

The State of Ontology Pattern Research

A Systematic Review of ISWC, ESWC and ASWC 2005–2009

Karl Hammar and Kurt Sandkuhl

School of Engineering at Jönköping University
P.O. Box 1026
551 11 Jönköping, Sweden
haka, saku@jth.hj.se

Abstract. While the use of patterns in computer science is well established, research into ontology design patterns is a fairly recent development. We believe that it is important to develop an overview of the state of research in this new field, in order to stake out possibilities for future research and in order to provide an introduction for researchers new to the topic. This paper presents a systematic literature review of all papers published at the three large semantic web conferences and their associated workshops in the last five years. Our findings indicate among other things that a lot of papers in this field are lacking in empirical validation, that ontology design patterns tend to be one of the main focuses of papers that mention them, and that although research on using patterns is being performed, studying patterns as artifacts of their own is less common.

Keywords: Ontology design patterns, literature review, content classification

1 Introduction

The use of patterns in computer science-related work already has some tradition. Since the seminal work on software design patterns by the “gang of four” [8], many pattern approaches have been proposed and found useful for capturing engineering practices or documenting organizational knowledge. Examples can be found for basically all phases of software engineering, such as analysis patterns [6], data model patterns [10], software architecture patterns [7, 4] or test patterns. The pattern idea has been adapted in other areas of computer science, such as workflow patterns [5], groupware patterns [12] or patterns for specific programming languages [1].

Ontology design patterns are a rather new development. Ontology design patterns in its current interpretation was introduced by Gangemi [9] and by Blomqvist and Sandkuhl [3]. Blomqvist defines the term as “*a set of ontological elements, structures or construction principles that solve a clearly defined particular modeling problem*” [2]. Despite quite a few well-defined ontology construction methods and a number of reusable ontologies offered on the Internet,

efficient ontology development continues to be a challenge, since this still requires both a lot of domain knowledge/experience and knowledge of the underlying logical theory. Ontology Design Patterns (ODP) are considered a promising contribution to this challenge. They are considered as encodings of best practices, which help to reduce the need for extensive experience when developing ontologies.

When a new research area is emerging, as is the case for ontology design patterns, both academia and industry are interested in the state of affairs, but for different reasons. Academia's interest often relates to identifying the open problems and most pressing challenges to be explored and solved. Industry's prime interest is often in new technologies promising better operational efficiency and effectiveness, or competitive advantages for particular products or services. In order to decide which technologies qualify, they need to understand the choices available, their advantages and disadvantages, and potential risks. This paper presents a literature survey intending to identify existing research in ontology design patterns, structure this research, and evaluate the way it has been tested and validated.

The remaining part of the paper is structured as follows: after an introduction to the research questions and description of study design (Sections 2 and 3 respectively), the paper presents some key results of the survey (Section 4). Section 5 discusses threats to validity and analyses the results, attempting to answer the research questions posed, and pointing out possibilities for future work.

2 Research Questions

The point of performing a literature survey is of course to get a handle on the current state of affairs within the topic area. But what does this really boil down to, what data should we gather and which questions should we ask? We turn to another information gathering profession for guidance, by adopting and adapting the Five Ws (and one H) formula for gathering information from the journalism field. This is a set of questions that need to be answered for any piece of reporting to be considered complete. The questions to pose are Who, What, When, Where, Why, and How.

Since, in the case of academic research, the Who and the Where often converge (few scholars work independently of research institutions) and since we do not want to perform reviews of individual scholars, the Who question is dropped from the list. The Why question is also left out - while it is indeed an important question to answer (why ontology pattern research is done at all), it seems out of scope of a survey such as this one. We are left with four questions to answer. We map these questions to concrete answerable research questions for the survey:

1. What kind of research on ontology patterns is being performed? (What)
2. How has research in the field developed over time? (When)
3. Where is ontology pattern research performed? (Where)
4. How is research in ontology patterns being done? (How?)

3 Method

The method used to perform the literature review is illustrated in Fig. 1. It consisted of first finding the relevant research papers (i.e. papers that discuss ontology patterns), then classifying and sorting them based on various metrics and measures in order to obtain the data needed to answer the research questions. The following sections discuss each of these steps in more detail.




Fig. 1. Method overview.

3.1 Paper Selection

For the purpose of this survey, we have studied conference and workshop proceedings of ISWC, ASWC and ESWC from the last five years (2005-2009). The reasons for choosing these delimitations are as follows:

- The conferences in question are three of the most well-established and prestigious semantic web conferences in the academia. Papers that have been accepted to them are therefore likely to be of a high quality and representative of the general direction in which the semantic web research field is heading.
- As mentioned in section 1, the current interpretation of ontology design patterns was introduced in [9] and [3], both of which were published in 2005. We thus consider this a natural starting year for the survey. 2009 was selected as the ending year for the simple reason that the 2010 conferences have not all been held at the time of writing.
- After initially performing the selection process detailed below with only the main conference proceedings as source, it was found that the number of papers returned were rather few. In order to gain more data, the scope was widened to include the workshop proceedings also, on the intuition that

the workshops while being narrower in focus may still be relevant markers of trends and directions in ontology pattern research (see for instance the WOP workshops that deal specifically with ontology patterns).

The proceedings were retrieved from various Internet databases including SpringerLink¹ and CEUR-WS². A total of 2462 papers were retrieved, divided into 861 conference papers and 1601 workshop papers. Unfortunately, not all workshop proceedings could be located and retrieved, leaving the following workshops' proceedings missing from the survey:

ISWC

SemUbiCare-2007

ESWC

Framework 6 Project Collaboration for the Future Semantic Web-2005

ASWC

All workshops from ASWC 2006

All workshops except DIST from ASWC 2008

All workshops from ASWC 2009

In order to find the subset of papers dealing with ontology patterns, all of these papers were subjected to a full-text search. All papers containing the phrases *ontology patterns* or *ontology **word** patterns* (**word** denoting any one single word) were selected for further analysis. Thereafter, in order to weed out false positives, papers mentioning patterns only in the reference list were removed from the dataset. This left us with a total of 54 papers, 19 from conferences and 35 from workshops.

3.2 Metadata Extraction

The following metadata posts were determined to be useful in answering the research questions:

- Publication year - necessary (but not sufficient) for answering RQ 2
- Author provenance (research institution) - sufficiently answers RQ 3
- Author count - necessary (but not sufficient) for answering RQ 4
- Institution count - necessary (but not sufficient) for answering RQ 4

These pieces of information were retrieved from the dataset. The publication year was retrieved automatically by way of database queries and web spidering scripts when downloading the source material, whereas the research institutions and institution/author counts were ascertained by studying the papers manually.

Research institutions were counted and ranked by the number of mentions they received in the headers of papers in the dataset. However, due to time constraints, further study of the specific organizational structure of the various

¹ <http://www.springerlink.com>

² <http://ceur-ws.org/>

institutions was not performed. Consequently, different departments of the same university were counted as belonging to the same unit, as were different branches of a company or other research organization. This gives us a course-grained view of what universities and organizations are involved in this type of research.

3.3 Data Extraction

Answering research questions 1 and 4 (the What and How questions) required more detailed reading and analysis of the papers. We determined that the What question could most accurately be answered by categorizing the papers based on how they contribute to or make use of ontology pattern research. This categorization is described later in this section and in Table 1. The How question, due to its rather fuzzy nature, is harder to answer in a concrete manner. We decided to focus on three measurable aspects: (1) how scientists cooperate, (2) how the theories they propose are validated and tested, and (3) how central ontology patterns are to the content of the studied papers (whether they are a core part of the performed research or merely used in passing). Aspect 1 can be ascertained by authorship metadata, see section 3.2. Aspects 2 and 3 are discussed in the following sections.

Table 1. Content categories and definitions.

Category	Definition
New pattern presentation	The paper presents a new ontology pattern.
Pattern usage method	The paper presents a new general approach or method of using patterns to achieve some goal(s).
Pattern creation method	The paper presents a new approach for isolating or creating ontology design patterns (including re-engineering from other knowledge representation forms).
Patterns used	The paper describes a case where patterns have been used for achieving some goal(s)
Evaluation	The paper focuses on evaluating patterns or pattern usage/creation methods
Pattern typology	The paper discusses or suggests different types or groupings of patterns.
Pattern features	The paper discusses specific features of ontology patterns. This includes features of not only the reusable reference implementation, but also the documentation and metadata associated with it.
Pattern identification	The paper deals with finding instances of patterns in an existing ontology.
Pattern languages	The paper discusses languages or formalisms for representing or displaying patterns.
Antipatterns	The paper concerns antipatterns or worst practices.
Not relevant	The paper is a false positive - it does not deal with ontology pattern research, but only mentions the term in passing.

Content Classification We read and tagged the papers with one or more labels categorizing how they related to or made use of ontology patterns. Since we did not know beforehand what categories would best describe the collected material, the labels used could not be decided ahead of time. Instead, the label were selected in an exploratory fashion as the readings of the papers progressed. However, to try to ensure as unbiased a classification as possible, each label was paired with a definition covering papers belonging to the category. The labels and definitions are listed in Table 1. The papers were studied twice, once before and once after deciding upon the categories.

In tagging the papers, we attempted to err on the side of caution. For instance, if a pattern is used or mentioned in a paper without reference to original creator, we have not assumed that the author of the paper has created this pattern for this particular project. Rather, new patterns or new methods need to be explicitly defined as new contributions before we add any of the relevant tags to the paper.

The papers that were classified as belonging to the *Not relevant* category were then pruned from the dataset, leaving 47 papers, of which 16 were conference papers and 31 workshop papers.

Validation Classification In order to survey how ontology pattern research is validated, two procedures were followed. To begin with, the papers were categorized according to in what manner validation or testing of the proposed ideas and theories had been performed. For this purpose, the following four categories were used:

- No validation - there is no mention of any validation of the ideas presented.
- Anecdotal validation - the paper mentions that the research has been validated by use in an experiment or in a project but it provides no detail on how this validation was performed or its results.
- Validation by example - one or more examples are presented in the text, validating the concepts presented in a theoretical manner.
- Empirical validation - some sort of experimental procedure or case study has been performed.

Each paper was assigned to one validation category only. In the cases where a paper matched more than one category, the category mapping to a higher level of validation was selected, i.e. empiricism trumps example which in turn trumps anecdote.

Note that there is a difference between examples used in passing in text to motivate some particular design choice or to illustrate syntax, and examples made explicit and presented as validation or test of theory. The latter is grounds for inclusion in the category *Validation by example*, the former is not. In the studied papers it was also common to take one example or use case and follow along with it from the beginning of the paper (where it occurs as problem description) to the end (presenting a solution to said problem). This is considered

to be thorough enough of an example for such a paper to be included in the *Validation by example* category.

Having categorized the papers by validation technique, a further study of validation quality was performed against the papers categorized as belonging to the *Empirical validation* group. For this study we made use of some of the metrics and corresponding measurement criteria developed in [11]:

- Context description quality - the degree to which the context (organization where the evaluation is performed, type of project, project goals, etc.) of the study is described.
- Study design description quality - the degree to which the design (the variables measured, the treatments, the control, etc.) of the study is described.
- Validity description quality - the degree to which the validity (claims of generalizability, sources of error etc.) of the study is discussed.

ODP Importance Classification In order to learn how important the use of or research into ontology patterns is to each particular paper, we studied in what section of the paper that patterns were mentioned. The intuition was that this information would give a crude indication as to whether ontology patterns were considered an essential core part of the research (warranting inclusion in the title or abstract) or not. Title inclusion indicate greater importance than abstract inclusion, in turn indicating greater importance than mere inclusion in the body text. For this measure, terms such as “Knowledge patterns”, “Semantic patterns”, or just plain “patterns” were also considered acceptable synonyms for ontology patterns, provided they were used in a manner indicating a link to ontologies and the use of such patterns in an ontology engineering context.

4 Results

The processes described in section 3 provided us with a large amount of data to analyze, a subset of which is presented in this chapter. The full dataset is too large to include in full, but is available on demand.

Table 2. ODP papers by year

Year	Conferences	Workshops
2005	2	2
2006	2	2
2007	7	3
2008	1	4
2009	4	20 ¹

¹ 15 of which are from WOP 2009

Table 3. Institutes by paper count

Institute name	Conference	Workshop
ISTC-CNR	4	2
U. of Economics, Prague	2	4
U. Politecnica de Madrid	1	5
U. of Karlsruhe	0	4
Jönköping U.	3	0
Salzburg Research GmbH	2	1
U. of Innsbruck	1	2
U. of Manchester	1	2
U. of Sheffield	1	2

Table 2 presents the number of papers in the dataset indexed by the year they were published, the idea being that this might give an indication as to whether research interest in the field is expanding. Table 3 lists the research organizations most often listed as affiliations of authors in the dataset (sorted by summarized publication count and alphabetically). The full list is considerably longer at 49 entries in total, but is for brevity here limited to organizations with three or more mentions.

Table 4. Classification of the reviewed papers’ connection to ODPs.

Classification	Conferences	Workshops
Antipatterns	0	2
Evaluation	0	2
New pattern presented	3	12
Pattern creation methods	1	1
Pattern features	1	5
Pattern identification	0	2
Pattern languages	1	4
Pattern usage method	4	11
Pattern typology	0	3
Patterns used	8	6

Table 4 contains the results of the process described in section 3.3, that is, the labels denoting categories of pattern-related research and the number of papers tagged with each such label. The results are divided into columns for the conference papers and workshop papers.

Table 5. Validation levels of reviewed papers.

Source	No validation	Anecdote	Example	Empiricism
Conferences	3	2	5	6
Workshops	3	2	19	7

Tables 5 and 6 present the results of the validation classification performed in section 3.3, i.e. how the results were validated and, in the case of empirical procedures being used for this purpose, how well the experiments or case studies were described. Table 7 presents the institution counts of the papers in the dataset. Table 8, finally, shows the result of the ODP importance classification in section 3.3, i.e. in what parts of the papers that the topic of patterns were addressed.

Table 6. Quality of empirical validations.

Quality indicator	Weak	Medium	Strong
<i>Conference papers</i>			
Context description	4	1	1
Study design description	0	3	3
Validity description	5	1	0
<i>Workshop papers</i>			
Context description	5	2	0
Study design description	1	3	3
Validity description	6	1	0

Table 7. Institution counts

Institutions	Conferences	Workshops
1	12	16
2	2	10
3	2	5

5 Analysis and Discussion

In this section we discuss some limitations to the generalizability of the survey and sources of error. We then analyze the data resulting from the survey, attempting to answer the research questions posed in section 2.

5.1 Sources of Error

There are of course limitations to the validity of this survey. To begin with, the selection of source material (ISWC/ASWC/ESWC 2005-2009) may give a limited view of the developments within the field as a whole. It is our belief that these three conferences are important and well-known enough for this selection to give a good overview of the general state of ontology pattern research, but we cannot guarantee that key influential work has not been presented elsewhere, thus falling below our radar. The reader should therefore keep this limitation of generalizability in mind when reading our results.

Furthermore, in the selection process described in section 3.1, we do not check for false negatives. This means that there may be papers within the source material which are not matched in the full text search and thus not selected for further analysis, even though they actually deal with ontology patterns. In order

Table 8. ODP importance classification of reviewed papers.

Group	Conferences	Workshops	Workshops (w/o WOP)
Title match	7	22	8
Abstract match	3	3	2
Body match	6	6	6

to study how common such false negatives are, a random set of discarded papers would need to be read manually. This is something that due to time constraints, we were unable to do. If the false negatives are evenly distributed over the various metrics that we have studied, the impact of missing these papers in the analysis might not be that great, but without further looking into this we cannot say either way.

Another possible error lies with the fact that we have been unable to gather all of the data from the ASWC workshops. This means that it is very possible that Asian universities are ranked lower than they actually deserve in the analysis of where research is done (RQ3).

Finally, in Section 3.3 there are at least two threats to validity. To begin with, we read and classify papers. This process is of course, by its very nature, prone to bias since judgement as to what categories a certain paper matches is not always entirely clear cut. We have tried to limit this threat as far as possible by defining fairly rigid categories, but the reader should keep in mind that this is not a 100 % objective measure by any means. Secondly, since we selected the categories on basis of the papers that we have read, it is likely that we have not exhausted the whole set of possible categorizations of ontology pattern research. In other words, there may be categories of research that are not present in our findings, for the very reason that they are not present in our data.

5.2 What kind of research on ontology patterns is being performed? (RQ1)

The results indicate that while patterns are used in various different ways in research and new patterns are being presented, there is quite a lot less work being done on how to formalize the creation and/or isolation of patterns. This, in our view, means one of two things. It could indicate that the best ways of creating and finding patterns have been established, and that there is thus little need for more research to be done in those areas. However, due to the youth of the field, we believe this to be less likely.

Instead, we believe it is more likely that the results indicate a lack of sufficient research in these areas. This situation could be problematic if it indicates that the patterns that are used are not firmly grounded in theory or practice. If one adds to this the fact that relatively little work is being done on pattern evaluation, the overall impression is that patterns are being presented and used as tools, but are not being sufficiently studied as artifacts of their own.

Another area that appears to be understudied is antipatterns, “worst practices” or common mistakes. We theorize that this may be because finding such antipatterns necessitates empirical study, which as we report in section 5.5, appears to be less common in the papers we have studied.

The distribution of research categories seems to be rather consistent between the set of conference papers and the set of workshop papers, with exception for the categories *Patterns used* and *Pattern creation methods*. The latter could be a simple statistical anomaly (as we’ve mentioned there are few papers dealing

with this topic, for which reason small discrepancies stand out more). The former seems a bit more notable however. We have no theory explaining this finding.

5.3 How has research in the field developed over time? (RQ2)

We can see that there is a much larger number of papers matching our search criteria published in 2009 than in 2005, in large part due to the first Workshop on Ontology Patterns being held in 2009. If we remove the WOP papers, we still see a growth in volume.

One interesting note in relation to RQ1 is that all of the work on pattern identification and pattern creation methods that we found was published in 2009. Furthermore, the papers dealing with pattern evaluation are all from 2008-2009, possibly indicating that researchers have noticed these gaps in theory and are attempting to do something about them.

Unfortunately both of these observations are based on so few papers and such a short period of time that it would be very difficult to claim them as a general trends or make predictions based on them.

5.4 Where is ontology pattern research performed? (RQ3)

Keeping in mind the possibility of errors in selection mentioned in section 5.1, our results indicate that ODP work is primarily taking place at European institutions. In Table 3 we can see that all of the top nine institutions publishing more than two papers are located in mainland Europe and the UK. However, even if we include all of the institutions that have two publications in the dataset, we still do not find any non-european organizations. As a matter of fact, out of a total of 49 institutions that had published, only four were located outside of Europe. Out of these four, three were based in the USA and one in New Zealand.

We can also see that while the dataset is dominated by research originating at universities, there are also a number of private corporations, research foundations, and other types of non-university organizations that work in the field. Out of the 49 institutions found, 17 were such non-university organizations (approximately 35 %). These proportions are slightly lower when taking into account the number of publications per institution, with the non-university organizations netting 31 % of all institutional mentions in the dataset.

5.5 How is research in ontology patterns being done? (RQ4)

Studying academic cooperation, one easily accessible metric is the number of authors per paper. However, author counts on their own can be misleading. In some academic cultures students' advisors get authorship, while in others they do not. We instead look at both author counts and institution counts which gives a better indication of actual research cooperation.

Out of a total of 47 papers, 19 (just over 40 %) list more than one affiliated institution and 7 (just under 15 %) list three institutions. These figures, however,

include papers written by only one author. If we look at the subset of 40 papers that were written by more than one author, the figures are 47.5 % and 17.5 % respectively. No matter which way you slice it, these numbers in our opinion indicate a quite healthy degree of cooperation between research institutions in the field.

Interestingly, there seems to be a difference between work published at workshops and work published at the main conferences - of the former, 51.6 % are credited to single institutions, whereas the of the latter, 75 % are. This possibly indicates that the prestige and/or difficulty associated with the higher barrier of entry to full conferences cause researchers to keep such papers “in-house” to a larger degree.

With regards to the validation and testing of ODP research, we find that there may be some work to be done. Nearly one third of the papers published at full conferences contain no validation or only anecdotal validation of the presented work. Another near-third validates the work via examples, but provide no real-world testing to ensure validity. For workshop papers, a smaller proportion of the papers have no validation or anecdotal validation, but on the other hand, a much larger proportion validate only through example. Of the papers that do contain empirical testing, it is uncommon to see discussions on the limits of validity of said testing. This situation may be somewhat problematic. Though not all types of research invite the opportunity to perform experiments or case studies, nor actually require them, we find quite a few papers that could have benefitted from a more thorough testing procedure.

Finally, looking at how central ODPs are to the content of the papers, we find that the most common situation is actually that patterns are mentioned already in the title (see Table 8), indicating that they are quite central to the papers. This remains the case even when we filter out the WOP papers which naturally skew the results. We have no explanation for this unexpected result. One possibility is that this indicates that ontology design patterns are primarily used within the ODP research community and thus written about by people who consider them to be important enough to warrant inclusion in the title.

5.6 Future Work

We can imagine followups to this study a few years down the line, studying the developments on a longer time scale, or using a wider set of source conferences (including for instance EKAW or KCAP). Also, an in-depth analysis of the individual papers could be performed, studying the citation counts of the papers and cross referencing this against the other metrics studied, in order to validate how well the latter map against the real world qualities that give a paper a high citation count.

Another aspect that we would like to explore when time and resources permit is the use of patterns within industry. It is certainly possible that there are patterns and best practices that are established outside of the academia, and if this is the case, it is not too far-fetched to assume that such patterns are more grounded in empirical knowledge than those discussed in this paper.

With regard to the more important question of future research opportunities in the ODP field we suggest that further work is needed on:

- methods of evaluating the efficiency and effectiveness of ODP,
- developing ODPs for particular usages,
- isolating ODPs from existing ontologies or other information artifacts.

We would also like to see, where possible, a greater focus on evaluating academic results in this field in an empirical manner. This would aid in firmly establishing the credibility and maturity of the ODP research field, which we believe to be necessary if ontology patterns are to be established as a viable solution to the complexity challenges of ontology development, in both the greater research community and in industry.

References

1. Beck, K.: Smalltalk Best Practice Patterns. Volume 1: Coding. Prentice Hall, Englewood Cliffs, NJ (1997)
2. Blomqvist, E.: Semi-automatic ontology construction based on patterns. Ph.D. thesis, Linköpings universitet (2009)
3. Blomqvist, E., Sandkuhl, K.: Patterns in ontology engineering: Classification of ontology patterns. In: Proceedings of the 7th International Conference on Enterprise Information Systems. pp. 413–416 (2005)
4. Buschmann, F., Henney, K., Schmidt, D.: Pattern-oriented software architecture: On patterns and pattern languages. John Wiley & Sons Inc (2007)
5. van Der Aalst, W., Ter Hofstede, A., Kiepuszewski, B., Barros, A.: Workflow patterns. *Distributed and parallel databases* 14(1), 5–51 (2003)
6. Fowler, M.: Analysis Patterns: reusable object models. Addison-Wesley (1997)
7. Fowler, M.: Patterns of enterprise application architecture. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA (2002)
8. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: Design patterns: elements of reusable object-oriented software. Addison-wesley Reading, MA (1995)
9. Gangemi, A.: Ontology design patterns for semantic web content. In: The Semantic Web–ISWC 2005. pp. 262–276. Springer (2005)
10. Hay, D.: Data Model Patterns: Conventions of Thought. Dorset House (1996)
11. Ivarsson, M., Gorschek, T.: Technology transfer decision support in requirements engineering research: a systematic review of REj. *Requirements engineering* 14(3), 155–175 (2009)
12. Schümmer, T., Lukosch, S.: Patterns for computer-mediated interaction. Wiley (2007)