Writing a survey

by Florian Daniel, 15 January 2020

Knowing the **state of the art** in a given domain is the most important precondition for making a new contribution to that domain.

With *state of the art* I mean the most recent snapshot of scientific knowledge, commercial products, practices or methods that are part of a well defined domain. A *domain* (or area) is a delimited field of research (e.g., business process management or artificial intelligence), business (e.g., enterprise resource planning software) or practice (e.g., developing software or designing a building). It can be *broad, coarse grained* (e.g., artificial intelligence) or *specific, fine grained* (e.g., similarity metrics for unsupervised machine learning). The granularity depends on the purpose of the survey and is chosen by you, the writer. Finally, with *contribution* I mean a scientific contribution, that is, new knowledge that can be added to the state of the art.

Misconception 1 (a survey only describes the state of the art and does not make any new contribution): If well done, a survey is an excellent instrument to (i) learn a lot about a given domain and (ii) communicate that knowledge to the reader. What is that knowledge? It's not just a list of items. It's also the insights that can be learned by analyzing those items, insights that would not be possible to obtain without placing all the items next to each other, organizing them into a meaningful structure and looking at them in a holistic fashion. See the next paragraph for examples.

So, what's the **purpose** of a survey? There may be multiple purposes, such as providing a reference or guidelines for practitioners or researchers, highlighting strengths and weaknesses of current approaches, identifying shortcomings or trends, or outlining new research or marketing opportunities.

How does it achieve that? By systematically analyzing the state of the art and comparing approaches, papers or products of interest. And what does *systematically* mean? It means using a system, i.e., a well-thought evaluation framework, to look at each approach, paper or product in the same, consistent (systematic) manner. It is important to treat all the objects of the analysis the same way and not to change focus, perspective or purpose from one to another. Why? Because doing so would not allow you to correctly and fairly compare the objects and, eventually, to draw meaningful conclusions. You would risk to draw arbitrary conclusions that are not backed by sufficient evidence in your survey. And this is of course what you would like to prevent.

Evaluation framework

How does an **evaluation framework** look like? An *evaluation framework* is the structured set of dimensions and attributes that can be used to compare the objects of the survey. For instance, if we were to compare operating systems of computers (e.g., Windows, MacOS, Linux, etc.), the evaluation framework would identify all those features that characterize any possible operating system (e.g., file management) and that the authors deem important for the purpose of their analysis.

Dimensions express domain concepts of interest (e.g., file management) and can typically be split hierarchically into finer-grained dimensions (e.g., file storage structure, file indexing logic, file search logic) an arranged in a tree structure.

Attributes are the leaves of the tree and represent the possible values or

manifestations of the dimensions with the lowest granularity. For example, file search could be supported as search by file name, by date and size, by file type, by content, etc.

Attention: as we are talking about writing a state of the art survey, which aims to capture the current state of a domain, the evaluation framework cannot contain dimensions or attributes that are not backed by at least one concrete example. In other words, it does not make sense to just invent attributes if it is not possible to name at least one example that implements it. The attribute would turn out to be useless for the purpose of the analysis and, therefore, be discarded in the very end.

Misconception 2 (to be credible, a survey must list all possibly known papers, products, practices that exist in our universe): Of course, in general the rule *the-more-the-better* applies also to surveys, but the key question is "more of what?" And this is important. The goal of a survey is not to list all possible papers, products, practices that are out there. That's anyway an impossible goal to achieve. The goal of the survey should instead be eliciting all relevant dimensions and attributes. These, if well described and explained, represent the real value of a survey as they tell the reader what's important of the domain and what not (if absent). As said before, it's however important that each the attributes be equipped with at least one paper, product, practice implementing it, possibly the most representative one.

The real challenge of a good survey is thus designing the best possible evaluation framework, which can then be used to compare the objects of the analysis among each other and to draw conclusions. There are essentially two approaches to the definition of the framework: top-down vs. bottom-up.

In a **top-down** definition of the framework the authors design the framework based on their own knowledge of the domain. For instance, if they are experts in operating systems, they will know which are they features that must be considered in order to best compare different operating systems and to identify strengths and weaknesses of each.

The **bottom-up** definition of the framework starts from the authors acknowledging that they do not yet have full knowledge of the domain (this I would say is the most prominent situation). Identifying the right dimensions and attributes in absence of complete prior knowledge requires, well, reading, reading and reading; perhaps also testing products, or using methods. The purpose of the reading is acquiring the necessary knowledge to be able to tell what should be included in the framework and what not. The more one reads the more relevant dimensions and attributes will pop up, e.g., because repeated multiple times or, to the contrary, because used as distinguishing feature.

Top-down frameworks thus express prior knowledge by the authors, and the survey aims to find concrete examples to back their claims. Bottom-up frameworks instead start from the examples and abstract them into dimensions and attributes. In practice, it is likely that reality will fall somewhere in between a pure top-down and a pure bottom-up definition of the framework, as the authors almost always have some prior knowledge of a domain, however never a complete one.

Example 1 on the next page provides a good example of how a concrete evaluation framework could look like in practice.

Example 1. Part of the evaluation framework used in a survey on quality control in crowdsourcing [1]. The hierarchical model is the result of a bottom-up analysis of which quality aspects have been studied in crowdsourcing. Dimensions are boxes with shadows and express high-level concerns of crowdsourcing; attributes are boxes without shadow and represent the concrete quality aspects that were found. For each of these attributes, the survey provides references to one/two/three examples for further inspection by the reader.

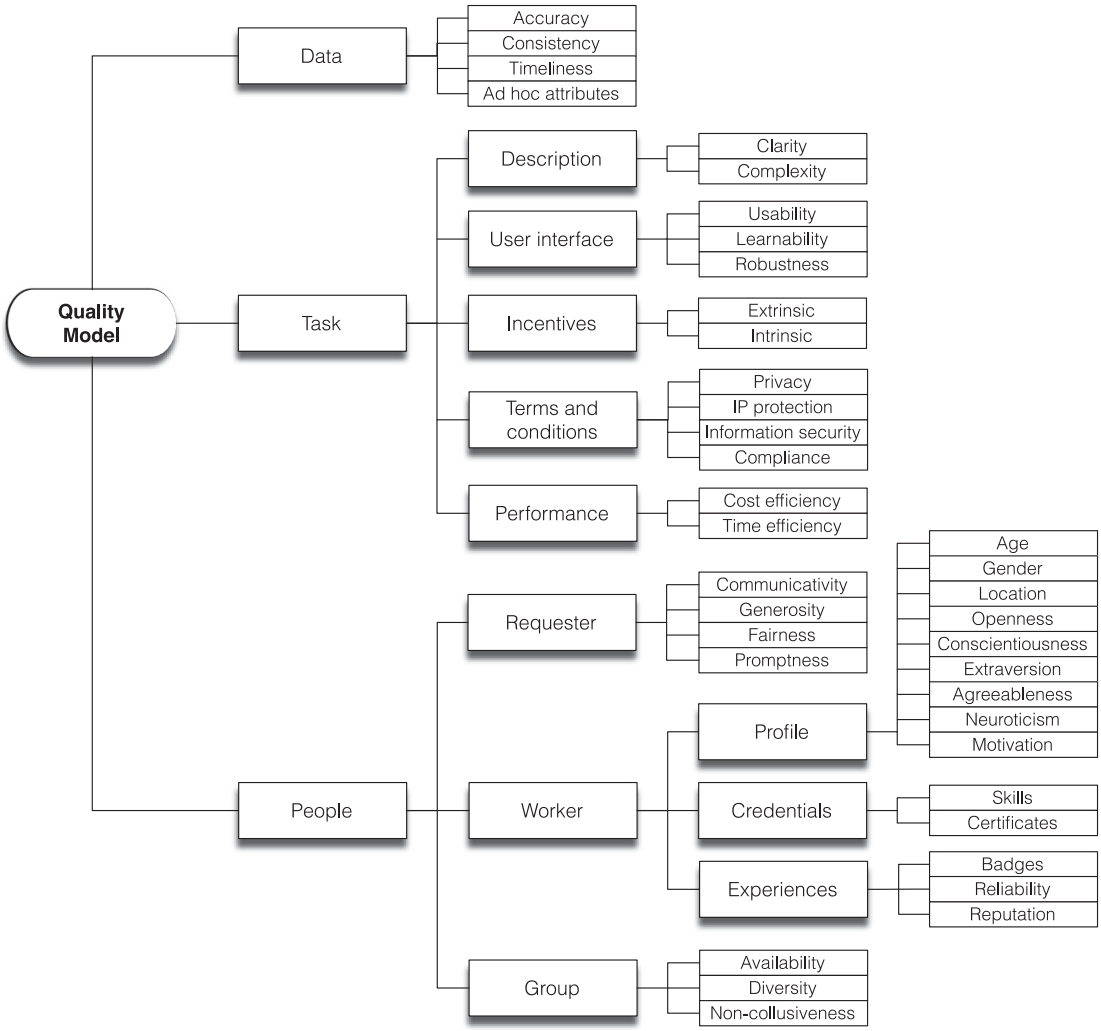

Conceptual model

Writing a good survey requires acquaintance with a set of concepts. Figure 1 illustrates a so-called conceptual model (a UML class diagram) bringing all necessary concepts under one hood:

- **Problem**: Each survey requires a clear problem statement, i.e., a sentence that clearly defines what the goal of the paper is. Here you define well the domain the survey studies. It is very important that you reader has the same understanding of what you are going to study, and why. If your reader does not know the exact focus of your work it is likely that he/she will be puzzled by what comes next in your analysis.
- **Source**: Being the survey an analysis of multiple sources (e.g., papers or product descriptions), it requires an explanation of which exact sources the survey uses to identify candidates for the reading. It is generally not enough to say "Google" or "the Internet." Very likely you search for papers in an online library. Typical libraries in computer science are:
  - ACM Digital Library https://dl.acm.org/

Figure 1. Conceptual model of a state-of-the-art survey with key concepts and relationships.

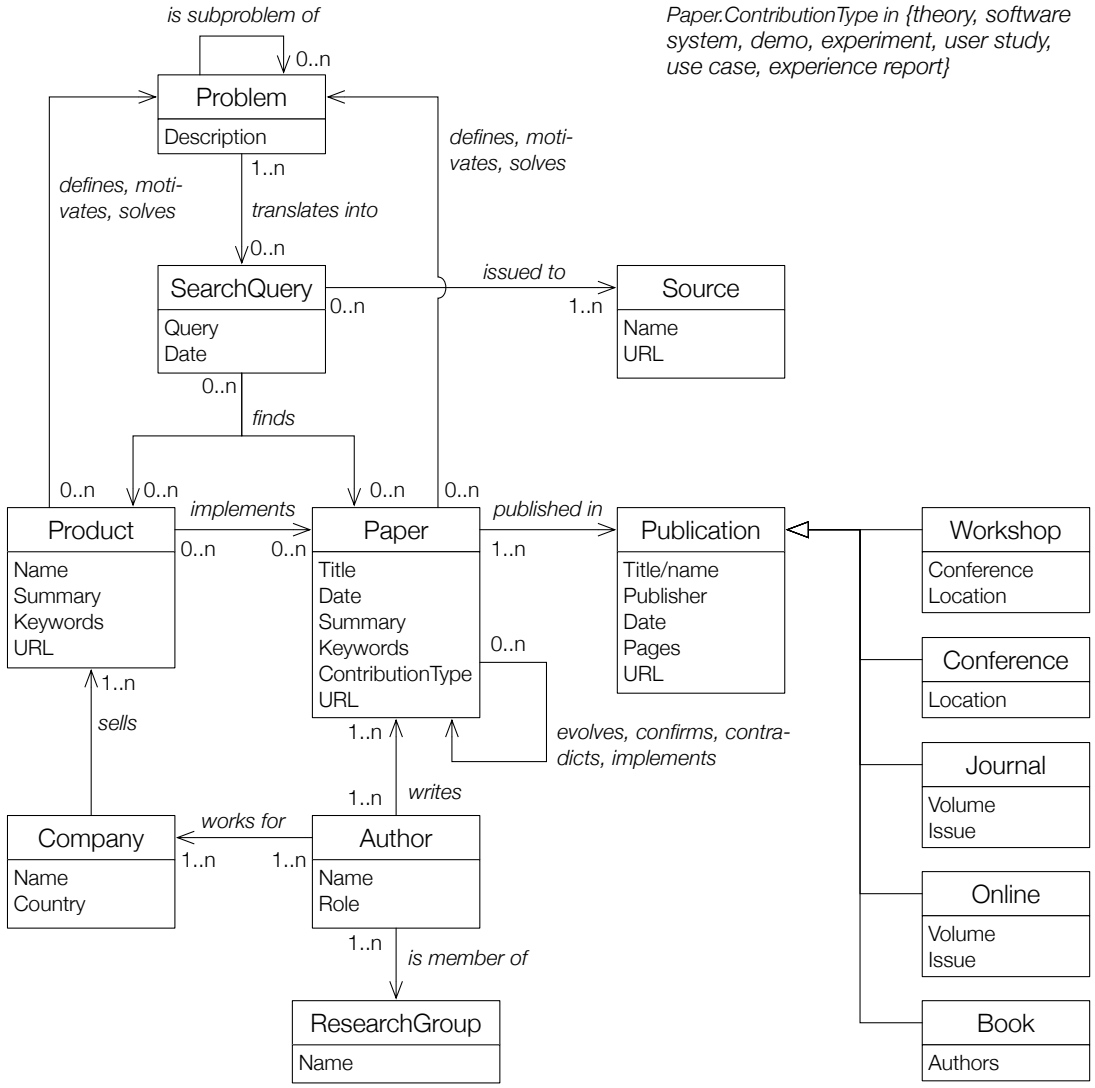

- IEEE Xplore
  http://ieeexplore.ieee.org/
- SpringerLink
  http://www.springerlink.com/
- Google Scholar
  http://scholar.google.com/

- **Search query**: Given your sources, in order to identify papers or documents to read you use some kind of query to interrogate the sources, typically a keyword query. Note down which keywords you use to find the documents to read. Doing so and telling the reader allows the reader to re-do a similar search and to find similar resources, e.g., for a verification of your claims. If you include documents that are not the result of a query but stem from your personal prior knowledge, state clearly which papers are included this way and why.
- **Paper**: Papers are typically the target of your search. These are were from you learn how to structure your evaluation framework and where you find examples for all attributes. Scientific papers are written by a set of *authors* (they may be part of a *research group*, which could be interesting if the group if known for good work in specific areas) and published in some *publication* venue, such as a workshop, conference, journal, book or online

resource. Properly referencing a paper usually asks for authors, title, name of publication, year, and page numbers.

- **Product**: Just like papers, you search may focus on or produce also a set of commercial products to review. You may get inspiration for the evaluation framework from products as well, but oftentimes the products are also used in the end of the survey to analyze how state-of-the-art products support the attributes identified in the framework. Properly referencing a product usually requires a name and URL, possibly also the name of the company producing it.

One important observation is due here: while the conceptual model in Figure 1 looks solid and stable, it describes only abstract concepts. Each survey must instantiate them in its own domain of analysis and, more importantly, iterate multiple times over them till the content of the survey can be considered stable. In fact, each search produces papers or documents to read, which in turn may come with new keywords or features that may lead to novel types of queries or new elements to add to the evaluation framework, respectively. The more iterations are done, the better the problem statement will be focused, the evaluation framework articulated, and the papers selected and classified.

Misconception 3 (in order to look scientific, it is enough to add citations to text): Just like with figures, citations are not just added as adornments of text. If cited, the reader must be able to understand why. For that, you must explain what the papers contribute to the state of the art (just like you must explain the figures you include in the text).

Survey structure

Now that you know the ingredients of a typical survey paper, the question is how to structure the survey. As usual in writing, there are no on-size-fits-all solutions, and each survey is a writing challenge on its own. However, the following sections will be there — perhaps in different order — in most good surveys:

1. **Introduction**: The intro provides the reader with the necessary context of the work and provides the problem statement of the article.
2. **Domain description**: The survey studies a specific domain that is limited by the assumptions and goals of the authors. Most readers are not familiar with that domain; that's the very reason why the actually read the survey, to learn about the domain. It is thus crucial to properly introduce the domain, provide the most important definitions and possibly refine the problem statement in light of the additional details now available to the reader.
3. **Evaluation framework**: Here you describe the structure of the framework you use to analyze the state of the art, the dimensions and the identified attributes. A good evaluation framework communicates excellent knowledge of the domain and anticipates details (the attributes) that raise the curiosity of the reader.
4. **Paper/product selection method**: After defining the framework, it is time to explain the dataset you use for your analysis. This dataset is the set of papers, products, documents and similar identified through the search. List the sources and queries used and explain the inclusion and exclusion criteria (e.g., select all papers that carry the searched for keywords in either their title or abstract, exclude those that are published only in workshops). Finally provide some descriptive statistics about your dataset, e.g., numbers of papers identified, excluded, analyzed.

5. **Body of analysis**: This is the part where you provide the main contribution of the survey, i.e., the description of the attributes with respective examples. This part may be structured similar to the evaluation framework and divided into sections and sub-sections, depending on the complexity of the work. The key here is providing the reader with a spectrum of which different attributes are there and explaining them to the non-expert, not just listing them along with some references.
6. **Overview of state of the art**: This part is optional and depends on the type of survey that is being written. It's essentially a "reality check." Let's say you construct your evaluation framework in a bottom-up fashion, read lots of paper and identify all relevant dimensions and attributes that characterize a domain. After explaining them well to the reader, here you take for instance a set of state-of-the-art products and compare them with the evaluation framework, that is, you tell for each of the products which of the attributes (features) it supports, and how. Doing so allows you (and the reader) to obtain an excellent picture of which product is the best according to which dimension, which dimensions are still underdeveloped, which over-engineered, and similar.
7. **Conclusion and outlook**: Of course, each article has a conclusion or final discussion, and a survey is no exception. Here you summarize the key findings of your analysis and possibly provide hints on new research challenges you identify (e.g., because your state of the art analysis has identified underdeveloped dimensions). When you draw your conclusions, it is important that such are also backed by evidence in the body of the article. Don't invent arbitrary ideas, just look at what your analysis tells you. And yes, you may add a paragraph or two in which you also express your very personal opinion. Just qualify those paragraphs as personal opinions.

Text to be extended…

References

1. F. Daniel, P. Kucherbaev, C. Cappiello, B. Benatallah, M. Allahbakhsh. Quality Control in Crowdsourcing: A Survey of Quality Attributes, Assessment Techniques and Assurance Actions. ACM Computing Surveys 51(1), Article No. 7, April 2018.