

The Researcher's Bible

Alan Bundy, Ben du Boulay, Jim Howe and Gordon Plotkin 1985

Including contributions by Graeme Ritchie and Peter Ross.

Thanks to [Anna Chekovsky](#), who has made [a Swedish translation of the Researcher's Bible](#), to Artem Delik, who has made [a Bulgarian translation](#), and to Valeria Aleksandrova, who has made [a Polish translation](#).

This version: 9 November 2004.

Abstract

Getting a Ph.D. or M.Phil is hard work. This document gives advice about various aspects of the task. Section 1 describes the problem - what is a thesis? Sections 2 and 3 describe some of the pitfalls and hurdles which students have encountered. Sections 4 and 5 advise about choosing and then executing a research project. Sections 6, 7 and 8 deal with two of the three R's: reading and writing. Section 9 describes the examination process for a research degree, and how to cope with it. Finally, section 10 gives some references to follow up.

1. What is a Thesis?

To get a Ph.D. or M.Phil. you must write a thesis and sit an oral examination, or vivas. The oral is generally used to ask for clarification of the thesis, so the main burden of assessment falls on the thesis. [Postgraduate Degree Programme Regulations 2014/15](#) give the precise requirements.

What is required from a Ph.D. thesis is described in [regulation 40](#). How original and significant must Ph.D. research be? It should be possible to distill from the thesis, a paper worthy of publication in a journal. This is not an infallible guide - refereeing standards are not always what they should be. The final decision must rest with the examiners.

What is required from an M.Phil. thesis is described [in regulation 43](#), but is rather vague. Most M.Phil. theses have been a record of research rather than a critical survey, but the latter would be possible. Again it should be possible to distill the essential message of the thesis into a short paper, but in this case publication in a journal would not be essential, but some form of publication is still desirable.

If you do not know what standards are expected in a journal paper or conference paper - read some! Read some theses too. Do not be intimidated by American theses. American Ph.D. students spend 5 or 6 years studying, compared to the British norm of 3 or 4 years.

2. Standard Pitfalls for Postgraduate Students

There are a series of standard traps lurking to catch postgraduate students, or anyone else, doing research for the first time. It is as well to be aware of these, then there is an outside chance of avoiding them. Some pitfalls are described below (I fell in most - AB).

2.1 Solving the World

Most students pick research goals which are far too ambitious. It is particularly easy in Informatics to underestimate the amount of work necessary to automate a task. Tasks that appear effortless to humans, reveal hidden depths when investigated in computational detail. Obviously the main burden of helping you choose a suitable topic will fall on your supervisor. In addition you should read the literature and talk to fellow workers to find out what the state of the art is. One good source of ideas is the further work sections of papers. Read the literature critically. Another good source is re-doing bad work, properly.

2.2 Manna from Heaven

Having chosen a topic, what do you do next? It is no good sitting in your room with a blank piece of paper and a pencil, expecting good ideas to come from above. What you can do is:

- Read the literature. Have projects similar to yours been tackled before, and were previous attempts successful or unsuccessful? What existing techniques can you borrow or adapt to your project? Do you need to adapt your project proposal to make it novel or feasible?
- Talk to people. Do not go away and hide. Do not be ashamed of your ideas. Other people's are sillier.
- Tackle a simplified version of your problem. Ask your supervisor for exercises, mini-projects, etc.
- Write down your ideas in a working paper. Imagine yourself explaining your ideas to someone. You will be amazed at how half-baked ideas take shape and errors are exposed or solved.
- Give a talk to a small group. This has a similar effect to writing down your ideas.

2.3 Computer Junky

Computers are very seductive. You can spend years at a terminal debugging your programs and tuning up the input/output routines. You get a satisfying sense of achievement when a bug is exposed or a nice output generated. This is illusory! Your program must be explainable at a higher level than code, for it to make a real contribution to knowledge. Try to plan your program theoretically before going to the terminal. If you must work some of it out at the terminal then rush away soon and work out the theory. If you find this hard, try to describe how it works: to a friend; in a paper or at a seminar. If people do not understand it is your fault - try harder.

2.5 Micawberism

Gathering experimental data can be fun and gives all the appearance of productive work. Make sure that you know what class of result you are attempting to establish with your experiments.

- Talk to people, explain what you think your experiment might show.
- Imagine the experiment finished and you have 'the data', what exactly would you do with it.

- Not only try out the experiment on one or two people first, but try out the analysis. Don't keep running experiments in the hope that something will turn up.

2.6 Ivory Tower

Single minded dedication to your topic is a good thing, but do not shut out the rest of the world completely. Keep in touch with the state of the art in related fields. Talk to other people about their research. Attend selected seminars and lectures. Set aside a part of the week for reading reviews and abstracts and skimming papers.

2.7 Misunderstood Genius

It is all too easy to believe that the reason why no one understands your ideas is because you are a genius and the others are all looneys and charlatans. There are alternative causes for misunderstanding that you should consider:

- **Love of Jargon:** Informatics is full of jargon: try to rephrase your ideas using ordinary English; try to rephrase your ideas in someone else's jargon. Do they come out any differently?
- **If I can do it, it's trivial:** Once you have seen the solution to a problem it appears trivial. Then it is tempting to say 'that's too easy, I'll try something else'. This is a non-terminating loop! Your solution won't be trivial to other people (probably it will be wrong or over-complex) and should anyway be used as a basis for further work. Motto: do the easiest thing first, then stand on the shoulders of these achievements and do the next easiest thing, etc. - a better infinite loop.
- **Love of complexity:** It is not a virtue to make a complicated program - it is just a nuisance to other people. Do it the simplest way you can. Occam was perfectly right.

2.8 Lost in Abstraction

To be worthwhile your research work should contribute to solve a hard problem in Informatics: e.g. making computers easier to use, smarter, more dependable, or better able to model natural systems. But to achieve anything you must tackle these abstract problems in a concrete situation, that is you must build a program that is easy to use, smart, dependable or a good model. Trying to tackle the problem in the abstract will only lead to paralysis and frustration.

2.9 Ambitious Paralysis

It is good to have high standards for your finished product but do not apply the same standards to your initial attempts, or you may never get started. Do something simple, then apply your standards to refine it into something worthwhile.

2.10 Methodology does not make a Thesis

Since Informatics is a relatively young field, and is interdisciplinary in nature, it does not have a single methodology. One of the difficult tasks that you face as an Informatics research student is the development, consciously or unconsciously, of a suitable approach to the problem(s) being tackled. In the course of evolving an appropriate methodology, you will encounter many other methodologies and philosophical positions, many of which will seem outrageous or hopelessly misguided. You will nevertheless find that these bizarre viewpoints have strong proponents,

perhaps at the next desk in your office. Hence, much of the formative period for your own methodology is spent having heated arguments with fellow researchers. Out of this struggle, your reading, your attendance at seminars, your debugging, and other hard work, will emerge your world view on Informatics and related philosophical issues. In later years, you will probably come to take this outlook for granted, perhaps modifying it occasionally in some way; however, it is quite likely to loom very large in your life during the period of your project, and when you come to write your thesis you may feel compelled to expand upon your philosophy of life at length. Restrain yourself - the examiners won't be all that interested. Give a brief summary of your methodological assumptions, giving references across to existing arguments or frameworks where appropriate, and confining yourself to the points which are essential to the understanding of the substance of your thesis. If your view is so wildly radical that you really need to spend fifty pages expounding it, it may be slightly suspect.

2.11 The Discovery Route is not a Justification

In the course of your project, you will come to certain beliefs about technical issues, some of which will be novel, and many of which will be rediscoveries (or new understandings) of established concepts. In presenting your thesis, it is important to distinguish between the justification (for instance, generality, efficiency, perspicuity, practicality) for some position or technique, and the route by which you happened to come to favour this idea (such as that it seemed similar to your ad hoc program, it was better than the theories you were taught as an undergraduate). The readers and examiners aren't particularly interested in reconstructing how you became convinced of an idea - they are interested in the general arguments in favour of the idea. When you have just become convinced of some point, your own discovery route will seem like the dominant reason for it, so you may need a cooling-off period before you can detach yourself sufficiently to write a reasoned support for the idea, particularly if it is your own idea as opposed to enthusiasm for someone else's.

3. Psychological Hurdles

Doing research shares the same psychological difficulties as other creative endeavours such as writing novels and plays or painting pictures. Some of these difficulties and their antidotes are set out below.

3.1 Mental Attitude

Part of a researcher's skill includes an appropriate mental attitude to his/her work. This can be learnt, if you know what you are aiming for and are determined enough. One of the main ingredients of this mental attitude is a belief in what you are doing. Do not be afraid of a little egotism! You must believe that the problem you are tackling is important and that your contribution to the solution is significant. Otherwise, how are you to generate the energy to see yourself through the long hours of hard work required? The first step in obtaining this self-assurance is to pick a research topic you can believe in (see section 4). Of course, you must not become so arrogant that you can no longer listen to criticism. You must be prepared to modify your ideas if they are wrong.

3.2 Research Impotence

For many people, research prowess is a kind of virility symbol. Lack of success at research is accompanied by the same feelings of inadequacy as sexual impotence and, like it, can be a self fulfilling prophecy. Doubts about your own ability can put you in a frame of mind where the dedication (Edison said that genius was 1% inspiration and 99% perspiration, and he should know.) and enthusiasm necessary to produce results evaporates. The way out of this vicious circle is to realise that research ability does not depend on some magic essence. It is a skill which can be learnt, like any other. You too can do original research by following the instructions in this guide.

3.3 Dealing with Criticism

We all find criticism hard to take, but some of us hide it better than others. If you are to make progress in your research you will have to learn to seek out criticism and take it into account. You will have to learn to differentiate between valid and invalid criticism. If you feel too close to the subject to decide, ask a friend for a second opinion. If the criticism is invalid, maybe the critic has misunderstood. Can you improve your explanation?

You are going to have to learn to take some knocks: rejections from journals; rough rides in question time. Take it with a smile. Learn what you can. Don't be tempted to give up - you are in good company. If you study the lives of famous scientists you will see that many of them had to endure very heavy criticism. In fact some of the best work is the product of personal feuds - each scientist busting to outdo the other. This is where your faith in yourself will be tested to the full.

3.4 Early Morning - Cold Start

Almost everybody finds it difficult to start work at the beginning of their working day, but once they have started, it is relatively easy to keep going. The remedy is twofold:

1. Make yourself a regular working schedule - and stick to it. It doesn't have to be 9-5, but there should be a definite time of day when you expect to start work. Otherwise, you will find yourself postponing the evil moment with endless, routine, domestic chores.
2. Make sure you leave some non-threatening, attractive task to do first thing. For instance, do not leave off writing the day before at the beginning of a new hard section. Leave something easy to start writing: a paragraph which is routine for you, a diagram to draw or a simple procedure to write.

3.5 Theorem Envy

You have chosen a new field where the research methodology has not yet been worked out. You may get a hankering for the methodology you were brought up on. For mathematicians this might be the longing to prove clean, clear theorems - theorem envy. For engineers this might be screwdriver envy, etc. Be wary! Only try to prove a theorem if it is clearly relevant to your overall purpose. For instance, proving the termination of a procedure you have found to be useful may well be relevant. Finding a procedure whose termination you can prove, but which is not otherwise interesting, is not relevant.

3.6 Fear of Exposure

You have a great idea and it only remains to test it by proving a theorem, writing a program, explaining it to a friend etc. But something is holding you back. You find it difficult to start work.

Could it be that you are secretly afraid that your idea is not so great after all? Hard experience has taught you that ideas that appear to be solutions to all your problems in the middle of the night, evaporate in the cold dawn. Courage! Research is always like this. Ten steps forward and nine steps back. The sooner you subject your idea to the acid test, the sooner you will discover its limitations and be ready for the next problem.

4. Choosing a Research Project

Your research project must fulfil the following criteria:

1. You must be enthusiastic about it.
2. Solving the problems it entails must be worthy of a Ph.D.
3. It must be within sight of the state of the art, i.e. it must be 'do-able' in three years.
4. There must be someone in the School willing to supervise it.

The importance of 1. cannot be overestimated. You are going to need all the enthusiasm you can muster to give you the perseverance and motivation to see you through what will be a hard, lonely and unstructured period. It will help if you choose to tackle a problem you consider of central importance (though you cannot expect to bite off more than a small chunk of it). It will also help if you choose an area which utilizes your already proven abilities, e.g. theoretical computing for mathematicians; computational linguistics for linguists. Beware of choosing an area new to you because of its superficial appeal. The gloss will soon wear off when you are faced with the hard grind necessary to get a basic grounding in it.

Having chosen the general area or problem you want to work on, you must try to define a specific project. This is where your supervisor comes in. Find a member of academic or research staff whose interests lie in this area and who is prepared to advise you. S/he may have some projects to suggest and will also be able to pass an opinion on the worthiness (2) and doability (3) of anything you suggest. On the whole, beginning students tend to underestimate the worthiness and overestimate the doability of projects - quite modest sounding projects prove harder than they look. So do listen to your supervisors advice and don't fall into 'solving the world', standard pitfall no. 2.1.

Get your supervisor to suggest some reading material. You will find suitable projects in the future work sections of papers and theses. It is good research methodology to continue working on a problem from where someone else left off. You may find some work you consider badly done - consider redoing it properly. You may be able to simplify someone else's program, relate it to other work or build a more powerful program.

Have a range of ideas on the boil. Try to construct a hierarchy of research goals. This imposes a structure on the work and also acts as a safety net when you find (inevitably) that you have attempted more than is possible in the time available.

Projects to avoid, because they lead to bad research, are programs which do a task without addressing any important issues and programs which are not based on previous work (also see the section on standard pitfalls).

5. Research Methodology

Informatics is a young science which draws on the methodologies of many fields and is gradually evolving its own methodology. See [the Informatics Research Methodologies course](#) for more discussion of these issues. This methodology supports a variety of approaches to your research project. For example, you might start by trying to build a theory of how some task might be automated, or by improving somebody else's theory, or you might try to rationally reconstruct someone else's work. The 'rational reconstruction' approach is often fruitful, since it is still regrettably often the case that Informatics research will focus on implementation and performance while leaving the assumptions and principles behind the work implicit and vague. But, however you start, get yourself a theory!

6. Writing Papers

Research papers are the major product of the School. They are the yardstick by which our individual and group progress and success are measured. They are therefore very important and you should expect to devote a large part of your research career to writing them. Writing papers is the main way of communicating with the rest of the Informatics world and it is also a good vehicle for clarifying and debugging your ideas.

As well as the dizzy heights of books, theses and journal papers, there are various lesser forms of writing. You should understand what these are so that you can make full use of them.

You should make writing a regular part of your life. Keep records of everything you do: notes of ideas you have; documentation of programs; lecture notes; notes on papers you read. These serve several purposes: an aid to your memory (you will be amazed at how quickly you forget); a vehicle for clarification (how often you will find that problems appear and are solved as you try to explain things to yourself and others) and as a starting point for a working paper. Make sure you write them legibly enough to read later and that you file them somewhere you can recover them. Recording and storing these notes electronically works well.

6.1 Informatics Technical Reports

The School has a [technical reports series](#) to which you are strongly encouraged to contribute. In particular, papers submitted to journals, conferences, etc should be included in this series. If you are asked to sign a copyright form by a publisher check it first with your local Service Manager to ensure that the School will retain the right to make your paper available electronically via our web pages.

6.2 Publishing Papers

When you and your supervisor think that you have something worth publishing externally you should submit a paper to a conference or journal. In preparing a paper for publication make sure that credit is given to everyone who has helped with its preparation, e.g. your supervisors and anyone else who has contributed ideas, others who have commented on the draft, and so on. Where a contribution is significant (for example, your supervisor's contribution) consider joint authorship. Remember to acknowledge sources of support such as source of your research studentship and related support for facilities used for the research and so on. If uncertain consult your supervisor about these points. Washington University in St Louis has [a policy on authorship](#) that reflects the scientific consensus on who should be the authors of a paper and what their rights and duties are.

A submitted paper will be vetted by several referees chosen by the journal editor. Do not be too downhearted if your paper is rejected - you will be in good company. Read the referees' comments carefully. Are they right or have they misjudged you? Is your rejection absolute, or have you been encouraged to resubmit after corrections or further work? Was your choice of journal appropriate? Consider submitting your paper elsewhere, but first take into account those criticisms you consider valid.

6.3 Conference Proceedings

A lesser form of publication is the proceedings of a conference. Conferences will often consider descriptions of work in progress. They will usually be refereed just like journal papers. Both papers and verbal presentations usually have strict length limits (from 5-15 pages and 10-30 minutes), so be prepared to be concise. Presenting a paper at a conference will be very valuable for you: you will get feedback from a wider audience than usual; you are more likely to meet people than a non-participant and you will find it easier to get funding to attend.

Advice about expressing your ideas to a large audience in plain English see [Orwell68].

7. Guide to Writing

During the course of your research project you will need to write many documents: a thesis proposal and thesis outline, research notes, records of papers you have read, conference and journal papers, and finally the thesis itself. A badly written thesis is not usually a cause for total failure, but can cause soul-destroying delays while it is rewritten and reexamined. Poor writing will also make it difficult for others to understand your work. It is, therefore, quite important that you learn to write well. This section contains some tips and rules to improve your writing. Nobody knows enough about good writing to do more than that. There is a good guide to style and presentation of scientific papers in [Booth75]. Helpful information about writing theses is given by [Parsons73]. Useful references for writing technical reports can be found in [Cooper64]. Advice about writing Computer Science papers is given in [Zobel 04]. The College Transferable Skills Programme runs a course on [How to Write an Informatics Research Paper](#) for Informatics postgraduate students, which you should consider attending.

There are no hard and fast rules of good writing, but if you are going to break one of the rules below you should have a good reason and do it deliberately, e.g. you want to overwhelm the funding agency with jargon rather than have them understand how little you actually achieved.

- Your paper should have a message, i.e. an argument that you are advancing, for which your research provides evidence. Make sure you know what this message is. Summarise it in a few words on paper or to a friend. Make sure the message is reflected in the title, abstract, introduction, conclusion and in the structure of the paper.
- Putting your case so that it can be understood is not enough - you must present it so that it cannot be misunderstood. Think of your audience as intelligent, but (a) ignorant and (b) given to wilful misunderstanding. Make sure that the key ideas are stated transparently, prominently and often. Do not tuck several important ideas into one sentence with a subtle use of adjectives. Do not assume that any key ideas are too obvious to say. Say what you are going to say, say it, and then say what you just said.
- Do not try to say too much in one paper. Stick to the main message and only include what is essential to that. Reserve the rest for another paper. A reader should get the main idea of the

paper from the first page. Long rambling introductions should be pruned ruthlessly.

- The basic framework for a scientific paper is: what claim/hypothesis am I making and what is the evidence for this claim.
- To keep the technical standard of paper uniform, have a particular reader in mind as you write.
- You do not have to start writing at the beginning. In particular, the introductory remarks are best written when you know what will follow. Start by describing the central idea, e.g. your main technique, procedure or proof. Now decide what your hypothetical reader has to know in order to understand this central idea and put this information into the introductory sections/chapters.
- Use worked examples to illustrate the description of a procedure. Do not use them as a substitute for that description.
- Clearly state what is new or better about what you have done. Make explicit comparisons with closely related work.
- If you find yourself using a long noun phrase to refer to the same entity or idea several times then you should probably define a new term. Do not define a new term unless you really need it.
- Learn to use a keyboard (all 9 fingers), a screen editor, a text formatter, a spelling corrector and a grammar corrector. Type your paper into a computer, either directly or from notes or from a handwritten manuscript. This will save time when it comes to alterations, corrections, etc. Run the finished product through spelling and grammar correcters.
- Ask several people to read the draft versions. Expect to spend time incorporating their suggestions into the text. If they did not understand it is your fault, not theirs. It is discourteous to ask anyone to reread a paper if you have not yet considered their previous comments. Draft theses should be read by supervisors, but should *not* be read by examiners. The remarks below are relevant to all writing, but are particularly addressed to thesis writing.
- Your thesis should not be a 'core-dump' of all you know about everything remotely related to the topic. Instead, there should be a single message, and you should carefully consider how each part of your thesis contributes to putting over this message. Remember that you are not writing specifically for your examiners. You should be addressing yourself to researchers following in your footsteps, who will be grateful for a good but relevant scene-setting and a clear argument. They will also be considering the state of knowledge at the time you were writing, which may be different from the state at the time they are reading it, and you should give sufficient detail to fix this without boring them rigid. It is also wise to remember that researchers around the world will also, implicitly at least, be judging the quality of the University and of the School when they read your work. Your examiners will be bearing this in mind even if you don't - so you should too.
- You can write your thesis top down, bottom up, or bi-directionally. Top down you start with some notes, and gradually unpack them into thesis chapters. Bottom up, you describe different aspects of what you have done, and then put these parts together to form the thesis. Neither of these approaches is very successful on its own. Top down tends not to work because your opinion as to what you have done changes as you unpack your ideas. Bottom

up produces a dog's dinner of unrelated snippets. A bi-directional combination is more successful.

- As you do your research you should write your ideas and results up as a series of notes and working papers. Some of these papers may be worthy of publication in a conference or journal. Collect these notes and papers into a single file (paper or electronic) entitled `thesis'. This is enough bottom-up work to start with. Now work top down.
- Build your thesis `message'. This should have the following properties.
 - It should consist of a few sentences, i.e. be of abstract length. - The sentences should form the steps of an argument. This argument is the message of your thesis.
 - Each sentence should outline the contents of some part, roughly a chapter, of your thesis.
 - The message should serve as a guide to the: title, abstract, summary, conclusion and the whole body of your thesis.
- The thesis message should help you in the following ways:
 - It should ensure that the parts of your thesis hang together in a coherent manner. It should suggest how to reorganise the various notes and papers in your `thesis' file so that they form an argument.
 - It should answer the questions `What have I done?' and `Why does this work deserve a degree?'. You should now know what to emphasize in the abstract, introduction, conclusion, title, etc.
 - It should answer questions like `What should be discussed in 'related work' ?'. In fact, you should know precisely what role each chapter is meant to play in the whole, i.e. what it is supposed to prove.
- The thesis message is short and easy to edit. You can play around with it until you get something you are happy with.
- Now you can go back to bottom up activity - reworking the existing material, and writing new material, to fulfil the demands of the `message'.

To give a flavour of the `message' described above, we give an example from the Ph.D. thesis of a famous AI researcher, Fr. Aloysius Hacker.

Example of a Thesis Message.

"The Computational Modelling of Religious Concepts"

by Fr. Aloysius Hacker

1. We apply ideas from Informatics to the understanding of religious concepts.
2. Previous attempts to explain religious concepts, e.g. the holy trinity and miracles, have often encountered philosophical problems.
3. These problems arose because the appropriate terminology was not available. Computational terminology often provides an appropriate analogy.
4. Although some problems still remain, e.g. free will,
5. We are seeing the beginning of a new, computational theology.

Each of these 5 points corresponds to one or two chapters of the thesis. Chapter 1 introduces the general notion of computer modelling and how it might be applied to religion by drawing analogies between computational concepts and religious ones to suggest consequences and non-consequences of religious positions, and hence debug some of the theological debate of the last two millenia.

Chapter 2 is 'related work'. It surveys the more important theological positions on a variety of 'problem' concepts, e.g. the holy trinity, miracles, free will, and points out the contradictions inherent in these positions.

Chapter 3 and 4 are the heart of the thesis. Chapter 3 draws an analogy between the trinity and trebly recursive functions, and uses this to resolve philosophical difficulties about God being both one and three entities, simultaneously.

Chapter 4 develops an extended analogy in which the universe is seen as a program for which God is the programmer, and in which miracles are seen as run time patches inserted during interruptions.

Chapter 5 is 'further work'. Outstanding problems are mentioned. There is a discussion of the problem of free will and possible computational accounts of it.

Chapter 6 is the conclusion. The results are summarised and the relative success of computational approaches to religious problems are summarised. The current work is seen as the humble beginnings of an important new approach to theology.

8. Guide to Reading

Staying in touch with related research is one of the main subgoals of obtaining a Ph.D. Some of the difficulties were raised at a Department of Artificial Intelligence 'research difficulties' meeting in the context of reading habits. Here is the relevant quote from the minutes of that meeting:

'Reading is difficult: The difficulty seems to depend on the stage of academic development. Initially it is hard to know what to read (many documents are unpublished), later reading becomes seductive and is used as an excuse to avoid research. Finally one lacks the time and patience to keep up with reading (and fears to find evidence that one's own work is second rate or that one is slipping behind).'

Clearly there are ways of staying in touch other than reading, but similar difficulties apply. One still has to maintain a proper balance between learning about other people's work and getting on with your own.

It may be helpful to think of the work of others as arranged in concentric circles around your own, where the relevance of the work decreases as you get further from the centre. For instance, if you were studying anaphoric reference, then the inner circles would consist of other work on anaphora; the middle circle would consist of work in natural language understanding and computational linguistics and the outer circle would contain other work in Informatics and linguistics. You can add extra circles to taste. Obviously, you can afford to spend less time keeping in touch with the work in the outer circle than that in the inner circle, so different study techniques are appropriate for the different circles.

8.1 Outer Circle

You can achieve an appropriate level of familiarity with the work in this circle by skimming papers or reading the abstracts. It is a good idea to set aside an hour each week for visiting the library (physically or electronically) to skim the latest arrivals. An alternative to skimming is attending conferences to listen to both the short presentations and the longer tutorial addresses. It is also very valuable to corner people in the coffee room or corridor and engage them in a short conversation about their latest ideas.

8.2 Middle Circle

Here you need to spend some more time. The methods described for outer circle are still applicable, but are not sufficient - you will also need to read some papers right through and engage in some longer conversations. You will want to read some more specialized textbooks and attend seminars etc. It is worthwhile keeping a record of papers you have read and some comments about them, otherwise the benefits derived from reading them will evaporate as your memory fades. It helps to write your literature survey in parallel with the rest of your research. Write a paragraph on each paper as you read it; this will save you re-reading it again when you come to write up your thesis.

8.3 Inner Circle

For a really deep understanding, reading a paper once is not sufficient. You should read it several times and get involved in it. Work through the examples. Set yourself some exercises. Get in touch with the author(s) about it. Talk or write to them with a list of queries and/or criticisms. One invaluable way to get a deep understanding of some work is to try to teach it to others. Offer a seminar, either formal or informal. You will need your own personal copy of papers you are making heavy use of. If you don't have one, photocopy someone else's.

When reading a paper you will find that you understand it better if you have a question in mind which you hope the paper will answer. The precise question will depend on the circumstances, but might be: What claim is being made? What is the evidence for this claim? Is it convincing? How can I use this work in my own research project? etc.

Finally don't be afraid to admit your ignorance by asking questions. Everybody feels sensitive about their areas of ignorance and in a field as multi-disciplinary as Informatics we all necessarily have wide areas of ignorance. People enjoy answering questions - it makes them feel important. You can usually get a far better feel for a piece of work by engaging in a discussion with someone who understands it than just by reading the paper alone.

9. The Examination of Theses

When you have written and rewritten your thesis to your supervisors satisfaction then you are ready to submit. Inform the College Office of your intention to submit. Make sure that your thesis is in accord with [regulation 41](#). Get two copies bound in the approved manner and take them to the College Office.

Your supervisor will suggest suitable internal and external examiners. They may consult you informally about the choice. The College will send your copies to the examiners. When the examiners are ready - and that could take several months - the internal examiner will arrange an oral examination or viva.

The viva is a question-answer session between you and your examiners, lasting several hours. Your supervisor may attend, as an observer, at the examiners' discretion. It will normally be in an office in the School; the external examiner (and possibly you) will travel up for the day. Dress is normal office wear and the occasion is fairly relaxed. Dress up a bit if it makes you feel more comfortable.

Before and after the viva the examiners have to submit reports to the College. The post-viva report is a joint one and contains a recommendation taken from [Regulation 3.10.4](#). Roughly speaking, the options are:

1. Accept the thesis as it stands.
2. Accept with minor alterations or with deficiencies that can be remedied without further study.
3. Accept the thesis, but not the oral, and examine the candidate further.
4. Reconsider after a further period of supervised study and resubmission.
5. Reconsider a Ph.D. as an M.Phil..
6. Reject.

You will usually be told the recommendation informally, with the understanding that it can be overturned by the College or Senate (and this is not unheard of).

Recommendations 1, 3 and 6 are very rare.

Recommendation 2 is to allow correction of errors that do not require further research. These can vary from minor typographical errors and spelling mistakes to major rewrites (there are actually two separate recommendations covering the two extremes). Typically, just the internal examiner will check that the thesis has been completely corrected and will then inform the College who will process your thesis and inform you of the outcome. This may take several months.

Recommendation 4 is to allow further research, usually requiring a major rewrite. You will have to rewrite, rebind and resubmit your thesis and go through the whole procedure again with the same examiners. This is your last chance. Recommendations 4 and 5 are not available the second time around.

Recommendation 5 is for theses which are not considered suitable for a Ph.D., but which are considered suitable for an M.Phil.. We are not supposed to say it is a consolation prize. You may or may not have to undertake further study and rewriting. You will have to get it rebound (in M.Phil. covers!), resubmit and have another viva.

The purpose of the viva is for the examiners to satisfy themselves that the thesis is acceptable as a Ph.D./M.Phil.. In particular, they will have raised various doubts in their pre-viva reports, which they must satisfy themselves about during the viva, and which they must discharge on the post-viva report. If they do not discharge these doubts in their post-viva reports then it is not unknown for the College to override their recommendations.

The examiners will ask you questions to try to satisfy their doubts. Because of time pressure, they often start with the most serious and/or most general questions. For instance, they might start by asking you to summarise in your own words what you consider to be the key contributions in the thesis. It is worth having a succinct answer ready to this one. You and your supervisors can try to anticipate other questions, but frequently the things you are most worried about have now been adequately covered in thesis, and the actual questions will surprise you. Thus it is better to have spent the previous night getting a good sleep, so that you are fresh and alert for the viva, than to have spent it rehearsing answers to questions that you will not be asked.

Do not ramble. Pay attention to the examiners questions and statements, and respond pertinently and succinctly. If the examiners can see that you are coherent, intelligent and aware of the issues in your field then they will be keen to award you your degree, and may be more prepared to overlook minor faults in the thesis.

Sitting a viva is a little like debugging a program. The thesis is the program, you are the programmer, the Ph.D./M.Phil. standards are the language syntax, and the examiners are the interpreter. During the viva you will get various error messages. These messages do not need to be taken at face value - they may be based on a misunderstanding - but they cannot be ignored. Assume that each error message will lead to some alteration in your thesis. Of course, you hope that this will only be a minor alteration, but do not let this hope blind you to the possibility that the problem is more fundamental. Do not get aggressive or defensive with your examiners. You cannot bludgeon or sweet-talk them into passing you, any more than you can force or persuade the computer to accept your buggy program. What you have to do is: clarify your own thinking, clear up any misunderstandings between you and your examiners, make sure you understand how to correct your thesis, and then correct it. The viva is a cooperative process. Your examiners want to pass you. Give them all the help they need.

10. References

- [Bligh72] Bligh D., *What's the use of lectures?*, Penguin, 1972
- [Booth75] Booth V., "Writing a scientific paper", in *Biochemical Society Transactions*, vol 3, 1975
- [Cooper64] Cooper B., *Writing Technical Reports*, Pelican, 1964.
- [Orwell68] Orwell G., *The Collected Essays*, Penguin, 1968.
- [Parsons73] Parsons C.J., *Theses and Project Work*, George Allen and Unwin, 1973.
- UK Research Integrity Office's [Code of Practice for Research](#)
- [Zobel 04] Zobel J., *Writing for Computer Science*, Springer, 2004.