Chapter 1

Introduction

Today many advances in all branches of science are result of contributions from computer scientists. Computers are essential because they permit a reliable storage and analysis of data, which make them particularly valuable in the field of medical imaging.
Most modern imaging devices embed a computer, which manages the data flow between the scanning module and the image screened by the clinician. Compared to the previous analogical processing, the digital processing is more flexible and reliable, and that means that a full new set of analysis tools are now applicable.
This thesis is within the field of medical imaging, and in particular it deals with the computer techniques to compare, build, combine and align images. We start with a brief example of how images are routinely used, and a specific type of improvement computers can make. Next, we review the chronological development of the thesis, and define a number of relevant terms. After review existing literature, we present the objectives pursued in this work.

1.1 Motivation

Although modern imaging devices provide increasing resolution, often the information is not employed to its full potential. The following example will explain this statement: a common procedure for patients suffering from Parkinson's disease is to be regularly examined in an MR device. The physician may check the image taken, and decide whether the patient is still in good condition, or request another examination, perhaps with another modality, to collect more information. Figure 1.1 presents two series of CT and MR slices of the same patient in a layout similar to the slices actually printed for examination in a photographic film.

After a number of examinations, the physician may take the decision to operate the patient. Prior to the operation, he or she requests another image scanning, similar to that of figure 1.2, and takes some measures used later, in the operation theatre. The whole operation procedure relies on the correctness of these measures: for the sake of achieving a minimal invasion, the tools are guided without direct vision, with

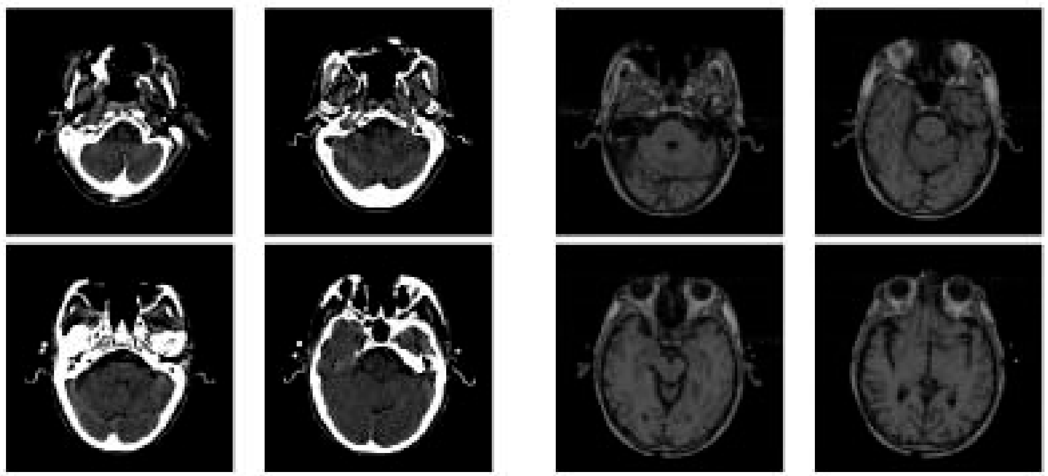

Figure 1.1: Physicians typically work with images printed in films in a layout similar to that of the figure. Selected slices are printed in separate sheets, a CT sequence (left) and an MR sequence (right).

the aid of some external coordinate frame. The purpose of the last pre-operative image is indeed to calibrate the coordinate frame, and therefore this is already a great achievement from medical images: before that, these minimally invasive operations were not available.

However, in the final operation stage as well as in the examinations, very little of the actual content of the images has been used. Although the images are created in a digital format in a computer, the current procedure is to print them in films, to be viewed with the help of a screening box. The films are kept together in the clinical history, to let the physician compare them at any time without the need of a computer. He or she mentally fuses the information from different sources, and states the evolution of the illness in the patient.

Although useful, given the technical conditions at many hospitals, this employment of the images discards relevant information. Sometimes, a point of view other than the orthogonal standard shows lesions hidden previously. Or maybe some of the slices have been wrongly considered uninteresting at the moment of the print, and afterwards are needed and not available any more. Even more interestingly, if images from different examinations could be aligned, that would make easier the assessment of changes in the tissues, perhaps with an automatic measure.

The first two issues can be labelled as a visualisation problem, which constitutes a separate research field in image processing. In our site we did not develop complex techniques, but even the simplest are significantly useful. Figure 1.3 shows the same information as 1.1 with a different presentation. Consecutive slices are presented as volumes in the computer screen, and the user interactively requests any desired view.

Also, images from both modalities have been aligned by means of an automatic procedure, and are presented simultaneously: all views from one set of images are shown together with their corresponding for the other set. The key point is, these

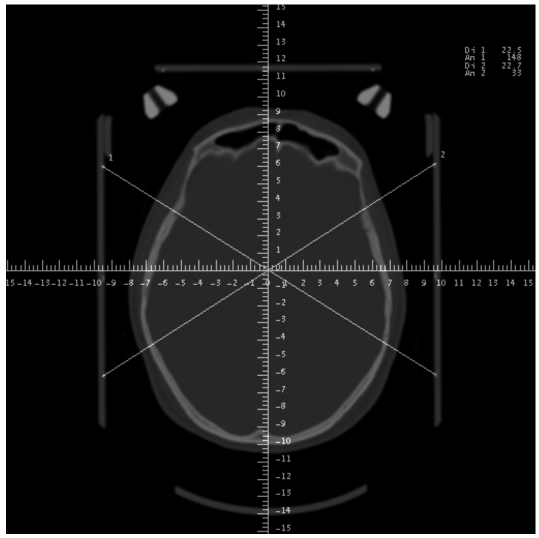

Figure 1.2: CT slice with some measures overdrawn. Results from this measures will be used later in the operating theatre.

capabilities imply far more than another presentation. The aid of the computer has represented such an advance that a new branch, called computer-assisted surgery (CAS) has arisen.

The issue presented above appears also in many image processing applications. It is called registration: whenever two images are to be compared, they need first to be geometrically aligned to make pixels correspond. This is not a trivial problem, for images may be very different in dimensions and content, and difficult to process because of their large volume. The thesis presented is mainly concerned to this subject.

1.2 A brief chronology

I started my studies in this field in 1995, when I was given the chance to join the *Unitat de Processament d'Imatges i Intel·ligència Artificial* (UPIIA) at *Universitat Autònoma de Barcelona* (UAB). Part of the deal was to work at the university as assistant lecturer. Although that has taken away a number of hours otherwise employed to research, I have enjoyed the contact with the real-life work of database systems, which was my teaching subject, and with my students.

Dr. Joan Serrat took the responsibility to direct my path in research. His first proposal was to extend the work done by my predecessor, Dr. Antonio López, in the field of creaseness measures: from 2–D images to 3–D volume images. Together with Felipe Lumbreras, I started working with medical images which, at that moment, were bigger than the computer memory space, and with algorithms which took hours

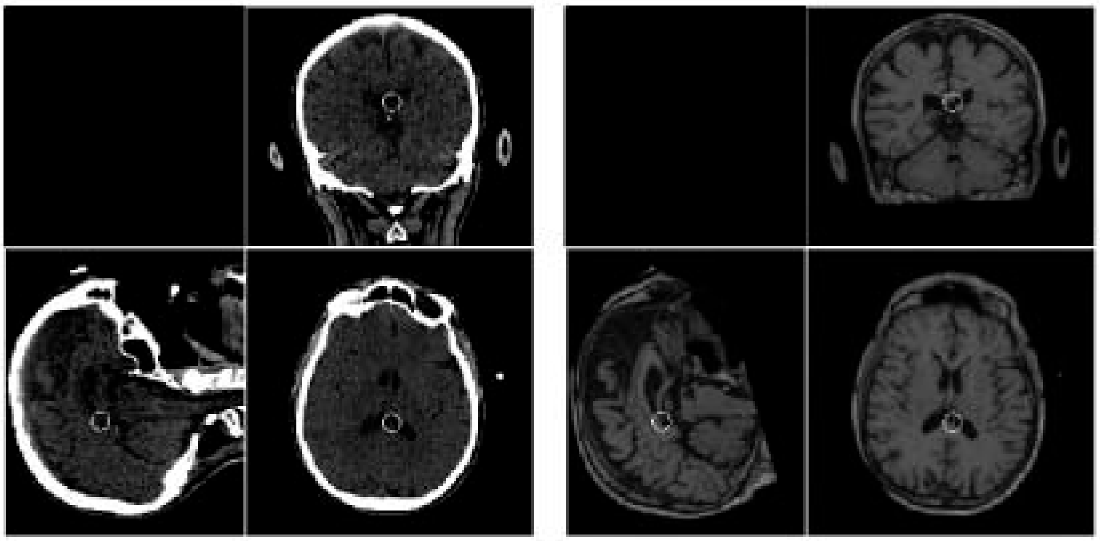

Figure 1.3: Stacks of slices may be seen as volumes. The circle marks corresponding points from the CT volume (left) to the MR volume (right). Also, this is the point where orthogonal views meet.

to complete.

The task was to implement a creaseness operator based on Rothe's definition, and since that required computation such as Gaussian products, derivatives and Fourier transforms, at the beginning we had to hire CPU time at a super-computation centre.

Nevertheless, the task was compelling because the major application was in the field of medical imaging. At that time I was for other reasons in contact with physicians, so the objective of using modern technologies to improve the quality of their work was highly motivating to me.

The first application of Rothe's creases, inspired by a paper by Dr. Petra van den Elsen, was to detect the skull in CT and MR images and then employ the results for the purpose of image registration. I finished the first algorithm seemingly capable of this task just in time to write the two-years compulsory report on the experimental work.

A major problem we had in our site was the lack of medical images. We had only two pairs, gently provided by Dr. van den Elsen, so our experiments lacked generality. Therefore, and also to start contact to the actual addressees of our work, we made intensive efforts to contact someone in a hospital interested in a collaboration. Dr. Joan Molet, at the Hospital de Sant Pau, showed from the beginning great interest in our results, and suggested many items which changed our vision of the problem. Part of this thesis, specially the chapter referring to ultrasound registration, and other papers not included here, are fruit of his collaboration and helpful assistance.

The initial registration algorithm was very slow (took hours, almost days to complete), and results were not very precise due to, amongst other factors, the limitation of the creaseness operator. Fortunately, Antonio took the charge to improved it, which constituted the core of his thesis, and let me concentrate on the alignment part. When that year was over, I could run it with reasonable confidence that things would work well, so I decided to submit the work to a conference, NMBIA98, to be hold in Glasgow, U.K. The reviews of the paper leaded to major improvements of the algorithm, and also originated the wish to widen my knowledge in medical imaging. The best way to achieve both was by means of a stage in some foreign research group.

We choose one of the groups leading the research at the field. The *CISG* group, in London, had just published several papers on the same subject, and had the advantage to be sited inside a hospital, the *Guy's Hospital Campus*. This meant I would experience day-to-day contact with a medical environment. Dr. Derek Hill, who was my tutor in the group, showed me a different, more experienced, point of view about medical imaging, and moreover the whole team made an ideal environment to exchange ideas and improve my methodology.

This improvement was tested with my first paper as first author to a journal, the *Journal of Electronic Imaging*, in a number devoted to selected papers from the NMBIA conference. At the same time, we participated in a program to validate the accuracy of the CT-MR registration methods, leaded by by Dr. Fitzpatrick at the Vanderbilt University. More than 100 pairs of images demanded emaciating tests and again some improvements to the original scheme. Chapter 2 gives a full report of the final method.

One of the aims of our team is that algorithms from our research ought to be of practical use in medicine, and in particular in Hospital de St. Pau. A routine operation in this site is the treatment of the Parkinson's disease, which employs stereotactic surgery to access some points in the brain. Part of the stereotactic surgery protocol is to take CT images of the patient wearing a metal box, and to measure several lengths in the resulting image in order to guide the tools during the operation. A small map with these lengths is then drawn to compute the coordinates in the stereotactic frame of the target point.

All these operations, carried out by an assistant in the short time between the image acquisition and the operation, can be easily modelled by a computer program. The first phase of this program constituted the final year project of a computer engineering student, Manuel Carballar. After its completion, we realised that the actual mathematical design had taken little of the effort, and that most of it had been put into reading the image format (DICOM) and the validation studies. We presented a paper on the subject [47], but for the sake of brevity the results are not included in this thesis.

The scheme we used for registering volume images is easily extendible to other modalities. The only requirement is that structures appearing in the images should be similar to creases. Veins and arteries in angiographies and, in general, ophthalmological images fulfil this requisite. Ophthalmology images are interesting from the image processing point of view because several modalities are available, some including time sequences, and vessels are not easy to segment because of the changes in intensity and contrast. Chapter 3 is devoted to the registration of ophthalmology images and the development of some quantitative measures to aid the physician.

During my stage in London I started studying how ecography images could be used in surgery. Ecographies are already used as a fast mean to check for the position

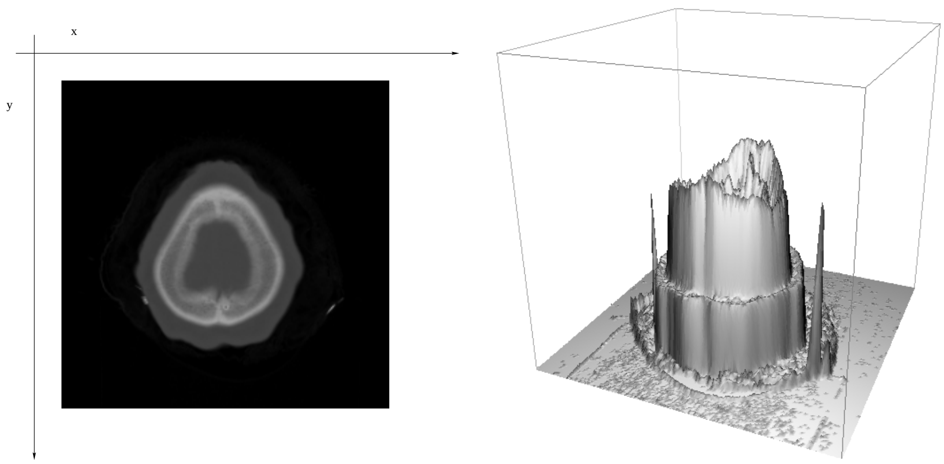

Figure 1.4: An image can be considered as a $\mathcal{N}^2$ function with real values, and depicted in a landscape form.

of underlying structures during the operating procedure, but our aim was to relate them automatically to pre-interventional images, with the aim of compensate the brain shifting. In addition to registration techniques similar to those employed for other images, we had to deal with other equally defying problems, such as tracking the probe or calibrating the coordinates system. The stage was over before I could complete the work, so we decided to continue it in our centre. Chapter 4 is devoted to this purpose.

1.3 A summary on registration

This section is an introduction to the body of image registration. We define some basic terms from image analysis in general and also present briefly the elements configuring a registration algorithm. The content of this section, although applied here to medical imaging, actually has a much broader scope.

When different images of the same object are taken, it is interesting to compare its contents. The images are different because of any combination of the following:

- the whole imaged object has moved. The patient inevitably moves for two different acquisitions, and even at the same acquisition series small movements are seen at the high resolution of the scanner.
- parts of the object have been displaced: the heard, the lungs and abdominal limbs move at different pace rates.
- the object itself has changed. A part of the image shows different content, e.g. when assessing the evolution of an illness.
- different devices have been employed. Patients are scanned under several modalities, depending on the requirements of the diagnosis.

Therefore, in general images will not be aligned, i.e., we will not know which pixel from an image corresponds to a given pixel of the other. This problem motivates the following definition:

to register two images is to find the transformation to align the contents of the images

A literature survey on registration algorithms included already in 1992 more than one hundred papers. They were classified in [25] according to four major components, explained graphically in figure 1.5:

feature space information from the images is used for matching

search space types of transformations we take into account.

search strategy given a non-satisfactory alignment, the method we use to choose the next transformation to try.

similarity metric measure of alignment for a transformation.

The similarity measure compares two images, called static and dynamic, for transformation parameters α:

$$\mathrm{M}(S, D; \alpha) := F(S, \mathcal{T}(\alpha, D)) \tag{1.1}$$

Usually only one of the two images, the dynamic, is iteratively transformed to achieve the highest alignment. The search strategy varies depending on many factors, but to summarise it is written as finding the value $\hat{M}$ such that:

$$\operatorname{Arg}\max_{\alpha} \; M(S, D; \alpha) = \hat{M} \tag{1.2}$$

Desirable properties of the measure M are[1]:

- **continuity**

$$\lim_{|\mathbf{h}| \to \mathbf{0}} M(\alpha + \mathbf{h}) = M(\alpha) \tag{1.3}$$

- **nil element** a non-empty image is most similar to itself with a nil transformation.

$$\operatorname{Arg}\max_{\alpha} \; M(S, S; \alpha) = \mathbf{0} \tag{1.4}$$

- **symmetry**

$$M(S, D) = M(D, S) \tag{1.5}$$

[1] For brevity only relevant arguments of M are explicitly written

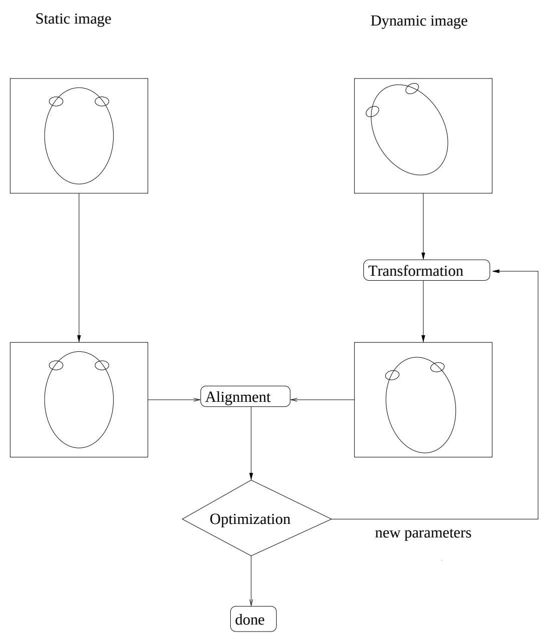

Figure 1.5: Scheme of a registration procedure. The dynamic image is transformed iteratively until it is properly aligned to the static image.

- **transitivity**

  given

$$\operatorname{Arg}\max_{\alpha} \; M(A,B) = \mathbf{T_{AB}}$$
$$\operatorname{Arg}\max_{\alpha} \; M(B,C) = \mathbf{T_{BC}}$$

  then

$$\operatorname{Arg}\max_{\alpha} \; M(A,C) = \mathbf{T_{AC}} \; = \; \mathbf{T_{AB}} \circ \mathbf{T_{BC}} \tag{1.6}$$

- **monotonicity**

  if

$$\operatorname{Arg}\max_{\alpha} \; M(S,D;\alpha) = \hat{M}$$

  then

$$M(S,D;\hat{M}) > M(S,D;\hat{M}+\bar{\epsilon}) > M(S,D;\hat{M}+2\bar{\epsilon}) \tag{1.7}$$

This is, M always decreases from the maximum value

- **invariant under transformations**

  if

  $$\operatorname{Arg}\max_{\alpha}\ M(S, D; \alpha) = \hat{M}$$

  then

  $$\operatorname{Arg}\max_{\alpha}\ M(\mathcal{T}(\beta, S), \mathcal{T}(\beta, D); \alpha) = \hat{M} \tag{1.8}$$

  This is, the aligning transformation is the same if both images are transformed.

A term related to registration is calibration:

to calibrate an instrument or an image is to find the transformation which relates it to an external coordinates system

For instance, one would need to calibrate an MR device when the pixel size happens to be different from the specified. Calibration is an important issue for compounding volume ultrasound images, and it will be developed in chapter 4.

This classification proved to be unable to hold the exponential growing of publications in this field. Another survey for medical image registration was published in 1993 [104], and further extended in 1998 [59]. In the latter, Maintz extends the number of criteria to nine, with the aim to clearly separate all the components and make the classification unique.

Our purpose here is not to give a full report of all variations, but rather to picture the relevant issues present in any registration algorithm. Following, we present of list of 7 items (numbered in bold), classified in groups to make it more readable.

- Image modality
  - **1** Monomodality/Multimodality
  - **2** Dimensions: 2–D/3–D + time
- Image contents
  - **3** Imaged area: head, abdomen, eyes.
  - **4** Patient: intrapatient, interpatient or patient to atlas.
- Algorithmic issues
  - **5** Transformation model: Rigid, affine, projective or curved; global or local.
  - **6** Comparison paradigm: intrinsic, extrinsic or manual.
  - **7** Optimisation method: hierarchical, heuristic; Simplex, Powel, Simulated Annealing