

Chapter 12

Reporting a Systematic Review

12.1	Planning reports	157
12.2	Writing reports	158
12.3	Validating reports	162

The final phase of the systematic review process is to document or report the study in ways that are suitable for the intended audiences. The context for this phase is illustrated in Figure 12.1.

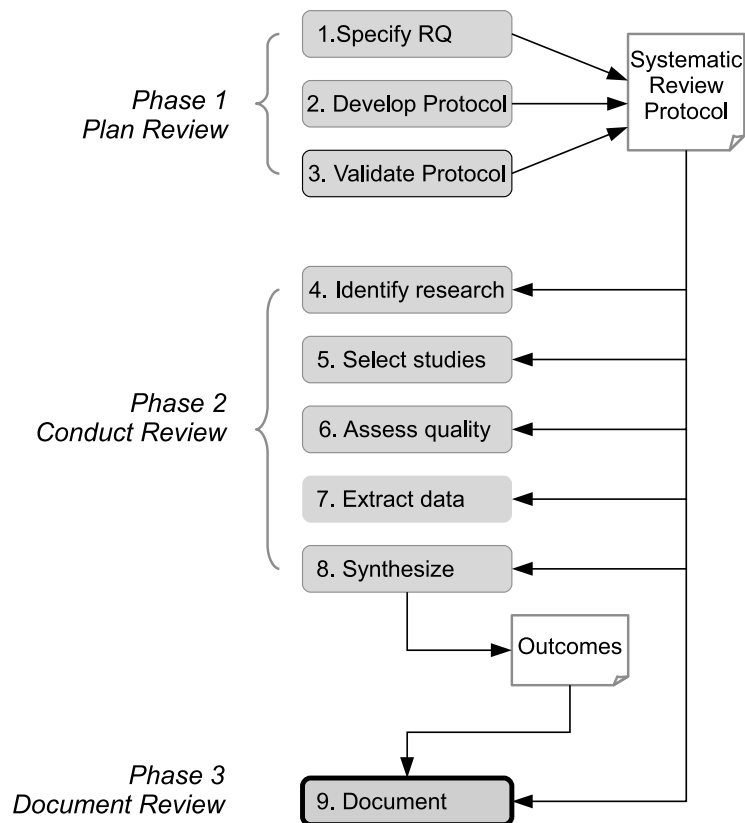


FIGURE 12.1: Reporting phase of the systematic review process.

Although, as we note below, the reporting of a systematic review may need to address a wider audience than would be expected for a conventional research paper, it still needs to answer the three key questions that Mary

Shaw identified in her tutorial on *Writing Good Software Engineering Research Papers* (Shaw 2003). These are:

- What, precisely, was your contribution [to your research field]?
- What is your new result?
- Why should the reader believe your result?

Indeed, one of her sub-questions expanding on the last of these does explicitly ask “what concrete evidence shows that your result satisfies your claims?”, a question that is directly relevant to this chapter. As Shaw observes, getting one’s work accepted for publication needs both interesting and sound results and also good clear communication of these to the reader. So far, we have been concentrating on how to obtain the results from a review (addressing the first two of the questions above), and now we need to consider how to communicate these results to others.

As we indicated in Section 4.4.9, a review can be reported in a number of ways. A journal paper, technical report or thesis chapter can usually provide full details of the process followed (including the reasons for the decisions taken and the approaches used for validation), together with the outcomes of each stage of a review plus the conclusions drawn from the results of a review. In contrast, conference papers (and sometimes journal papers too) are usually limited in length (often to 8 or 10 pages) and so can only cover some of the aspects of a review. And of course the use of these can be combined, with full details being given in a Technical Report, and the key aspects perhaps summarised as a conference or journal paper. Certainly, publication in some form in a refereed venue can contribute towards persuading others of the validity of your results.

So, in this chapter we consider three steps that need to be performed as part of the reporting phase.

1. *Planning* reports—which involves specifying the possible audiences and deciding what sort of document would suit their needs. Ideally, planning of the final reports will be initiated during the preparation of a review protocol.
2. *Writing* the actual reports.
3. *Validating* the reports—which involves asking internal, and possibly external, reviewers to assess the quality of reports.

The wider process of dissemination and publicising any important outcomes is outside the scope of this book. You can find a discussion of the issues involved in dissemination in Petticrew & Roberts (2006), and some aspects of this are also addressed in Chapter 14.

12.1 Planning reports

The nature of a systematic review means that the reporting of its outcomes is very likely to target a wider set of readers than would be the situation for a conventional research study, aimed simply at other researchers in a specific topic area, or in the case of a PhD thesis, the examiners. For example, Booth et al. (2012), identify six other groups of stakeholders that might be interested in medical and sociological research comprising: research funders, policy makers, practitioners, the research community, the media, and the general public.

While software engineering research should also be relevant to the same categories of stakeholder, the low levels of awareness usually exhibited by the media and the public means that, for the present, software engineering attracts little interest from these groups. Nevertheless, it is important that systematic reviewers make their results accessible to practitioners as well as the research community. In terms of policy making, the results may also be of relevance to standards organisations, and to authors of software engineering guidelines and text books. In contrast, mapping studies are usually of benefit only to the research community.

In planning therefore, it is important to identify the groups most likely to be interested in any outcomes (obviously, at this stage, you don't know what the actual results will be). Doing so may also provide an element of feedback to the design of the review. If changes to the research question(s) could provide outcomes likely to be of greater interest to one or more group, then this is the stage at which it should still be feasible to make such changes.

As already noted, to ensure that systematic review and mapping study reports are of benefit to the research community, they will need to be published in refereed journals and conference papers. All the details of the review process, data extraction and analysis/synthesis will need to be reported, including the rationale and criteria for excluding any 'marginal' papers (and hence, it is necessary to ensure that these details are all recorded while conducting the review—which in turn will also affect the planning for data extraction). In particular, the references for all the primary studies need to be provided, as well the data extracted from each paper, including quality data and data related to the specified research questions. In some cases, this can mean that the size of the resulting report exceeds journal and conference size guidelines. In such cases, the report should be supported by ancillary information such as a technical report and/or a database, if possible, held in an online repository. Planning should therefore take account of the likely reporting forms and, possibly, of the scale of these.

One way to make the results of a systematic review more accessible to other interested parties, particularly practitioners, is to produce shorter versions of the report concentrating on the practical implications of the results. These

shorter versions can be directed to software engineering ‘magazines’ (as long as you do not violate originality requirements). Again, the readers of shorter versions of the report need be able to easily find a copy of the full report and be able to access any ancillary material.

12.2 Writing reports

The use of systematic reviews and, to a certain extent, of mapping studies, are advocated because they adhere to a rigorous methodology. So any resulting report should demonstrate clearly that you have used an appropriate systematic review process, and also that you have used it rigorously. In particular, a report should show:

- *Traceability*—providing the reader with a clear link from the research questions to the data needed to answer the questions; from the data to the data analysis; and finally from the data analysis to the answers to the questions and the study conclusions. This is relatively straightforward for quantitative systematic reviews and mapping studies, but may be much more difficult to demonstrate in the case of qualitative systematic reviews.
- *Repeatability*—ensuring that the methodology is defined clearly and in sufficient detail that other researchers could replicate it. This does not mean that other researchers would obtain exactly the same search outcomes and results. Time differences, at the very least, would make the results of searches different. Furthermore, differences are quite likely to occur for qualitative systematic reviews where researchers often use different synthesis methods which can result in different conclusions. Nonetheless, if using the same basic protocol, researchers should get broadly similar search results and be in a position to identify and investigate any divergences in conclusions.

Booth et al. note that many people prefer the term ‘replicability’ to avoid implying claims for laboratory-like repetitions. (We discuss the issues relating to replication of primary studies in Chapter 21, where we also note that a ‘differentiated’ replication of a study may be useful for determining the boundaries or scope of an effect.)

The structure of a report/paper is usually fairly well-defined. While there may be variations, we would suggest that the following outline is one that addresses most of the above needs.

1. The *abstract*. The role of the abstract is to aid *selection*, by providing the reader (or search engine) with enough information to suggest relevance.

As noted earlier, we are advocates of structured abstracts (Budgen, Kitchenham, Charters, Turner, Brereton & Linkman 2008, Budgen, Burn & Kitchenham 2011), not least because their use encourages an author to include relevant information.

2. *Introduction.* A major role for this is to set the context, and following Mary Shaw's criteria, to make clear why this study is a useful contribution to a particular research field and why a systematic review is appropriate. So this is where we usually pose the research questions too.
3. *Background.* Usually this relates to the topic of the review, expanding on the description provided in the introduction, and where relevant, providing information about previous studies or reviews (whether expert reviews or systematic ones) and their contributions.
4. *Method.* This is where the core elements of the research protocol should be included in the paper, justifying the choice of the type of review, and the plan for its conduct, as well as the rationale for any other choices involved.
5. *Conduct.* This section is usually used to highlight any *divergences* from the plan, as well as to provide information about how well the team agreed about such issues as inclusion and exclusion (including providing kappa (κ) values to indicate the level of agreement where appropriate).
6. *Results.* This section usually describes the outcomes of searching and inclusion/exclusion, as well as of data extraction. As a section, this needs to be factual and thorough, leaving most of the interpretation for later sections.
7. *Analysis.* The outcomes from the synthesis process are described here. For a mapping study this may largely consist of tabulation and grouping. Some ideas about other forms of representation to use when reporting synthesis are covered in Chapters 9–11.
8. *Discussion.* This section is where the outcomes from analysis are considered within the wider context, and as such, has a large interpretive element. This is also where we assess the limitations of our study, through a discussion of the threats to validity.
9. *Conclusion.* This should seek to address how well the research questions have been answered, and what the answers are. This section provides the 'take home' message of a paper and so it is important for it to be concise and well focused, building upon the *Results*, *Analysis* and *Discussion* sections as appropriate.

Further ideas about how to report the outcomes from a systematic review may be obtained by consulting the PRISMA guidelines discussed in Section 7.2.1.

The way in which a report is written and presented is important too. There are many good textbooks about how to write technical reports and papers and so we have not tried to cover this aspect in any depth. However, there are some recurring issues that we have observed when refereeing or reviewing reports of systematic reviews, and so these are briefly discussed below.

Use diagrams and tables. Many authors illustrate the process involved in conducting the early stages of a review with a diagram that shows how this was conducted and the number of studies being retained at each stage of searching and inclusion/exclusion. Usually this can provide a good visual summary for the reader, and makes it easier to write the description of the process, since this can be ‘written around’ the figure. Figure 12.2 provides an example of such a figure, and is based upon one that we ourselves used when reporting a systematic review of empirical studies of the UML (Budgen, Burn, Brereton, Kitchenham & Pretorius 2011). We are not suggesting that this is the only way to structure such a figure, simply that they can play a useful role.

Tables can also provide useful summaries of complex processes. For example, when searching multiple sources for primary studies, tabulated results can illustrate the numbers found from each source, and also how many of these were unique (that is, not also found from other sources). Table 12.1, reproduced by permission of the Institution of Engineering & Technology, illustrates this using some of the values from a study we did into how reproducible systematic reviews were, published as (Kitchenham et al. 2012).

TABLE 12.1: Example of Tabulation: Papers Found at Different Stages

Papers papers	Digital Libraries	Scopus	Additional papers found from references	Additional papers from previous search	Duplicated reports
Search strings	1480	1275			
After initial screening	160	94	22		
Unit testing papers after 2nd screen	39	10	2	8	10
Regression papers after 2nd screen	25	6	2	2	9
Total testing primary papers	64	16	4	10	19

Take care with tenses. One of the arguments for producing a full and thorough research protocol is that it then helps when writing the final report.

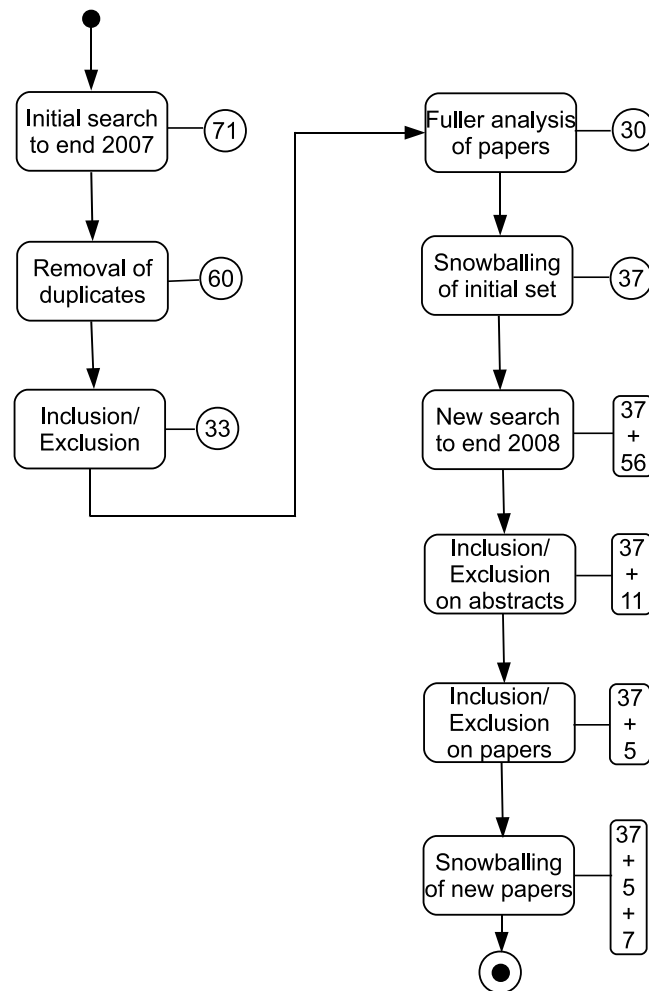


FIGURE 12.2: Example of a graphical model for the selection process. Reproduced with permission.

This is indeed very true. However, this is not just a matter of ‘cut and paste’, and one reason for this is that much of a protocol is likely to be written using future tense (what we are going to do), whereas a report is largely written in past tense (what we did). So yes, by all means reuse the protocol, but do take care to edit it so that the tense for the resulting report is consistent.

Avoid inventing new terms. Computing in general is a discipline that is rather apt to reinvent its wheels, and authors of computing (and hence software engineering) papers are often prone to try to introduce new terms to describe what they do. We strongly recommend that you avoid doing so when writing a report about a systematic review, largely because:

- the terminology used in systematic reviewing has been around for quite some time now (especially for the analysis elements), and should be quite adequate to describe what you are doing;

- using ‘standard’ terminology makes it easier for others to find your work (which after all, is usually what you want them to do of course)—it is not unheard of for systematic reviewers to be found complaining about failure by the authors of primary studies to use consistent terminology, while demonstrating just such a failure on their own part!

Conduct a review of the report. Because a systematic review can involve many people, different sections of the report may well have been written by different authors—and this sometimes shows in terms of differences of style, terminology and even grammar and tenses. Given that there are some marked differences between American and British English, using a mix of these can also be confusing¹. So, once the report is drafted, an important task is that of an editorial review, carried out by one or two people, with the aim of making the paper more consistent and (one hopes) readable. The purpose of this type of review is editorial rather than technical, and we address the latter role in the next section. Also, where the paper is written in English, and none of the authors is a native-English speaker, we suggest that it may be useful to get a native speaker to go through the report, both to check grammar, and maybe also to revise any idioms that might not have translated well. (As we ourselves are essentially mono-lingual, we do admire the ability of others to write reports in another language, but as referees, we are also aware that even good papers can have some strange phrasing in places.)

12.3 Validating reports

Editorial aspects apart, all report authors have a responsibility to read and review a report, with the aim of ensuring that the following situation is true.

- The research questions are clearly specified and fully answered.
- The research methodology is fully and correctly reported.
- There is traceability from the research questions to data collection, data synthesis and conclusions.

¹The phrase ‘two nations separated by a common language’ can be very true. As an illustration, fortunately unlikely to appear when reporting a systematic review, consider the use of the word ‘momentarily’. To Americans, this means ‘in a moment’, while to the British it means ‘for a moment’. So, consider the effect on British passengers of American air-crew announcing that “we will be landing momentarily”!

- All the tables and figures used to present the results are correct and internally consistent.
- In the case of systematic reviews, the conclusions are written clearly and are targeted both at researchers and also practitioners.

If possible, reports of systematic reviews should be independently reviewed. Within a research group (for example, a university department), colleagues may be willing to undertake independent reviews of reports on a quid-pro-quo basis.