Chapter 16. The writing process

Table of Contents

The role of writing in research

Writing up throughout research

The most important and lasting output of any research are the publications that document the ideas and results. Despite its vital nature, both novice and experienced researchers frequently delay writing

up their work because they find writing difficult, and because they lack confidence and are afraid to expose their work to critical review.

It is best to make writing up an integral part of any research activity from the very beginning. Not only will this mean that style improves with practice, but equally importantly, it also gives important deliverables along the way. By documenting what is done and what is learnt throughout the research process, you ensure that no work is ever wasted. Much time spent in writing is in fact time spent tackling a problem away from the computer. This is not tedious at all it is extremely valuable. When writing up research one often discovers improvements, alternatives or problems which were not apparent before. This new vision is fed right back into the system/solution being described, and the final product substantially better as a result of reporting on the work.

The Research Diary

A diary keeping track of your actions, questions, problems, attempts that failed and solutions that worked is an invaluable tool for organizing yourself, recording your progress and helping you subsequently write articles on your work. Write down speculations, problems, possible solutions, random ideas, references to follow up, notes on your readings, interesting quotes, paper outlines, and so forth in the same diary as they occur to you. Read through this regularly to see what occurs frequently and how such ideas are evolving, to avoid solving the same or a similar problem more than once over (and avoid repeating what proved fruitless), and to have encouraging evidence that you have not been idle and are indeed making progress.

Benefits of keeping a research diary:

- it provides a reminder of ideas, experiences and actions which would otherwise be forgotten later when you come to report your research
- it gives practice in writing and makes writing a natural part of the research process, rather than something seen as a hurdle to avoid as long as possible
- it is a source of comfort and support, particularly if you write down not only problems and solutions but also your feelings of frustration, fear or lethargy
- it is encouraging evidence that progress is being made, and that you have not been idle

Follow-up notes

You should also keep a collection of follow-up notes each describes a problem or idea that occurred to you while focused on some other task, which could not be considered at the time but which someone could investigate in the future. If you do not jot these down (about a page or two on each) they will be forgotten altogether or remembered incorrectly.

Meeting minutes

Another form of ongoing writing involves note-taking at research meetings. Ideally, you should have regular meetings with your supervisor as well as regular meeting with the entire research group working in your field of interest. These notes document progress and decisions taken, and provide a quick way of remembering the context from one meeting to the next. Meetings themselves are important, to make sure you are on the right track, to discover any problems or relevant research of which you were unaware, and to share ideas with others to learn from their feedback so don't be afraid to speak out, even when you are unsure if your thinking is sound (in fact, especially then!)

Writing to get feedback

A successful researcher should actively seek criticism. This is hard for us all, and particularly so if your personality is not suited to this, but remember that it is a part of the research process, and a key one at that. You will learn resilience to destructive comments and how to listen carefully to constructive, valid

criticisms. In order to invite criticism, you must present your ideas to others. Much of this will be verbal make a point of attending conferences, seminars and other events where academics or practitioners get together, and talking to people about your work. Give presentations yourself wherever you can, and if necessary ask your supervisor to arrange a forum where you can do so.

You can also write pre-publication-stage papers to show people in order to gain feedback. Take the same care with writing these as you would if you were submitting the paper for publication, as it is just as important to make your points clearly to the reader. You should therefore write papers regularly during your research, either to give to others directly or as preparation for a presentation. You will see that explaining your work often leads you to detect where your own thinking is faulty or incomplete. Feedback and criticism naturally can result in finding better ways of solving problems, can show up errors early and so reduce wasted time, and can point out aspects which you have failed to notice yourself. It also shows up which parts of your research are the most interesting or innovative, and which are the hardest to convince people of, and as a result you will be better focused thereafter. Feedback from readers of pre-publication writing will also improve the quality of your subsequent written and oral presentation.

Giving feedback to others

Giving feedback to others is also an important ingredient of research. It sharpens your critical thinking skills, enhances your confidence and exposure, and increases the likelihood of others giving you feedback in turn. In addition to active involvement at research group meetings, conference talks and seminars, it is very helpful to volunteer to review papers for a journal or conference programme committee. If you are too inexperienced to do so personally, you should tell reviewers in your department that you are keen to help them. A written review or critique should be structured, preferably from high-level comments down to lowest-level nit-picking criticisms relating to e.g. grammatical and typographical errors (and of course nit-picking should be avoided altogether unless the paper will be considered for publication). High-level comments give ones overall impression of the paper, suggest improved organization and presentation of the content, mention related work that has been omitted, and ask questions about alternative approaches, additional supporting evidence, future extensions and the like. To be constructive, try wherever possible to suggest an improvement rather than simply state what does not work.

Getting started

The best way to get started is to simply aim to get it written down, without regard for style or level of detail. Later on, you can improve the structure and language, add details, figures, discussion, examples and so on the important part is to get over the hurdles where you cannot face starting or fail to make progress because you are overly concerned about writing perfectly. Start with the part you consider easiest, to help you get going. If you encounter parts where you struggle to explain the work, jot down rough notes to yourself about what should be covered and keep going. The more you write the easier it becomes, and you will find it less difficult to tackle those sections when you return to them.

Overcoming writers block

Procrastination is common, particularly at the start of a writing task. Generally, once the process is set in motion, it is much easier to continue, to extend and refine what already exists. Some activities to get you going at the outset or when struggling to proceed, are:

- make notes on interviews, observations, experiments being performed
- make notes from your research diary
- make notes on any recent discussions with colleagues, team members and others
- draft the outline of a chapter or section, or produce a mindmap
- talk to someone about something that needs to be written

- work on a diagram, table, reference list or bibliography
- talk to a tape recorder or Dictaphone
- imagine you have to teach a class of students what you have done
- describe the work in the simplest possible way, as if you were talking to a lay person

If you can, enlist the help of friends and colleagues who can offer support, chart your progress and give feedback on rough drafts. The best such helper is a fellow student who is writing their dissertation at the same time you can reciprocate, and indeed benefit from seeing how they present their work (identifying through reading their thesis both techniques that work well and mistakes that you should guard against yourself).

Overcoming fears when it is done

When you reach the stage of writing your final report, it is often the case that your work seems obvious and insignificant, and this diffidence can cause you to delay the writing unnecessarily. Remember that your perspective arises from your familiarity and experience with the piece of research (which has increased enormously since you started), and your work will be more interesting to others than to yourself because of this. Cast your mind back to your original plans and ideas, and this will usually show clearly that the work you are now presenting is not obvious at all.

Writing is hard work, but it is the most important product of your research, and its rewards are substantial.

Writing as a team

If you are writing as part of a team, it is important to decide on a team leader from the outset. The group should write side-by-side, taking turns to write and revise different sections, and discuss the task as much as possible. With this approach, differences of opinion can be accommodated by agreeing to disagree for the present and return to the issue later in the process.

Obtaining assistance

Many universities have units or individuals who help students with writing. If you are unhappy about the writing you are doing, or if you are not writing in your home language, then you should visit such a facility to get help and advice. If there is more than one such person, try to make an appointment with the one whose background is most suited to your research field. Before arriving, prepare an outline of the document and two copies of the current draft. One copy should be kept clean, the other used to highlight the key phrase of each major section, and to annotate where there are concerns about specific pieces of writing that seem unsatisfactory. It is important to take notes when meeting with the helper, particularly about what needs to be done next, otherwise it can be very difficult to remember specific points raised after the visit, when writing is resumed. A similar approach should be used when students meet with their supervisors during the writing-up stage of the research.

Sections and Kinds of document

Document structure

A report, thesis or paper must include:

- the context in which the research was done (the study area/problem, motivation behind the project, assumptions made, and how the work relates to existing literature)
- the key issues, ideas or questions (introduced at the outset, forming the thread that ties together everything that follows, and addressed in the conclusion to demonstrate what has been learnt)

- the approach used, any experiments conducted and an honest evaluation of results
- a good structure and adequate linkages to make it easy for readers to find their way through the document, understand how sections relate to each other, and be able to start reading at any section and still understand the material.

The broad structure of the document follows from the above skeleton context, questions addressed, approach, experimentation, results. In addition, an abstract upfront summarises the work; the key points are presented in the introduction before being developed in the body; and the main results are brought together in the conclusion.

The Introduction

The introduction is the most important part of any writing, as few readers will read more than this, particularly if you don't grab their attention here. This means you should get to the point directly at the very start of your introduction. Many papers will instead start with grand sentences describing the topic or research context. To those who are knowledgeable in the field, this is a waste; to the rest, it is partly or completely incomprehensible. Even if you manage to briefly define the field in a way that researchers outside this field can understand it, this will not help if the rest of the paper is such that these readers cannot follow the material anyway.

Make it clear from the outset whether or not the system you describe has been implemented and used. The reader has a right to know from the beginning, and certainly this information must not be omitted entirely. If the work has been implemented, describe its usage as precisely as possible in what follows: who has been using it, for what purpose, and with what results. If this usage shows that the idea has not worked out as expected, this is generally just as publishable readers can learn from such an experience; negative results also contribute knowledge to the field.

Establish carefully, from the start, the assumptions and constraints on which the work is based, and how sensitive it is to these. Discuss why these assumptions are realistic. For example, if the paper describes a design which has not been implemented, convince the reader that it has taken into account the true characteristics of components used; a theoretical paper must similarly show its assumptions reflect reality.

The Conclusion

Conclusions should integrate the results of a paper and bring out the most significant ideas or results of the work. Be careful not to repeat much of your introduction in the conclusion rather give perspective that was not possible until the paper had been read, and focus on what lessons your readers should take away with them. This is the reason the work is published, so state this clearly to ensure that people benefit from the paper, and don't repeat what you have done. Give the context or scope in which these lessons can be applied, by explaining to what extent your ideas are generalisable, and making clear what assumptions or constraints you have incorporated. The conclusion is also a good place to mention or discuss conjectures, wish lists and open problems.

References

When you develop the content of your paper, you will naturally cite papers you describe or mention along the way. These are then gathered together and form the first draft of your References section, most of which probably come from that part of the paper where you compare your research with related work. Remember that your references are there to ensure that any reader, including one less knowledgeable in the field, is able to fully understand your work without you having to describe everything in detail. Papers with ideas that you have built on, should be added to your references for this reason; it also enables you to briefly acknowledge your debt in the body of your article.

Make sure that you reference the most accessible paper on the work you are citing journal or conference papers rather than departmental technical reports - and avoid personal communication or unpublished

paper references wherever possible. Researchers should read and cite authoritative and recent sources, and avoid unreliable sources such as unrefereed articles or Internet publications.

Some documents require references, others a bibliography and, more rarely, there are still others that require both. References provide details of the works you have mentioned in your text; bibliographies list all relevant material you have consulted during your research (and are hence a super-set of the References).

The Abstract

The abstract should be a summary of the technical content of the paper, not a prose table of contents as is often the case (e.g. This paper reviews and presents a solution based on . We describe how this has been implemented, and give performance results for the system, concluding with several directions for future work). State the assumptions on which the work is based, the idea/approach used, and the results. Focus only on this; leave out discussion and argument, and avoid the passive voice. As many people read only abstracts, these should be as specific and precise as possible. Write the abstract first, so that this focus is there from the beginning, to keep you from straying into irrelevant descriptions in the rest of the paper.

The extended abstract

An extended abstract is not a lengthy abstract, but a very short paper. Its purpose is to communicate your ideas clearly and concisely so that others can quickly understand the essentials of your research; it is also often used to judge your work (e.g. by funding bodies, conference organizers, potential supervisors or collaborators, etc.)

The extended abstract should thus summarize the research, its context, how it was executed, and the results. It should motivate and explain your contribution, comparing it to related work and showing key results in a way that is easy to assimilate, such as graphs and figures. Do not cover anything that isn't essential to understanding the main ideas that you wish to convey (e.g. future work or insignificant implementation details).

The thesis

Every institution has standards and conventions expected of their Masters and PhD theses, and it is naturally important to obtain these guidelines early on, and also to look at examples of (good) theses completed in the recent past. Generally a thesis comprises at least the following sections: Abstract (an executive summary), Introduction (research goals/questions, with motivation, and an outline of the thesis content that follows), Literature review, Project description (covering some or all of the following: theory, algorithms, designs, implementation, experiments, results), Conclusion (a summary relating back to the goals/questions set in the introduction, giving your main findings and an indication of future work).

How to write well

Preparatory steps before writing

Most research is part of a pipeline building on what others have done before, and in turn hopefully being built upon by researchers in the future. As such you have a duty and a responsibility to document your work well.

When making a presentation it is important to know your audience and tailor your talk accordingly; it is equally important to know your community of readers when you are writing. This is established from what you know of the book, journal or conference where you are submitting your work, and can require some extra effort to discover, if you are not familiar with that publication or forum. To write appropriately for your readers, you must know how much knowledge they are likely to have of the

general subject area as well as the specific topic you are addressing, and whether they have a special interest in it (e.g. some journals emphasise theory and others practice/experience). At the very least, authors must take into account how well the problem they are tackling is known to their readers, and if it is considered important (if only a few have this opinion, the document will have to convince the others early on).

It is helpful to look at review forms when writing your own papers, particularly if you (or your supervisor or colleague) have in the past conducted reviews for the very journal or conference series where you are hoping to publish. Naturally you should also carefully study the guidelines for that particular journal or the call for papers of the conference you are interested in.

Re-read papers which you noted as being easy to follow, particularly if you are finding some material difficult to put across clearly and concisely yourself

Both the research and the writing tasks are greatly facilitated if these highly daunting objectives are broken down into smaller task of more manageable size. The sub-tasks should not be the same size or level of difficulty, so that one is more likely to find the best thing to tackle according to the time and energy of the moment. If you are feeling tired or demotivated, you can have the satisfaction of getting all or some of a simple part out of the way choosing a straight-forward section to write ; when you are energetic and enthusiastic you can choose to tackle the more challenging sections.

The three types of IT publication

There are three basic types of Computer Science paper:

1. engineering papers describing an implemented system in its entirety, or parts of such a system, to illustrate a specific theme or idea

2. concept papers describing an unimplemented system that uses a novel technique or idea

3. theoretical papers covering formal aspects of computing, performance modeling, security verification and the like

Implementations are extremely expensive, so an engineering paper reporting on an implementation of an existing- and still unproven - concept paper is very valuable and deserving of publication. This is equally true if the implementation confirms or refutes the claims of the concept paper. Reviewers are more skeptical of concept papers, so the standard expected of ideas in such a paper is much higher. If there is any doubt, reviewers will recommend implementing and using the proposed system before submitting the work for publication.

What to cover

One of the main differences between a good and a poor presentation of a piece of work is an appropriate focus. Avoid the temptation to write down something because you know it (when giving background to the research) or have done it (when describing your contribution). Cover only what is necessary to understand your main idea, and give just enough detail for your reader to follow. Compare and relate your research to existing work, to demonstrate your contribution and also to help the reader understand your goals and your approach better. For example, mere references, as in The system is based on the model of A[1] using the language of B[2], with the algorithm of C[3] as extended by D[4] employed for . are not adequate. Poor description of related work may not be this blatant, but many papers do frustrate readers by omitting one or two sentences that could have saved the reader from find and reading the referenced work.

If your paper describes a system, decide at the outset whether you are presenting it in its entirety or just one aspect, and then write the entire paper accordingly. This will serve you much better than an article which tries to do both, changing between overview and focused writing as it proceeds.

A paper that simply describes a system/solution/task is far less informative than one that also makes mention of ideas and avenues of exploration that proved fruitless or inferior. You should also include

alternatives considered by not explored. It is therefore better to outline the path followed during your research, rather than simply presenting the final product as if you went directly to this solution.

Writing Style

It is not easy to describe the work you have done briefly yet adequately this requires good writing as well as a good structure/organization to the document, with appropriate figures to refer to. Structure your paper carefully, and follow normal practice for the type of paper and publication, where appropriate. For example, most theoretical papers have a format comprising definition, lemma, theorem, example, corollary.

Avoid forward references (as described later in section N). If you are forced to use them, give at least an outline of the term/idea to be covered later, so that the reader can continue with some notion of what this is, even if the details have not yet been given.

Visual presentation

Make sure that the visual impact of your paper is not displeasing, as this will give an initial impression of sloppiness which reviewers will then associate with your work as well as its presentation. Consider using latex and bibtex for your documents, which are particularly geared towards scientific writing. Bibtex is a useful way of organizing your references, since latex will automatically lookup references needed in a document from among all the references you have collected. If you are not good at layout and visual presentation, find a colleague who is and ask them for help.

Tables, charts and graphs can tell a story more concisely and strikingly than textual description, and are convenient ways of showing large amounts of information. Tables are more precise than graphs and charts, and encourage readers to interpret the information for themselves. Charts and graphs are more visually appealing and are a way of bringing a specific point home clearly - charts particularly effective for highlighting differences in values when the number of values is relatively small, and graphs more easily showing many values and more readily showing trends. Tables, charts and graphs should be clearly labeled the axes named and units of measure given, multiple bars or lines distinguished and labeled separately, and the entire figure given an appropriate caption. Effective documents use simple writing. They contain interesting examples to illustrate ideas. If program code, models or mathematics is used, this is also explained in plain English. If your first language is not the language of the conference or journal, it is essential to find a suitable person for whom this is their native tongue, and have him/her improve your writing with you. On the other hand, if the conference/journal language is also your own home language, remember that non-native speakers will be reviewing and reading your paper, so keep to simple sentences and plain language.

Tables

The rows and columns of a table should be given in a logical order, such as alphabetic order, chronological order, increasing/decreasing order, ordered from parts to wholes/totals, etc. Where it makes sense to do so, include totals, averages and medians. Lay the information out with independent variables in different rows and dependent variables as the different columns, and if the table is long or has logical subsections, leave blank lines between rows where appropriate. Tables are mostly used to present numeric data or results, but can occasionally be used for text entries, as long as this is not an oversimplified representation of the information.

Charts

Simpler bar charts are generally easier to read than stacked bars, and 3-dimensional bars are best avoided altogether. If stacked bar charts are used, it is best to keep the darkest shades at the bottom and the lightest on top, as people tend to over-estimate the size of the dark areas, and use numbers or connecting lines to clarify relative proportions. In general, numbers at the end of each bar help to make the chart more precise and easily readable. Bar charts are usually read from left to right (or from top to bottom, if bars are arranged to lie horizontally one above the other). You can use this fact to

order the bars in a way that highlights the point you are trying to make e.g. to show an increase or a decrease. Pie charts are usually less informative than bar charts, and particularly ineffective when there are many segments or when some segments are particularly small. If you must use a pie chart to specifically show proportions, then arrange the segments clockwise in a logical order (e.g. from largest to smallest), and explode or make darkest the segment which is most significant (if applicable).

Graphs

Graphs can take the form of lines, connected points, or area plots (which represent proportions by shading the spaces between the lines). For area plots, as for stacked bar charts, it is best to have the darkest shades at the bottom and the lightest on top. Area plots are harder to interpret than conventional line graphs, because all but the bottom value are visually affected by the lines below them (i.e. one needs to focus on the area between the lines, but it is natural to look at the line itself instead). If several graphs are shown on the same set of axes it is important to clearly distinguish the individual lines. This is particularly important with graphs having more than three lines, as these are difficult to read if the lines cross each other. The choice of axis for a multi-line graph, as for a bar chart, can affect the readers interpretation in subtle ways for example if percentages rather than absolute values are plotted, differences will be accentuated and appear greater than they really are. If the axis range is small, differences are also accentuated: for example, a difference of 10 looks small if the Y-axis ranges from say 0 to 100, but looks significantly large if the Y-axis ranges from say 80 to 100. Hence the Y-axis should be carefully selected to show what is scientifically significant in the given context.

Preparing to submit a paper

Writing is an iterative process. As new sections are written, others already drafted are revisited and refined - in the light of the new context, by incorporating what has been learnt by the writer since that section was last tackled, because the text is being reorganized, to reduce the length, etc. As with any cyclic process, it is important to recognize when to stop.

When examining drafts, put yourself in the shoes of a reviewer and ask if there is undue emphasis or de-emphasis of any aspect, or any worrying omission; and write down the main strengths and weaknesses of the paper. Consider changing the paper to address these, or stating in the paper why you have not done so.

If your document is too long, read it critically for clauses, sentences and sections that are not needed by the reader at all or that could be summarized briefly instead; delete references and quotations which are not essential; and replace lengthy descriptions by diagrams or tables (with brief explanations) where possible.

- A checklist such as the one below can be used to determine whether a draft is ready for submission:
- title page contains all required information
- number of pages is commensurate with expectations
- appropriate citations have been included and all are listed in the references or bibliography in full detail
- spelling and grammar have been checked throughout the document
- the text adheres to the requirements set out by the institution or journal
- each chapter and section is complete and covers the appropriate material

To check your spelling and grammar, make sure you use only complete sentences, and avoid jargon and slang. A spelling checker and word processor, and a friend who is not a Computer Scientist, can help with this, particularly as they will not be distracted from their task by the content (as you are likely to be). Be consistent in your style throughout your document, for example do not use the third person almost everywhere and then suddenly write a section in the first person (etc.)

Timing your submission

For successful publication, your paper should be:

- Technically sound
- Correctly timed (the work is complete, rather than almost complete)
- A clear and focused description of your work
- With results adequately and convincingly demonstrated
- Making a sufficiently important/significant contribution that others can build upon
- As interesting and thought-provoking as possible

Make sure that you do not submit your work for publication prematurely. While it is tempting to submit to a conference work that you expect to have finished by the time the conference takes place, the reviewers of your work cannot take your word for this, and will generally advocate that you resubmit when the research has indeed been completed. Under no circumstances should you submit a paper that is incomplete, with notes such as to be included in the final draft.

Presenting evidence and argument

Convincing readers of your results

The key question that readers and reviewers of your paper will ask is whether you have convinced them of the validity of your conclusions. If not, your paper will not lead to others taking up your ideas and using or building upon them, and hence your work will not form part of the research pipeline that advances the subject. Make sure you give enough evidence in the way of arguments, analysis, surveys, measurements, proofs or evaluations. Try to avoid omitting issues whose absence is likely to concern your reader. Where you have chosen between alternative design or implementation options, explain how and why you did so; and in general remember to discuss not only what has been done but also why. Mentioning alternatives that were not implemented or investigated can in itself be useful to readers by suggesting to them possibilities for future work. One of the best ways to convince people of your conclusions is to run experiments or case studies to confirm your claims (whichever is most appropriate to your task or innovation). Any activities you conducted to convince yourself that your research goal has successfully been achieved should also be described. You should also explain why you chose particular evaluation methods and approaches, in the same way as you give reasons for other aspects of your research.

Claims and evidence

Every paper or thesis has at least one major claim, the point of the research being presented. This is the core of the piece of writing, and should have the following properties:

1. It is a conclusion reached, rather than a statement on what has been done (e.g. < X > outperforms < Y > rather than we compared < X > and < Y >)
2. It is new and contestable rather than unsurprising (if readers accept the claim immediately, without puzzling over why it should be the case or whether it is true, they will not bother to read much further)
3. It is precise and specific (e.g. Under conditions < Z >, < X > outperforms < Y > because < W >)

If a claim has been explicitly stated as a specific conclusion/result, it naturally provides a number of pointers as to where evidence is needed and what type of evidence is required. Evidence must be accurate, and if any data or measurements are of questionable quality, this should be explained.

The experiments conducted should be adequate for the topic of a suitable size and sufficiently representative of the area or population studied. For example, enough users should have participated in the study, or databases/workloads of suitable size tested, etc. In presenting evidence, be as precise as you can sensibly be for your particular context, replacing fuzzy language such as a great deal, some, frequently etc. with more concrete evidence like 60% (or whatever value applies, given with appropriate precision).

The place of argument in research

All research requires some argument, if only to convince others that your idea/approach is sound and beneficial. Argument is also needed to illustrate why your conclusions follow from your results, why your survey/case study/experiment uses appropriate measurement indicators, samples and analyses; why you interpret your observations in a particular way, and so forth.

The structure of an argument

Most people find arguments arising in the course of a conversation easy to handle, and doing so in writing should not be much more difficult. The main differences are the degree of confidence required (as writing will hopefully be critiqued by many) and the need for an author to cover both sides of the argument. A typical argument involves the following: on the one hand, a claim, with evidence and reasons to support it, and with qualifications indicating under which conditions it holds; on the other hand, questions and rebuttals from those unconvinced or unclear of the extent of the claim. An author should cover as much as possible of both aspects in order to convince the reader.

An argument in a piece of writing therefore has the following ingredients:

- A claim
- Evidence to support the claim
- An explanation why the claim follows from the evidence
- Qualifications indicating the limits/conditions under which the claim holds
- Reservations, questions and rebuttals that could come to mind, with responses that address these

Arguing your case

Arguments begin by stating premises and using them to derive new facts. Premises come from facts and definitions established in the literature (quote the source so that skeptics can find details), from the conclusions of previous arguments, and from observations made in the course of the research. Examples are widely used to strengthen arguments an argument that something exists needs just one example, a complex argument may also have just one or two examples to illustrate at the same time as clarifying what is meant; arguing that some phenomenon is common or generally causes another, requires a number of diverse examples to convince the reader.

The reasons why the evidence supports and justifies the claim should be given explicitly in a sense, evidence is much like a sub-claim which must itself be explained and justified by further evidence. A collection of quotations from other documents and numbers from experimental results does not constitute evidence unless the author critically analyses and explains their meaning and interpretation. The point of a graph, table or chart should be explicitly given so that the significance of the data presented there is clear.

To cover the other side of the argument, i.e. to address concerns the reader may have, an author can start with their own earlier ideas, describing those which were tried and explaining why they were rejected. Awareness of how a claim contradicts existing belief or practice is essential knowledge, and this aspect must be addressed directly. The evidence should be read critically from an outsiders perspective to identify any alternatives or problems a reader might see; each of these should then be

described and refuted or conceded. Some typical objections: the author is oversimplifying the problem, counter-examples haven't been considered, etc.

Criteria for publishable work

Is your work publishable?

The primary purpose of publication is to present innovative and important contributions to the subject. One way of demonstrating that your work has these properties is to point out who will benefit from your research, and how. Another is to explain how your work differs from similar or related papers already published. A paper must have focus, originality and significance to be worthy of publication.

Focus

For each original idea you believe your paper is contributing, describe it in a single paragraph, to make sure that the contribution is clearly and explicitly stated. This paragraph/paragraphs should then be placed upfront, in the abstract, which people read first. If you cannot describe the idea in a few sentences, then you should not write the paper until you understand it better. Develop the idea further, formulate it in more detail, apply it in case studies or experiments, or modify it in some way until you are able to capture it clearly and concisely.

Originality

The most important criterion for publication is that a paper contains at least one original idea. To ensure this is true of your research, you must be sufficiently familiar with existing knowledge in the field. The most common failing of rejected papers is that the system/solution/theory/model/language/approach they describe is isomorphic to an existing one described in the literature. Therefore, you must first and foremost make sure that your literature search has been thorough.

Look critically at your bibliography: if it comprises mainly seminal works, check that a recent paper has not already published something that you are repeating; or if you reference mainly recent work, check that your idea was not proposed a long time ago. And when you come to convince the reader that published works differ significantly from your own, make sure that you have actually read the papers you reference (rather than merely heard of them). Then you can explain the differences explicitly and thoroughly, and can argue more convincingly particularly for those readers who actually have read what you reference!

Significance

If your idea is original and can be succinctly explained in this way, the next question to ask yourself is how much of what still remains to be said is simply a variation on existing work. If the answer is very little, consider thinking more creatively and extending your research before submitting the paper. If you are writing a paper where the main contribution is an extension or variation of existing work, ask yourself honestly whether the difference you describe is sufficiently significant to truly warrant reading all the pages you are filling. Alternatively, consider publishing the idea as a short note, short paper or even a poster. It may be that you are planning a journal paper when your work would be more suitable for a conference presentation. Journal articles are typically longer than those in conference proceedings, as they should be the more significant contributions to the field, whereas conferences aim at disseminating new ideas widely and rapidly.

Concerns about your contribution

If you encounter a paper which seems to present what you are aiming to do yourself, remember it is virtually impossible to be exactly identical. It may be tackling a similar problem/task using a similar approach. Discoveries of such papers are common in research (particularly if the work is very recently

published, or is not easily accessible). Add the paper to your literature review, and critically evaluate it with a view to improving your own project. Use the paper to compare and contrast with our own research, and study with interest where it confirms or refutes your own ideas.

Another concern that often arises when writing up research is the feeling that the ideas are little more than common sense, unworthy of publication because they are obvious or uninteresting. To overcome this problem it is best to get another's opinion, and to review early entries in your research diary to convince yourself that the contribution only appears straight forward due to familiarity with the project.