

A Beginner's Guide to Research

David Vernon

Carnegie Mellon University Africa

www.vernon.eu

Revision 7.1
January 9, 2024

Table of Contents

1. Getting Started	1
2. Moving Along	4
3. Reading	7
4. Writing	9
4.1 Good Writing Discipline	9
4.2 Good Writing Style	9
5. Writing a Research Proposal	13
6. Writing a Literature Survey	15
7. Writing Scientific Papers	16
8. Ph.D. and M.Sc. Theses	21
9. Looking Forward	23
Acknowledgements	23
Further Reading	24
Appendix I Some Rules of English Usage	i
Appendix II Writing a Literature Survey	viii
Appendix III Typical Structure of a Thesis	x

1. Getting Started

Research is Difficult



Writing a thesis is hard work and getting started is perhaps the most difficult part. A research degree focuses on you becoming an expert in a particular topic and adding to the body of scientific and engineering knowledge on that topic. However, research is also concerned with learning, and especially with learning to work independently and being able to develop your own understanding of any given topic. Starting with very little, where do you begin? The following are some pointers on how to get started.

Understand the Problem

Your supervisor will have provided you with an outline description of your research topic. To begin with, and to test your understanding of this outline, you should try to expand on the problem statement. In your own words, try describing informally exactly what the final system should do, what data it will take as input, what data it will produce as output, and how these input and output processes are accomplished. Next, try to say exactly how the input data is transformed in order to produce the required output. Describe why this problem is relevant. Say why it is important to solve it. What are the consequences of finding a solution? What does solving this problem mean?

Start Reading

Your supervisor will suggest some initial reading: journal papers, conference papers, book chapters. Read them all.

Start Writing



Arm yourself with a pen and write a short synopsis of every paper and book chapter you read. Write down the key message (one or two sentences) and the main contribution (one or two paragraphs). It is also worth writing down one or two quotes from the paper if they provide some important insight into the topic. This is hard work. Don't underestimate it but do be aware of how important it is to do it. The very act of writing helps crystallize ideas. Remember: summarize every paper you read. Don't be tempted to copy the paper abstract: the point of the exercise is for you to express your understanding of the paper in your own words. This is a necessary part of the learning experience. You won't learn if you don't write it down in your own words.

Start Talking

Talk to your supervisor: ask questions, try to explain your ideas, tell him what you are finding hard (and easy). Talk to other students: explain to them what the goal of your work is, what the central problem is, how you are going to try to solve it, how other people have tried to solve it, what difficulties they encountered, why their approach isn't as good as the one you are trying to use. This is hard, but it gets easier with practice. Now, try to write down these ideas. Use a pen and paper first. Type the notes later. Give a short presentation to your group.

Start Listening



Go to seminars. Go to talks being given by other students. Go to all group meetings. Research is a social activity. There are exceptions, of course: Andrew Wiles spent over seven years working on Fermat's Last Theorem without telling anyone ... but you aren't Andrew Wiles! Even so, it's worth reading Simon Singh's account of the Wiles's dedication and commitment on the web.¹ Every Ph.D. and M.Sc. student needs a little of this type of dedication.

Start Simple

Carve up the problem. Is there a simpler version you can tackle? If so, do that first, and then move on to the more complicated version later. In other words, try to structure your research goals so that there is (a) an essential but fairly easily achieved goal, (b) a desirable but somewhat harder or risky goal, and (c) the ideal goal, something that would represent a real breakthrough. This gives you both a way of structuring your work and a safety net if everything doesn't go as expected (and it won't).

Work Hard



Remember: genius is 1% inspiration and 99% perspiration. You can't get a Ph.D. or M.Sc. without a lot of hard work (early mornings, late nights, frustration, fatigue, and sometimes depression). That's normal! You are not alone. It is hard work but it's worth it.

Formalize

As early as possible, you need to try to formalize the problem you are working on. You won't be able to do this straight away because you will need to get a grounding in the background theory first. But you should at least make an attempt, so that, after having

¹ <http://www.simonsingh.net/books/fermats-last-theorem/who-is-andrew-wiles/>

read around the topic, the next time you try you will see how much progress you have made.

Learn About Tools



As you study for your degree and pursue your research, you will become an expert in two different domains: the problem domain and the solution domain. The problem domain refers to the theoretical issues by which we can model and solve the problem. This is the primary focus of research. At the same time, you also need to know how to handle the tools that will enable you to implement the solution. These might include configuration tools, compilers, libraries, programming languages, operating systems, application programming interface (API) definitions, class libraries, toolkits, networking infrastructures and protocols, hardware platforms, input devices (data acquisition), output devices (control systems). These tools belong to the solution domain and they give you an opportunity to learn other things in parallel to your formal research in the problem domain. Typically, they are more straightforward to deal with so will give you a chance to feel you are making some progress early on.

2. Moving Along

Make Notes



The best way to organize your research is to keep a logbook of all work in progress. Get a notebook and write everything you do on the project into this logbook every day. Every thought and observation you have on your project should go into this book, along with notes of meetings with your supervisor, results, theoretical developments, calculations, references, anything that is relevant to your work. Again: the act of writing helps crystalize ideas.

Believe In Your Own Ability

Because there is so much uncertainty associated with research, you need to have a lot of self-belief to do it. At the outset, there is a lot you don't know. You don't understand the problem, you don't know how other people have approached it, you don't know how successful they were, you don't know if you will be able to solve it yourself, and you don't know how long it will take. All that is certain is that you have to spend many hours every day, often with no reward, chipping away at the problem, hoping for progress. To get through this, you need to believe in your own ability. You need to constantly remind yourself that you can do it. It will take hard work, determination, and sometimes a little help from others, but you can do it.

Believe In Your Problem

You must also believe that the topic you are working on is worthwhile. If you are not convinced that solving this problem matters, don't even think about working on it. Find another topic. You should also enjoy working on it: it should represent a challenge that motivates you to get out of bed each morning (well, maybe not every morning, but most mornings!) There is no substitute for enthusiasm in a research degree, and you will need lots of it to get you through. So, if you find yourself working on something that simply does not fire your imagination, consider switching topic. Talk to your supervisor and do it as soon as possible. Don't delay.

Practice, Practice, Practice

Research is a skill and can be learned. However, just like every other skill, you need to practice it to get better. You can't improve just by thinking about it. Research is not a spectator event: it is a contact sport.

Framing the Research Question

Don't be afraid to re-scope your research as you learn more about it. Your supervisor will help guide you in this. If it turns out that the problem is much more difficult than it seemed at first, it needs to be redefined. A Ph.D. or M.Sc. thesis typically will ask and attempt to answer a hard question, what is commonly referred to as the *research question*. Make sure it answers convincingly. If it doesn't, either you need to generate a better answer or ask a slightly different question.

Jump Start Your Day



Getting started first thing in the morning is often difficult so, when you stop the night before, set yourself a task for the morning that is easy and even enjoyable. For example, you might update your bibliography, comment some code, or draw a diagram. That way, you will reduce the temptation to avoid starting work early by doing other 'important' tasks such as checking email or browsing the web. Take a lead from Ernest Hemmingway:

"Mr. Hemingway works on a strict schedule that produces an average of 500 to 1,000 words a day. 'I start in at seven in the morning,' he says, 'and I always quit when I'm going good, so that I'll be able to pick right up again the next day.'"²

Dealing with Criticism

"Remember, though, if we can't criticize ourselves, someone else will save us the trouble."

Drew McDermott, Artificial Intelligence Meets Natural Stupidity, ACM SIGART Newsletter, No 57, pp. 4-9, April 1976.

You will inevitably have to take a lot of criticism during your time as a research student. Being wrong and being ignorant³ goes with the territory. Remember too that you can't learn without mistakes: learning means improving, and improvement implies that you start out being less capable. Criticism is just a way of pointing this out so don't take offence. Criticism is a positive act, not a negative one. Criticism allows you to learn and to improve and that's what you are here for. Most criticism will be valid, but some may not be. Learn to trust your own judgment. If you think the criticism is invalid, say so, but be prepared to defend your position cogently, rationally, and quietly. Most invalid criticism is caused by a misunderstanding. This means you may not have explained your

² From "On the Books", Roger Bourne Linscott, 1946, reprinted in *Conversations with Ernest Hemingway*, Matthew J. Bruccoli, Ed., University of Mississippi, 1986.

³ Ignorant does not mean stupid. It means lacking knowledge.

ideas clearly in the first place. Try to explain them better and then see if the same criticism is forthcoming.

Don't Give Up



Even when everything looks black and your problem looks impossible, keep on going. However, give yourself lots of mini-holidays: take ten-minute breaks away from the daily grind (that's ten minutes, not ten days!). Don't keep thinking about a problem all the time.

Be Prepared For Inspiration

Frequently, the solution you are looking for will come to you when you least expect it and often when you are not thinking about the problem. So always keep pen and paper handy, just in case!


Get Yourself a Theory!

The great power of science and engineering is that it allows us to predict how systems will behave. To be able to predict something, though, you must have a model: an abstract formulation. This is how science and engineering works. Mathematics is the bed-rock of engineering so if at all possible try to describe your problem in mathematical terms. If this isn't possible, at least try to formalize the problem and use rigorous arguments: identify your axioms and state your hypothesis clearly and in a manner that can be refuted (falsified is the more correct term). You can never prove a hypothesis (only a theorem) but any worthwhile hypothesis will have a clearly stated test that will show it is incorrect.

Get Yourself a Benchmark

The hallmark of good engineering is to assess the system's performance and compare it to that of other similar systems. Ideally, you should identify some quantitative metric by which to compare the systems, since numbers are the best and perhaps the only way to objectively describe performance. Quite often, we use statistical measures as our comparative metric, e.g. the mean and standard deviation of some performance measure when the system is subjected to a large variety of input parameters and conditions.

Check & Compare



3. Reading

Recursive Reading



Your supervisor will have given you an initial list of papers to read. Get copies and read them. They will refer to other papers. Get the main papers cited and read these. This is a recursive procedure. Eventually, you will begin to see closure and you won't be encountering many new papers. At that point, you have covered the topic in depth.

Different Levels of Reading

Not all the articles you collect will be equally relevant or important. Consequently, it's not efficient to give each of them the same attention. But it is not easy to know how relevant it is until you read it. So how do you proceed? To solve this dilemma, you should know that there are three levels of reading: shallow reading, focused reading, and deep reading.

Shallow Reading

Some papers will be mainly background material providing the context to the research. The topics are relevant to your work but not directly related. Often it is sufficient to read just the abstract, introduction, and conclusion, and then write a one or two line summary of the main issue addressed.

Focussed Reading

Other papers are directly relevant and provide, for example, alternative ideas on the topic of your research. These papers should be read thoroughly, perhaps twice. You should aim to write a one-paragraph summary after you have read it. These papers will typically go in your literature survey.

Deep Reading

Finally, some papers, typically 10 to 15 for an M.Sc. and 15-30 for a Ph.D., will be absolutely central to your work and you will need to read them several times to really understand them. You may need to work through some examples and you will probably have to refer to other papers or texts to understand some of the concepts described. You should write a 1-2 page summary of the content of these papers. If they contain mathematical results, you should include these too, explaining each term and the importance of the results. After many careful readings, you should know as much

about the topic as the author. A good test of your understanding of a paper is to see if you can give a short presentation on it and explain it to other people in your group.

Build a Bibliography



For everything you read, insert the full citation of the paper or book chapter in your bibliography so that you can refer to it in your subsequent writing.

Make sure you keep a full citation index, i.e., you must record exactly where every article you copy comes from. Typically, you need to record the title of the article, the authors, the name of the magazine/journal/book, the volume and number of the journal or magazine, and the page numbers. If it's a chapter in a book and the author of the chapter is different from the editor of the book, you need to record both sets of names.

Variety

Note that the 'reading in' phase of the project can last quite a long time (there's a lot of reading *and writing* to be done) and it can help to overlap with some of the other early tasks, such as learning about the solution domain.

Reading Means Writing

Finally, it is very important that you realize that, in order to fully understand anything that you read, you must write it up in your own words. If you can't express or speak about a given idea, then you haven't truly understood it in any useful way. This is so important that it's worth saying a third time:

Writing is an Essential Part of Understanding

This brings us nicely to the next section.

4. Writing

Writing is not easy. There are two aspects to this: the discipline of writing and the style of writing. Discipline refers to the physical act of writing: the process of assembling ideas and getting them down on paper. Style refers to elegance and simplicity: the power of your writing to communicate an idea.

4.1 Good Writing Discipline

- Keep records. Write notes on papers you read, notes on software you develop, notes on ideas you have, notes on tests you run. Writing these notes serves several purposes. It helps to crystallize ideas and very often clarify them. It helps the learning process and it makes sure you don't lose or forget any important points. It also acts as a basis for subsequent writing: for reports and papers.

- Make writing a way of life. You should allocate a large proportion of your day to writing. Writing should be an integral part of your working day.

- Use pen and paper. Write things down long-hand. Later on, write these notes up more neatly and in a more organized fashion. Once you get good at this, you can go straight from long-hand notes to typed document, but it's very helpful at the beginning to first create an intermediate long-hand version.

4.2 Good Writing Style[†]

Effective writing is difficult. It takes practice and a willingness to revise your work, many times. One of the best ways to learn how to write well is to read. Almost any reading material will do, as long as it's well written. For technical writing, you should at least read several previous theses, conference papers, journal papers, magazine articles. The popular scientific press, e.g., *Scientific American* or *New Scientist*, employs a particularly simple and effective form of written expression. Try to emulate their style.

Why is writing well so difficult? As we have already noted, the goal of writing is to convey a message to the reader. Writing and reading are sequential processes. Therefore, you have to

[†] Several of these guidelines are adapted from Strunk and White 1979.

construct the meaning of your message, piece by piece, in a linear time-line. However, the meaning you intend to convey may emerge from many sources, not all related in a nice orderly fashion. This creates a problem for the writer: how to order the messages contained in each sentence effectively. The job of a writer is to make the sequence of pieces as mutually-relevant as possible so that the story or message builds naturally, each piece adding to the previous one. When the pieces are presented out of order (*e.g.* parenthetical expressions in the wrong place, or two related sentences split by a third) the impact on the reader is lessened. And remember the golden rules of good writing: keep to the point and keep it simple.

The following are some pointers to help you in your task.

- Use short sentences and make sure the sentences are complete.

A complete sentence has a subject followed (usually) by a verb, and then an object. For example: “Cognitive systems can adapt to changes in the environment.” Of course, we can add other words to enhance the descriptiveness and richness of the sentence: “Cognitive systems can adapt autonomously to unexpected changes in the environment”, and we can even include additional phrases (which will normally have a subject-verb-object structure of their own). For example: “Cognitive systems can adapt autonomously to unexpected changes in the environment, especially those that the designer did not anticipate.”

Remember that, if you remove all of the extra supporting words, you should be left with a valid sentence. It’s a good idea to check all your sentences this way.

- Good writing strikes a balance between short sentences and longer more descriptive ones. Just as in spoken communication, the full stops mean pauses: too many pauses and the text sounds disconnected, too few and it can be hard to follow the story line. Strike a balance but favour brevity over complexity.

- ❑ Make the paragraph your unit of construction. Each paragraph should bind one or more sentences about a specific subject or idea. If the subject or idea changes, start a new paragraph.
- ❑ Omit unnecessary words. They distract the reader. Don't write "This is a system the performance of which is very adaptive". Instead, write "This is an adaptive system".
- ❑ Write in a way that comes naturally. Speak the sentence. If it sounds correct, trust your ear, and use the sentence. If it sounds unnatural, rewrite it.
- ❑ Avoid fancy words; they don't impress anyone.
- ❑ Be clear in your expression. Write clearly. Write concisely. Write with a purpose. You have a message to convey so make sure you know what it is and keep it in focus throughout. Don't stray from the key point you are making. Present your argument in a structured manner so that what comes first depends as little as possible on what comes later, and so that what comes later builds on and adds to what you have just stated. If the idea you are trying to convey is getting lost in a sea of words and phrases, draw a line through the sentence and start again.
- ❑ Don't take short-cuts. Explain what you mean. Don't leave the reader to struggle trying to figure out the real meaning of your carefully constructed but complicated sentence. He or she may conclude there is none. Explain all acronyms the first time you use them.
- ❑ Use consistent names. Refer to each significant entity (algorithm, concept, device) using the same word everywhere.
- ❑ To distinguish one-to-one relationships from n-to-m relationships, refer to each entity in the singular, not the plural.

- ❑ Let information flow. In each sentence, move your reader from familiar information to new information.
- ❑ Place material you want to emphasize at the end of the sentence.
- ❑ Be brief. This isn't easy. Blaise Pascal, the French mathematician, once remarked: "I have made this [letter] longer, because I have not had the time to make it shorter".⁴
- ❑ If you use pictures and diagrams (and you should), make sure each one has a self-contained explanatory caption. Never refer to a picture or diagram in the main text without saying what it is. For example, never say "Figure 2.3 shows the results of the noise test" and then carry on to another topic. Help the reader. Summarize the content of the figure in a short sentence: "Figure 2.3 shows the results of the noise test. These results demonstrate the robustness of the system to Gaussian noise with a standard deviation of 2.3 or less." If you have copied the figure from a book or article you must cite the source.
- ❑ Don't be afraid to write a far-from-perfect first draft; its value is that it provides you a place to start and the basis for improvement.
- ❑ Revise and rewrite. It is highly unlikely that you will manage to find the best way of expressing an idea with your first attempt. Nonetheless, make that attempt and then be prepared to revise it, repeatedly.



Remember, if you want to learn how to write a good thesis or paper, you need to do two things: you need to read other good theses or papers and you need to practise your own writing.

Appendix I provides further details on some of the principal rules of English usage.

⁴ Blaise Pascal, *Lettres Provinciales*, Lettre XVI, 1657.

5 Writing a Research Proposal

The Research Question

In Section 2, we said that a Ph.D. or M.Sc. thesis typically asks and attempts to answer a research question. A good thesis answers it convincingly. This research question is typically one element of what every researcher, from aspiring graduate student to seasoned professional, will write at some point: a research proposal.

The Goal

The goal of a research proposal is to identify the problem that you wish to solve, to state clearly why this problem is important, and to highlight the reasons why finding a solution is challenging. The research question asks: how can we overcome these challenges? A research thesis provides the answer, in full or in part.

Not Just an Interesting Idea

It is not enough just to ask the research question. A convincing research proposal also demonstrates you that you understand the extent of the challenge. The best way to do this is to include a brief survey of the approaches that others have taken when attempting to solve your chosen problem, possibly identifying weaknesses in these approaches. You need to demonstrate that you are conversant with the field of research. It doesn't have to be a detailed literature survey (see Section 6 Writing a Literature Survey) but it must show some familiarity with topic. Ideally, the proposal should identify a plausible solution strategy.

The Heilmeier Catechism

The questions in The Heilmeier Catechism⁵ capture the key points addressed by a good research proposal.

1. "What are you trying to do? Articulate your objectives using absolutely no jargon."
2. "How is it done today, and what are the limits of current practice?"
3. "What is new in your approach and why do you think it will be successful?"
4. "Who cares? If you are successful, what difference will it make?"
5. "What are the mid-term and final "exams" to check for success?"

⁵ The Heilmeier Catechism, named after George H. Heilmeier, a former director of DARPA (1975-1977), is a set of questions crafted to help DARPA officials evaluate proposed research programs; see <https://www.darpa.mil/work-with-us/heilmeier-catechism>.

Proposal Structure

We can recast these questions a little differently under the following headings.

Goals of the Research Project.

Review of Current Approaches.

Proposed Solution.

Novelty and Significance.

Anticipated Results.

Evaluation Strategy and Metrics of Success.

It can help the reader to adapt these titles to include some text that makes them specific to your research project.

Critical Analysis

As you will appreciate by now, a good research proposal requires some preliminary research and a good deal of critical analysis. It is not just a suggestion for an avenue of enquiry, no matter how interesting or exciting it might appear.

Brevity, Brevity, Brevity

A good research proposal is short. This is not good news: writing a concise, clear, compelling proposal takes a lot of time and effort, with many rewrites before it is done. Writing a short one takes even more effort. If you are not convinced of this, re-read Section 4.

To Summarize ...

A research proposal is not a statement of interest in studying some topic. It is a clear statement of a specific problem, an explanation of its importance, and a description of a plausible solution, an indication of anticipated results, and an explanation of the novelty and significance of the results. The description does not have to be complete: the research program is the exercise in completing it, but it must have sufficient detail to be convincing, its argument must be compelling, and the proposed strategy must be plausible. Together, they must make a case that the research program is likely to lead to results that represent an improvement over those achieved by others who have tackled the problem.

6 Writing a Literature Survey

Survey the Field



One of the hallmarks of good research, and an essential element of a research thesis, is an understanding of how others have approached the problem you are tackling. You need to develop a deep understanding of the theoretical basis of their techniques, the assumptions they make, and, to a lesser extent, the tools and methodologies they use. The combined understanding of the nature of a problem and the spectrum of possible approaches is encapsulated in what is known as the literature survey. This is a well-structured synthesis that collects together all the relevant ideas, organizes them, and presents each of them in turn, highlighting their strengths and weaknesses.

It is not trivial to create this synthesis. It will require many attempts and many re-writes. You must start with your own short summary of each paper you have read. Then you try to organize the ideas, identifying different classes of topics and relating them together in the form of a taxonomy, i.e., a hierarchical classification tree.

This taxonomy provides you with a way to structure the literature survey, typically by effecting a breadth-first traversal of the taxonomy tree, using the material in the cited papers as examples. If this doesn't alarm you, it should. But don't worry: Appendix II provides more detail on the process involved in conducting and writing a literature survey.

Why is the literature survey so important? There are three reasons. First, it provides the essential background for your thesis: it allows you to critically appraise the state of the art in the field you are conducting research, thereby establish your mastery of the field. Second, it allows you to identify the gap in existing knowledge that is encapsulated in your research question (refer back to Section 2 Moving Along: Framing the Research Question) and make a compelling argument that this gap actually exists. Without this, you could be wasting your time by trying to fill a non-existent gap or a gap that the research community has decided is not worth filling. Third, it helps you identify the tools (mathematical, analytical, software, data collection, data presentation, and style of argument) that you will need to use to answer your research question.

7 Writing Scientific Papers

The Importance of Writing Papers



The ultimate test of M.Sc. or Ph.D. research is whether or not it is worthy of publication in a journal or in the proceedings of a conference. Consequently, research papers are the primary output of a research degree, not software. Software may be needed to validate the ideas in the paper and software may be the whole point of developing the idea, but the contribution to knowledge is the idea itself, encapsulated in a paper. Indeed, it should be possible to re-write or re-generate the software based on the information contained in the paper.

Know Your Reader

Assume the person reading your paper is intelligent but not knowledgeable (i.e., they don't know everything about the topic so part of the role of the paper is to bring them up to speed quickly). Assume the person reading your paper misunderstands things easily so make sure the argument is really clear. Make it easy for them to say: 'nice idea, good model, great validation; yes, the community would like to know about this'. If someone doesn't understand your paper, assume it is your fault, not theirs. Find out where they got lost and improve it.

Structure

A paper should begin with a statement describing a claim or hypothesis. It should then provide an argument to support that claim or hypothesis. This will require you to provide the context for the claim (e.g. other people's work and alternative approaches), to state its relevance or importance, to offer a model of the subject you are investigating, and to provide either some theoretical or empirical evidence that the model is valid. You should also provide an assessment of how well it works.

The Thread of an Argument

The way you present your ideas in your paper (or a chapter in your thesis) is critical to the impact you make in your writing. Remember that you are trying to convey a message to the reader and, since this message is likely to be quite complicated, you must assist the reader by making your points clearly and in a logical order. As you write, you reveal increasingly more of your message, but doing so in a way which builds on what you have already said. This is what we mean by a logical structure: breaking up the message into a linear sequence of messages which follow naturally one from the other and

which lead to a definite and clear conclusion. The golden rule is to construct your message incrementally, piece by piece, building on ideas you have already developed. Ideally, you build on ideas you have just introduced in the previous sentence. If you need to build on ideas that were introduced in the previous paragraph or a few pages back, you will have to provide the reader with a short reminder of what these ideas are. We call this type of incrementally-constructed message the thread of an argument. You need to keep the thread as cohesive as possible: threading sequentially from one idea to another, each following logically from the preceding one. Making repeated reference to ideas that were introduced much earlier (or, worse, not at all) leads to a tangle that will result in a very confused reader. Similarly, introducing a new idea without warning only serves to distract and further confuse the reader. Finally, trying to weave several threads together at the same time will cause the reader great difficulty unless it is done very skillfully.

Don't Start at the Beginning

Although papers are intended to be read from beginning to end, and the argument should flow linearly from beginning to end, it isn't always best to write the paper in that order. Often, it is a good idea to start by describing the technique and write the introduction later, once you've established the core message, and possibly after you have drawn your conclusions. The abstract should always be written last, once everything is clear and in place.

Do Your Best and then Improve It

Be prepared to write, and re-write, many times. It can take up to ten attempts just to get a good first draft of a paper. Once you do have a good draft, ask other people to read it. Do not ask people to read early drafts: it's impolite and says to them "I think my time is more valuable than yours so I'm not going to bother giving you my best attempt". Give them your best and thank them for their time. Under no circumstances use your supervisor as a proof-reader to correct mistakes or get hints on how to improve structure before you are certain that the paper can't be improved! You will probably be wrong, but that should be your goal.

Provide Results

Qualitative results are usually not as convincing as quantitative results so it's a good idea to try to identify a metric for the performance of your technique or system and then measure how well it performs using this metric. One metric is good, more than

one is better. Ideally, compare your system to others using the same metric. Sometimes this will mean re-implementing other people's work, but increasingly today you can get open-source versions of standard approaches to save you this effort. These provide an excellent benchmark against which to judge the value of your own contribution.

Citations

In a research paper (and a research thesis), everything you say must be substantiated. There are two ways you can substantiate something: one is to provide a citation of a publication which corroborates or supports the statement you are making, the second is to provide either theoretical or empirical evidence to support it. It is important to understand why citations are so important. The key issue is that they are references to published material which, therefore, has been subject to some form of peer review (usually rigorous, occasionally not). This means that the statements or claims have been judged to be legitimate by a group of people, not just the author. This elevates the statement from being mere personal opinion of the author to some level of mutually-agreed knowledge. Avoid unsubstantiated claims or statements in your thesis. Either provide a citation⁶ or provide evidence. If you can't provide either, then provide a compelling argument in support of your claim.

There are other reasons to cite published material in a research paper or a research thesis. One is to acknowledge the source of your ideas: even the most original thoughts are based on the work of others. Another is to demonstrate that you have done the research and are familiar with the literature in your area.

Don't cite a reference unless you have read it! You should never just copy the citations you find in someone else's papers to support your statement without first consulting them to make sure they say what the author claims they say. You would be surprised how often they don't.

Quotations

Never copy other people's writing, even if you change it slightly. You must re-express it in your own words. If you must use someone else's writing, put it in quotation marks "...". And, of

⁶ Sometimes one citation is not enough. The more the evidence in the citation veers from fact into opinion, the more citations you will need.

course, you should cite the source of the quotation.

If the grammar or spelling in the quotation is wrong, don't correct it. Quote it exactly as it is but put *[sic]* after the error; e.g. "How r *[sic]* you?". In general, you should change nothing in the quotation. There are two exceptions.

The first is where you replace a contiguous group of words with ellipsis. For example, instead of "computational attention, in its most general form, is a pre-requisite for action selection" you could write "computational attention ... is a pre-requisite for action selection".

The second is where you insert a word to help its comprehensibility. You can do this by inserting the word in square brackets [and]. For example, instead of "the lowest level is the most difficult to implement" you could write "the lowest level [of the stack] is the most difficult to implement".

Reviews



Journal papers, and most conference papers, are subject to peer review. This means that two to four referees read the paper and provide a critical assessment of its contents. These reviewers are normally experts in the field. You need to convince them that there is merit in your work. This is hard: they are very busy people and don't have the patience and motivation to read poorly written text, no matter how important the issue. Make it easy for them.

Submitting a paper for review can be a terrifying experience. You are asking for the reviewers' approval but there is a strong possibility of rejection. Nobody likes rejections. It is one of the things that human-beings fear most in life. If the paper is rejected, accept the rejection gracefully and learn from it. Was it rejected because you explained things badly? Was it because there was a flaw in your argument? Was there something wrong with the way you stated the problem, with your theoretical development, with the model? Did you provide enough evidence of the validity of the model, e.g., by providing quantitative tests? How did you establish the robustness, generality, or limitations of the technique?

Review your paper with these criteria in mind before submitting it.

**Choose Your Forum
Carefully**

Quite often, journals and conferences have different standards, with journals requiring a more substantial contribution to knowledge. This is not always the case: it can be just as difficult to publish in some top-flight conferences as it is to publish in a journal. In any case, choose your conference and journal carefully. Make sure there is a good match between the subject matter of a journal or conference and the topic of your paper.

Try, Try, Try Again

Get used to rejection. If your paper is rejected, don't give up. Accept the criticism, take on board the reviewers' comments, and improve the paper. Then submit it somewhere else.



8. Ph.D. and M.Sc. Theses

Length In general, a Ph.D. thesis will be 60,000 to 80,000 words in length, spread over six to eight chapters. In addition, there will be a bibliography and appendices. An M.Sc. thesis will be shorter; typically 60% to 75% of a Ph.D. thesis. Note that these figures represent typical upper bounds, not targets. It may well be that you can express your thesis more succinctly (and you should certainly try).

Structure Because all research projects involve the exercise of a fairly standard research methodology, you can, if you wish, adopt a relatively standard structure. Note, however, that this does not mean that you just simply have to fill in the gaps in a general thesis template: a standard structure simply provides you with a place to start as you begin to design the final structure and content of your thesis. You will still have to do quite a lot of work to make it fit your own project. A typical outline of a research thesis is provided in Appendix III.

You should design your own thesis to the level of detail given in the proposed structure, adapting it to your own particular needs. Note that you should do this *before* you start writing anything. It's just like designing a piece of software: the sooner you start coding, the longer it will take (and the more likely it is to be wrong). Try to achieve modularity and independence amongst your chapters and sections. At the same time, remember you are trying to convey a convincing message to the reader. Again, it's like telling a good story: you have to keep the reader interested and he has to be able to follow the storyline (which means there has to be one: make sure there is). Keep the thread of the story going continuously, from section to section, and from chapter to chapter by providing link sentences or paragraphs. At the end of a chapter, for example, remind the reader of the important messages, tell her or him why these messages are important, and then say what you need to look at next, and why, in order to continue with the story. That's your cue for the next chapter.

Standards The research described in a Ph.D. thesis should be of sufficient quality and depth to allow you to write a paper that would be

accepted for publication in a journal. This should be your target. The research described in an M.Sc. thesis should have a good chance of being accepted for a relevant conference.

It is very easy to pick a research goal that is too ambitious. Many students do this. One of the most important jobs of your supervisor is to help you avoid this danger so check and listen to the advice. Remember: a Ph.D. degree is not a Nobel Prize. Your work has to be good; it doesn't have to be revolutionary or world-changing.

Your thesis must clearly demonstrate your ability to *assimilate*, *synthesize*, and *critically appraise* all material relevant to your research project.

Your main opportunity to display your talents at assimilation and synthesis comes when you describe the background material you read, the requirements, specification, and design phases of your work. Synthesis means that you must write the text yourself, expressing your understanding of the material in your own words. Resist the temptation, no matter how strong, to copy sentences or paragraphs (or whole sections) from other books or articles. Copying is not synthesis and it demonstrates neither your assimilation of material nor your understanding of it. If you do come across a sentence or paragraph which is so good that it just has to be used, then use it and do so as a direct quotation, providing a reference to the original source.

It is extremely important that you also assess your own work critically, *i.e.*, with objectivity and with a view to seeing how it could be improved. Such honest criticism does not mean you will be judged harshly; in fact, it is likely that you will be given credit for your openness and transparency. Typically, you exercise your talents of critical appraisal at the end of the thesis in a discussion chapter. However, you can also exercise it throughout the thesis wherever it seems appropriate. Note that this exercise of critical appraisal is different from the testing processes of verification, validation, and evaluation, which refer to the functionality of the system you have designed. The critique you write applies less to the system and more to the overall objectives, methodologies, and findings of the thesis.

9. Looking Forward

Congratulations on having completed the guide. At this point, you might be wondering if all this research is worth it. Why bother? Why do all this work? Here's why. Along the path of a research degree, you grow. You become able to do things – hard things – that you could only dream of doing before: developing a new model or algorithm of your own, learning how to master a new technique, seeing simplicity in a complex equation, forming your own view and being able to defend its validity. But these are the little rewards that accompany the process. The big reward comes after the degree is complete and after the papers have been published. This is when you realize that you yourself have changed and that you now have the ability to tackle more or less any problem, with complete confidence. The unknown becomes a challenge, instead of an obstacle, and you are now equipped to meet this challenge: to overcome the unknown and create knowledge. Acquiring this ability is the true reward of a research degree.



Acknowledgements

Much of the advice on research methodology is based on [Bundy *et al.* 2004] while the pointers on good writing have their roots in [Strunk and White 2000].

Thanks go to Nicholas Moore for his many helpful comments, including the suggestion that the text would benefit from some relevant illustrations and that it needed a more encouraging conclusion.

Thanks also go to Tim Brown for his useful suggestions.

The illustrations were sourced at www.free-graphics.com.

Further Reading

S. Marshall & N. Green. "Your PhD Companion", How to Books Ltd., Oxford, UK, 2007.

P. Dunleavy, *Authoring a PhD: How to Plan, Write and Finish a Doctoral Thesis or Dissertation*, Palgrave Macmillan, 2003.

A. Bundy, B. du Boulay, J. Howe, G. Plotkin, *The Researcher's Bible*, Version 2004, <http://homepages.inf.ed.ac.uk/bundy/how-tos/resbible.html>

A. Sloman. "Notes on Presenting Theses", Version May 2004, <http://www.cs.bham.ac.uk/research/projects/poplog/teach/theses>

W. Strunk & E.B. White. *The Elements of Style*, Macmillan, 4th Edition, 2000.

Lynne Truss. *Eats, Shoots & Leaves: The Zero Tolerance Approach to Punctuation*. Gotham, 2004.

J. de Bruin and B. Hertz, *Project Management for PhDs*, Eleven International Publishing, The Hague (2010).

B. E. Lovitts, "The Transition to Independent Research: Who Makes It, Who Doesn't, and Why", *The Journal of Higher Education*, Vol. 79, No. 3, pp. 296-325 (2008).

Appendix I

Some Rules of English Usage

For the purposes of this short guide, we will list common rules and guidelines under five headings: spelling, punctuation, capitalization & typefaces, grammar, and word choice. The list is far from complete and is only intended to help you understand some of the issues in English usage that distinguish good writing from bad. For a much more complete treatment of these issues, you should consult other resources such as the Online Writing Lab (OWL) at <http://owl.english.purdue.edu/owl/>

Spelling

There is usually just one way to spell a word correctly, but some words do have variations. Most of the variations arise from which form of English you have adopted in your writing. In general, you should use standard British English rather than American English. Consequently, you should use a British English dictionary, such as the Oxford English Dictionary, not an American English dictionary, such as Webster's English Dictionary, to check your spelling.

Since you will use LaTeX to typeset most of your technical writing, you should consult a dictionary frequently as LaTeX does not have a built-in spelling checker. This a good habit to acquire. Automated spelling checkers can make you lazy and cannot always be trusted.

Some words, such as *organize*, can end in either –ize or –ise. The –ize ending is not an Americanism, although it is standard in American English. Both –ize and –ise are acceptable in British English, with the –ise ending having its roots in French-derived spelling. Use the –ize version.

Other words can be spelled in two legitimate ways, but each has a different meaning. For example, *practise* and *practice*. The former is a verb and the latter is a noun. So, you practise programming but your programming needs more practice. In general, the s version of a word is the verb and the c version is the noun: compare *license* and *licence*.

Punctuation

There are many rules of punctuation. We won't go into all of them here. However, some are ignored so often that it's worth highlighting them. For an accessible and entertaining introduction to the rules of punctuation, you should read *Eats, Shoots and Leaves* by Lynn Truss.

You always put a space after a punctuation mark. Normally, you put two spaces after a punctuation mark that ends a sentence. Good typesetting tools, such as LaTeX, will properly adjust the spacing for you, as long as you remember to put at least one space (or a new line) after the punctuation mark.

Many people are confused about when you should use a comma, a semi-colon (;) and a colon (:). To help with this, remember that punctuation marks were invented to help people read writing without having to depend completely on the grammar to break the text up into meaningful chunks. When reading, many people internally voice what they are reading and the punctuation marks provide pauses of varying length. This is also the case when reading aloud. In general, the pause associated with a comma is short, with a semi-colon longer, a colon longer still, and a full stop provides the longest pause.

A comma is normally used to designate a break between two clauses in a sentence. For example, *programs are simple to understand, assuming you know the language*. Commas are often used in pairs to separate out an explanatory or incidental clause in a sentence. For example, *programming, or coding, can be very rewarding*. These pairs of commas are called *parenthetical commas*. The omission of the second of a pair of parenthetical commas is one of the most common mistakes in writing.

A semi-colon separates two relatively independent clauses in a sentence. For example, *C++ can be a difficult language to learn; the same is true of other object-oriented languages such as Java*.

A colon also separates clauses in a sentence and it is usually used to introduce a list of items. For example, *there are three important factors in good writing: clarity, clarity, and clarity*.

A full-stop terminates a sentence. An exclamation mark (!) does the same, but you should use it only in very exceptional circumstances, if at all.

The rules for using a dash often confuse people. There are several types of dash, the three most common of which are the hyphen (-), the en dash (–), and the em dash (—). The hyphen is not really a dash at all but is the punctuation mark used to join two parts of a word which has been broken in two by the typesetter, typically at the end of a line. The en dash, which is the same length as the letter n, is used to join two words to form, for example, a compound adjective; for example, *easily-understood writing requires hard work and careful thought*. It is also used to designate a range; for example, *pages 101–120*. The em dash is normally used in the same way as a colon (one em dash) or a pair of parenthetical commas (two em dashes). The em dash is the same length as the letter m. You would use the em dash if you wanted to achieve greater emphasis, such as when a clause that is being separated out by the em dash expresses a distinctly different idea or thought. For example, *punctuation is confusing — many new things are — but it is also essential to good writing*. In general, the em dash can be used where a full stop is too strong and a comma too weak. One final point about the em dash: we said that you always place a space after a punctuation mark. Some style manuals advise setting the em dash closed with no space on either side of it— like this one— although it is common practice to set it open with spaces on either side.

In LaTeX the hyphen is set by using the minus character, the en dash by using two minus characters together (--), and the em dash by using three minus characters together (---).

Footnote markers,⁷ typically indicated by a superscript, are placed after the punctuation mark, not before.

Capitalization and Typefaces

There are quite strict rules for setting the titles of articles and section heading. In general, you write every word with the first letter in capitals, except prepositions (e.g. in, with, on) and conjunctions (and, or). If you have a hyphenated word such as a composite noun or composite adjective, you write the first letter of the second word in lower-case. For example: *An Analysis of Easily-understood Technical Writing on Humanoid Robotics*. Note that you do not use a full stop at the end of the title or section heading.

Text should be typeset in a normal typeface. In the body of the text, never use bold-face for emphasis; if you want to distinguish or emphasize a word or phrase in a sentence, set it in italics. Use italicized emphasis very sparingly. Do not use single or double quotation marks for emphasis; reserve them for quotations. When you do use them, be careful to use opening marks at the beginning of a sentence and closing marks at the end. In LaTeX, the opening quotation mark is the ` character (top-left of the keyboard) and the closing quotation mark is the ‘ character (bottom-left of the keyboard). Write two mark together to achieve the double quotation mark: `` or “.

There are special rules for typesetting mathematical formulae and you should consult a style manual on the matter. In general, though, scalar variables are typeset in italic typeface with lower-case letters, vectors are set in bold italic typeface with lower-case letters, and matrices with italic upper-case letters.

Foreign words are written in italics. For example, you should write

The *a posteriori* probability can be computed using Bayes' theorem.

Note in the above example we put the apostrophe after Bayes. This is not because Bayes is the plural of Baye (in which case, it would indeed be the correct thing to do) but because in this instance we are forming the possessive of Bayes which is a person's name. In particular, it is an old and well-known name. If it were a modern name, such as Des, you would form the possessive just as you would with any other word by writing Des's (as in *Wow! Look at Des's car!*).

The rule for writing foreign words in italics also extends to writing abbreviations of foreign words in italics, such as *e.g.*, (for example), *etc.* (and the rest of such things), and *et al.* (and others). Note

⁷ A footnote is a useful way of providing essential additional information without disrupting the flow of the narrative. A reader can safely skip over it on the first reading, and return to it later, when time permits.

that *et al.* is used when citing a reference involving three or more authors: e.g. Moore *et al.* 2008 have shown that precise writing is laborious and time-consuming.

One final point about italics: if you have put paragraph, sentence, or even part of a sentence in italics for emphasis, and some of the words included would normally be written in italics, you then write them in normal roman type. For example:

The student quoted from Moore et al. 2008 in an effort to impress the examiners at his viva voce examination.

Grammar

Our goal in this section is not to cover all the rules of English grammar, but to highlight some of the principal rules for the construction of valid sentences and to flag some of the most common mistakes people make.

Structure: subject-verb-object

Almost all sentences have a basic structure with three parts: the subject, the verb (sometimes referred to as the predicate), and the object. The subject is what does the action, the verb is the action, and the object is the target of the action. For example, in the sentence *the student wrote a report*, the subject is the 'student', the verb is 'wrote', and the object is the 'report'. Of course, many sentences are more complicated than this because we often add in adjectives, adverbs, parenthetical expressions, or additional clauses that amplify or qualify the sentence. Now here is the key point: no matter how complicated a sentence, you should always be able to reduce it to this simple structure if you remove the additional words and clauses. You should check all sentences you write by doing this.

Structure: transitive and intransitive verbs

In the previous example, *the student wrote a report*, the action denoted by the verb is done directly to the object (by the subject). The verb is a transitive verb: it takes a direct object. Not all verbs are transitive; some are intransitive. This means that they don't take a direct object – they take an indirect object. Practically speaking, this means you need a preposition before the object. For example, *the student spoke to his supervisor*. Here, the verb to speak (past tense: spoke) is intransitive. You wouldn't say *the student spoke his supervisor* – you need the proposition 'to' before the object. You need to be careful when choosing the proposition. Depending on the verb, only some prepositions will be appropriate.

Relative pronouns: which and that

Consider the following two sentences.

Robots that have several degrees of freedom are hard to program.

Robots which have several degrees of freedom are hard to program.

These two sentences have different meanings depending on the relative pronouns *which* and *that*. This issue confuses most people. When do you use *which* in a sentence and when do you use *that*? You use *that* when you want the clause that follows to be a qualification on the noun to which it refers in the preceding clause. In other words, *that* constrains or restricts meaning. You use *which* when you want the clause that follows to be an amplification on the noun to which it refers in the preceding clause. In other words, *which* adds to or augments meaning. In the examples here, the first sentence says only robots that have several degrees of freedom are hard to program (the rest are easy); the second sentence says that all robots are hard to program and asserts that robots in general have several degrees of freedom. These are two quite different meanings. Choose your relative pronoun carefully. Before leaving this point, note that the clarity of second sentence would be much enhanced by using parenthetical commas, as follows.

Robots, which have several degrees of freedom, are hard to program.

Matching Number

The number of the verb in a sentence or clause must agree with the number of the subject of a sentence for clause. For example:

The interface will display information that are needed by the operator.

This should be:

The interface will display information that is needed by the operator.

In this case both *information* and *is* are in the singular. Every sentence you write should be checked to ensure that the number of the subject and the verb agree. This means that you must also check that the sentence has a valid subject and a valid verb.

Word-choice

Finding the Correct Word

In English, different words convey different shades of similar meaning. It is often difficult to find exactly the right word to convey the meaning you want to express. For example, don't use the word *colour* to describe a region in an image if you actually mean the *brightness* (i.e. the intensity), or the *contrast* (the ratio of the minimum and maximum brightness), or *saturation* (the amount of white present). If you are not absolutely sure that the meaning of a word is exactly what you want, look it up in a dictionary. If you are not looking up several words every day, you are probably not writing

effectively or precisely. Remember: precise writing implies logical construction of the argument and use of exactly the right word. It is a very time-consuming exercise.

Collocation: Pairing Verbs and Nouns

The problem of finding the right word gets worse when you realize that you cannot pair words arbitrarily. That is, if you choose one word as the subject or object in a sentence, then there is a limited number of verbs you can use with it. This works in reverse too: if you choose a verb, then there will be restricted number of nouns that will match well. This is called collocation.

For example, don't write *the process specifies a level of priority*. Instead, write *the process assigns a level of priority*.

Similarly, don't write *Chapter 2 assigns the requirements of the system*. Instead, write *Chapter 2 specifies the requirements of the system*.

Consistent Constructs in the Same Sentence

Consider the following sentence.

Fusing infrared and visible spectrum images helps not only in detecting defects but also to locate them precisely.

Can you see what's wrong? While the meaning is clear, the sentence sounds cumbersome. This is because there are two different grammatical constructs: the participle (*detecting*) and the infinitive (*to locate*). Both are correct but the sentence would sound better if either one or the other type of construct is used in both cases:

Fusing infrared and visible spectrum images helps not only in detecting defects but also in locating them precisely.

Fusing infrared and visible spectrum images helps not only to detect defects but also to locate them precisely.

Which of these two constructs you choose is a matter of personal style. The second form is probably better though as it is more direct and therefore has greater impact.

Link Words

Link words are sometimes used at the beginning of a sentence to make a logical connection to the preceding sentence. For example:

Robots are ugly machines. However, the iCub humanoid has a functional cover that give it an aesthetically-attractive appearance.

Here, the word *However* indicates that the second sentence is in some way going to contradict what the first sentence said. Unfortunately, for inexperienced writers, it is easy to use the wrong link word at the start of a sentence (common examples include *Moreover*, *Thus*, *Furthermore*, *Hence*, ...). You should avoid using these link words unless you are absolutely sure they make sense and add something to the meaning of the two sentences they are coupling.

Appendix II

Writing a Literature Survey

Conducting a literature survey and writing it up in a cogent form is a daunting task. The following are some hints to help you get started.

First, let's establish what a literature survey is. It is a structured synthesis of all the published material that is relevant to your research question (refer again to Section 2 Moving Along: Framing the Research Question) and your attempt to answer it. This means that it collects together all the relevant ideas, concepts, models, formalisms, techniques, and results that have been published by others in their attempt to address your research question. The word "all" in the preceding sentence is important: your survey must exhibit comprehensive coverage.

However, a literature survey is far more than a mere collection of references and summaries of these references. It must also bring some order to that collection: some form of organizing structure. This is often referred to as a taxonomy, i.e., a structured classification of the various approaches that have been published. Usually, but not always, this structure is hierarchical. At the top of the hierarchy (or the root of the hierarchy tree) there is the topic itself. The next level might have the main alternative approaches that can be deployed, and at the next level, there is a further sub-classification of each approach, and so on, until all the approaches that have been taken are covered.

But there is more. While the taxonomy, or tree of possible approaches, is important, and demonstrates your knowledge of the field, a literature survey should also demonstrate your understanding of each topic. In other words, it should be a synthesis of the topics, expressed as a narrative that critically appraises each of the approaches by highlighting their similarities and differences, and discussing their relative strengths and weaknesses. This is perhaps the hardest aspect of the literature survey. Because it will take the form of a narrative, i.e., a sequence of connected statements leading the reader through the various approaches, you need to decide the order in which to treat each one. One way of approaching this is to view it as a breadth-first traversal of the hierarchical taxonomy tree, i.e., visiting each approach in turn, beginning at the root (i.e., the research topics itself), followed by each node at the first level, i.e., the initial decomposition of approaches, then each node at the next level, appropriately grouped according to the parent approach, and so on. You shouldn't view this as a strict algorithm, but as a way of making your initial attempt. As you write the syntheses, you may find different topics align according to other attributes, other than their parent topic, and you should make use of this in your synthesis.

Recognize that it will require many attempts and many re-writes, polishing your narrative on each pass. You should start with your own short summary of each paper you have read as you try to

achieve complete coverage of the topic through recursive reading (refer again to Section 3 Reading). That provides the input for the second stage: organizing all the various topics, possibly creating a taxonomy of approaches. The leaf nodes of the hierarchical taxonomy tree will be the specific published papers describing the work of other researchers. Finally, move on to creating the synthesis, possibly by traversing each node in the taxonomy tree, but being alert to other ways of highlighting similarities between approaches. You will certainly have to make more than one pass at this, probably several. Don't despair: writing is hard, but, like any skill, it gets easier with practice.

Appendix III

Typical Structure of a Thesis

Title Page

- Specific title of the thesis (e.g., “Multi-stage Learning in Biomimetic Search and Rescue Robots”)
- General Title (i.e., “Final Year Project Report”)
- Degree (e.g., Ph.D., M.Sc., B. Sc.)
- Author (name and student identification number)
- Institution (i.e., Carnegie Mellon University Africa)
- Supervisor
- Date

Abstract

- What is the subject matter of the thesis: what did you do?
- Motivation: why is it important?
- Significance: what contribution does the thesis make?
- The abstract should be approximately 200 words long. It normally takes at least ten revisions to achieve a good abstract.
- The abstract should be written after the thesis has been completed.

Table of Contents

- Chapters
- Sections

Acknowledgements

- Help from friends, colleagues, and staff.
- Support from sponsor
- Support from Parents, etc

Chapter 1. Introduction & Overview

Chapter 2. Literature Survey

Chapter 3. Theoretical Foundations: Background Material

Chapter 4. Formal Model: Theoretical Development (use additional chapters if necessary)

Chapter 5. Algorithmic Considerations

Chapter 6. Implementation Issues

Chapter 7. Evaluation

Chapter 8. Discussion & Critical Appraisal

References

Appendices

- Key Software listings
- Mechanical schematics
- Mathematical proofs