Response to Recommendations Report

Roy S. Schestowitz

July 27, 2009

My revision involves changes affecting the text as indicated below and the new version of the thesis is about 40 pages longer. It comprises additional parts and rephrasing of sentences which add clarifications. Particularly, changes were made based on each pertinent point raised by the examiners in their personal reports and notes. Certain sections were scrapped and rewritten (mostly in Chapter 5 where experiments were redone), whereas others were introduced to cover new experiments that defend my thesis (e.g. 1D experiments and additional examples in 3D). Some changes are larger than others and some were not explicitly asked for, but they improve clarity and help ensure that explanations are more comprehensive, and ideally complete.

For the examiners' convenience, this is a response to their main report. It shows which portions of the new version of the thesis address which request/suggestion from the examiners. I shall tackle changes proposed in the report by referencing parts which address each of them in turn. Quotes are indicated by bold fonts and in order to improve readability, these quotes contain the entire report as-is, only split apart. In order to make it easier to get a clear impression from the document as to what exactly

has changed, each chapter is accompanied by a very brief summary of what I have done.

Chapter 1 Summary: No major change. Typographical and grammatical errors corrected.

Chapter 2 Summary: No major changes, with the exception of 2 parts about feature-based registration methods that were added (Subsection 2.3.1). Various organisational changes were also applied so that there is improved separation between different strands of methods. Many references were added (almost 20 in total since the previous submission).

Chapter 2: Non-Rigid Registration

The review of non-rigid registration methods should encompass feature-based registration methods.

This point was addressed in two parts. Subsection 2.3.1 on "Objective Function" (page 37) now incorporates an introduction to feature-based registration methods, which are explained in more detail at the end of this subsection (page 41 to page 43), just before "Optimisation" (Subsection 2.3.2).

Chapter 3 Summary: Typographical and grammatical errors corrected. Explanation about the covariance matrix was added to Subsection 3.2.3 (page 59)

Chapter 4 Summary: Section 4.2 on information theory was added and the old Section 4.5 removed (Section 4.5 used to be about experiments that were prematurely and too scarcely explained). Section 4.2 describes MDL and Section 4.3 gives corresponding mathematical descriptions. The connection to the work of Kotcheff was made clearer. Typographical and grammatical errors were also corrected.

Chapter 4: MDL Shape Models

This chapter mostly provides an important section of background for the research presented in later chapters. However, key components of the MDL (minimum description length) approach are not described (the MDL objective function itself, for example).

Section 4.2 (page 79) was added to provide background on information theory. Therein, an introduction to Shannon's entropy and to MDL is included as well. In Section 4.3 (page 82), a detailed explanation of the MDL formulation is given and its components are described. A formal definition of the covariance matrix was added to Chapter 3 (Subsection 3.2.3, page 59) to make derivations more complete.

The methods developed for image segmentation in chapter 5 are inspired by the MDL method, but in practical terms owe more to the earlier work of

Kotcheff. It is important at that stage to have a clear idea of which ideas are being carried forward from the shape domain and which are not.

The explanations now place greater emphasis on ideas which are encountered in Chapter 5. Better distinction is also made between the idea of MDL and the work of Kotcheff. A relationship between those two is described mathematically in page 88.

Overall the description of the MDL method (sections 4.1 and 4.3) is sparse, and its underlying idea would be rather opaque to a reader who was not thoroughly familiar with it already.

I believe that the revisions made (see above) not only give the necessary background information but also inform in greater depth how the MDL formulation is derived and used.

More formal information about the work of Taylor and Kotcheff was also added, along with its relationship to MDL and to my work (see Section 4.3 and Section 4.4, for example, in page 82 onwards). Later sections were modified to reflect on that.

Technical explanations were preceded by more verbal descriptions for those who are not already familiar with the underlying idea.

Section 4.5 describes some extensions to the MDL

optimisation for shape. While this element of the work is rather peripheral to the main focus on the thesis, it is still important to describe it clearly, particularly as a similar approach is adopted later in chapter 5 on registration.

Two extensions are described: use of subsets and varying optimiser tolerance. The description needs to make clear exactly what experiments were carried out here (how large were the subsets, what was the optimisation scheme...) and to make explicit what the results were.

This section about extensions was removed because it had been put in the wrong place. Some new practical work was undertaken to achieve the same thing in the context of appearance models in Chapter 5. This includes similar experiments but for full appearance models with detailed explanations and results.

In the current version we see one graph (figure 4.8) which shows something improving in a rather uneven fashion followed by the assertion that improvements were made.

This relates to the section above, which was removed. It covered extensions.

Chapter 5 Summary: Extensively rewritten, includes new material on NRR experiments, and a lot of work in 2D (see Section 5.5, which runs from page 129 to page 142). The chapter no longer includes material related to older and poorly explained experiments, it is better divided structurally, it contains improved explanation of the NRR method, and it presents new experiments that explore shape/intensity weighting, optimiser tolerance, and subsets. Experiments in 2D provide better separation between my own work and that of colleagues.

Chapter 5: Model-based Registration

This entire chapter is very hard to follow.

The chapter was significantly restructured and old experiments discarded in favour of large new ones that I performed especially for this final revision. I then explained them in greater detail, as well as provided results for. Additional work on 2D registration was added to better demonstrate my personal contributions and clearer illustrations presented as figures.

The methods and results were made better connected to the new text in Chapter 4. Additional references were added to attribute related work.

There seems to be a mixture of granularity in the exposition, such that some very preliminary experiments are presented (such as those that give rise to figures 5.2 and 5.3) at the same level as more substantive experiments, with little guidance to the

reader.

There is now a clear distinction which starts with an introduction to the problem and an explanation of the shape component, as seen in Figure 5.2 (page 100) for example. No experiments are presented at the early stage; instead, they are quite consistently presented in the dedicated sections where experiments are systematically explained and their results shown (sections 5.3, 5.4, and to a lesser degree 5.5).

The "model-based" objective function (equation 5.1) is not fully described (what is δ?). This equation is referred to in the context of the MDL function in shape, but is in fact related to the earlier approach of Kotcheff.

Terms like δ are properly described now (page 105) and they are put in a suitable context, relating to the clearer background which Chapter 4 provides after extensive revision. The relationship between MDL and Kotcheff's method was made a lot clearer and explanations less ambiguous. This is fully described in Subsection 5.2.3 (page 104).

This in itself is fine, but it should be made clear what is being modelled in this case, particularly as we subsequently (section 5.3.3) have the (unsubstantiated) assertion that some observed unwanted behaviour of the optimisation is due to the approx-

imation to the MDL function that is actually used.

Subsection 5.2.3 on the objective function was extended so that it includes all the information required. Paragraphs were also changed to explain how it relates to the work of Kotcheff.

The results of the registration experiments are presented in a confusing way. The registration process is represented in figures 5.10 and 5.11 by the progressive reduction of the "score". However, as the score is simply the minimand, it is not helpful to know that it is minimised.

These results were left out altogether because they did not satisfy a high standard in terms of explanation. These were replaced entirely by new experiments that are described in Section 5.3 (page 111). The section covers large-scale experiments that explore shape/intensity weighting, optimiser tolerance, and subsets.

The experiments need to show improvement of registration with reduction of the score. Since the 1-D images used are all synthetic, it should be possible to do this.

Figure 5.18 (page 127) shows reduction of the score as a function of iterations. Figure 5.13 (page 120) and Figure 5.15 (page 123 at the bottom) show how the distance of

the correct solution varies as a function of optimiser tolerance and as a function of set size, respectively. Figure 5.17 (page 125) shows early results that are later improved, as shown in Figure 5.19 (page 128). That latter figure shows the final results by dissecting the models before and after NRR.

In particular figure 5.12 shows a comparison of a number of registration methods in terms of the score. The diagram is presumably intended to show the author's method is superior to the others, but in fact indicates that the "model-based" method behaves dramatically differently from methods that would be expected to do quite well.

The results from these old experiments were not comparable for reasons I shall explain in a moment. They were therefore discarded and proper experiments put in their place (see Section 5.3, which starts in page 111).

These results also need to be presented in terms of registration accuracy.

The new results (e.g. Figure 5.11 in page 118, Figure 5.13 in page 120, and Figure 5.15 in page 123) include a measure of distance from the correct solution, which for this synthetic data is always known.

It is not clear in any of the approaches (including the author's) what optimisation method (as distinct from objective function) was used. This seems to have been the MATLAB default, but we are not told what it is. Was the method used for all registration methods compared? Was it appropriated for all of them?

The new text (see page 119) makes it explicit and clear that a general-purpose Nelder-Mead optimiser was used throughout all the experiments.

There is a mention on page 105 of Taboo search, but it is left unclear whether this was implemented, and if so with what result.

Taboo search was not implemented, so this portion of the text has been removed.

It seems quite likely that in a properly conducted comparison, the "model-based" method will not emerge showing any great advantage over some others. While disappointing, such a result should be presented honestly with discussion of the rationale behind the approach and the significance of a negative result.

These old experiments were conducted back in 2003/4 using deficient algorithms that unfortunately flattened the bumps. Thus, they were not helpful in showing anything

meaningful or significant. The figure should sensibly be omitted.

New experiments (mostly in Section 5.3) show that this problem of flattening has been resolved since these old experiments were performed.

Figures 5.16 and 5.17 are particularly confusing. They can only be understood in the knowledge that the "shape" element of the 1-D appearance model refers to the spatial transformation between the input image and the output image. This is not stated explicitly. There needs to be a very clear description of what constitutes a "shape" and "texture" model in this case.

In rewriting of this chapter, great attention was paid (as seen in explanations) to the need to clarify what "shape" means in this context. Unhelpful figures, which depicted curves with insufficient reasoning, were removed and text rewritten to prevent unnecessary confusion. Background can be found in Section 3.3 (page 63) and shape in these experiments is explained in Subsection 5.2.1 (page 98), with Figure 5.2 (page 100) and Figure 5.7 (page 110) to illustrative this visually.

There should also be some discussion of the "spiky" nature of the registered models on the bottom line of figure 5.17.

This is now Figure 5.19 (page 128). An explanation was added to the end of Section 5.4 (page 129).

In rewriting, care should be taken that figures show clearly what is intended, and have clear captions explaining the contents.

As major changes were applied, many of these old figures got dropped along with their captions. Most of these figures showed the value of the objective function declining (owing to optimisation), but lacked proper explanation of the corresponding experiments. The new figures and the flow of information are now considerably improved.

Section 5.4.2 refers to using registration for 2D model construction. This work is of a different character from the rest of Chapter 5, and refers to material, some of which is the author's work and some the work of others.

The big new section, Section 5.5 (page 129 to page 142), embodies this work and adds a considerable amount of material which demonstrates my contribution and makes it a lot clearer to see what I achieved.

It would be possible to present more quantitative results here, but it is very important to distinguish explicitly the author's contribution from that of collaborators.

In general, a good deal of the text in the thesis was changed to eliminate cases where my personal work was wrongly re-

ferred to as "we" or "our" (usually just meaning the reader and I). A lot more technical work was included in this section to share results that I reached. Although quantitative results were obtained and published in papers from the group, I was not involved in this, so this is not included.

On page 115 the registration is described as using the "MDL framework". Be clear what registration metric is used here and how it relates to the method evaluated on 1-D images.

There was a slip in the text. The metric used was mutual information, as the corresponding figure, Figure 5.21 in page 132, explicitly stated. The text was changed to correct this mistake.

It may be appropriate to place this material in a separate chapter.

Given the reorganisation of the chapter and the improved flow of argument, I believe this now fits where it is.

Chapter 6 Summary: The method is discussed in Subsection 6.2.3 where I defend it from criticism of a potential issue arising from misinterpretation of the results. No other substantial changes were made.

Chapter 6: Assessment of Non-Rigid Registration

This chapter (And the succeeding ones) is based on the premise that it is possible to have a measure that is monotonic with misregistration, but is not a measure of misregistrations, provides a basis for validation. It can be argued that it doesn't. There should be some discussion of this point.

Subsection 6.2.3 (page 151) now embodies a discussion of this point.

Chapter 7 Summary: The explanations were improved in the sense that warp parameters are now described more fully and there is better information about the methods. The decision to focus on Specificity is also defended, with a table that illustrates possibilities one faces.

Chapter 7: Validation Methodology

An editorial note: Figure 7.1 belongs more naturally in chapter 6.

The image (Figure 6.5) was moved where it belongs and it can now be found in page 155.

In section 7.1.2 there should be some discussion of whether the synthetic warps introduced are valid in the context of the distortions one might expect among real images. Why were the specific values of d chosen?

The second paragraph in this chapter was extended to include an important point about these distortions. Several new paragraphs were added to Subsection 7.1.2 (page 161) where this choice is described and defended. Values of d are also better explained.

Specificity and generalisation are used as distinct measures of registration quality. However, it is pointed out in chapter 6 that neither alone is indicative of a "good" model; they require to be op-

timised together. There should be some discussion of the validity of using them as separate measures.

Section 7.3 (page 171) is used to present a discussion of these points.

Chapter 8 Summary: The explanations were made better connected to the remainder of the thesis, a new important graph was added as required, and a new section demonstrates my work on 3D NRR.

Chapter 8: Application to Evaluation of NRR

In Section 8.1.3 the groupwise registration method is described as a MDL formulation. Here, as previously, the precise method, and its relation to that described in Chapter 5, need to be made quite clear.

Chapters 4 and 5 were redone to include a detailed explanation of the MDL (and MDL-inspired) formulations, so Subsection 8.1.3 (page 177) was augmented to connect this new work with previous explanations of MDL.

Figure 8.1 shows how the specificity varies with the number of modes used in the model. There should be a similar figure addressing generalisation and some discussion of the nature of the relationship shown.

The complete data for this experiment was found and then used to produce the corresponding graph for generalisation. Section 8.2 (starts in page 178) was extended to include a discussion of the new figure (Figure 8.1 in page 179) and its relationship to others.

It transpired in the thesis that some initial experiments were carried out with 3D image sets. This work, albeit preliminary should be included in the thesis.

Section 8.3 (page 181 to page 185) was added to cover my preliminary work on 3D NRR.

Chapter 9 Summary: Changes were made to shed light on my personal contribution and Subsection 9.2.1 now explains a lot better the work on KL divergence.

Chapter 9: Further Work

Here your work needs to be separated from the work of others. Explain how the method of KL divergence (not your work) is related to your development.

The Chapter was rewritten where appropriate to better differentiate my work from that of others. Subsection 9.2.1 (page 187) was extended to explain the work of Twining *et al.* It also explains how it relates to my work.

Chapter 10 Summary: No changes made. I may have dealt with very minor changes.

In addition to the above there are a number of detailed textual or typographical corrections that can be made. These are annotated on the examiners' copies of the thesis.

Done.