

The Researcher's Bible

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

DAI Teaching Paper No: 4 (revised July 1989)

Last updated on 24 July 1989 by Peter Ross.
Includes contributions by Graeme Ritchie.

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

Contents

1	What is a Thesis?	4
2	The Programme of Postgraduate Study	5
3	Standard Pitfalls for Postgraduate Students	6
3.1	Solving the World	6
3.2	Manna from Heaven	6
3.3	Computer Bum	7
3.4	Yet Another Language	7
3.5	Micawberism	7
3.6	Ivory Tower	7
3.7	Misunderstood Genius	7
3.8	Lost in Abstraction	8
3.9	Ambitious Paralysis	8
3.10	Methodology does not make a thesis	8
3.11	The discovery route is not a justification	9
4	Psychological Hurdles	9
4.1	Mental Attitude	9
4.2	Research Impotence	9
4.3	Dealing with Criticism	10
4.4	Early Morning – Cold Start	10
4.5	Theorem Envy	10
4.6	Fear of Exposure	11
5	Choosing a Research Project	11
6	Research Methodology	12
6.1	Stage 1	12
6.2	Stage 2	12
6.3	Stage 3	12
6.4	Stage 4	12
6.5	Stage 5	12
6.6	Stage 6	13
6.7	Stage 7	13
7	Writing Papers	13
7.1	Departmental Research Papers	13
7.2	Departmental Teaching Papers (formerly Occasional Papers)	13
7.3	Departmental Software Papers (formerly Technical Papers)	14
7.4	Departmental Discussion Papers	14
7.5	Departmental Working Papers	14
7.6	Publishing Papers	14
7.7	Conference Proceedings	14
8	Guide to Writing	15

9	Guide to Reading	18
9.1	Outer Circle	19
9.2	Middle Circle	19
9.3	Inner Circle	19
10	The Examination of Theses	19
11	Journals Publishing AI Material	21
11.1	General AI Journals	21
11.2	Other Journals Accepting AI Material	22
11.3	Newsletters	27
11.4	Abstracts	28
11.5	Regular Conferences	28
12	References	29

The Researchers' Bible

by

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

This version: February 17, 1995

Abstract

Getting a Ph.D. or M.Phil is hard work. This document gives advice about various aspects of the task. Section 1 describes the problem – what is a thesis? Section 2 sets out the formal requirements of gaining a thesis. Sections 3 and 4 describe some of the pitfalls and hurdles which students have encountered. Sections 5 and 6 advise about choosing 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 process for a research degree, and how to cope with it. Finally, section 11 lists journals which accept A.I. material.

1 What is a Thesis?

To get a Ph.D. or M.Phil. you must write a thesis and sit an oral examination (there are no written examinations for postgraduate work in the AI Department). The oral is generally used to ask for clarification of the thesis, so the main burden of assessment falls on the thesis.

The definitions of an adequate thesis, at least as far as Edinburgh University is concerned, are given in the University Calendar: reg. 3.2.6 and reg. 3.3.7 of the section on postgraduate study. We reproduce them below:

Ph.D. Thesis “The grounds for the award of the degree of Ph.D. are:

- (i) The candidate must have demonstrated by the presentation of a thesis and by performance at an oral examination (unless this is exceptionally waived) that he is capable of pursuing original research in his field of study, relating particular research projects to the general body of knowledge in the field, and presenting the results of his researches in a critical and scholarly way.
- (ii) The thesis must be an original work making a significant contribution to knowledge in or understanding of the field of study and containing material worthy of publication; show adequate knowledge of the field of study and relevant literature; show the exercise of critical judgment with regard to both the candidate's work and that of other scholars in the same general field; contain material which presents a unified body of work such as could reasonably be achieved on the basis of three years postgraduate study and research; be satisfactory in its literary presentation, give full and adequate references and have a coherent structure understandable to a scholar in the same general field with regard to intentions, background, methods and conclusion.

- (iii) The thesis may not exceed 100,000 words save that, exceptionally, on the recommendation of the supervisor, the Committee may grant permission to exceed the stated length on the ground that such extension is required for adequate treatment of the thesis topic.”

M.Phil. Thesis “The grounds for the award of the degree of M Phil or M Litt [or LL M] are:

- (i) The candidate has demonstrated by the presentation of a thesis and by written and/or oral examination that he has acquired an advanced level of knowledge and understanding in his field of study, is capable of relating knowledge of particular topics to the broader field of study involved and of presenting such knowledge in a critical and scholarly way.
- (ii) The thesis must be an original work comprising a satisfactory record of research undertaken by the candidate, or a satisfactory critical survey of knowledge in the approved field of study; show competence in the appropriate method of research and/or an adequate knowledge of the field of study; exhibit independence of approach or presentation; be satisfactory in literary presentation and include adequate references.
- (iii) The thesis may not exceed 60,000 words save as specified in Section 6.1 and save that, exceptionally, on the recommendation of the supervisor, permission may be granted by the Committee to exceed the stated length on the ground that such extension is required for adequate treatment of the topic of the thesis.’

How original and significant must Ph.D. research be? The phrase ‘containing material worthy of presentation’ suggests a simple rule of thumb. It should be possible to distil from the thesis, a paper worthy of publication in a journal. This is not an infallible guide – refereeing standards are not always what they should be. The final decision must rest with the examiners.

The definition of an M.Phil. thesis is even less helpful. In the AI Department we do not require a written examination. Most M.Phil. theses have been a record of research rather than a critical survey, but the latter would be possible. Again it should be possible to distil the essential message of the thesis into a short paper, but in this case publication in a journal would not be expected and a departmental working paper would suffice.

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 studying, compared to the British norm of 3 or 4 years.

2 The Programme of Postgraduate Study

At Edinburgh University, the minimum prescribed period for a Ph.D. is 33 months full-time or 45 months part-time, and for an M.Phil. is 21 months full-time or 33 months part-time. During this time the student usually must be resident in Edinburgh. Students

are normally expected to complete Ph.D.s in three years and M.Phil.s in two years. Studentships usually cover only these minimum periods. The maximum permitted periods of study are 5 years for a Ph.D. and 4 years for an M.Phil..

In the past, few postgraduate students in Artificial Intelligence have submitted their thesis in these minimum periods. In an effort to reduce the time taken, each student's work will be organised in the way described in Departmental Notice 16. This sets out the requirements that a postgraduate student must meet at various stages. The requirements change a little from time to time and so are not included in this document.

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 (I fell in most – AB).

3.1 Solving the World

Most students pick research goals which are far too ambitious. This is particularly easy in AI. So many tasks which humans find easy turn out to be really hard to model. Obviously the main burden of helping you choose a suitable topic will fall on your supervisor. In addition you should read the literature and talk to fellow workers to find out what the state of the art is. One good source of ideas is the further work sections of papers. Read the literature critically. Another good source is re-doing bad work, properly.

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) Read the literature. Read with a question in mind, e.g. What is wrong with this? How can I use this? etc.
- (b) Talk to people. Do not go away and hide. Do not be ashamed of your ideas. Other people's are sillier.
- (c) Tackle a simplified version of your problem. Ask your supervisor for exercises, mini-projects, etc.
- (d) 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.
- (e) Give a talk to a small group. Same effect as (d).

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 program must be explainable at a higher level than code, for it to make a real contribution to knowledge. Try to plan your program theoretically before going to the terminal. If you must work some of it out at the terminal then rush away soon and work out the theory. If you find this hard, try to describe how it works: to a friend; in a paper or at a seminar. If people do not understand it is your fault – try harder.

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. Writing a useful, new language requires an encyclopaedic knowledge of AI needs and experience of systems programming. 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 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** AI is full of jargon: try to rephrase your ideas using ordinary English; try to rephrase your ideas in someone else's jargon. Do they come out any 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 make a complicated program – it is just a nuisance to other people. Do it the simplest way you can. Occam was perfectly right.

3.8 Lost in Abstraction

To be worthwhile your research work should be aimed at understanding some major property of intelligence, e.g. controlling search, representing knowledge, learning. But to achieve anything you must tackle the abstract property in a concrete situation, that is you must build a program to do some task that requires search to be controlled, knowledge to be represented, knowledge to be learnt. Trying to tackle the problem in the abstract will only lead to paralysis and frustration.

3.9 Ambitious Paralysis

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

3.10 Methodology does not make a thesis

Since AI 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 an AI 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 views which may originate from (for example) Dreyfus, Schank, Fodor, Minsky, or which may even be their concoction. Out of this struggle, your reading, your attendance at seminars, your debugging, and other hard work, will emerge your world view on AI 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.

3.11 The discovery route is not a justification

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

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? ¹ The first step in obtaining this self-assurance is to pick a research topic you can believe in (see section 5). 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

¹Edison said that genius was 1% inspiration and 99% perspiration, and he should know.

enthusiasm necessary to produce results evaporates. The way out of this vicious circle is to realise that research ability does not depend on some magic essence. It is a skill which can be learnt, like any other. You too can do original research by following the instructions in this 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:

1. Make yourself a regular working schedule – and stick to it. It doesn't have to be 9–5, but there should be a definite time of day when you expect to start work. Otherwise, you will find yourself postponing the evil moment with endless, routine, domestic chores.
2. Make sure you leave some non-threatening, attractive task to do first thing. For instance, do not leave off writing the day before at the beginning of a new hard section. Leave something easy to start writing: a paragraph which is routine for you 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 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. Beware of choosing an area new to you because of its superficial appeal. The gloss will soon wear off when you are faced with the hard grind necessary to get a basic grounding in it.

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

6 Research Methodology

There are many possible methodologies, but you should have one. Many start from different beginnings but merge later. For example, you might start by analysing some aspect of human performance, or somebody else's theory to account for it; or, you might start by trying to rationally reconstruct someone else's theory as embodied in their reported program, and look for the strengths and weaknesses of it. The 'rational reconstruction' approach is often fruitful, since it is still regrettably often the case that a published research paper will concentrate on implementation and performance while only hinting at the assumptions and principles behind the work reported. But, however you start, get yourself a theory!

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. Use existing AI mechanisms wherever appropriate. Showing that a problem can be solved with an existing mechanism is also a research achievement!

6.3 Stage 3

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

6.4 Stage 4

When you are satisfied that your proposed program is stable, choose the programming language which fits your needs closest and implement your program. Debug it on the scenarios in stage 1 and 3.

6.5 Stage 5

Find some examples you have not previously considered and run your program on them. Modify your program until it is robust – i.e. runs on a wide range of examples and does

not collapse ignominiously every time you input a new one.

6.6 Stage 6

Evaluate your program by testing it on some examples. Keep statistics on its success/failure and analyse any failures that occur. Present these findings in your thesis.

6.7 Stage 7

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 product of the 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 AI 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. The Department of AI at Edinburgh runs various different series of research memoranda: these are described separately in Departmental Notice 8, which is the definitive statement about the purpose of each series. An informal statement, not necessarily up to date, is given below.

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. I personally (AB) find it very useful to type them into the computer (learn about troff, \LaTeX or – more unusually – \TeX).

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 Teaching Papers (formerly Occasional Papers)

These are for educational material, e.g. lecture notes and programming primers etc.

7.3 Departmental Software Papers (formerly Technical Papers)

These are software produced in the Department, eg programming primers, manuals and editing facilities, etc. It also includes descriptions of special hardware/software built in the Department, e.g. the LOGO turtle, robots etc.

7.4 Departmental Discussion Papers

These are for papers for internal circulation and discussion. Also includes student projects (MSc or AI 1/2) and thesis proposals.

7.5 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 a load of c**p then the sooner other people can see them and let you know the better for you.

7.6 Publishing Papers

Each paper must be approved by 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 one from the list in section 11. 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. These will be vetted by several referees chosen by the journal editor. 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.

7.7 Conference Proceedings

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

There is a good guide to style and presentation of scientific papers in [Booth75]. Helpful information about writing theses is given by [Parsons73]. Both these documents can be found in the Forrest Hill Library.

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. The standard of writing for journals, conferences and theses is not high, and has been 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. Further useful references for writing technical reports can be found in [Cooper64] and for expressing your ideas to a large audience in plain English see [Orwell68].

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

- Your paper should have a message, i.e. an argument that you are advancing, for which your research provides evidence. Make sure you know what this message is. Summarise it in a few words on paper or to a friend. Make sure the message is reflected in the title, abstract, introduction, conclusion and in the structure of the paper.
- Putting your case so that it can be understood is not enough – you must present it so that it *cannot* be *misunderstood*. Think of your audience as intelligent, but (a) ignorant and (b) given to wilful misunderstanding. Make sure that the key ideas are stated transparently, prominently and often. Do not tuck several important ideas into one sentence with a subtle use of adjectives. Do not assume that any key ideas are too obvious to say. Say what you are going to say, say it, and then say what you just said.
- Do not try to say too much in one paper. Stick to the main message and only include what is essential to that. Reserve the rest for another paper. A reader should get the main idea of the paper from the first page. Long rambling introductions should be pruned ruthlessly.
- The basic framework for a scientific paper is: what is the problem, what did you use to tackle it, what results followed.
- To keep the technical standard of paper uniform, have a particular reader in mind as you write.
- You do not have to start writing at the beginning. In particular, the introductory remarks are best written when you know what will follow. Start by describing the central idea, e.g. your main technique, procedure or proof. Now decide what your hypothetical reader *has to know* in order to understand this central idea and put this information into the introductory sections/chapters.

- Use worked examples to illustrate the description of a procedure. Do not use them as a substitute for that description.
- Clearly state what is new or better about what you have done. Make explicit comparisons with closely related work.
- If you find yourself using a long noun phrase to refer to the same entity or idea several times then you should probably define a new term. Do not define a new term unless you really need it.
- Learn to use a keyboard (all 9 fingers), a screen editor, a text formatter and a spelling corrector. Type your paper into a computer, either directly or from notes or from a handwritten manuscript. This will save time when it comes to alterations, corrections, etc. Run the finished product through a spelling corrector.
- 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.
 - It should consist of a few sentences, i.e. be of abstract length.
 - The sentences should form the steps of an argument. This argument is the message of your thesis.
 - Each sentence should outline the contents of some part, roughly a chapter, of your thesis.
 - The message should serve as a guide to the: title, abstract, summary, conclusion and the whole body of your thesis.
- The thesis message should help you in the following ways:
 - It should ensure that the parts of your thesis hang together in a coherent manner. It should suggest how to reorganise the various notes and papers in your ‘thesis’ file so that they form an argument.
 - It should answer the questions ‘What have I done?’ and ‘Why does this work deserve a degree?’. You should now know what to emphasize in the abstract, introduction, conclusion, title, etc.
 - It should answer questions like ‘What should be discussed in ‘related work’?’. In fact, you should know precisely what role each chapter is meant to play in the whole, i.e. what it is supposed to prove.
- The thesis message is short and easy to edit. You can play around with it until you get something you are happy with.
- Now you can go back to bottom up activity – reworking the existing material, and writing new material, to fulfil the demands of the ‘message’.

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

Example.

“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 for which God is the programmer, and in which miracles are seen as run time patches inserted during interruptions.

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

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

9 Guide to Reading

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

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

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

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

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. AI abstracting has improved recently. Section 11.4 contains a list of publishers and journals which produce AI 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. Background information about AI is to be found in the general purpose textbooks, e.g. [Boden77], [Winston77], [Bundy78], and of course there are the PG introductory lectures.

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. It is worthwhile keeping a record of papers you have read and some comments about them, otherwise the benefits derived from reading them will evaporate as your memory fades. Some people find a card indexing scheme valuable here.

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. Talk or write to him 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?'.

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

10 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 send them to the faculty office.

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*.

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. Accept the thesis, but not the oral, and examine the candidate further.
4. Reconsider after a further period of supervised study and resubmission.
5. Reconsider a Ph.D. as an M.Phil..
6. Reject.

Verdicts 1, 3 and 6 are very rare. You will usually be told the recommendation informally, with the understanding that it can be overturned by the faculty or senate (and this is not unheard of).

Verdict 2 is to allow correction of typographical errors, spelling mistakes, minor rewrites, etc. You should make the alterations on the bound copies of the thesis, pasting in new pages if necessary. The internal examiner will check that the thesis has been completely corrected and will then inform the faculty who will process your thesis and inform you of the verdict. This may take several months.

Verdict 4 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 with the same examiners. This is your last chance. Verdicts 4 and 5 are not available the second time around.

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

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

doubts in their post-viva reports then it is not unknown for the faculty to override their recommendations.

The examiners will 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.Phil. standards are the language syntax, and the examiners are the interpreter. During the viva you will get various error messages. These messages do not need to be taken at face value - they may be based on a misunderstanding - but they cannot be ignored. Assume that each error message will lead to some alteration in your thesis. Of course, you hope that this will only be a minor alteration, but do not let this hope blind you to the possibility that the problem is more fundamental. Do not get aggressive or defensive with your examiners. You cannot bludgeon or 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.

11 Journals Publishing AI Material

The following list consists of journals known to publish AI papers. The information on each journal is arranged in the order: title of journal; name and address of most appropriate editor for sending AI material to; the type of AI material which this journal accepts. Further additions to this list are always welcomed. Information in the above format should be sent to Peter Ross, ideally via electronic mail.

11.1 General AI Journals

Artificial Intelligence

Bobrow D.

Xerox Parc,

3333 Coyote Hill Rd,

Palo Alto, Ca. USA.

General Artificial Intelligence.

Cognitive Science

Norman D.
Department of Psychology,
University of California, San Diego,
La Jolla, CA 92093, USA.
General AI, Psychology and Language.

Computational Intelligence

Cerconne, N. and McCalla, G.
Department of Computing Science,
Simon Fraser University,
Burnaby, British Columbia,
Canada, V5A 1S6.
General Artificial Intelligence.

Journal of Experimental and Theoretical Artificial Intelligence

Dietrich E. and Fields C.,
Computing Research Laboratory,
Box 30001/3CRL,
New Mexico State University,
Las Cruces, NM 88003-0001,
USA.
*Short papers on experimental, theoretical or methodological issues.
Deliberately short turn-around time of three months.*

11.2 Other Journals Accepting AI Material**International Journal of Approximate Reasoning**

Bezdek J.C.,
Dept. of Computer Science,
University of South Carolina,
Columbia, SC 29208, USA
*Approximate reasoning in the design of artificially intelligent systems,
eg robotics, computer vision, control processes, expert systems, database
management, information retrieval, medical computing and NLP.*

Applied Artificial Intelligence

Trapp R.
Austrian Institute for Artificial Intelligence,
University of Vienna, Austria.
*Applications of AI, evaluations of existing systems and tools,
user experience, economic, social and cultural impacts of AI.*

Journal of Artificial Intelligence in Education

Marks G.H.,
Association for Computers in Mathematics and Science Teaching,
P.O. Box 60730,
Phoenix, Arizona AZ 85082
*Intelligent tutoring systems, learning environments, computational models of
learning and instruction, AI programming environments for educational use.*

**Artificial Intelligence for Engineering Design,
Analysis and Manufacturing (AI EDAM)**

Dym C.L.,
Dept. of Civil Engineering,
University of Massachusetts at Amherst,
Amherst, Massachusetts

*Particularly interested in the uses of AI in support of planning,
design, finite element analysis, simulation, spatial reasoning and
graphics, process planning, optimisation and manufacturing*

Artificial Intelligence and Society

Gill K.S.,
Director SEAKE Centre,
Faculty of Information Technology,
Brighton Polytechnic,
Brighton, UK.

*Major Societal and AI perspectives, detailed case studies, current research
and applications, surveys, monitoring and forecasting of AI and IT research.*

Journal of the Association of Computing Machinery

Fischer M.J.
Department of Computer Science,
Yale University,
PO Box 2158
New Haven, CT 06520, USA.

Theory of Computation and Mathematical type, AI papers (eg Theorem Proving).

Journal of Automated Reasoning

Wos, L.
Argonne National Lab.,
Math. and Computer Science Division,
9700 S. Cass Avenue,
Argonne IL 60440, USA.

Automatic Theorem Proving, Inference, Expert Systems.

Cognition

Mehler J.
Laboratoire de Psychologie
54 Blvd Raspall*
F-75006 Paris, France.

AI of direct interest to Cognitive Psychologists.

Cognitive Psychology

Hunt E.
Department of Psychology,
University of Washington,
Seattle, Washington 98195, USA.

AI of direct interest to Cognitive Psychologists.

Communications and Cognition

Heefer, A. and Vervenne, D.

Blandijnberg 2,
B-9000 Gent,
Belgium.

General AI with a philosophical flavour.

Journal of Computer and System Sciences

Blun E.K.

Department of Mathematics,
University of Southern California,
Los Angeles, CA 90089, USA.

Data and Knowledge Engineering

van de Riet, R.

Dept. of Math. and Computer Science,
Free University,
1081 HU Amsterdam,
The Netherlands.

Knowledge representation and expert systems with a database flavour.

Expert Systems

Evison Look H. Learned Information,
Besselsleigh Road,
Abingdon, Oxon, OX13 6LG.

Expert systems, knowledge engineering.

International Journal of Expert Systems Research and Applications

Harandi M.T.,

Dept. of Computer Science,
University of Illinois at Urbana-Champaign
1304 West Springfield Avenue,
Urbana, IL 61801, USA.

Theoretical or Practical issues dealing with anything that relates to expert systems, current or future.

Future Computing Systems

Shaw M.L., Sugeno M.

Department of Computer Science,
York University,
4700 Keele Street,
Downsview,
Ontario, Canada, M3J 1P3.

Future Generation Computer Systems

Aiso H., Kuo F., Raulefs P.

University of Kaiserslautern,
Fachbereich Informatik, Postfach 3049,
D-6750 Kaiserslautern, FRG.

General AI with IKBS and AI applications flavour.

Journal of Intelligent Systems

George F.H., Johnson L., Wright M., Lovejoy A.

A.A. Publishing Services,
39 Oakleigh Gardens,
London N20 9AB, UK.

*Interdisciplinary collaboration between computer scientists, psychologists,
electronic engineers, neuroscientists and philosophers*

Knowledge Acquisition

Gasines B.R.

Centre for person-Computer Studies,
University of Calgary,
Calgary, Canada

All aspects of knowledge acquisition

The Knowledge Engineering Review

Fox J.

Imperial Cancer Research Fund,
Lincoln's Inn Fields,
London WC2 3PX, England.

Publication of BCS specialist interest group on expert systems.

Language and Cognitive Processes

Tyler L.

MRC Applied Psychology Unit,
15 Chaucer Road,
Cambridge CB2 2EF, UK.

Mental Processes and Representations involving language use.

Journal of Logic Programming

Lassez J.-L.,

IBM Thomas J. Watson Research Centre,
Yorktown Heights, New York NY 10598

Logic Programming.

Machine Learning

Langley P.

Dept. of Information & Computer Science
University of California,
Irvine, CA 92717, USA.

Machine Learning.

International Journal of Man-Machine Studies

Gaines B. and Hill D.

Department of Computer Science,
York University,
4700 Keele Street,
Downsview, Ontario,
Canada, M3J 1P3.

CAI, NLU, Speech, Bionics

Education and Computing

Levrat B.
University of Geneva,
24 Rue General-Dufour,
1211 Geneva 4,
Switzerland.
Computing and education.

IEEE Pattern Analysis & Machine Intelligence

Pavlidis T.
Rm. 2C-519,
Bell Labs.,
600 Mountain Ave.,
Murray Hill, NJ 07974, USA.
Computer Vision, Pattern Recognition.

Journal of Pragmatics

Mey J.L and Haberland H.
AI of direct interest to NLU, discourse analysis.

IEEE Journal of Robotics and Automation

Bekey G.A
The Robotics Institute,
School of Engineering,
University of Southern California,
Los Angeles, CA 90089, USA.
Theory and applications in robot dynamics and control; robot languages; robot vision; robot locomotion; management of multi-robot systems; simulation of robot and manufacturing systems; motion planning and task planning.

Journal of Intelligent and Robotic Systems

Spyros G. Tzafestas
National Technical University,
Division of Computer Science,
Department of Electrical Engineering,
15773 Zographou, Athens, Greece.
Almost equally divided between specific robotic interests and all other AI interests; wide-ranging. "System intelligence" is the linking theme.

Journal of Robotic Systems

Beni G. and Hackwood S.
University of California,
Santa Barbara, CA
Robotics Systems, Planning, Dynamics, vision, etc.

International Journal of Robotics Research

Brady J.M and Paul R.
Artificial Intelligence Laboratory,
Massachusetts Institute of Technology,
545 Technology Square,
Cambridge, MA 02139, USA.
Robotics Systems, Planning, Dynamics, vision, etc.

Journal of Symbolic Computation

Buchberger B.
Johannes-Kepler-Universitat,
Institut fur Mathematik,
A-4040 Linz,
Austria.
Computer Algebra, Automatic Theorem Proving and Algorithmic Geometry.

Knowledge-based Systems

Sawyer H.
Butterworth Scientific Ltd.
PO Box 63
Westbury House
Bury Street
Guildford, Surrey GU2 5BH
Design and building of knowledge-based systems, HCI

11.3 Newsletters

AISB Quarterly

Sharples M.
Cognitive Studies Programme,
University of Sussex,
Brighton, BN1 9QN.
*Short general interest articles, conference reports, titles of recent papers etc.
Not technical articles.*

AAAI Magazine

Englemore B.
Teknowledge Inc.,
525 University Ave,
Palo Alto, CA 94301.
Similar to AISB Quarterly, but including technical articles.

SIGART Newsletter

Price K.
Image Processing Institute,
University of Southern California,
Los Angeles Ca 90007
Similar to AISB Quarterly, but abstracts of papers as well as titles.

11.4 Abstracts

The following list mentions some of the abstracting publications. New ones are started almost monthly at present. They provide abstracts of papers in other journals, but do not accept papers directly.

ACM – *Computing Review*

Basil Blackwell – *Artificial Intelligence Abstracts*

INSPEC Marketing – *Key Abstracts (Artificial Intelligence)*

SIGART

Turing Institute – *Abstracts in Artificial Intelligence*

11.5 Regular Conferences

IJCAI (International Joint Conference on Artificial Intelligence)

Biennial on odd years. The major AI conference.

ECAI (European Conference on Artificial Intelligence)

Biennial on even years. Major European conference. Succeeded British-based AISB conferences.

AAAI (American Association for Artificial Intelligence)

Biennial on even years. Major US conference.

AISB (Artificial Intelligence and the Study of Behaviour)

Biennial on odd years. Major British conference.

CSCSI (Canadian Society for Computational Studies of Intelligence)

Biennial on even years. Major Canadian conference.

IFIP (International Federation for Information Processing)

Major computer conference. Has AI section.

Meeting of the ACL (Association for Computational Linguistics)

Annual Meeting. Main international forum for the exchange of research results in Natural Language Processing.

COLING (International Conference on Computational Linguistics)

Biennial on even years in venues between Europe, North America and Japan.

Meeting of the European Chapter of ACL

Biennial on odd years, complement the COLING sequence of conferences.

BCS Expert Systems Group

Annual in December in the UK. Mix of research and commercial interests.

12 References

- [Bligh72] Bligh D., *What's the use of lectures?*, Penguin, 1972
- [Boden77] Boden M., *Artificial Intelligence and Natural Man*, Harvester Press, 1977
- [Booth75] Booth V., "Writing a scientific paper", in *Biochemical Society Transactions*, vol 3, 1975
- [Bundy78] Bundy A., R.M.Burstall, S.Weir and R.M.Young, "To Artificial Intelligence", Dept. of AI Occasional Paper 3, 1978
- [Cooper64] Cooper B., *Writing Technical Reports*, Pelican, 1964
- [Orwell68] Orwell G., *The Collected Essays*, Penguin, 1968
- [Parsons73] Parsons C.J., *Theses and Project Work*, George Allen and Unwin, 1973
- [Winston77] Winston P., *Artificial Intelligence*, Addison-Wesley, 1977