Supplementary Report of Research Progress - Revised

R. S. Schestowitz*
Research Student
Imaging Science and Biomedical Engineering
Stopford Building
Victoria University of Manchester
United Kingdom

17th July 2004

Overview

Interpretation of images is a rudimentary task which can be tackled by various different approaches. One such approach involves *modelling* objects which are expected to appear in images. Having modelled objects in some form or another, it is then possible to use this model to analyse new images describing similar objects.

Another problem which is associated with interpretation is image registration. Registration of images is an essential step which enables one to compare images at greater ease. It does so by *aligning* an arbitrary number of images so that homologous elements described in the image eventually overlap. This alignment of images is performed by transforming the space of images. This means that each image pixels is projected onto a new location and this results in images which are deformed. The ultimate aim is to deform all images until they all appear alike.

An existing method of solving the registration problem is to choose one common reference (or target) frame. There is a need to define a target according to which images are transformed, yet

*Contact: sch@danielsorogon.com
Electronic version: http://www.danielsorogon.com/Webmaster/Research/Progress/Forms/Supplements/form5sup.pdf

the flawed approach defines only one such target. The choice of this target affects the results, hence there is no unique solution to the problem.

The study we conduct addresses this issue by expressing the registration task in term of the entire set of images. We transform all the images at the same time and evaluate the transformation using a model representative of all the images. We define a principled cost function for this model and the images it encapsulates. Using an optimiser, we then strive to minimise that cost. Since we optimise over a single representation that accounts for all images, we obtain one unique solution.

This research work unifies several different strands. It takes a method which is commonly used to analyse images and indirectly uses it to guide image registration. What we also achieve as a result of the process is a description of transformations from which observed image variation can be derived. Essentially this produces one mean image and operations that deform it to assimilate to plausible, yet unseen images. Even more usefully, by registering images, correspondences can be identified automatically. This resolves the need for manual landmark which is otherwise necessary for the construction of statistical models. Such models can regenerate images similar to the ones which were registered; they can also be used to perform various measurements over objects in the images.

Experiments and Milestones

Environment for Experimental Studies

From the very start, a tool was needed to conduct experiments in a systematic manner. The problem we needed to solve was that of simple data registration so 1 dimensional data needed to be generated and then analysed in sensible ways, before, throughout and after registration.

We used MATLAB to construct a package that can be rapidly built and tested. On top of this package, a graphical user interface was laid and results were displayed in the form of hyper-text. As development of this package progressed, it was decided to name it Autonomous Appearance-based Registration Test-bed (AART). As the name implies, this is primarily a flexible environment in which registration tasks can be performed. Of particular interest is registration which is based on appearance of images. This appearance is described by the means of a model and the process of registration is intended to be free of user intervention, hence it is autonomous.

As it presently stands, AART is a stable tool that has a great number of run-time options. By setting these options, new experiments can be quickly conducted and results returned in visual form, as well as in textual form.

Initial Exploration of the Problem

One of our aims was to benchmark different registration methods and come up with comparative results which highlight the up- and down-sides of each method. We had a particular interest in the underlying behaviour of each method and the quality of registration as evaluated by a model of appearance. Results are shown in Figure 1.

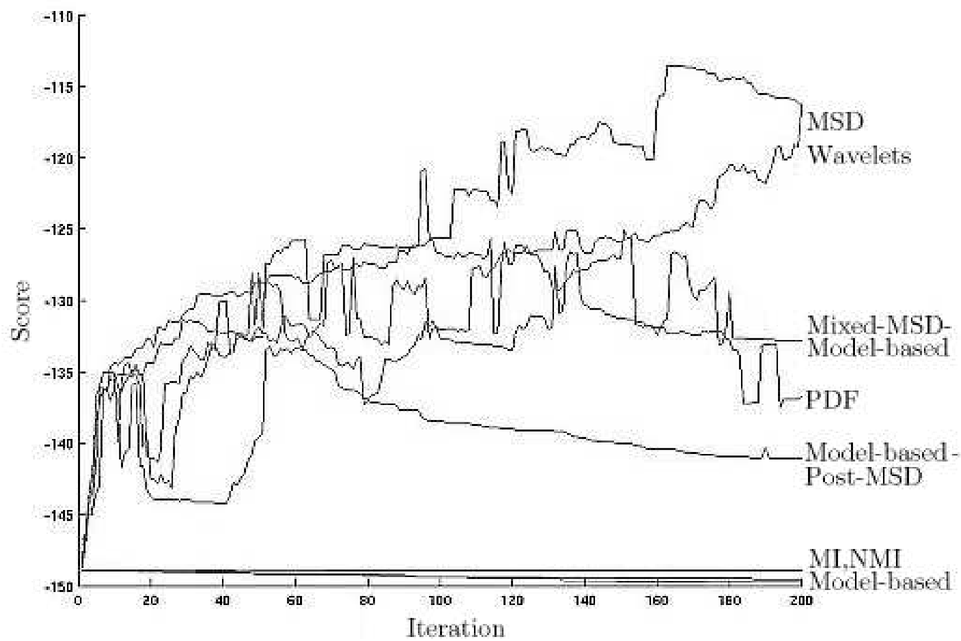

Figure 1: A comparative analysis of different objective functions. It illustrates that the model complexity decreases only for our newly-proposed objective functions. The Y-Axis value is an indicator of model compactness.

Months later we discovered that the registration method which we had proposed could become ever more successful. Up to a certain point in time, we were simply unable to get decent results. It was revealed that the transformations applied were restricted to remain small in extent. The problem was resolved by changing this restriction term, whereupon larger, more radical transformation were permissibly applied and a good solution was shortly approached. The issue of speed (or efficiency) remained a worrying factor. It had to be addressed in order to make a our registration method more practical in 2-D (and potentially an even greater number of dimensions).

Speeding up Convergence

The objective we are after is the minimisation of some cost – a cost is associated with the model and the images which we transform. Convergence of the algorithm is declared once that cost can no longer be reduced. We have found that very much computational effort is spent on refining existing transformations even when the solution is yet far away. As it turns out, at early stages of the optimisation, coarse changes to the images should suffice.

Amongst the more important developments of this project, we managed to find a way of substantially decreasing the amount of time which is required to register sets of 1-D data. This was done primarily by tweaking the optimisation which is involved in the process.

We decided to aim for a different optimiser tolerance depending on our advancement towards the correct solution. By doing do, little time should be spent on obtaining lower costs at the early stages of registration. This rational observation motivated us to implement a similar algorithm in a separate domain[1].

Dealing with Representation Errors

As we attempted to identify the point of convergence for our algorithm, we discovered a clear flaw with the cost we had defined. It turned out that the way in which we evaluate models neglects to account for small artifacts which must be encapsulated in this model. More crucially, these small artifacts which are left-out residuals need to form part of the model cost. In their absence, our objective function was able to drift away, thereby hiding vital structures in images. Not only was the result of registration poorer due to an improper model cost, but also it was impossible to contend that we ever reach one unique solution.

The Next Steps

The last unsolved problem, which was listed in the previous section, is being addressed at present. It has also led us to realisation of a fundamental problem with the way in which we describe models and, more importantly, our definition of model cost.

Having solved some of the problems that we had encountered, we plan to extend our method to deal with 2-D data and then

[1]This domain is the selection of landmarks in shapes for the construction of statistical models of shape.

preferably apply it to bio-medical images. This will make a comparison between our results and results produced by other methods possible. We will then be able to argue that our method indeed does what it was set to do and that the one unique solution we get is most preferable.