Chapter 3. Conducting a literature review

Table of Contents

What is a literature review?

A literature review summarises and evaluates published material in a subject in order to present current knowledge in a coherent and understandable way. Its purpose is to provide an integrated, organized overview of significant literature published on a topic. A literature review can be an end in itself, or a preface to primary research, or a requisite part of a research or grant proposal, or a chapter in a thesis. If it is a first step in achieving some other research goal, this research must be guided to a large extent by what you find in the literature survey – viz. what problem to tackle, how to tackle it, and how to publish your findings. A literature review is not just a description of existing work in some field – it is defined by the guiding concept of your research objective.

A literature review makes a contribution to the discipline by synthesizing existing knowledge into a framework against which a piece of research can be evaluated. It is a self-contained work, and facilitates finding out more about the ideas and results it describes by providing a good list of references to the sources from which it was compiled.

Outline of the literature review process

Reviewing the literature in a topic requires three skills: information gathering, critical appraisal and coherent writing. The first is the ability to scan the literature efficiently, read documents effectively and record relevant material; the second is the ability to analyse studies to discover precisely what they contribute to our understanding of a subject while also being able to detect bias, flaws and limitations where these exist. The third skill requires creativity in finding a framework and structure for the knowledge acquired, and the ability to present this in a clear, logical way.

A literature review requires finding some starting point(s), reading and noting what you find, and changing your list of things to read accordingly – repeating this process until you conclude that you have covered what is currently available. You then need to find an appropriate way of presenting the information you have gathered, and write the review in your own words, with references to the articles and texts you have read. Of course a literature review can never be complete, because new knowledge is continually being published. Therefore it is important to keep abreast of the field even after you move on to the subsequent steps in your research project, and to continue looking for, reading and noting new ideas as they appear in the literature. To keep reading throughout your career is particularly vital in a field such as IT which changes so rapidly. Keep up with the latest issues of the key journals in your field as they arrive, and remember that Computer Science conferences are an excellent way of tracking what others are doing. In cases where new results appear after you needed them in your

own work, it is helpful to note the date they became available and what stage you had then reached in your research.

Where to start a literature review

The best way to start any literature survey is to ask those actively doing research in the field, and also to ask a librarian. They can advise you which are the best sources to use and how to search them effectively.

If your research concerns a particular product, process or existing study, then that will clearly give you a starting point; without this, finding some good initial references can be time-consuming and frustrating. Bibliographic databases, CDs of journal abstracts, citation indexes, textbooks, journals and the Internet can lead you to many initial references. Appropriate search terms are needed to uncover the most suitable ones. Write down a list of key words and phrases describing what you are looking for, and include synonyms, alternative names and word variations. If your area of interest is broad or unclear, it is helpful to look over recent copies of journals in that field (e.g. ACM Transactions on) to get a better idea of what you are wanting and which search terms are most likely to lead to this. If you find references which appear promising but cannot be found in your library or on the Internet, contact the authors personally as well as their institutions, and try an inter-library loans service. You will be surprised how an email request for a paper can not just lead to a convenient electronic copy being sent to you, but also result in a valuable contact who will give you good insight and advice on the problem you are researching. Of course, if the author in question is a well-known expert in the field, such a relationship is far less likely – nonetheless, even experts will generally send you a paper by email.

If you are unfamiliar with the subject area, don't begin with bibliographic databases and research articles, as such papers have a narrow focus. You must first obtain a good background in the subject; it will help you to conduct a better search as well as ensuring that you understand the material you read better than you otherwise would have. This can be done by using secondary sources, namely handbooks, textbooks, review articles and encyclopedias. Secondary sources summarise and integrate information on a subject, drawing from primary sources already published; often they also give bibliographies as well as references to major papers in the field.

Many researchers advocate using a pack of index cards (3" x 5") to control their search, (and others use a database package). The top card contains the Main Research Question you are addressing. This question will keep changing as your search proceeds, and you may well need to clip together several such cards due to many changes (keep them, do not discard old versions). The remaining cards are headed Main Topic – each has one major topic to be addressed, and you should produce as many as you can think of. Keep these in alphabetical order. Start searching for the first topic in catalogs, journals, CDs, bibliographic databases, the Internet etc. When you find relevant documents, look at the subjects associated with them and use these to expand or contract your search. Annotate your index cards with comments such as "nothing!" or "too much".

Explicitly-related literature is directly related to your topic; implicitly-related material is not, but concerns work that might be applicable to your study. If you use the latter in order to produce a literature review, you will have to be creative in tying together seemingly disparate pieces of work and relating them to your research problem. If you cannot find enough explicitly-related work and find yourself settling for implicitly-related papers, you should probably change your research goal. Fortunately, searching such implicitly-related work often gives you good ideas as to what you might tackle instead –starting with a good background in the area. Moreover, concrete examples of similar work show what research methods are typically used and how results are presented for publication in that field.

Citations

When paper B references paper A, this is called a citation of A by B. While the number of references that a paper like B makes is finite, easy to determine and does not change, the number of citations of a paper like A is continually growing, potentially infinite, and difficult to ascertain. The number

of citations that an article has is an excellent indicator of its importance. The number of citations to specific journals, conference proceedings or monographs also has a strong influence on the number of libraries that stock it. Citation indexes like the Science Citation Index show the patterns of citations into and out of journals. Online citation indexes include Citebase (the Open Citations Project) [http://citebase.eprints.org]. When you are conducting a literature search, use a citation index to see how significant a work is and how long-lasting its influence; when you are publishing, use a citation index to see how important the journal is. In addition, many countries like South Africa have rating systems for journals based on their citations, and these are a good way of easily discovering the impact and importance of the different journals in your field. A study of citation indexes shows that papers generally cite papers of similar impact, and that very few papers receive a large number of citations – these are key papers which no literature survey should miss.

Literature searches on the Internet

Start with a well-known search engine such as Google or yahoo, or a well-known bibliographic database, and compile an appropriate Boolean expression involving search terms you are interested in and those you are not interested in. If you are unsuccessful and receive too many or too few results, you will need to change your search terms and expression, and possibly also the search-engine or database you are using. Ask your librarians for help with search terms and expressions.

There is an increasing number of on-line journals. Thousands of articles, abstracts, reviews, books, theses, etc. are available electronically on the Web – so you should keep your initial search as narrow as possible.

Some helpful Computer Science websites are the CiteSeer citation index [http://citeseer.nj.nec.com/cs], which has electronic copies of the vast majority of publications, bibliographies such as that at http://liinwww.ira.uka.de/bibliography/index.html, online contents pages of journals such as those of the IEEE Computer Society [http://www.computer.org/publications/dlib], technical report libraries (e.g. the Networked Computer Science Technical Reference Library [http://www.ncstrl.org]), digital libraries available at your institution (e.g. the ACM Digital Library at UCT), as well as the websites of relevant Computer Science departments and individuals.

Reading throughout research

The basis of good research is finding and reading good related work. One of your most important secondary research goals is to produce a bibliography and review of the topic you are studying. But to keep reading after this is fundamental to success as a researcher, especially in computing where many advances are continually occurring.

Even if you need to study a paper thoroughly, it is helpful to first skim the entire article to get an overview on which you can “hang” sections when you subsequently read them more carefully the second time. Many people find it easier to make notes as they read, but be careful of writing so much detail that your reading takes too long.

Some articles in Computer Science are very hard to read. It can help to read them with your supervisor, research group, or others in your department who are interested in the same piece of work. It will be much slower but will be enormously helpful in raising everyone’s understanding. Larger research groups typically find regular get-togethers to discuss papers very beneficial, not only because it facilitates keeping up with the literature but also because the group will usually look more creatively at questions raised by a paper than would a single individual.

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 one’s 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.

Reading Efficiently

Naturally it is impossible to read all the relevant background material for your project, in all pertinent books, journals, websites, etc. The key is selectivity, not only in choosing your sources but also in how much of each you read. At the start of your literature search, you are likely to come across a mix of highly relevant, less relevant and irrelevant information. At this stage it is best to read papers quickly to get an overall picture and an idea of how your specific interest fits this. Decide whether it is worthwhile reading the entire document, and if so in how much detail to study it. As your search progresses, you will find yourself considering few texts that are not truly relevant, but it is still worthwhile using the same approach.

Can you take a book relevant to your research that you have not read before, read it in 5 minutes, and have notes covering what you want from it at the end of that time? This requires practice, focus and discipline. A technique used by many researchers is to skim the material, looking mainly at headings and figures, and to use this quick preview to think of questions the text is likely to answer for you - then look for those answers. Another common approach is to look at the table of contents to identify which chapters or sections are useful. Read the introduction and conclusion of these chapters, then look for sub-headings indicating relevant parts and for key points highlighted in the text itself. The first and last paragraphs of a section generally encapsulate the points that section is covering. If there are specific systems, people, keywords etc. of interest, look for them in the index. The references are another way of discovering very quickly what the paper or chapter is about. After the first few minutes you will have an idea whether the book (or paper) requires more detailed reading, and of which parts. In the case of a book you should rarely be reading more than a quarter of its contents.

Making notes on your reading

Having read a seemingly relevant paper, ask yourself whether or not you should be skeptical of its claims. Not all papers are equally sound – the status of the source (e.g. refereed journal through to unknown web site) and of the authors (ranging from guru in the field to completely unknown) gives you an initial idea of how cautious you should be. The article itself should be read critically, rather than accepting at face value what authors profess to have discovered.

When you read a document, ask yourself questions, such as:

- What is the contribution of this work to the field and how does it compare with and relate to papers you have already read?
- What references does it cite?
- What useful examples does it contain?
- What can I learn about writing from this paper – is it easy or hard to read and why?
- What techniques and research methods have been applied by the authors?
- What is the author’s thesis, i.e. what is s/he trying to convince you of?
- How does the author go about convincing you of this?
- Does the author succeed in convincing you?
- Does the paper explain how the work differs from other work in the field? If not, can you do so?

- Does the author suggest future work to follow on what has been reported? What new ideas are generated by the research described?

- Does the paper raise any questions in your mind that it does not answer, e.g. relating to potential problems with their idea or the ability to apply it in a different context or the possibility of generalizing the results?

Make notes of interesting answers. Summarise the main ideas and the approach taken. If the paper presents a new way of doing things or a new way of thinking about something, note and describe this too.

Notes help considerably if you return to re-read the paper later or when you come to include it in your literature review. Whereas note-taking delays reading and can seem tedious, remember that it is extremely difficult to do this well, so see it as a challenge and watch yourself improve with experience.

The amount of notes you make should be tailored according to the degree of relevance of the specific piece. You should also note any potentially useful research or presentation techniques, which you might adopt yourself. Use the article to improve your literature search itself, by altering your search terms and making a note of new references they contain (and why they might be useful to you). Take note too of references to documents you have encountered before, since the more frequently you come across references to the same paper, the more important it is to read it. If your reading raises a question to which you are unable to find an answer in the literature or in your department, it is worthwhile contacting the author(s) by email with a clear, concise question phrased in such a way that it can be answered in a single sentence. Once again, such email contact can result in a great deal of help; in other cases, even busy researchers often find the time to send you a reference where the answer to your question can be found.

With each good article that you read, you need to keep a complete reference to it for your bibliography and to make notes in your own words of its content and how it might relate to your own work. It is advisable to use bibtex to keep an electronic copy of your reference list, this is convenient for your future writing and also for sharing your bibliography with others. Many people add extra fields to their bibtex entries, such as keywords, or where they obtained the paper.

Structure of a literature review

A good review organizes material well and includes critical thought. It typically has the following structure: the introduction describes the problem area and its significance, the author's reason for doing the review, and the organization of the review document. This is followed by a background section explaining any underlying principles, explaining criteria used to categorise, analyse and compare the works, and also defining and justifying the scope of the review. Your findings are then presented – first, a description of current knowledge, starting with general aspects and gradually focusing on specific research hypotheses, purposes or questions.

It is usually helpful to divide the material into categories, and to explain how each work differs from others in its category. Authors of academic texts often create new taxonomies to describe their research findings, but this is often embedded deeply within the document, and inexperienced readers may not be able to "stand back" sufficiently from the content to see the organization of knowledge. The author of a literature review frequently creates his/her own taxonomy of the topic, based on historical development or similarities of approach and so on. In your review you should introduce this taxonomy first, and then for each category in turn explain what it is and why it is important, before describing works in that category and showing how they fit into this category.

Conclude with a brief summary (keep it at two or three sentences), an indication of any trends, conflicts or themes detected, and preferably some conclusion reached – such as a statement of gaps identified or issues yet to be adequately addressed (ideally this in turn establishes the need for the research undertaken). Make sure that your summary is an accurate summary of what you have actually presented in the review. It should summarise major contributions of significant works and maintain the focus established in the introduction.

The review should read as an essay rather than a list of paper summaries, and focus on the findings and contributions of the works reviewed. Most initial drafts comprise a sequence of text summaries, particularly as the writer may have insufficient confidence at this stage to include his/her own opinions and knowledge. Start to improve on this in subsequent drafts by linking these summaries – e.g. comparing their differences or highlighting their similarities or simply indicating their relationship to each other. Later, when you are sufficiently familiar with the field and have done some work in the area yourself, add your own understanding and evaluation of the material. The value of a review is how it places each work in context and relates it to other work. Where applicable, it should resolve conflicts among seemingly contradictory studies and point the way forward for further research. In the end, your review should be a description of your understanding of the topic, supported by information gleaned in the papers you have read.

If you are producing an annotated bibliography, you will need to briefly summarise each item, but should still follow themes and do some critical assessment of the works. Group the items into sections, with a paragraph introducing each section, and use this to assist in comparing and relating works in the same section.

How to critique a paper

To critically analyse the literature you review, follow through a set of concepts and questions, comparing works to each other according to how they deal with these concepts and questions; and evaluate the papers you describe, discussing their strengths and weaknesses. To do so, ask the following questions about the documents you include: has the author clearly formulated a problem/ issue and established its significance? Has the author evaluated the literature relevant to this problem/ issue? How reliable and valid are the measurement and conclusions? How does it contribute to our understanding of the topic and how can it be useful in practice? Does it have any limitations and could it have been approached more effectively in a different way ? How does it relate to my own research goal?

To assess a paper, you need to know what research question or goal it is addressing, what methods were used to do so, and what results it claims. Look at the evidence or support on which these claims are based and evaluate their soundness and completeness. Are the claims reasonable? Were the test cases all merely “toy” examples? Are significant aspects missing from the paper? Also consider how the work fits in with other work in the field, and what follow up steps are suggested. To understand a paper properly, you need to fully understand the problem, the alternatives possible and the choices made, the assumptions behind the solution selected and whether these are reasonable, the conclusions and their justification, and the potential for extending or scaling the solution.

In evaluating each document, ask yourself if the author’s arguments are supported by evidence such as case studies or statistics, if relevant information has been ignored in order to prove the author’s point, if the work contributes in any significant way to our understanding of the topic. Criticism should be balanced and justified, set forth for the sake of improved systems or understanding. Any aspect of existing work can be criticized – assumptions, arguments, methodologies, results, interpretation. To challenge an argument, one can challenge one of its premises, or challenge the fact that the premises together imply the conclusion. Examples can be challenged if they are not representative of the claim being made; the context, constraints or exceptions to a premise can be used to refute or weaken an argument; the implication that the claim follows the premises can be disputed by means of a counter-example; and so on.

Writing the literature review

It is helpful to read other literature reviews before commencing your own, to see how they organize material and manage the transition from one topic to the next. Once you have completed your review, ask as many friends and colleagues as possible to read it; their comments will help improve the review by indicating what is unclear or difficult to follow.

Avoid detailed descriptions (after all, the original paper is available to the reader of your review). In general the amount of space devoted to a single work should be proportional to its significance in the

field. Limit the number of direct quotations, as these are often too detailed and seldom unaffected by being read out of (their original) context. Too many quotations means a variety of writing styles is being used, which is disconcerting for the reader and also raises the question of whether the authors of the review have in fact failed to understand the concept sufficiently well to explain it in their own words.

Reporting verbs are used when talking about texts Thompson and Ye (1991) distinguish between those denoting what the author has done (positive examples: points out, emphasizes, notes, subscribes to; negative examples: questions, attacks, dismisses, disputes, opposes, questions, rejects) from those denoting the review writer's stance (factive examples: acknowledges, demonstrates, proves, throws light on, identifies; counter-factive examples: confuses, ignores, disregards, neglects); non-factive examples: believes, claims, proposes, uses). Factive verbs are used where the review writer believes the author to be correct; counter-factive verbs are used where the author is considered incorrect; non-factive express no opinion on the author's accuracy or inaccuracy.

For a discussion on writing style and technique, see the chapter on Writing later in the course.