

The Researchers' Bible

(How to survive a PhD)

by

Alan Bundy, Ben du Boulay, Jim Howe and Gordon Plotkin
University of Edinburgh

This version: February 17, 1995

Modified for Computer Science at TCD by Vance Gledhill, 1999

Contents

1	What is a Thesis?	4
2	Choosing a Research Project	5
3	Standard Pitfalls for Postgraduate Students	6
	3.1 Solving the World	6
	3.2 Manna from Heaven	6
	3.3 Computer Bum	6
	3.4 Yet Another Language	6
	3.5 Micawberism	7
	3.6 Ivory Tower	7
	3.7 Misunderstood Genius	7
	3.8 Lost in Abstraction	7
	3.9 Ambition Paralysis	8
	3.10 Methodology does not make a thesis	8
	3.11 The discovery route is not a justification	8
	3.12 Sidetracks	8
4	Psychological Hurdles	9
	4.1 Mental Attitude	9
	4.2 Research Impotence	9
	4.3 Dealing with Criticism	9
	4.4 Early Morning -- Cold Start	10
	4.5 Theorem Envy	10
	4.6 Fear of Exposure	10
5	Seminars	10
	5.1 Departmental	10
	5.2 Group/Unit Seminars	10
6	Research Methodology	11
	6.1 Stage 1	11
	6.2 Stage 2	11
	6.3 Stage 3	11
	6.4 Stage 4	11
	6.5 Stage 5	11
	6.6 Stage 6	11
	6.7 Stage 7	11

7	Writing Papers	12
	7.1 Departmental Research Papers	12
	7.2 Departmental Technical Papers	12
	7.3 Departmental Working Papers	12
	7.4 Personal Notes	12
	7.6 Publishing Papers	13
	7.7 Conference Proceedings	13
8	Guide to Writing	13
9	Guide to Reading	17
	9.1 Outer Circle	17
	9.2 Middle Circle	17
	9.3 Inner Circle	17
10	Progress	18
11	The Examination Process	18
	The Viva Voce	19

The Researchers' Bible

by

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

University of Edinburgh

This version: February 17, 1995

Modified for Computer Science at TCD by Vance Gledhill, 1999

Abstract

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

1 What is a Thesis?

To get a Ph.D. or MSc. you must write a thesis. A PhD also requires an oral examination (*viva voce*). The oral is generally used to ask for clarification of the thesis, so the main burden of assessment falls on the thesis.

The requirements for submission in Trinity College are contained in the University Calendar and are quite brief. We reproduce them below:

A thesis submitted for a Master's degree must show evidence of independent enquiry and/or originality in either conclusions or method.

A doctoral thesis must show evidence of independent enquiry, originality in the methods used and/or in the conclusions drawn and must make an appreciable new contribution to knowledge in the candidates field.

A thesis must be a candidate's own work.

A 'viva voce' examination is a normal requirement of the examination of a PhD candidate.

How original and significant must Ph.D. research be? The phrase 'make an appreciable new contribution to knowledge' suggests a simple rule of thumb. It should be possible to distil from the thesis, papers worthy of publication in a journal. This is not an infallible guide -- refereeing standards are not always what they should be. Your supervisor will help clarify this notion of 'new

knowledge', and a review of some past theses will also help - particularly if they are close to your planned topic. The final decision must rest with the examiners.

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

2 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 department willing to supervise it.

The importance of (1) cannot be overestimated. You are going to need all the enthusiasm you can raise 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. mathematical reasoning for mathematicians; natural language for linguists; complex systems for development, or hardware design. 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 do-ability (3) of anything you suggest. On the whole, beginning students tend to underestimate the worthiness and overestimate the do-ability 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 number 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 the 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 one finds (inevitably) that one has attempted more than is possible in the thesis.

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).

3 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 (and most of us have fallen for most of them).

3.1 Solving the World

Most students pick research goals which are far too ambitious. 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, which contained the germ of a good idea.

3.2 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:

- (a) Complete an initial electronic literature search in your chosen topic/field. Use your librarians to help in this search - they have been trained to assist in retrieving research publications and enjoy the professional challenge. Your librarian can be a major aid over the years of your study.
- (b) Use electronic sources - papers published on the WWW and through subscriptions to relevant News Groups.
- (c) Read the literature. Read with a question in mind, e.g. What is wrong with this? How can I use this? etc
- (d) Talk to people. Do not go away and hide. Do not be ashamed of your ideas. Other people's are sillier.
- (e) Tackle a simplified version of your problem. Ask your supervisor for exercises, mini-projects, etc.
- (f) 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 bugs are exposed.
- (g) Give a seminar to your research group. Same effect as (f).

3.3 Computer Bum

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 work 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 computer. Undertake some of Einstein's 'thought experiments'. If you must work some of it out at the keyboard 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. Try to explain your ideas to a non-computer person. The necessary simplification will expose the framework of your ideas.

3.4 Yet Another Language

A terminal case of ‘computer bum’ is to get involved in writing yet another programming language. Of course the existing languages do not offer exactly what you need for your project, but that is no excuse for writing another one. You can usually find a reasonable candidate and add to it what you need. No one will use the language you write -- not even you! You will have spent all your time on the language and none on the project you started with. If you really believe that existing languages are inadequate then write a paper on it, carefully describing the deficiencies. If you do a good enough demolition job, enhanced languages will spring up overnight.

3.5 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.

- (a) Talk to people, explain what you think your experiment might show.
- (b) Imagine the experiment finished and you have ‘the data’, what exactly would you do with it.
- (c) Not only try out the experiment on one or two people first, but also try out the analysis. Don’t keep running experiments in the hope that something will turn up.

3.6 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.

3.7 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. Try to rephrase your ideas using ordinary English; try to rephrase your ideas in someone else’s jargon. Do they come out any different?
- 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 shoulders and do the next easiest thing etc. -- a better infinite loop.
- Love of complexity. It is not a virtue to build a complicated solution -- it is just a nuisance to other people. Do it the simplest way you can. Occam was perfectly right. If you look back on scientific publications of 50 years ago you will see how complex ideas have been understood and simplified into school texts.

3.8 Ambition 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 worthy.

3.9 Methodology does not make a thesis

Since Computer Science is a relatively young field, and is somewhat interdisciplinary in nature, it does not have one received framework for research, or one well-defined methodology. One of the (difficult) tasks that you face as a 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 violent arguments with fellow researchers who are promoting alternative views. Out of this struggle, your reading, your attendance at seminars, your debugging, and other hard work, will emerge your world view on your area 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.

3.10 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. 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.

3.11 Sidetracks

You will be tempted from time to time to take diversions from your core work because you feel that you have made a significant discovery that will change the world. Beware of the sidetrack - it can consume months of fruitless distraction. On the other hand, these sidetracks sometimes become the basis for significant new directions in your work. Present these ideas to your supervisor or in a

seminar as it might be fruitful. Beware the sidetrack after the first 12 months and plan to pick it up when you have finished (note it in your day-book).

3.12 The Development Trap

Remember that research is not development, although development will be necessary in solving your problem. If you stray too far into the big D, you will probably be passed by industrial R&D and find your years of work will be released as a product before you start to write-up.

4 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.

4.1 Mental Attitude

Part of a researcher's skill includes an appropriate mental 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 2). 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.

4.2 Research Impotence

For many people, research prowess is yet another virility symbol. Lack of success at research is accompanied by the same feelings of inadequacy and impotence as post-mature virginity. Like sexual impotence, research impotence is a self-fulfilling prophecy. Doubts about your own ability can put you in a frame of mind where the dedication and enthusiasm necessary to produce results evaporates. (Edison said that genius was 1% inspiration and 99% perspiration, and he should know.) 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 pamphlet.

4.3 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.

4.4 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:

- 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.
- 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 or a diagram to draw.

4.5 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.

4.6 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.

5 Seminars

Research seminars are a vital part of your research training and a great source of ideas to validate your own work or to keep abreast of recent developments. If you are presenting, a seminar provides a goal to complete a unit of work.

5.1 Departmental Seminars

These provide opportunities for staff or visitors to present their most recent work. This should provide a model of how they go about exploring their field and containing a piece of work in a presentation. You will form your own views on the quality of different seminars and will use that as a basis for your own style when you are given the opportunity to make a major presentation as your

work develops. Notice that a good scientist or engineer is clear headed and avoids unnecessary jargon. Their seminars are simply constructed and you can feel the enthusiasm of the researcher.

5.2 Group or Unit Seminars

These are meant as working seminars with your colleagues within a research area. They should expose work in progress and to obtain constructive criticism of the work. You should seek to make a presentation every 6 weeks or so - it reminds you that you should be progressing and provides you with a useful sounding board for your ideas. Remember that no ideas are 'out of bounds' here and your peers can provide useful guidance on similar work or can constructively criticise your ideas or methods. You should also ensure that you contribute in discussion to these seminars where possible, with positive suggestions.

6 Research Methodology

There are many possible methodologies, but you should have one. Many start from different beginnings but merge later.

Here is one example of a methodology (contributed by AB).

6.1 Stage 1

Think of a scenario -- i.e. a sample output which would show that your computer program was exhibiting the ability you want it to model. In mathematical reasoning this scenario might be a proof; in natural language a sample dialogue; in vision the recognition of a scene, etc.

6.2 Stage 2

Hypothesise what processes might achieve such a scenario. Outline the procedures and data structures that might be involved. Try to make these as general as possible. See the problems you encounter as examples of general problems. Do not use ad hoc mechanisms except to overcome problems that are not central to the issue you are addressing.

6.3 Stage 3

Think of further scenarios. See whether your proposed program or system could cope with them. Use them to refine; generalize; extend and debug it.

6.4 Stage 4

When you are satisfied that your proposed solution or project is stable, choose the programming language that fits your needs closest and implement your program.

6.5 Stage 5

Do 'thought experiments' before subjecting your problem/solution to a program. Use a pencil and paper to see if your 'mental' solution might work before subjecting it to a coding solution.

6.6 Stage 6

Describe your program using language independent of your particular implementation. Try to draw out any new techniques. Compare them to previous techniques in your area. If time permits, apply

your technique to other areas. N.B. Stages 4 and 5 will take longer than you think -- years not months -- so leave plenty of time!

7 Writing Papers

Research papers are the major research product of a Department. 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 scientific 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.

When examining a PhD thesis, I (VG) have a simple rule of thumb that says if the candidate has published several papers in quality refereed journals, then much of my work has been done. The review process for a journal paper indicates that peers in the field view this work to be a significant contribution. My task as examiner is then largely concerned with style and to see that the thesis represents a unification of ideas. This, of itself, is a good reason to publish as you progress as a PhD student.

7.1 Departmental Research Papers

These are papers submitted for publication in a journal, conference or book. Papers may be promoted to this category from the other categories after submission. This arrangement encourages external publication of the Department's work and to provides advance copies of published papers to other researchers.

7.2 Departmental Technical Papers

These are formal internal papers which report on the conclusion of a piece of research

7.3 Departmental Working Papers

These are for descriptions of research work, either finished or in progress. You should not be afraid of presenting ideas that are not fully developed. Make writing up your ideas in a working paper a regular habit. If your ideas are off the planet then the sooner other people can see them and let you know the better for you.

7.4 Personal Notes

Each research student is encouraged to maintain a 'day book'; a diary of ideas or results. It should be a hard covered book always at your elbow in which you enter ideas, results or the outcome of interesting discussions. It is your personal record and you will be surprised how valuable it is when you come to write more formally.

7.5 Publishing Papers

Each paper should be discussed with your supervisor before it is published internally. When you and your supervisor think that you have something worth publishing externally you should submit a paper to a journal. Choose an appropriate journal and write in its style. In preparing the 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 research supervisor about these points.

You will submit several copies of the paper to the journal. Several referees chosen by the journal editor will vet these. Do not be too downhearted if it is rejected -- you will be in good company. Read the referees comments carefully. Are they right or have they misjudged you? Was your choice of journal appropriate? Consider submitting your paper elsewhere, but first take into account those criticisms you consider valid.

7.6 Conference Proceedings

Another form of publication is the proceedings of a conference. In Computer Science, some conference proceedings have equivalent status to papers published in journals. Conferences will sometimes 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

8 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. If the standard of writing for journals, conferences and theses is not high, then it will be a major cause of rejection in all three cases. 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 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 willful 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. You know yourself, when asked what a paper is about, you usually answer with a single idea (and a quick appraisal of its quality). 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 is the problem, what did you use to tackle it, what results followed (introduction, materials and methods, results, and discussion.)
- 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.
- 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, may be read by internal examiners and may not be read by external 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 department 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. Bottom up produces a dogs 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 magnetic) 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.
 1. It should consist of a few sentences, i.e. be of abstract length.
 2. The sentences should form the steps of an argument. This argument is the message of your thesis.
 3. 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:
 1. 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.
 2. 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.

3. It should answer questions like ‘What should be discussed in ‘related work’?’.
 4. 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, (AB) gives an example from the Ph.D. thesis of a famous AI researcher, Fr. Aloysius Hacker.

Example.

“The Computational Modelling of Religious Concepts”
by
Fr. Aloysius Hacker

1. We apply ideas from Computer Science 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 or 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.

9 Guide to Reading

Staying in touch with related research is one of the main subgoals of obtaining a Ph.D.

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 finding, 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 computational 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.

9.1 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 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.

9.2 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. You should keep a record of papers you have read in the middle and inner circles. Otherwise the benefits derived from reading them will evaporate as your memory fades. Establish a bibliographic computer file which contains the full reference (in the form that you will publish) plus your own abstract of the essential features in relation to your work. You will be able to do key word searches on your abstracts when writing up and include the references in your bibliography.

9.3 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 about it. Email 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: 'Can I use this technique in my program?'; 'How does he tackle the X problem?'; 'Is this acceptable as a journal article?'; 'How can I present this idea to my class?'

Students may wish to form a 'Journal Club' where the graduate students meet on a weekly/monthly basis and one of the group dissect a paper from some point of view. This builds critical/analytical skills, informs colleagues of the field and improves presentation techniques.

Finally don't be afraid to admit your ignorance by asking questions. Everybody feels sensitive about their 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 by reading the paper alone.

10 Progress

Impatience with yourself and your progress is healthy provided that it doesn't get your mind into a state of anxiety so that it actually impedes your work or encourages you to cut corners.

It is a useful exercise to set a plan and to establish milestones to be achieved. It is not a disaster to miss these milestones or change the plan - this is research after all. However the fact that you have a plan ensures that you have thought about the steps necessary to demonstrate or prove your thesis. Your supervisor may wish you to complete a 6 monthly report to mark progress. I (VG) have always found these a useful tool, but no more than that. A sample of a 6 monthly progress form is attached.

Trinity College requires that the progress of research students be "reviewed each year, and the supervisor and/or head of department may test the student in whatever way is appropriate. If unsatisfactory progress is not being made the student's registration should be terminated". The 6 monthly reports support the review process.

11 The Examination of Theses

When you have written and rewritten your thesis to your supervisors satisfaction then you are ready to submit. Inform the faculty office of your intention to submit. Make sure that your thesis is in accord with the guidelines given in the calendar. Get two copies bound in the approved manner and presented at the Graduate Studies Office, accompanied by one copy of the abstract (see University Calendar, Part 2) .

Your supervisor and head of department will choose suitable internal and external examiners. They may consult you informally about the choice. The faculty 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 voce*.

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 Department; 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 faculty. The post-viva report is a joint one and contains a recommended verdict. The verdict is roughly one of the following, according to a set list which may vary a little from time to time:

1. Accept the thesis as it stands.
2. Accept with minor alterations.
3. Refer the thesis for significant revision.
4. Consider the award of a lower degree.
5. Failure

(See University of Dublin Calendar, Part 2)

Verdicts 1 and 5 are very rare.

Verdict 2 is to allow correction of typographical errors, spelling mistakes, minor rewrites, etc. Corrections required by examiners must be made before reports are submitted to Council. Minor corrections should be completed within two months of the candidate being informed of what corrections are required.

Verdict 3 is to allow a major rewrite with or without further research. You will have to rewrite, rebind and resubmit your thesis and go through the whole procedure again. Revision should be completed within 6 months. This is your last chance. Verdicts 4 and 5 are not available the second time around.

Verdict 4 is for theses which are not considered suitable for a Ph.D., but which are considered suitable for an M.Sc.

The Viva Voce

The examiners may ask you questions to try and 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 question 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.Sc. 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 mis-understanding -- 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 sweettalk 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.

Progress Report

Name: _____

Date: _____

Achievements over the past 6 months:

Goals for the next 6 months:

Name 2 or 3 key papers that you have read in the past 6 months:

Supervisors Comments:

Progress: Good Satisfactory Fair Inadequate