans the full content of the head, while the other cover only central slices. As a consequence, the overlapping area is small compared to the full content and the alignment function will be easily mis-leaded into a local maximum.

Still more troublesome is the question of whether the segmented surfaces do actually represent the same structure. For some MR modalities, e.g. MR–PD, the segmented surfaces match consistently over all the image. But for others, and specially for MR T1, the skull at higher slices does not show valleys as clearly as in the example. The matching needs then to rely on other areas, and expect the mismatching not to influence the final result.

The correlation of two images as a measure of alignment is well suited for this purpose: it discards non common features and it is very fast to compute. More properties of the correlations are examined in section 2.3.

Van den Elsen and Maintz evaluated the performance of several creaseness operators [56]: $L_{\mathbf{vv}}$, $L_{\mathbf{pp}}$ and $L_{\mathbf{qq}}$, revisited in section 2.2.1. They studied their output for a set of images, including some generated synthetically. They concluded that some of them could not be suitable, but only for a number of theoretical cases not found very often in real images and they could be used for registration purposes.

Nevertheless, these operators were not included in the accuracy study at Vanderbilt's. They used only edge operators, studied in [58], to extract the edges of the head.

According to our experiment, the skin may not be such a good choice as the skull for registration. It is not undeformable, may not measure the same width for all modalities and its signal is less clear compared to the bone. For these reasons, we decided further investigate the use of the creaseness operator for registration.

Yet another reason to continue our research was that creases are not seldom employed in the image registration literature, in spite of their being a recurrent visual feature: veins and arteries in medical imaging and roads and rivers in satellite images. Since the whole registration process relies on the accuracy of the segmentation, we devoted large energies to develop the creaseness operators. Section 2.2 compares a number of standard definitions of creaseness to our own.

Compared to van den Elsen's, our method presents the following novelties:

- our 3–D creaseness (ridge and valleyness) operator is able to extract neatly the skull from both images, with almost no spurious response (section 2.2).

- an initial alignment is computed in the z–axis by means of a comparison of principal axis of the two stacks of images. The comparison, robust and very fast, is suitable for the initialisation purposes.

- we have added surface matching using chamfer distances as in [27], with the purpose of improving robustness.

- we have investigated experimentally the performance of our method for a number choices of parameters (section 2.5.2) and several MR modalities for a set of patient studies. Consistently, our evaluation of error is more complete and quantitative.

In next section we briefly describe two classical creaseness operators, $L_{\mathbf{vv}}$ and the level-curve curvature κ, and introduce new operators with better properties. We describe our hierarchical optimisation method in section 2.4. Afterwards, we shortly reviewed the mutual information paradigm in previous chapter, section 1.3.5, we compare results from a mutual information method in section 2.5.3. Results referring to the submission to the Vanderbilt database are discussed in section 2.5.4, and we present the final conclusions in 2.6.

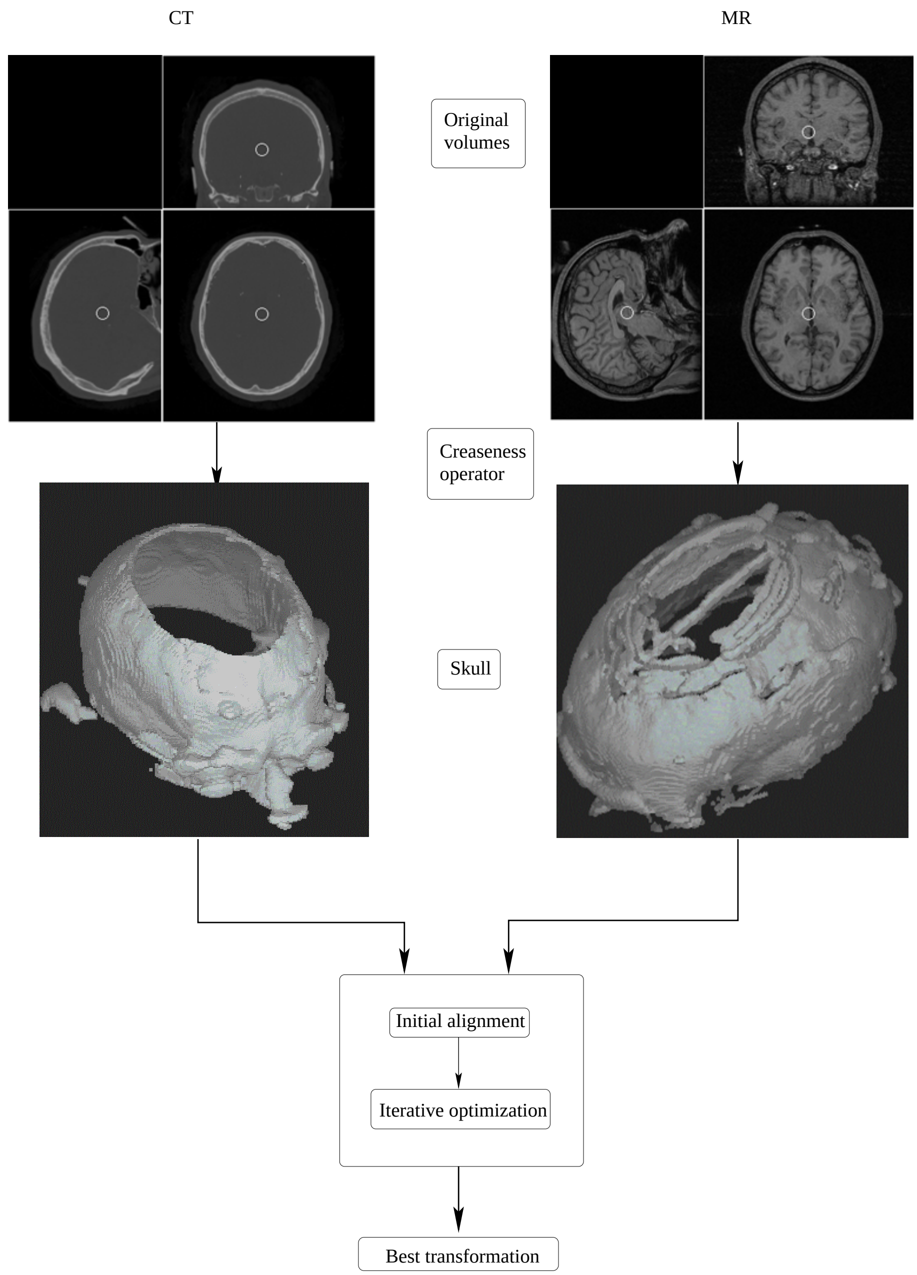

Figure 2.4: Scheme of our registration algorithm: we extract the skull of the original volumes by means of a creaseness operator. The resulting creaseness images, presented here as surfaces for the sake of visualisation, are after-wards aligned in two steps: rough alignment and hierarchical optimisation.