A Framework for Building Deformable Atlases

T.F.Cootes, C.J. Taylor, C.Twining, N.Thacker ...

March 14, 2002

Abstract

There is considerable overlap between recent work on non-rigid image registration and that on building and using statistical models of shape and appearance. Non-linear registration can be considered as a special case of matching a deformable model to a new image, in which the model is simply generated from a single image. This document describes a framework which is an attempt to unify both fields. In particular it describes the components required for a successful system, the properties they should have and how they would interact. It is hoped that such a description will act as a focus for research in the field, as currently the best way of achieving the desired properties of the components is not known.

1 Introduction

Recently there have been considerable advances in Non-Rigid Registration (NRR) and in deformable model matching; two fields are drawing together. Simplistically, NRR can be considered as a special case of matching a deformable model to a new image, in which the model is simply generated from a single image. Conversely, statistical models of shape and appearance require accurately corresponded training data. In early work the correspondences were defined manually but recently it has been shown that NRR can be used to obtain such matches. In addition, recent work on automatically building statistical models from boundaries and surfaces by treating the process as a large optimisation problem has shown considerable promise, and gives insights into many of the core problems in combining NRR and statistical model building.

Important recent results include the following

- Under the right circumstances, NRR can provide accurate correspondences between image pairs
- It is difficult to ensure that the deformation does not tear or fold
- Mutual information like measures are effective when corresponding images of differing modalities
- Corresponding *groups* of images can be treated as an optimisation process

- The choice of objective function in such a process is critical
- Explicit representation of the class of warps, chosen so as never to fold or tear, is possible in certain cases, and extremely useful.

An obvious question is: If NRR is becoming increasingly reliable, why would we need statistical models of appearance? Answers include

- Statistical models allow comparison of groups of images, not just pairs
- By applying constraints learnt from the models, the matching process can be made much more efficient and reliable.

A natural objective in bringing the fields together is to build statistically valid deformable models (or atlases) of anatomy which can be matched to new images rapidly and accurately.

2 Building A Deformable Atlas

We require an algorithm with the following behaviour

Input Two or more images (+ optional annotation)

Output Description of correspondences across set

The correspondences (or their parametric description) can be used to build a statistical model of shape and appearance.

The approach can be summarised as an attempt to find the group correspondences which give rise to a model which is in some sense optimal. In practise, correspondences should be defined implicitly, by determining the parameters of a warp field from the mean of a model to each individual image.

In order to construct such an algorithm, we require the following components.

2.1 Required Components

- A parameterised class of warps (deformations), $\mathbf{x}_{new} = W(\mathbf{x}, \phi)$.
- A representation of constraints (points, curves, surfaces etc)
- An objective function $F()$
- An optimisation strategy
- Methods of selecting model complexity to match the data
- Statistical modelling methods

The final atlas will contain (at the least), a statistical model of the allowed deformations, a statistical model of the "mean" shaped image (eg the texture in a mean co-ordinate frame) and a statistical model of the relationship between shape and texture (if appropriate). Annotations of the mean image can then be propogated to new images by matching the model, in the usual way.

2.2 Algorithm Outline

Assume we have N images I_i. Assume $F_2(I_1, I_2)$ measures the difference between two images, and $F_N(\{I_i\})$ is a measure of the compactness of a group of N images. [1]

- Select one image as the initial estimate of the mean, I_m
- Perform affine registration to initialise each ϕ_i to minimise $F_2(W(I_i, \phi_i), I_m)$
- Repeat
  - For $i = 1..N$, modify each ϕ_i in turn to minimise $F_N(\{W(I_i, \phi_i)\})$
- Until convergence

The aim is to find the deformations which, when applied to a group of images, lead to as similar a set of deformed images as possible (where 'similar' is measured using a suitable objective function).

2.3 Statistical Deformation Model

Having found such a set of deformations, with parameters $\{\phi_i\}$, we can build a statistical model of the deformation parameters,

$$\phi = f_\phi(\mathbf{b}) \tag{1}$$

Notice that we explicitly model the *parameters* of the deformation, and not the deformation grid itself. This ensures that any new deformation generated by the model, $W(I, f_\phi(\mathbf{b}))$, is itself a diffeomorphism (assuming we can generate a suitable class of warps). This is not guaranteed by existing shape models or deformation field models.

3 Properties of Transformations

There are two cases which we will consider

- Topology preserving transforms (diffeomorphisms)
- Transformations in which topology is preserved almost everywhere, but there are a finite number of local intrusions or removals (eg a tumour appearing in a brain, or bowel gas in the abdomen).

[1] Here we use our notation rather loosely, to give a flavour of the algorithm. Later we will be more precise.

3.1 Topology preserving transforms

We require representations of warps which will not 'tear' of fold space - structure must not appear or disappear. In practise this means that the warps must be *diffeomorphisms*. A diffeomorphism is a function $f : R^n \rightarrow R^n$ which is infinitely differentiable and for which an inverse exists.

In particular, we would like a parametric basis for such diffeomorphic warps, so that we can easily represent them and perform statistics on sets of diffeomorphisms. By representing each warp by its parameters, we can combine warps in ways which are guaranteed to produce diffeomorphisms. Thus, although $f_1(\mathbf{x}) = 0.5(W(\mathbf{x} : \phi_1) + W(\mathbf{x} : \phi 2))$ is not necessarily a diffeomorphism, $f_2(\mathbf{x}) = W(\mathbf{x} : 0.5(\phi_1 + \phi_2))$ is.

Ideally, we would like a parameterised set of functions $W(\mathbf{x} : \phi)$ with the following properties

1. $W(\mathbf{x} : \phi)$ is a diffeomorphism
2. $W(\mathbf{x} : 0) = \mathbf{x}$ (identity exists and is represented by zero vector)
3. $\forall \phi_1$ there exists ϕ_2 such that $W(\mathbf{x} : \phi_2) = W^{-1}(\mathbf{x} : \phi_1)$ (inverse exists in class)
4. If $W(\mathbf{x} : \phi)$ is a diffeomorphism for $\phi = \phi_1$ and $\phi = \phi_2$ then $W(\mathbf{x} : \alpha\phi_1 + (1 - \alpha)\phi_2)$ is a diffeomorphism when $0 < \alpha < 1$.
5. Small changes to ϕ lead to small changes in the diffeomorphism, and vice-versa.

Ideally we'd like to have a parametric function which uniquely generates all possible diffeomorphisms. This appears to be difficult, if not impossible, to achieve.

Desirable additional properties include

- The effects of individual parameters are 'localised' (useful for efficient image matching)
- The diffeomorphism can be decomposed into a set of coarse to fine scale components (useful for image matching)

3.2 Occasionally discontinuous transforms

We need to be able to allow for the intrusion or absence of structures when mapping from one image to another. For instance, one image may have bowel gas appearing in places where a second image does not (or vice versa).

It may be possible to represent these as a composition of a global diffeomorphism together with functions which locally create/destroy structure.

There's plenty of room for research here.

4 Properties of the Objective Function

We need a method of comparing the quality of match of one image warped to another. Existing metrics include sum of squared errors, dot products and various MDL and mutual-information type measures.

We would like the objective function to have the following properties:

- Works for groups as well as pairs
- Amenable to error propogation
- Invariant to appropriate things (scaling of images in size and intensity etc)
- Should be robust to outliers, or simple to make robust
- Efficient to evaluate

5 Optimisation Strategy

In the simplest case, given a parameterised family of warps and an objective function, we can use a general purpose optimiser (eg Downhill Simplex) to do the matching. Though generally reliable, this may be slow. Work on AAM's suggests that careful study of the problem can lead to optimisation strategies which are tuned to the application and can be very fast.

When matching one image to another, coarse to fine methods have been shown to be efficient and robust - we would be wise to build in such a strategy from the start.

A key question when matching a pair of images is how complex a deformation should be allowed? One can probably achieve a perfect match given a complex enough transformation, but this is likely to be overfitting the data somewhat. Techniques drawn from the model-selection literature will be needed to determine what level of model complexity is justified given the data (and its uncertainty).

An important issue is determining the error bounds on the deformation that is estimated - without such bounds one cannot make reliable decisions.

6 Evaluation

We'll need some of this.

In particular

- Sets of training images (for building models)
- Sets of test images (for testing performance)
- Ground truth annotation (including estimate of uncertainty in annotation)
- All of the above for various modalities, both 2D and 3D

7 Conclusion

This document proposes a general framework for combining non-linear registration algorithms with statistical model building techniques in order to automatically construct deformable anatomical models.

It currently is more in the form of a wish list, as we do not yet know how to achieve several of the key steps (eg representing deformations, forming an objective function, estimating the uncertainty of the result etc).

However, it is hoped that it will encourage work to be done on the key unsolved problems, eventually turning the dream into reality.