



Writing a Literature Review

Contents

The literature review
Why write one?
Process
Choosing a structure
The bottom line
References

The literature review

A *literature review* is a survey of reputable, scholarly work that has been published in a relatively specific topic area. A good literature review is more than just a catalog of vaguely-related papers: rather, it is a synthesis, with commentary, of a significant body of related work. It exposes similarities and differences; draws parallels; points out issues, concerns, and open research questions; offers criticism; and draws conclusions.

You will find literature reviews everywhere: in the introductory sections of conference and journal papers, in survey-oriented journals such as *ACM Computing Surveys*, as academic term papers, and as the second chapter of most Master's and Doctoral theses. Literature surveys may range from a few paragraphs to hundreds of pages. The odds are good that you will have to write several of them during your time as a graduate student.

Since literature surveys are so common, much has been written about them. In addition to this document (which provides a personal perspective), you might want to look at [1] and [2]. For those in search of further guidance, a list of references on the subject of literature reviews can be found at [3].

Why write one?

You typically write a literature review for one (or more) of three reasons:

1. to educate yourself regarding a particular field, frequently so to assist you in identifying gaps in the literature which represent potential research topics;
2. to demonstrate to someone else (your course instructor or your thesis examination committee) that you have educated yourself; or
3. to educate others by providing a synthesis, overview, or framework for understanding the existing literature.

How you structure your review and what you choose to include will be at least partly determined by the reasons you are writing it, so make sure you understand them before you start.

Process

The major activities in writing a literature review include:

- determining the scope of the review;
- finding, reading, selecting and rejecting articles and books to include;
- making notes;
- determining what facts, issues, themes, *etc.*, are relevant;
- making inferences, identifying lacunae, and drawing conclusions;

- determining the gross and detailed structure of the review; and
- actually writing the review itself.

You will notice that I called these "major activities" rather than "steps". While they are presented in an order that may appear logical, in fact these activities are deeply intertwined. For example, it is impossible to determine the scope of the review without having selected the source material; conversely it is impossible to determine what source material is relevant without knowing the scope of the review. As with all the other seeming paradoxes in academic work (and life itself) we cope by doing a little bit of everything, in an iterative fashion, and muddling through.

Determining a scope for the review as early as possible is important. Scopes tend to "creep"; that is, to get broader as you discover more potentially-related material. Don't be afraid to alter your scope as you proceed, but make sure it stays small enough to meet whatever your objectives are in conducting the review.

One frequently asked question is "how much source material do I need". The answer depends on the scope and nature of your review. For a doctoral dissertation depth survey, you may include hundreds of papers. For a term paper in a graduate course, seven to ten *good* papers is frequently enough. One obvious point related to scope and number of papers particularly bears emphasizing: choosing which source material to *leave out* of your review is every bit as important as choosing what to include. You have limited time and space: no matter what the purpose of your review, make sure you use it wisely.

There are several good suggestions in [2] about what kinds of questions you need to ask yourself as you study your source material.

Choosing a structure

Your literature review needs some kind of organizing principle. In my experience, once you understand the literature and have found a structure, writing the review itself is fairly straightforward. Possible structures include:

the project-by-project summary

In some fields, the interesting work happens at the level of projects or systems developed by particular research groups. Each of these will have been documented in one or more publications. In a project-by-project or system-by-system summary, each section of your paper talks about one project or system. You will also need an introduction, plus a section or more at the end that draws together the common themes that occur across the projects.

the historical review

An historical review is similar to a project-by-project summary, except that the main organizing principle is time rather than the project. If the subject you are reviewing really does have an interesting historical dimension, this can work well; however, many areas do not---so be careful.

the thematic survey

A thematic survey is organized around the important *ideas* in a particular area. The survey explains what each of those ideas are, and then indicates how each of the projects, systems or publications that you have referenced contributed to the development of the idea. Of course, the hard part is figuring out what the important ideas are! You can do this only by reading your reference material attentively and identifying the important problems, discoveries, designs, experiments, incremental advances, and so on.

Each of these have advantages and disadvantages, which are discussed at some length in [1]. (In [1] the annotated bibliography is also listed as a type of literature review; however, in my view an annotated bibliography is more of a possibly-communal note-taking activity, and not a true literature review at all.)

For an example of a system-by-system summary, see section 3 of [4]; to see essentially the same material in a shorter thematic summary form, see section 2.1 of [5]. The system-by-system summary in [4] allows for significant detail about each project, but inter-project links and evaluation are relegated to occasional comments and the discussion section. Conversely, the thematic survey in [5] provides significantly less detail about each system, but provides a much more coherent overview of the important ideas in the area. In my own (admittedly somewhat biased) opinion both of these surveys are valuable, and each meets the goals for which it was written---the survey in [4] was written to educate others and has been used as a resource in some graduate level courses; [5] was written to convince my doctoral committee that I had done my homework, and they did give me the degree!

The bottom line

For course term papers and chapter 2 of your thesis, what you really want to write is a **thematic survey**. Writing a historical or project/system-based review is often acceptable, but is really a week substitute. Why? You can't write a good thematic survey well unless you have deeply understood your topic area---and if you *do* write a good one, you will have conclusively demonstrated your deep understanding to whoever is evaluating your work.

— Greg Phillips (<http://last3.in/greg.html>)

References

- [1] Saul Greenberg *et al.* How to write a literature review (<http://pages.cpsc.ucalgary.ca/~saul/wiki/pmwiki.php/Chapter1/HowToWriteALiteratureReview>). Saul Greenberg's Chapter 1 Club Wiki, University of Calgary. Version of December 23, 2006, at 04:06 PM.
- [2] Dena Taylor and Margaret Proctor. The Literature Review: A Few Tips On Conducting It (<http://www.writing.utoronto.ca/advice/specific-types-of-writing/literature-review>). University of Toronto Writing Support Centre. 2007.
- [3] The Literature Review (<http://www.deakin.edu.au/library/findout/research/litrev.php>). Deakin University Library. Version of 9th October 2006.
- [4] W.G. Phillips. Architectures for Synchronous Groupware (<http://symbiosis.rmc.ca/pub/arch-for-sync-groupware-tr1999-425.html>). Technical Report 1999-425, Queen's University, Kingston, Ontario, Canada. May 1999.
- [5] W.G. Phillips. The Workspace Model: Dynamic Distribution of Interactive Systems (<http://symbiosis.rmc.ca/pub/phillips-workspace-2006.html>). PhD thesis, Queen's University, Kingston, Ontario, Canada, July 2006.

From <http://last3.in> © Greg Phillips. Unlimited use is permitted within the Government of Canada. All other use is governed by the Creative Commons by-nc 3.0 license; see <http://last3.in/colophon.html> for details.