Registration Experiments

22nd November 2004

1 Notes on Optimisation Method

When it comes to optimisation, in the non-rigid case, only downhill search ever reaches an end. Simplex which is the default optimisation method and also Powell do not appear to ever stop.

I have performed some experiments and what follows is a summary and some conclusions.

2 Experiments

2.1 Studying the Behaviour of longer NRR

Experiment 20112004-4

Registration comprised:

- 3 iterations of translations
- 3 iterations of scaling
- 3 iterations of affine transformation
- 91 iterations of NRR using downhill search, 4 knot-points in each dimension

Results show that improvements are first made to the skull and the ventricles. At the later stages, some improvements are made to CSF and the discrepancy looks like an almost consistent brain. The duration was under an hour on my machine.

2.2 Studying the Effects of Pyramiding

Experiment 20112004-5

The rigid and affine transformations remained as before. At the stage of NRR, there was a gradual increase from 1 to 5 knot-points in each dimension; that's a total of 5 iteration.

Another 5 iteration were performed in the same manner but at a higher resolution (level 1, down (or up) from level 2).

The duration was about 15 minutes for the 20 or so iterations in total. The results were not as good as in the previous case. CSF in particular was still fuzzy. This is not surprising considering the fact that 10 iterations of NRR were used here, compared with 90 in the previous case.

2.3 W_s Weighting Experiments

I then changed the 3 stages at the start to work in a coarse-to-fine approach. This was to make sure that they get close to the 'best' rigid and affine transformations. I tried to investigate the weighting between shape and texture. I started with 0.5 and stuffed a zero at the right side of the decimal point. NRR now comprised 10 iteration with 3 knot-points along each dimension.

- 20112004-6: weight = 0.5
- 20112004-7: weight = 0.05
- 20112004-8: weight = 0.005
- 20112004-9: weight = 0.0005
- 20112004-10: weight = 0.00005
- 20112004-11: weight = 0.000005
- 20112004-12: weight = 0.0000005

Videos depicting a sequence of the results above suggested that we should try and see what happens at weight > 0.5. There appears to be better alignment when the weight is set to high values[1]. Here are some of the next experiments:

- 20112004-13: weight = 0.9 (trying higher values which result in better alignment, according to visual appraisal)
- 20112004-14: weight = 2
- 20112004-15: weight = 10

Perhaps this is obvious, but values above 1 do not seem to affect registration. The results do not surpass these which were obtained at weight = 0.9. They appear almost identical, in fact, but they are not. There is a difference between weight = 2 and weight = 10 and it seems like a small 'vibration' in the brain. weight = 0.9 and weight = 2 give identical results, maybe because they are the same order of magnitude. The same behaviour exactly is observed for weight = 0.9 and weight = 0.5. Comparison of the results is in 21112004-1.

2.4 Back to Rigid and Affine

Experiment 21112004-2

As weight at 0.5 gave reasonable results, it was decided to stick with it for the time being and investigate another aspect of registration.

It was so far the case that 9 stages were spent optimising rigid and affine transformation. It was worth checking how expensive this optimisation was and identify stages that were of little use.

2.4.1 Translation

The first stage brought significant improvements:

[1]**Reminder:** this is based on registration which takes 15 minutes or so for 2 images.

```
pairwise_stage: {
  warper: translation
  warp_penalty_fn: zero { }
  region_picker: all { }
  levels: 1 2
  optimisation_method: simplex
  param_tol: 0.001
  use_exhaustive_search:  false
  n_per_dim: 5
  optimise_many_warps:  false
}
```

The next stage made very little improvement which is subjective.The only difference was that of levels to:

```
levels: 0 1
```

The next stage was identical to its predecessor and it appeared to align one side of the brain at the expense of another. The overall conclusion is that the first stage only (was default) should suffice. Going to a finer resolution does not result in much better results of translation.

2.4.2 Scale

There were 3 stages as before and the levels reduced from [1 2] in the first stage to [0 1] in the second and third. Judging by the results, the first stage makes great improvements, whereas animations of the second and third stages show just tiny vibrations around the skull.

The conclusion is that the far more significant scaling stage is the first and the latter are negligible. I must point out that it will be interesting to try treating scale and translation not as independent. It will be interesting alternating between these two stages, e.g. performing translation, scaling, then translation again. At present, the first stage of translation does the best it can given a known discrepancy in scale.

2.4.3 Affine

Levels were lowered down from [2 2] to [1 1] and finally [0 0] throughout the three stages. The first stage makes significant improvements, quite unsurprisingly. The subsequent stage **does** make improvements as well. It's refining the overlap of the brain edges. The third stage makes noticeable differences, but it is hard to see if these are beneficial. It would be interesting to look at the difference in objective function value.

2.5 Optimising Scale and Translation "Together"

Experiment 21112004-3

The previous set of experiments suggested returning to translation one scale was 'fixed'. The effects were investigated by performing 6 stages: translation, scaling, translation, scaling, translation and finally scaling.

The results appear **very interesting**.

After the first 2 iterations, which is what normally gets done, improvement continue to be made. These improvements are quite significant too.

To find out how long such improvements can be made, a long such experiment was performed, namely:

Experiment 21112004-4

Same as above, but 18 stages rather than just 6. The result is very interesting to look at. Improvement are made during the first 4 iterations or so and then the brain 'wobbles' between two states (optimal scale versus optimal translation). It still seems that before that stage of 'wobbling', actual improvements are made. It was decided to try the same procedure on two completely different brain data which leads to the next experiment.

Experiment 21112004-5

Two very different brains were chosen. In this case, after the first 2 stages, the 'wobbling' state was reached. This would imply that optimising translation and scale 'together' might not always be helpful.

2.6 Multi-scale in NRR

Experiment 21112004-6

The optimisation regime remained the same for consistency. Also, it the older couple of images which were under investigation, not the different ones mentioned in experiment 21112004-5. To find out how increase in resolution affects change, the following experiment was performed:

Stage 1: Translation

Stage 2: Scaling

Stage 3: Translation

Stage 4: Scaling

Stage 5: Affine - coarse

Stage 6: Affine - finer

Stage 7: NRR - coarse - 5 iterations

Stage 8: NRR - finer - 5 iterations

Evaluation of each of the stages, judged by the results:

Stage	Significant
1	Improvement
2	Significant
3	Noticeable
4	Small
5	Noticeable
6	Small/neutral
7	Small/noticeable
8	Small

It was of course taken into consideration that bigger changes will be made at the start.

2.7 NRR in More Depth

Experiment 22112004-1

To see the effect of NRR above in a bit more detail, the next experiment was performed. It would also be interesting to see the pace of improvement as a function of scale (resolution). So, a multi-scale approach was again used: 3 levels, 10 iterations in level 2, then 2 iterations in level 1, and 1 iteration at the finest level. There would be 13 frames of NRR to look at and the experiment would be relatively long.

It is only the last (NRR) stages that are of interest here: many beneficial changes are made in the first 5 iterations of NRR, the next 5 iterations (still at coarse resolution) make small changes to only one small region of the brain. At a lower resolution (level 1), there is change yet it is not very visible to the eye if it is a good change. The same goes for the last stage at level 0. It seems to almost make things worse.

Experiment 22112004-2

As above, but 6 iterations at level 1, and 6 at level 0. This experiment was run to reassure that conclusion were not made at haste. At level 1 there are little 'vibrations' that don't affect the non-reference image much. At level 0 there is a little more change, but again there is only a mere movement, up to a pixel or two in extent.

The conclusion is that going to finer resolutions will consume a great deal of time and produce little gain. There is significant change, but not necessarily good change.