

Recovering from a Decade: A Systematic Mapping of Information Retrieval Approaches to Software Traceability

Markus Borg

markus.borg@cs.lth.se

Dept. of Computer Science

Lund University, Sweden

Per Runeson

per.runeson@cs.lth.se

Dept. of Computer Science

Lund University, Sweden

Anders Ardö

anders.ardo@eit.lth.se

Dept. of Electrical and Information Technology

Lund University, Sweden

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Abstract Context: Engineers in large-scale software development have to manage large amounts of information, spread across many artifacts. Several researchers have proposed expressing retrieval of trace links among artifacts, i.e. trace recovery, as an Information Retrieval (IR) problem. **Objective:** The objective of this study is to produce a map of work on IR-based trace recovery, with a particular focus on previous evaluations and strength of evidence. **Method:** We conducted a systematic mapping of IR-based trace recovery. **Results:** Of the 79 publications classified, a majority applied algebraic IR models. While a set of studies on students indicate that IR-based trace recovery tools support certain work tasks, most previous studies do not go beyond reporting precision and recall of candidate trace links from evaluations using datasets containing less than 500 artifacts. **Conclusions:** Our review identified a need of industrial case studies. Furthermore, we conclude that the overall quality of reporting should be improved regarding both context and tool details, measures reported, and use of IR terminology. Finally, based on our empirical findings, we present suggestions on how to advance research on IR-based trace recovery.

DRAFT

Recovering from a Decade: A Systematic Mapping of Information Retrieval Approaches to Software Traceability

Markus Borg · Per Runeson · Anders Ardö

Received: date / Accepted: date

1 Introduction

The successful evolution of software systems involves concise and quick access to information. However, information overload plagues software engineers, as large amounts of formal and informal information is continuously produced and modified (Olsson, 2002; Cleland-Huang, Chang, and Christensen, 2003). Inevitably, especially in large-scale projects, this leads to a challenging information landscape, that includes, apart from the source code itself, requirements specifications of various abstraction levels, test case descriptions, defect reports, manuals, and the like. The state-of-practice approach to structure such information is to organize artifacts in databases, e.g. document management systems, requirements databases, and code repositories, and to manually maintain trace links (Gotel and Finkelstein, 1994; Huffman Hayes, Dekhtyar, and Sundaram, 2006*)¹. With access to trace information, engineers can more efficiently perform work tasks such as impact analysis, identification of reusable artifacts, and requirements validation (Antoniol et al., 2002*; Winkler and Pilgrim, 2010). Furthermore, research has identified lack of traceability to be a major contributing factor in project overruns and failures (Gotel and Finkelstein, 1994; Dömges and Pohl, 1998; Cleland-Huang, Chang, and Christensen, 2003). Moreover, as traceability plays a role in software verification, safety standards such as ISO 26262 (International Organization

Markus Borg
Dept. of Computer Science
Lund University, Sweden
E-mail: markus.borg@cs.lth.se

Per Runeson
Dept. of Computer Science
Lund University, Sweden
E-mail: per.runeson@cs.lth.se

Anders Ardö
Dept. of Electrical and Information Technology
Lund University, Sweden
E-mail: anders.ardo@eit.lth.se

¹ We use an asterisk (*) to distinguish primary publications in the systematic mapping from general references.

1 for Standardization, 2011) for the automotive industry, and IEC 61511 (International
2 Electrotechnical Commission, 2003) for the process industry, mandate maintenance of
3 traceability information (Katta and Stålhane, 2011), as does the CMMI process im-
4 provement model (Carnegie Mellon Software Engineering Institute, 2010). However,
5 manually maintaining trace links is an approach that does not scale (Heindl and Biff,
6 2005). In addition, the dynamics of software development makes it tedious and error-
7 prone (Dömges and Pohl, 1998; Huffman Hayes, Dekhtyar, and Sundaram, 2006*; Fa-
8 lessi, Cantone, and Canfora, 2010).

9
10 As a consequence, engineers would benefit from additional means of dealing with
11 information seeking and retrieval, to navigate effectively the heterogeneous information
12 landscape of software development projects. Several researchers have claimed it feasible
13 to treat traceability as an information retrieval (IR) problem (Antoniol et al., 2002*;
14 Marcus and Maletic, 2003; De Lucia et al., 2004*; Huffman Hayes, Dekhtyar, and
15 Sundaram, 2006*; Lormans and Deursen, 2006*). Also, other studies have reported
16 that the use of semi-automated trace recovery reduces human effort when performing
17 requirements tracing (Huffman Hayes, Dekhtyar, and Sundaram, 2006*; Natt och Dag,
18 Thelin, and Regnell, 2006*; De Lucia et al., 2006b; De Lucia et al., 2007*; De Lucia,
19 Oliveto, and Tortora, 2009*[a]). The IR approach builds upon the assumption that
20 if engineers refer to the same aspects of the system, similar language is used across
21 different software artifacts. Thus, tools suggest trace links based on Natural Language
22 (NL) content. During the first decade of the millennium, substantial research effort has
23 been spent on tailoring, applying, and evaluating IR techniques to software engineering,
24 but we found that a comprehensive overview of the field is missing. Such a secondary
25 analysis would provide an evidence based foundation for future research, and advise
26 industry practice (Kitchenham and Charters, 2007). As such, the gathered empirical
27 evidence could be used to validate, and possibly intensify, the recent calls for future
28 research by the traceability research community (Gotel et al., 2012), organized by
29 the Center of Excellence for Software Traceability (CoEST)². Furthermore, it could
30 assess the recent claims that applying more advanced IR models does not improve
31 results (Oliveto et al., 2010*; Falessi, Cantone, and Canfora, 2010).

32 We have conducted a Systematic Mapping (SM) study (Kitchenham, Budgen, and
33 Brereton, 2011; Petersen et al., 2008) that clusters publications on IR-based trace
34 recovery. SMs and Systematic Literature Reviews (SLR) are primarily distinguished
35 by their driving Research Questions (RQ) (Kitchenham, Budgen, and Brereton, 2011),
36 i.e. an SM identifies research gaps and clusters evidence to direct future research, while
37 an SLR synthesizes empirical evidence on a specific RQ. The rigor of the methodologies
38 is a key asset in ensuring a comprehensive collection of published evidence. We define
39 our overall goals of this SM in three RQs:

- 40
41 RQ1 Which IR models and enhancement strategies have been most frequently applied
42 to perform trace recovery among NL software artifacts?
43 RQ2 Which types of NL software artifacts have been most frequently linked in IR-
44 based trace recovery studies?
45 RQ3 How strong is the evidence, wrt. degree of realism in the evaluations, of IR-based
46 trace recovery?

47
48 This paper is organized as follows. Section 2 contains a thorough definition of
49 the IR terminology we refer to throughout this paper, and a description of how IR

50 ² www.coest.org
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

tools can be used in a trace recovery process. Section 3 presents related work, i.e. the history of IR-based trace recovery, and related secondary and methodological studies. Section 4 describes how the SM was conducted. Section 5 shows the results from the study. Section 6 discusses our research questions based on the results. Finally, Section 7 presents a summary of our contributions and suggests directions for future research.

2 Background

This section presents fundamentals of IR, and how tools implementing IR models can be used in a trace recovery process.

2.1 IR background and terminology

As the study identified variations in use of terminology, this section defines the terminology used in this study (summarized in Table 1), which is aligned with recently redefined terms (Cleland-Huang, Gotel, and Zisman, 2012). We use the following IR definition: “*information retrieval is finding material (usually documents) of an unstructured nature (usually text) that satisfies an information need from within large collections (usually stored on computers)*” (Manning, Raghavan, and Schütze, 2008). If a retrieved document satisfies such a need, we consider it relevant. We solely consider text retrieval in the study, yet we follow convention and refer to it as IR. In our interpretation, the starting point is that any approach that retrieves documents relevant to a query qualifies as IR. The terms Natural Language Processing (NLP) and Linguistic Engineering (LE) are used in a subset of the mapped publications of this study, even if they refer to the same IR techniques. We consider NLP and LE to be equivalent and borrow two definitions from Liddy (2001): “*NL text is text written in a language used by humans to communicate to one another*”, and “*NLP is a range of computational techniques for analyzing and representing NL text*”. As a result, IR (referring to a process solving a problem) and NLP (referring to a set of techniques) are overlapping. In contrast to the decision by Falessi, Cantone, and Canfora (2010) to consistently apply the term NLP, we choose to use IR in this study, as we prefer to focus on the process rather than the techniques. While trace recovery truly deals with solutions targeting NL text, we prefer to primarily consider it as a problem of satisfying an information need.

Furthermore, a “*software artifact is any piece of information, a final or intermediate work product, which is produced and maintained during software development*” (Kruchten, 2004), e.g. requirements, design documents, source code, test specifications, manuals, and defect reports. To improve readability, we refer to such pieces of information only as ‘artifacts’. Regarding traceability, we use two recent definitions: “*traceability is the potential for traces to be established and used*” and “*trace recovery is an approach to create trace links after the artifacts that they associate have been generated or manipulated*” (Cleland-Huang, Gotel, and Zisman, 2012). In the literature, the trace recovery process is referred to in heterogeneous ways including traceability link recovery, inter-document correlation, document dependency/similarity detection, and document consolidation. We refer to all such approaches as *trace recovery*, and also use the term *links* without differentiating between dependencies, relations and similarities between artifacts.

1 In line with previous research, we use the term *dataset* to refer to the set of artifacts
2 that is used as input in evaluations and *preprocessing* to refer to all processing of NL
3 text before the IR models (discussed next) are applied (Baeza-Yates and Ribeiro-Neto,
4 2011), e.g. stop word removal, stemming and identifier (ID) splitting names expressed
5 in CamelCase (i.e. identifiers named according to the coding convention to capitalize
6 the first character in every word) or identifiers named according to the `under_score`
7 convention. *Feature selection* is the process of selecting a subset of terms to represent a
8 document, in an attempt to decrease the size of the effective vocabulary and to remove
9 noise (Manning, Raghavan, and Schutze, 2008).

10 To support trace recovery, several IR models have been applied. Since we identified
11 contradicting interpretations of what is considered a model, weighting scheme, and
12 similarity measure, we briefly present our understanding of the IR field. IR models
13 often apply the *Bag-of-Words* (BoW) model, a simplifying assumption that represents
14 a document as an unordered collection of words, disregarding word order (Manning,
15 Raghavan, and Schutze, 2008). Most existing IR models can be classified as either al-
16 gebraic or probabilistic, depending on how relevance between queries and documents
17 is measured. In algebraic IR models, relevance is assumed to be correlated with sim-
18 ilarity (Zhai, 2007). The most well-known algebraic model is the commonly applied
19 *Vector Space Model* (VSM) (Salton, Wong, and Yang, 1975), which due to its many
20 variation points acts as a framework for retrieval. Common to all variations of VSM
21 is that both documents and queries are represented as vectors in a high-dimensional
22 space (every term, after preprocessing, in the document collection constitutes a dimen-
23 sion) and that similarities are calculated between vectors using some distance function.
24 Individual terms are not equally meaningful in characterizing documents, thus they are
25 weighted accordingly. Term weights can be both binary (i.e. existing or non-existing)
26 and raw (i.e. based on term frequency) but usually some variant of *Term Frequency-*
27 *Inverse Document Frequency* (TF-IDF) weighting is applied. TF-IDF is used to weight
28 a term based on the length of the document and the frequency of the term, both in the
29 document and in the entire document collection (Singhal, 2001). Regarding *similarity*
30 *measures*, the cosine similarity (calculated as the cosine of the angle between vectors)
31 is dominating in IR-based trace recovery using algebraic models, but also Dice's co-
32 efficient and the Jaccard index (Manning, Raghavan, and Schutze, 2008) have been
33 applied. In an attempt to reduce the noise of NL (such as synonymy and polysemy),
34 *Latent Semantic Indexing* (LSI) was introduced (Deerwester et al., 1990). LSI reduces
35 the dimensions of the vector space, finding semi-dimensions using singular value decom-
36 position. The new dimensions are no longer individual terms, but concepts represented
37 as combinations of terms. In the VSM, *relevance feedback* (i.e. improving the query
38 based on human judgement of partial search results, followed by re-executing an im-
39 proved search query) is typically achieved by updating the query vector (Zhai, 2007).
40 In IR-based trace recovery, this is commonly implemented using the Standard Rocchio
41 method (Rocchio, 1971). The method adjusts the query vector toward the centroid vec-
42 tor of the relevant documents, and away from the centroid vector of the non-relevant
43 documents.

44
45 In probabilistic retrieval, relevance between a query and a document is estimated
46 by probabilistic models. The IR is expressed as a classification problem, documents
47 being either relevant or non-relevant (Singhal, 2001). Documents are then ranked ac-
48 cording to their probability of being relevant (Maron and Kuhns, 1960), referred to
49 as the probabilistic ranking principle (Robertson, 1977). In trace recovery, the *Binary*
50 *Independence Retrieval* model (BIM) (Robertson and Jones, 1976) was first applied
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 to establish links. BIM naïvely assumes that terms are independently distributed, and
2 essentially applies the Naïve Bayes classifier for document ranking (Lewis, 1998). Differ-
3 ent weighting schemes have been explored to improve results, and currently the *BM25*
4 weighting used in the non-binary Okapi system (Robertson and Zaragoza, 2009) con-
5 stitutes state-of-the-art.
6

7 Another category of probabilistic retrieval is based on the model of an inference
8 process in a *Probabilistic Inference Network* (PIN) (Turtle and Croft, 1991). In an in-
9 ference network, relevance is modeled by the uncertainty associated with inferring the
10 query from the document (Zhai, 2007). Inference networks can embed most other IR
11 models, which simplifies the combining of approaches. In its simplest implementation, a
12 document instantiates every term with a certain strength and multiple terms accumu-
13 late to a numerical score for a document given each specific query. Relevance feedback
14 is possible also for BIM and PIN retrieval (Zhai, 2007), but we have not identified any
15 such attempts within the trace recovery research.

16 In the last years, another subset of probabilistic IR models has been applied to
17 trace recovery. *Statistical Language Models* (LM) estimate an LM for each document,
18 then documents are ranked based on the probability that the LM of a document would
19 generate the terms of the query (Ponte and Croft, 1998). A refinement of simple LMs,
20 topic models, describes documents as a mixture over topics. Each individual topic is
21 then characterized by an LM (Zhai, 2008). In trace recovery research, studies applying
22 the four topic models *Probabilistic Latent Semantic Indexing* (PLSI) (Hofman, 2001),
23 *Latent Dirichlet Allocation* (LDA) (Blei, Ng, and Jordan, 2003), *Correlated Topic Model*
24 (CTM) (Blei and Lafferty, 2007) and *Relational Topic Model* (RTM) (Chang and Blei,
25 2010) have been conducted. To measure the distance between LMs, where documents
26 and queries are represented as stochastic variables, several different measures of dis-
27 tributional similarity exist, such as the *Jensen-Shannon divergence* (JS). To the best
28 of our knowledge, the only implementation of relevance feedback in LM-based trace
29 recovery was based on the *Mixture Model method* (Zhai and Lafferty, 2001).
30

31 Several attempts are made to improve an IR model, in this paper referred to as
32 *enhancement strategies*. Apart from the already described relevance feedback, one of
33 the most common approaches in IR is to introduce a *thesaurus*. A thesaurus is a tool
34 for vocabulary control, typically defined for a specific subject area, such as art or bi-
35 ology, formally organized so that *a priori* relationships between concepts are made
36 explicit (Aitchison, Bawden, and Gilchrist, 2000). Standard usage of a thesaurus is to
37 provide an IR system with *preferred and non-preferred terms*, restricted vocabularies of
38 terms that the IR system is allowed to, or not allowed to, use for indexing and search-
39 ing, and *semantic relations*, relations between terms such as synonymy and hyponymy.
40 Another enhancement strategy in IR is *phrasing*, an approach to exceed indexing ac-
41 cording to the BoW model (Croft, Turtle, and Lewis, 1991). A phrase is a sequence
42 of two or more words, expected to be more accurate in representing document content
43 than independent words. Detecting phrases for indexing can be done using either a
44 statistical analysis of term frequency and co-occurrence, or by a syntactical approach,
45 i.e. analyzing grammatical structure using a parts-of-speech tagger. Yet another en-
46 hancement strategy is *clustering*, based on the hypothesis that “documents relevant
47 to a query tend to be more similar to each other than to irrelevant documents and
48 hence are likely to be clustered together” (Charikar et al., 1997). Clustering can be
49 used for different purposes, e.g. presenting additional search results or to structure the
50 presentation of search results.
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Retrieval Models			Misc.		
Algebraic models	Probabilistic models	Statistical language models	Weighting schemes	Similarity measures / distance functions	Enhancement strategies
Vector Space Model (VSM)	Binary Independence Model (BIM)	Language Model (LM)	Binary	Cosine similarity	Relevance feedback
Latent Semantic Indexing (LSI)	Probabilistic Inference Network (PIN)	Probabilistic Latent Semantic Indexing (PLSI)	Raw	Dice's coefficient	Thesaurus
	Best Match 25 (BM25) ^a	Latent Dirichlet Allocation (LDA)	Term Frequency Inverse Document Frequency (TFIDF)	Jaccard index	Phrasing
		Correlated Topics Model (CTM)	Best Match 25 (BM25) ^a	Jensen-Shannon divergence (JS)	Clustering
		Relational Topics Model (RTM)			

^a Okapi BM25 is used to refer both to a non-binary probabilistic model, and its weighting scheme.

Table 1 A summary of fundamental IR terms applied in trace recovery. Note that only the vertical organization carries a meaning.

Finally, a number of measures used to evaluate IR tools have been defined. Accuracy of a set of search results is primarily measured by the standard IR-measures *precision* (the fraction of retrieved instances that are relevant), *recall* (the fraction of relevant instances that are retrieved) and *F-measure* (harmonic mean of precision and recall, possibly weighted to favour one over another) (Baeza-Yates and Ribeiro-Neto, 2011). Precision and recall values (P-R values) are typically reported pairwise or as precision and recall curves (P-R curves). Two other set-based measures, originating from the traceability community, are *Recovery Effort Index* (REI) (Antoniol et al., 2002*) and *Selectivity* (Sundaram et al., 2010*). *Secondary measures* aim to go further than comparing sets of search results, and also consider their internal ranking. Two standard IR measures are *Mean Average Precision* (MAP) of precision scores for a query (Manning, Raghavan, and Schütze, 2008), and *Discounted Cumulative Gain* (DCG) (Järvelin and Kekäläinen, 2000) (a graded relevance scale based on the position of a document among search results). To address this matter in the specific application of trace recovery, Sundaram et al. (2010*) proposed *DiffAR*, *DiffMR*, and *Lag* to assess the quality of retrieved candidate links.

2.2 IR-based support in a trace recovery process

As the candidate trace links generated by state-of-the-art IR-based trace recovery typically are too inaccurate, the current tools are proposed to be used in a semi-automatic

process. De Lucia et al. (2012) describe this process as a sequence of four key steps, where the fourth step requires human judgement. Although steps 2 and 3 mainly apply to algebraic IR models, also other IR models can be described by a similar sequential process flow. The four steps are:

1. document parsing, extraction, and pre-processing
2. corpus indexing with an IR method
3. ranked list generation
4. analysis of candidate links

In the first step, the artifacts in the targeted information space are processed and represented as a set of documents at a given granularity level, e.g. sections, class files or individual requirements. In the second step, for algebraic IR models, features from the set of documents are extracted and weighted to create an index. When also the query has been indexed in the same way, the output from step 2 is used to calculate similarities between artifacts to rank candidate trace links accordingly. In the final step, these candidate trace links are provided to an engineer for examination. Typically, the engineer then reviews the candidate source and target artifacts of every candidate trace link, and determines whether the link should be confirmed or not. Consequently, the final outcome of the process of IR-based trace recovery is based on human judgment. Concrete examples, put in work task contexts, are presented in Section 3.4.

A number of publications propose advice for engineers working with candidate trace links. De Lucia et al. have suggested that an engineer should iteratively decrease the similarity threshold, and stop considering candidate trace links when the fraction of incorrect links get too high (De Lucia, Oliveto, and Sgueglia, 2006*; De Lucia, Oliveto, and Tortora, 2008*). Based on an experiment with student subjects, they concluded that an incremental approach in general both improves the accuracy and reduces the effort involved in a tracing task supported by IR-based trace recovery. Furthermore, they report that the subjects preferred working in an incremental manner. Working incrementally with candidate trace links can to some subjects also be an intuitive approach. In a previous experiment by Borg and Pfahl (2011*), several subjects described such an approach to deal with tool output, even without explicit instructions. Coverage analysis is another strategy proposed by De Lucia, Oliveto, and Tortora (2009*[b]), intended to follow up on the step of iteratively decreasing the similarity threshold. By analyzing the confirmed candidate trace links, i.e. conducting a coverage analysis, De Lucia et al. suggest that engineers should focus on tracing artifacts that have few trace links. Also, in an experiment with students, they demonstrated that an engineer working according to this strategy recovers more correct trace links.

3 Related work

This section presents a chronological overview of IR-based trace recovery, previous overviews of the field, and related work on advancing empirical evaluations of IR-based trace recovery.

3.1 A brief history of IR-based trace recovery

Tool support for the linking process of NL artifacts has been explored by researchers since at least the early 1990s. Pioneering work was done in the LESD project (Lin-

1 guistic Engineering for Software Design) by Borillo et al. (1992), in which a tool suite
2 analyzing NL requirements was developed. The tool suite parsed NL requirements
3 to build semantic representations, and used artificial intelligence approaches to help
4 engineers establish trace links between individual requirements (Bras and Toussaint,
5 1993). Apart from analyzing relations between artifacts, the tools evaluated consis-
6 tency, completeness, verifiability and modifiability (Castell et al., 1994). In 1998, a
7 study by Fiutem and Antoniol (1998) presented a recovery process to bridge the gap
8 between design and code, based on edit distances between artifacts. They coined the
9 term “traceability recovery”, and Antoniol et al. published several papers on the topic.
10 Also, they were the first to clearly express identification of trace links as an IR prob-
11 lem (Antoniol et al., 2000). Their milestone work from 2002 compared two standard
12 IR models, probabilistic retrieval using the BIM and the VSM (Antoniol et al., 2002*).
13 Simultaneously, in the late 1990s, Park et al. (2000*) worked on tracing dependencies
14 between artifacts using a sliding window combined with syntactic parsing. Similarities
15 between sentences were calculated using cosine similarities.

16
17 During the first decade of the new millennium, several research groups advanced
18 IR-based trace recovery. Natt och Dag et al. (2002*) did research on requirement depen-
19 dencies in the dynamic environment of market-driven requirements engineering. They
20 developed the tool ReqSimile, implementing trace recovery based on the VSM, and later
21 evaluated it in a controlled experiment (Natt och Dag, Thelin, and Regnell, 2006*). A
22 publication by Marcus and Maletic (2003), the second most cited article in the field,
23 constitutes a technical milestone in IR-based trace recovery. They introduced Latent
24 Semantic Indexing (LSI) to recover trace links between source code and NL documenta-
25 tion, a technique that has been used by multiple researchers since. Huffman Hayes et
26 al. (2004*) enhanced VSM retrieval with relevance feedback and introduced secondary
27 performance metrics. From early on, their research had a human-oriented perspective,
28 aimed at supporting V&V activities at NASA using their tool RETRO (Huffman Hayes
29 et al., 2007*).

30 De Lucia et al. (2005*) have conducted work focused on empirically evaluating
31 LSI-based trace recovery in their document management system ADAMS. They have
32 advanced the empirical foundation by conducting a series of controlled experiments
33 and case studies with student subjects (De Lucia et al., 2006*; De Lucia et al., 2007*;
34 De Lucia, Oliveto, and Tortora, 2009*[a]). Cleland-Huang and colleagues have pub-
35 lished several studies on IR-based trace recovery. They introduced probabilistic trace
36 recovery using a PIN-based retrieval model, implemented in their tool Poirot (Lin et
37 al., 2006). Much of their work has focused on improving the accuracy of their tool by
38 enhancements such as: applying a thesaurus to deal with synonymy (Settimi et al.,
39 2004*), extraction of key phrases (Zou, Settimi, and Cleland-Huang, 2010*), and us-
40 ing a project glossary to weight the most important terms higher (Zou, Settimi, and
41 Cleland-Huang, 2010*).

42 Recent work on IR-based trace recovery has, with various results, gone beyond the
43 traditional models for information retrieval. In particular, trace recovery supported
44 by probabilistic topic models has been explored by several researchers. Dekhtyar et
45 al. (2007*) combined several IR models using a voting scheme, including the prob-
46 abilistic topic model Latent Dirichlet Allocation (LDA). Parvathy, Vasudevan, and
47 Balakrishnan (2008*) proposed using the Correlated Topic Model (CTM), and Geth-
48 ers et al. (2011*) suggested using Relational Topic Model (RTM). Abadi, Nisenson, and
49 Simionovici (2008*) proposed using Probabilistic Latent Semantic Indexing (PLSI) and
50 utilizing two concepts based on information theory, Sufficient Dimensionality Reduction
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

(SDR) and Jensen-Shannon Divergence (JS). Capobianco et al. (2009*[b]) proposed representing NL artifacts as B-splines and calculating similarities as distances between them on the Cartesian plane. Sultanov and Huffman Hayes (2010*) implemented trace recovery using a swarm technique, an approach in which a non-centralized group of non-intelligent self-organized agents perform work that, when combined, enables conclusions to be drawn.

3.2 Previous overviews on IR-based trace recovery

In the beginning of 2012, a textbook on software traceability edited by Cleland-Huang, Gotel, and Zisman (2012) was published. Presenting software traceability from several perspectives, the book contains a chapter authored by De Lucia et al. (2012) specifically dedicated to IR-based trace recovery. In the chapter, the authors thoroughly present an overview of the field including references to the most important work. Also, the chapter constitutes a good introduction for readers new to the approach, as it describes the basics of IR models. Consequently, the book chapter by De Lucia et al. is closely related to our work. However, our work has different characteristics. First, De Lucia et al.'s work has more the character of a textbook, including enough background material on IR, as well as examples of applications in context, to introduce readers to the field as a stand-alone piece of work. Our systematic mapping on the other hand, is not intended as an introduction to the field of IR-based trace recovery, but requires extensive pre-understanding. Second, while De Lucia et al. report a large set of references to previous work, the method used to identify previous publications is not reported. Our work instead follows the established guidelines for SMs (Kitchenham and Charters, 2007), and reports from every phase of the study in a detailed protocol.

Furthermore, basically every publication on IR-based trace recovery contains some information on previous research in the field. Another good example of a summary of the field was provided by De Lucia, Oliveto, and Tortora (2009*[a]). Even though the summary was not the primary contribution of the publication, they chronologically described the development, presented 15 trace recovery methods and 5 tool implementations. They compared underlying IR models, enhancing strategies, evaluation methodologies and types of recovered links. However, regarding both methodological rigor and depth of the analysis, it is not a complete SM. De Lucia, Fasano, and Oliveto (2008) have also surveyed proposed approaches to traceability management for impact analysis. They discussed previous work based on a conceptual framework by Bianchi, Fasolino, and Visaggio (2000), consisting of the three traceability dimensions: type of links, source of information to derive links, and their internal representation. Apart from IR-based methods, the survey by De Lucia et al. contains both rule-based and data mining-based trace recovery. Also Binkley and Lawrie (2010) have presented a survey of IR-based trace recovery as part of an overview of applications of IR in software engineering. They concluded that the main focus of the research has been to improve the accuracy of candidate links wrt. P-R values, and that LSI has been the most popular IR model. However, they also report that no IR model has been reported as superior for trace recovery. While our work is similar to previous work, our review is more structured and goes deeper with a more narrow scope.

Another set of publications has presented taxonomies on IR techniques in software engineering. In an attempt to harmonize the terminology of the IR applications, Canfora and Cerulo (2006*) presented a taxonomy of IR models. However, their surveyed

1 IR applications are not explicitly focusing on software engineering. Furthermore, their
2 proposed taxonomy does not cover recent IR models identified in our study, and the
3 subdivision into ‘representation’ and ‘reasoning’ poorly serves our intentions. Falessi,
4 Cantone, and Canfora (2010) recently published a comprehensive taxonomy of IR tech-
5 niques available to identify equivalent requirements. They adopted the term variation
6 point from Software Product Line Engineering (Pohl, Bockle, and Linden, 2005), to
7 stress the fact that an IR solution is a combination of different, often orthogonal, design
8 choices. They consider an IR solution to consist of a combination of algebraic model,
9 term extraction, weighting scheme and similarity metric. Finally, they conducted an
10 empirical study of various combinations and concluded that simple approaches yielded
11 the most accurate results on their dataset. We share their view on variation points,
12 but fail to apply it since our mapping study is limited by what previous publications
13 report on IR-based trace recovery. Also, their proposed taxonomy only covers algebraic
14 IR models, excluding other models (most importantly, the entire family of probabilistic
15 retrieval).

16
17 Concept location (a.k.a. feature location) is a research topic that overlaps trace
18 recovery. It can be seen as the first step of a change impact analysis process (Marcus
19 et al., 2004). Given a concept (or feature) that is to be modified, the initial information
20 need of a developer is to locate the part of the source code where it is embedded. Clearly,
21 this information need could be fulfilled by utilizing IR. However, we distinguish the
22 topics by considering concept location to be more query-oriented (Gay et al., 2009).
23 Furthermore, whereas trace recovery typically is evaluated by linking n artifacts to m
24 other artifacts, evaluations of concept location tend to focus on n queries targeting a
25 document set of m source code artifacts (where $n \ll m$), as for example in the study
26 by Torchiano and Ricca (2010). Also, while it is often argued that trace recovery
27 should retrieve trace links with a high recall, the goal of concept location is mainly to
28 retrieve one single location in the code with high precision. Dit et al. (2011) recently
29 published a literature review on feature location.

30 31 32 33 3.3 Related contributions to the empirical study of IR-based trace recovery

34
35 A number of previous publications have aimed at structuring or advancing the re-
36 search on IR-based trace recovery, and are thus closely related to our study. An early
37 attempt to advance reporting and conducting of empirical experiments was published
38 by Huffman Hayes and Dekhtyar (2005a). Their experimental framework describes the
39 four phases: *definition*, *planning*, *realization* and *interpretation*. In addition, they used
40 their framework to characterize previous publications. Unfortunately, the framework
41 has not been applied frequently and the quality of the reporting of empirical evalua-
42 tions varies greatly (Borg, Wnuk, and Pfahl, 2012). Huffman Hayes, Dekhtyar, and
43 Sundaram (2006*) also presented the distinction between studies of methods (are the
44 tools capable of providing accurate results fast?) and studies of human analysts (how
45 do humans use the tool output?). Furthermore, they proposed assessing the accuracy
46 of tool output according to quality intervals named ‘acceptable’, ‘good’, and ‘excellent’,
47 based on Huffman Hayes’ industrial experience of working with traceability matrices
48 of various qualities. Huffman Hayes et al.’s quality levels were defined to represent the
49 effort that would be required by an engineer to vet an entire candidate traceability
50 matrix.
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65



Fig. 1 The Integrated Cognitive Research Framework by Ingwersen and Järvelin (2005), a framework for IR evaluations in context.

Considering empirical evaluations, we extend the classifications proposed by Huffman Hayes, Dekhtyar, and Sundaram (2006*) by an adapted version of the *Integrated Cognitive Research Framework* by Ingwersen and Järvelin (2005). Their work aimed at extending the de-facto standard of IR evaluation, the *Laboratory Model of IR Evaluation*, developed in the Cranfield tests in the 60s (Cleverdon, 1991), challenged for its unrealistic lack of user involvement (Kekäläinen and Järvelin, 2002). Ingwersen and Järvelin argued that IR is always evaluated in a context, referred to the innermost context as “the cave of IR evaluation”, and proposed a framework consisting of four integrated contexts (see Figure 1). We have adapted their framework to a four-level context taxonomy, tailored for IR-based trace recovery, to classify in which contexts previous evaluations have been conducted, see Table 2. Also, we add a dimension of study environments (university, proprietary, and open source environment), as presented in Figure 12 in Section 5. For more information on the context taxonomy, we refer to our original publication (Borg, Runeson, and Brodén, 2012).

In the field of IR-based trace recovery, the empirical evaluations are termed very differently by different authors. Some call them ‘experiments’, others ‘case studies’, and yet others only ‘studies’. We use the following definitions, which are established in the field of software engineering.

Case study in software engineering is an empirical enquiry that draws on multiple sources of evidence to investigate one instance (or a small number of instances) of a contemporary software engineering phenomenon within its real-life context, especially when the boundary between phenomenon and context cannot be clearly specified. (Runeson et al., 2012)

Experiment (or controlled experiment) in software engineering is an empirical enquiry that manipulates one factor or variable of the studied setting. Based in randomization, different treatments are applied to or by different subjects, while keeping other variables constant, and measuring the effects on outcome variables. In human-oriented experiments, humans apply different treatments to objects, while in technology-oriented experiments, different technical treatments are applied to different objects. (Wohlin et al., 2012)

Empirical evaluations of IR-based trace recovery may be classified as case studies, if they evaluate the use of, e.g. IR-based trace recovery tools in a complex software engineering environment, where it is not clear whether the tool is the main factor or other factors are at play. These are typically level 4 studies in our taxonomy, see Table 2. Human-oriented controlled experiments may evaluate human performance when using

Level 1: Retrieval context	The most simplified context, referred to as “the cave of IR evaluation”. A strict retrieval context, performance is evaluated wrt. the accuracy of a set of search results. Quantitative studies dominate.	Precision, recall, F-measure	Experiments on benchmarks, possibly with simulated feedback
Level 2: Seeking context	A first step towards realistic applications of the tool, “drifting outside the cave”. A seeking context with a focus on how the human finds relevant information in what was retrieved by the system. Quantitative studies dominate.	Secondary measures. General IR: MAP, DCG. Traceability specific: Lag, DiffAR, DiffMR.	Experiments on benchmarks, possibly with simulated feedback
Level 3: Work task context	Humans complete real tasks, but in an in-vitro setting. Goal of evaluation is to assess the casual effect of an IR tool when completing a task. A mix of quantitative and qualitative studies.	Time spent on task and quality of work.	Controlled experiments with human subjects.
Level 4: Project context	Evaluations in a social-organizational context. The IR tool is studied when used by engineers within the full complexity of an in-vivo setting. Qualitative studies dominate.	User satisfaction, tool usage	Case studies

Table 2 A context taxonomy of IR-based trace recovery evaluations. Level 1 is technology-oriented, and level 3 and 4 are human-oriented. Level 2 typically has a mixed focus.

two different IR-tools in an artificial (in vitro) or well-controlled real (in vivo) environment, typically at level 3 of the taxonomy. The stochastic variation is here primarily assumed to be in the human behavior, although there of course are interactions between the human behavior, the artifacts and the tools. Technology-oriented controlled experiments evaluate tool performance on different artifacts, without human intervention, corresponding to levels 1 and 2 in our taxonomy. The variation factor is here the artifacts, and hence the technology-oriented experiment may be seen as benchmarking studies, where one technique is compared to another technique, using the same artifacts, or the performance of one technique is compared for multiple different artifacts.

The validity of the datasets used as input in evaluations in IR-based trace recovery is frequently discussed in the literature. Also, two recent publications primarily address this issue. Ali, Guéhéneuc, and Antoniol (2012) present a literature review on characteristics of artifacts reported to impact trace recovery evaluations, e.g. ambiguous and vague requirements, and the quality of source code identifiers. Ali et al. extracted P-R values from eight previous trace recovery evaluations, not limited to IR-based trace recovery, and show that the same techniques generate candidate trace links of very different accuracy across datasets. They conclude that research targeting only recovery methods in isolation is not expected to lead to any major breakthroughs, instead they suggest that factors impacting the input artifacts should be better controlled. Borg, Wnuk, and Pfahl (2012) recently highlighted that a majority of previous evaluations of IR-based trace recovery have been conducted using artifacts developed by students. The authors explored this potential validity threat in a survey of the traceability community. Their results indicate that while most authors consider artifacts originating from student projects to be only partly representative to industrial artifacts, few respondents explicitly validated them before using them as experimental input.

3.4 Precision and recall evaluation styles for technology-oriented trace recovery

In the primary publications, two principally different styles to report output from technology-oriented experiments have been used, i.e. presentation of P-R values from evaluations in the retrieval and seeking contexts. A number of publications, including the pioneering work by Antoniol et al. (2002*), used the traditional style from the ad hoc retrieval task organized by the Text REtrieval Conference (TREC) (Voorhees, 2005), driving large-scale evaluations of IR. In this style, a number of queries are executed on a document set, and each query results in a ranked list of search results (cf. (a) in Figure 2). The accuracy of the IR system is then calculated as an average of the precision and recall over the queries. For example, in Antoniol et al.'s evaluation, source code files were used as queries and the document set consisted of individual manual pages. We refer to this reporting style as *query-based* evaluation. This setup evaluates the IR problem: “given this trace artifact, to which other trace artifacts should trace links be established?” The IR problem is reformulated for each trace artifact used as a query, and the results can be presented as a P-R curve displaying the average accuracy of candidate trace links over n queries. This reporting style shows how accurately an IR-based trace recovery tool supports a work task that requires single on-demand tracing efforts (a.k.a. reactive tracing or just-in-time tracing), e.g. establishing traces as part of an impact analysis work task (Antoniol et al., 2002*; Li et al., 2008*; Borg and Pfahl, 2011*).

In the other type of reporting style used in the primary publications, documents of different types are compared to each other, and the result from the similarity- or probability-based retrieval is reported as one single ranked list of candidate trace links. This can be interpreted as the IR problem: “among all these possible trace links, which trace links should be established?” Thus, the outcome is a candidate traceability matrix. We refer to this reporting style as *matrix-based* evaluation. The candidate traceability matrix can be compared to a gold standard, and the accuracy (i.e. overlap between the matrices) can be presented as a P-R curve, as shown in b) in Figure 2. This evaluation setup has been used in several primary publications to assess the accuracy of candidate traceability matrices generated by IR-based trace recovery tools. Also, Huffman Hayes, Dekhtyar, and Sundaram (2006*) defined the quality intervals described in Section 3.3 to support this evaluation style.

Apart from the principally different meaning of reported P-R values, the primary publications also differ by which sets of P-R values are reported. Precision and recall are set-based measures, and the accuracy of a set of candidate trace links (or candidate trace matrix) depends on which links are considered the tool output. Apart from the traditional way of reporting precision at fixed levels of recall, further described in Section 4.3, different strategies for selecting subsets of candidate trace links have been proposed. Such heuristics can be used by engineers working with IR-based trace recovery tools, and several primary publications report corresponding P-R values. We refer to these different approaches to consider subsets of ranked candidate trace links as *cut-off strategies*. Example cut-off strategies include: *Constant cut point*, a fixed number of the top-ranked trace links are selected, e.g. 5, 10, or 50. *Variable cut point*, a fixed percentage of the total number of candidate trace links is selected, e.g. 5% or 10%. *Constant threshold*, all candidate trace links representing similarities (or probabilities) above a specified threshold is selected, e.g. above a cosine similarity of 0.7. *Variable threshold*, a new similarity interval is defined by the minimum and maximum similarities (i.e. similarities are normalized against the highest and lowest similarities),

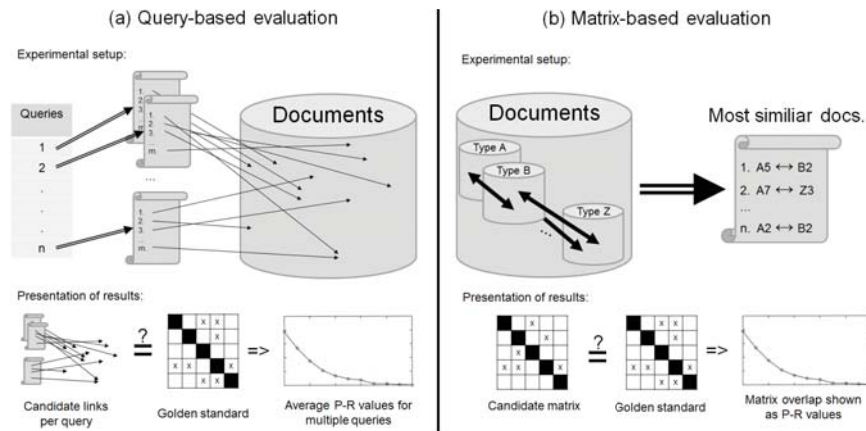


Fig. 2 Query-based evaluation vs. matrix-based evaluation of IR-based trace recovery.

and either a percentage of the candidate trace links are selected or a new constant threshold is introduced.

The choice of what subset of candidate trace links to represent by P-R values reflects the cut-off strategy an imagined engineer could use when working with the tool output. However, which strategy results in the most accurate subset of trace links depends on the specific case evaluated. Moreover, in reality it is possible that engineers would not be consistent in how they work with candidate trace links. As a consequence of the many possible ways to report P-R values, the primary publications view output from IR-based trace recovery tools from rather different perspectives. For work tasks supported by a separate list of candidate trace links per source artifact, there are indications that human subjects seldom consider more than 10 candidate trace links (Borg and Pfahl, 2011*), in line with what is commonplace to present as a ‘pages-worth’ output of major search engines such as Google, Bing and Yahoo. On the other hand, when an IR-based trace recovery tool is used to generate a candidate traceability matrix over an entire information space, considering only the first 10 candidate links would obviously be insufficient, as there would likely be thousands of correct trace links to recover. However, regardless of reporting style, the number of candidate trace links a P-R value represents is important in any evaluation of IR-based trace recovery tools, since a human is intended to vet the output.

The two inherently different use cases of an IR-based trace recovery tool, reflected by the split into matrix-based and query-based evaluations, also call for different evaluations regarding cut-off strategies. The first main use case of an IR-based trace recovery tool is when one or more candidate trace links from a specific artifact are requested by an engineer. For example, as part of a formal change impact analysis, a software engineer might need to specify which test cases to execute to verify that a defect report has been properly resolved. This example is close to the general definition of IR, “to find documents that satisfy an information need from within large collections”. If the database of test cases contains overlapping test cases, it is possible that the engineer needs to report just one suitable test case. In this case precision is more important than recall, and it is fundamental that the tool presents few false positives among the

1 top candidate trace links. Evaluating the performance of the IR-based trace recovery
2 tool using constant cut points is suitable.

3
4 The second main use case of an IR-based trace recovery tool is to generate an
5 entire set of trace links, i.e. a candidate traceability matrix. For instance, a traceability
6 matrix needs to be established between n functional requirements and m test case
7 descriptions during software evolution of a legacy system. If the number of artifacts is
8 $n + m$, the number of possible trace links (i.e. the number of pair-wise comparisons
9 needed) to consider is $n * m$, a number that quickly becomes infeasible for manual
10 work. An engineer can in such cases use the output from an IR-based trace recovery
11 tool as a starting point. The generation of a traceability matrix corresponds to running
12 multiple simultaneous queries in a general IR system, and typically recall is favored
13 over precision. There is no natural equivalent to this use case in the general IR domain.
14 Furthermore, when generating an entire traceability matrix, it is improbable that the
15 total number of correct trace links is known *a priori*, and consequently constant cut-
16 points are less meaningful. A naïve cut-off strategy is to instead simply use a constant
17 similarity threshold such as the cosine similarity 0.7. More promising cut-off strategies
18 are based on variable thresholds or incremental approaches, as described in Section 2.2.
19 Typically, the accuracies of traceability matrices generated by IR-based trace recovery
20 tools are evaluated *a posteriori*, by analyzing how precision varies for different levels
21 of recall.

22 23 4 Method

24
25 The overall goal of this study was to form a comprehensive overview of the existing re-
26 search on IR-based trace recovery. To achieve this objective, we systematically collected
27 empirical evidence to answer research questions characteristic for an SM (Kitchenham
28 and Charters, 2007; Petersen et al., 2008). The study was conducted in the following
29 distinct steps, (i) development of the review protocol, (ii) selection of publications,
30 (iii) data extraction and mapping of publications, which were partly iterated and each
31 of them was validated.

32 33 4.1 Protocol development

34
35 Following the established guidelines for secondary studies in software engineering (Kitchen-
36 ham and Charters, 2007), we iteratively developed a review protocol in consensus
37 meetings between the authors. The protocol defined the research questions (stated in
38 Section 1), the search strategy (described in Section 4.2), the inclusion/exclusion cri-
39 teria (presented in Table 3), and the classification scheme used for the data extraction
40 (described in Section 4.3). The extracted data were organized in a tabular format to
41 support comparison across studies. Evidence was summarized per category, and com-
42 monalities and differences between studies were investigated. Also, the review protocol
43 specified the use of Zotero³ as the reference management system, to simplify general
44 tasks such as sorting, searching and removal of duplicates. An important deviation from
45 the terminology used in the guidelines is that we distinguish between *primary publi-*
46 *cations* (i.e. included units of publication) and *primary studies* (i.e. included pieces of
47 empirical evidence), since a number of publications report multiple studies.

48
49
50 ³ www.zotero.org

	Inclusion criteria	Rationale/comments
I1	Publication available in English in full text	We assumed that all relevant publications would be available in English.
I2	Publication is a peer-reviewed piece of software engineering work	As a quality assurance, we did not include technical reports, master theses etc.
I3	Publication contains empirical results (case study, experiment, survey etc.) of IR-based trace recovery where natural language artifacts are either source or target	Defined our main scope based on our RQs. Publication should clearly link artifacts, thus we excluded tools supporting a broader sense of program understanding such as COCONUT (De Lucia et al., 2006a). Also, the approach should treat the linking as an IR problem. However, we excluded solutions exclusively extracting specific character sequences in NL text, such as work on Mozilla defect reports (Ayari et al., 2007).
	Exclusion criteria	Rationale/comments
E1	Answer is no to I1, I2 or I3	
E2	Publication proposes one of the following approaches to recover trace links, rather than IR: a) rule-based extraction b) ontology-based extraction c) machine learning approaches that require supervised learning d) dynamic/execution analysis	We included only publications that are deployable in an industrial setting with limited effort. Thus, we limited our study to techniques that require nothing but unstructured NL text as input. Other approaches could arguably be applied to perform IR, but are too different to fit our scope. Excluded approaches include: rules (Egyed and Grunbacher, 2002; Spanoudakis et al., 2004), ontologies (Assawamekin, Sunetnanta, and Pluempitwiriyaew, 2010), supervised machine learning (Spanoudakis, d'Avila-Garcez, and Zisman, 2003), semantic networks (Lindvall et al., 2009), and dynamic analysis (Eisenbarth, Koschke, and Simon, 2003).
E3	Article explicitly targets one of the following topics, instead of trace recovery: a) concept/feature location b) duplicate/clone detection c) code clustering d) class cohesion e) cross cutting concerns/aspect mining	We excluded both concept location and duplicate detection since it deals with different problems, even if some studies apply IR models. Excluded publications include: duplicate detection of defects (Runeson, Alexandersson, and Nyholm, 2007), detection of equivalent requirements (Falessi, Cantone, and Canfora, 2010), and concept location (Marcus et al., 2004). We explicitly added the topics code clustering, class cohesion, and cross cutting concerns to clarify our scope.

Table 3 Inclusion/exclusion criteria applied in our study. The rightmost column motivates our decisions.

Table 3 states our inclusion/exclusion criteria, along with rationales and examples. A number of general decisions accompanied the criteria:

- Empirical results presented in several articles, we only included from the most extensive publication. Examples of excluded publications include pioneering work later extended to journal publications, the most notable being work by Antoniol et al. (2000) and Marcus and Maletic (2003). However, we included publications describing all *independent replications* (deliberate variations of one or more major aspects), and *dependent replications* (same or very similar experimental setups) by other researchers (Shull et al., 2008).
- Our study included publications that apply techniques in E2a–d in Table 3, but use an IR model as benchmark. In such cases, we included the IR benchmark, and noted possible complementary approaches as enhancements. An example is work using probabilistic retrieval enhanced by machine learning from existing trace links (Di and Zhang, 2009*).

- We included approaches that use structured NL as input, i.e. source code or tabular data, but treat the information as unstructured. Instead, we considered any attempts to utilize document structure as enhancements.
- Our study only included linking between software artifacts, i.e. artifacts that are produced and maintained during development (Kruchten, 2004). Thus, we excluded linking approaches to entities such as e-mails (Bacchelli, Lanza, and Robbes, 2010) and tacit knowledge (Stone and Sawyer, 2006; Huffman Hayes, Antoniol, and Guéhéneuc, 2008).
- We excluded studies evaluating trace recovery in which neither the source nor the target artifacts dominantly represent information as NL text. Excluded publications comprise linking source code to test code (Van Rompaey and Demeyer, 2009), and work linking source code to text expressed in specific modelling notation (Antoniol et al., 1999; Cleland-Huang, Marrero, and Berenbach, 2008).

4.2 Selection of publications

The systematic identification of publications consisted of two main phases: (i) development of a gold standard of primary publications, and (ii) a search string that retrieves them, and a systematic search for publications, as shown in Figure 3. In the first phase, a set of publications was identified through exploratory searching, mainly by snowball sampling from a subset of an informal literature review. The most frequently recurring publication fora were then scanned for additional publications. This activity resulted in 59 publications, which was deemed our gold standard⁴. The first phase led to an understanding of the terminology used in the field, and made it possible to develop valid search terms.

The second step of the first phase consisted of iterative development of the search string. Together with a librarian at the department, we repeatedly evaluated our search string using combined searches in the Inspec/Compendex databases. Fifty-five papers in the gold standard were available in those databases. We considered the search string good enough when it resulted in 224 unique hits with 80% recall and 20% precision when searching for the gold standard, i.e. 44 of the 55 primary publications plus 176 additional publications were retrieved.

The final search string was composed of four parts connected with ANDs, specifying the *activity*, *objects*, *domain*, and *approach* respectively.

```
(traceability OR "requirements tracing" OR "requirements trace" OR
"trace retrieval")
AND
(requirement* OR specification* OR document OR documents OR
design OR code OR test OR tests OR defect* OR artefact* OR
artifact* OR link OR links)
AND
(software OR program OR source OR analyst)
AND
("information retrieval" OR IR OR linguistic OR lexical OR
semantic OR NLP OR recovery OR retrieval)
```

The search string was first applied to the four databases supporting export of search results to BibTeX format, as presented in Table 4. The resulting 581 papers

⁴ The gold standard was not considered the end goal of our study, but was the target during the iterative development of the search string described next.

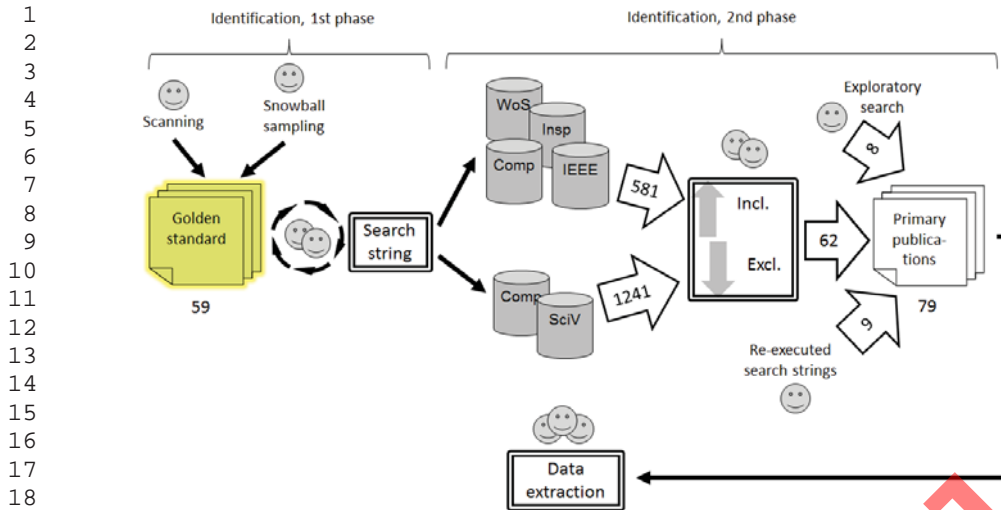


Fig. 3 Overview of the publication selection phase. Smileys show the number of people involved in a step, while double frames represent a validation. Numbers refer to number of publications.

Primary Databases	Search options	#Search results
Inspec	Title+abstract, no auto-stem	194
Compendex	Title+abstract, no auto-stem	143
IEEE Explore	All fields	136
Web of Science	Title+abstract+keywords	108
Secondary Databases	Search options	#Search results
ACM Digital Library	All fields, auto-stem	1038
SciVerse Hub Beta	Science Direct+SCOPUS	203

Table 4 Search options used in databases, and the number of search results.

were merged in Zotero. After manual removal of duplicates, 281 unique publications remained. This result equals 91% recall and 18% precision compared to the gold standard. The publications were filtered by our inclusion/exclusion criteria, as shown in Figure 3, and specified in Section 4.1. Borderline articles were discussed in a joint session of the first two authors. Our inclusion/exclusion criteria were validated by having the last two authors compare 10% of the 581 papers retrieved from the primary databases. The comparison resulted in a free-marginal multi-rater kappa of 0.85 (Randolph, 2005), which constitutes a substantial inter-rater agreement.

As the next step, we applied the search string to two databases without BibTeX export support. One of them, ACM Digital Library, automatically stemmed the search terms, resulting in more than 1000 search results. The inclusion/exclusion criteria were then applied to the total 1241 publications. This step extended our primary studies by 13 publications, after duplicate removal, and application of inclusion/exclusion criteria, 10 identified in ACM Digital Library and 3 from SciVerse.

As the last step of our publication selection phase, we again conducted exploratory searching. Based on our new understanding of the domain, we scanned the top publication fora and the most published scholars for missed publications. As a last complement,

1 we searched for publications using Google Scholar. In total, this last phase identified
2 8 additional publications. Thus, the systematic database search generated 89% of the
3 total number of primary publications, which is in accordance with expectations from
4 the validation of the search string.
5

6 As a final validation step, we visualized the selection of the 70 primary publications
7 using REVIS, a tool developed to support SLRs based on visual text mining (Felizardo
8 et al., 2011). REVIS takes a set of primary publications in an extended BibTeX format
9 and, as presented in Figure 4, visualizes the set as a document map (a), edge bundles
10 (b), and a citation network for the document set (c). While REVIS was developed to
11 support the entire mapping process, we solely used the tool as a means to visually
12 validate our selection of publications.

13 In Figure 4, every node represents a publication, and a black outline distinguishes
14 primary publications (in c), not only primary publications are visualized). In a), the
15 document map, similarity of the language used in title and abstract is presented, calcu-
16 lated using the VSM and cosine similarities. In the clustering, only absolute distances
17 between publications carry a meaning. The arrows point out Antoniol et al.’s publica-
18 tion from 2002 (Antoniol et al., 2002*), the most cited publication on IR-based trace
19 recovery. The closest publications in a) are also authored by Antoniol et al. (Anto-
20 niol et al., 1999*; Antoniol et al., 2000*). An analysis of a) showed that publications
21 sharing many co-authors tend to congregate. As an example, all primary publications
22 authored by De Lucia et al. (De Lucia et al., 2004*; De Lucia et al., 2005*; De Lucia
23 et al., 2006*; De Lucia et al., 2007*; De Lucia, Oliveto, and Sgueglia, 2006*; De Lu-
24 cia, Oliveto, and Tortora, 2009*[a]; De Lucia, Oliveto, and Tortora, 2008*; De Lucia,
25 Oliveto, and Tortora, 2009*[b]), Capobianco et al. (Capobianco et al., 2009*[b]; Capo-
26 bianco et al., 2009*[a]), and Oliveto et al. (Oliveto et al., 2010*) are found within the
27 rectangle. No single outlier stands out, indicating that none of the primary publications
28 uses a very different language.
29

30 In b), the internal reference structure of the primary studies is shown, displayed
31 by edges connecting primary publications in the outer circle. Analyzing the citations
32 between the primary publications shows one outlier, just below the arrow. The publi-
33 cation by Park et al. (2000*), describing work conducted concurrently with Antoniol
34 et al. (2002*), has not been cited by any primary publications. This questioned the
35 inclusion of the work by Park et al., but as it meets our inclusion/exclusion criteria
36 described in Section 4.1, we decided to keep it.

37 Finally, in c), the total citation network of the primary studies is presented. Re-
38 garding common citations in total, again Park et al. (2000*) is an outlier, shown as I in
39 c). The two other salient data points, II and III, are both authored by Natt och Dag et
40 al. (Natt och Dag et al., 2004*; Natt och Dag, Thelin, and Regnell, 2006*). However,
41 according to our inclusion/exclusion criteria, there is no doubt that they should be
42 among the primary publications. Thus, in December 2011, we concluded the set of 70
43 primary publications.

44 However, as IR-based trace recovery is an active research field, several new studies
45 were published while this publication was in submission. To catch up with the latest
46 research, we re-executed the search string in the databases listed in Table 4 in June
47 2012, to catch up with publications from the second half of 2011. This step resulted
48 in 9 additional publications, increasing the number of primary publications to 79. In
49 the rest of this paper, we refer to the original 70 publications as the “core primary
50 publications”, and the 79 publications as just the “primary publications”.
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

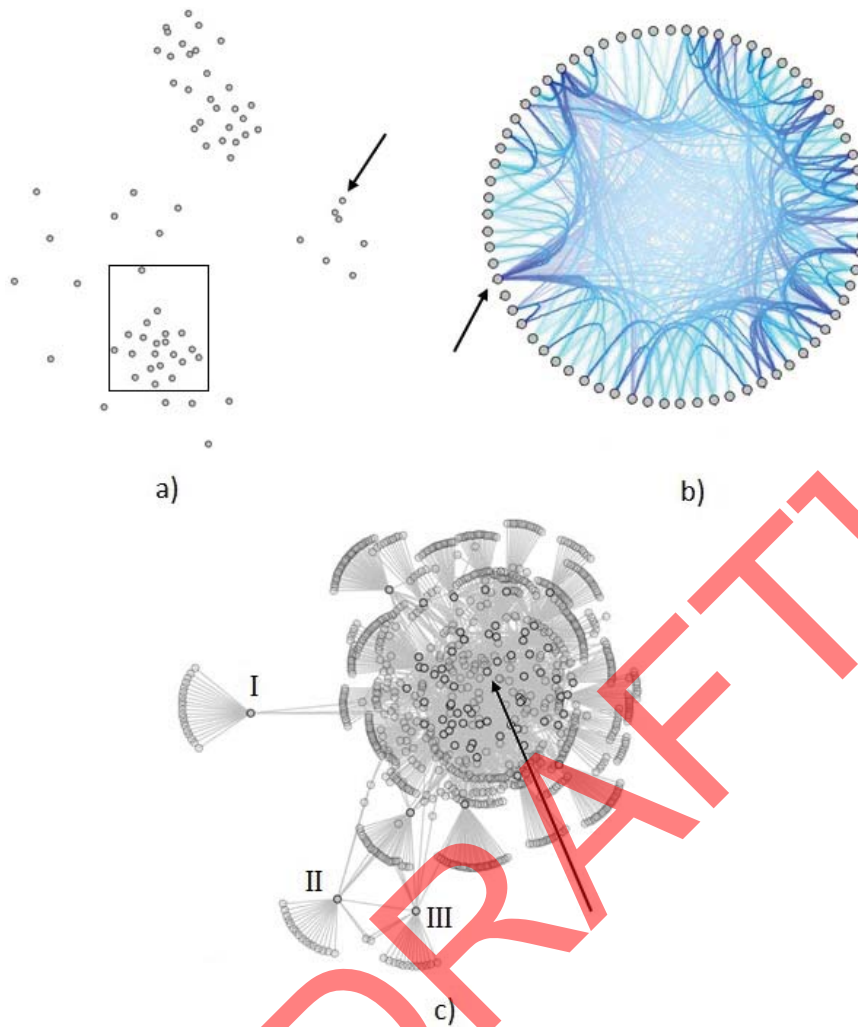


Fig. 4 Visualization of core primary publications. a) document map, shows similarities in language among the core primary publications. b) edge bundle, displays citations among the core primary publications. c) citation network, shows shared citations among the core primary publications.

4.3 Data extraction and mapping

During the stage of the study, data was extracted from the primary publications according to the pre-defined extraction form of the review protocol. We extracted general information (title, authors, affiliation, publication forum, citations), details about the applied IR approach (IR model applied, selection and weighting of features, enhancements) and information about the empirical evaluation (types of artifacts linked, size

and origin of dataset, research methodology, context of IR evaluation, results of evaluation).

The extraction process was validated by the second and third authors, working on a 30% sample of the core primary publications. Half the sample, 15% of the core primary publications, was used to validate extraction of IR details. The other half was used by the other author to validate empirical details. As expected, the validation process showed that the data extraction activity, and the qualitative analysis inherent in that work, inevitably leads to some deviating interpretations. Classifying according to the four levels of IR contexts, which was validated for the entire 30% sample, showed the least consensus. This divergence, and other minor discrepancies detected, were discussed until an agreement was found and followed for the rest of the primary publications. Regarding the IR contexts in particular, we adopted an inclusive strategy, typically selecting the higher levels for borderline publications.

4.4 Threats to validity

Threats to the validity of the mapping study are analyzed with respect to construct validity, reliability, internal validity and external validity (Runeson et al., 2012). Particularly, we report deviations from the study guidelines (Kitchenham and Charters, 2007).

Construct validity concerns the relation between measures used in the study and the theories in which the research questions are grounded. In this study, this concerns the identification of papers, which is inherently qualitative and dependent on the coherence of the terminology in the field. To mitigate this threat, we took the following actions. The search string we used was validated using a golden set of publications, and we executed it in six different publication databases. Furthermore, our subsequent exploratory search further improved our publication coverage. A single researcher applied the inclusion/exclusion criteria, although, as a validation proposed by Kitchenham and Charters (2007), another researcher justified 10% of the search results from the primary databases. There is a risk that the specific terms of the search string related to ‘activity’ (e.g. “requirements tracing”) and ‘objects’ cause a bias toward both requirements research and publications with technical focus. However, the golden set of publications was established by a broad scanning of related work, using both searching and browsing, and was not restricted to specific search terms.

An important threat to *reliability* concerns whether other researchers would come to the same conclusions based on the publications we selected. The major threat is the extraction of data, as mainly qualitative synthesis was applied, a method that involves interpretation. A single researcher extracted data from the primary publications, and the other two researchers reviewed the process, as suggested by Brereton et al. (2007). As a validation, both the reviewers individually repeated the data extraction on a 15% sample of the core primary publications. Another reliability threat is that we present qualitative results with quantitative figures. Thus, the conclusions we draw might depend on the data we decided to visualize; however, the primary studies are publicly available, allowing others to validate our conclusions. Furthermore, as our study contains no formal meta-analysis, no sensitivity analysis was conducted, neither was publication bias explored explicitly.

External validity refers to generalization from this study. In general, the external validity of a SM is strong, as the key idea is to aggregate as much as possible of the

Publication forum	#Publications
International Requirements Engineering Conference	9
International Conference on Automated Software Engineering	7
International Conference on Program Comprehension	6
International Workshop on Traceability in Emerging Forms of Software Engineering	6
Working Conference on Reverse Engineering	5
Empirical Software Engineering	4
International Conference on Software Engineering	4
International Conference on Software Maintenance	4
Other publication fora (two or fewer publications)	34

Table 5 Top publication channels for IR-based trace recovery.

available literature. Also, our research scope is tight (cf. the inclusion/exclusion criteria in Table 3), and we do not claim that our map applies to other applications of IR in software engineering. Thus, the threats to external validity are minor. Furthermore, as the review protocol is presented in detail in Section 4, other researchers can judge the trustworthiness of the results in relation to the search strategy, inclusion/exclusion criteria, and the applied data extraction. Finally, *internal validity* concerns confounding factors that can affect the causal relationship between the treatment and the outcome. However, as our mapping study does not investigate casual relationships, and only relies on descriptive statistics, this threat is minimal.

5 Results

Following the method defined in Section 4.2, we identified 79 primary publications. Most of the publications were published in conferences or workshops (67 of 79, 85%), while twelve (15%) were published in scientific journals. Table 5 presents the top publication channels for IR-based trace recovery, showing that it spans several research topics. Figure 5 depicts the number of primary publications per year, starting from Antoniol et al.’s pioneering work from 1999. Almost 150 authors have contributed to the 79 primary publications, on average writing 2.2 of the articles. The top five authors have on average authored 14 of the primary publications, and are in total included as authors in 53% of the articles. Thus, a wide variety of researchers have been involved in IR-based trace recovery, but there is a group of a few well-published authors. More details and statistics about the primary publications are available in Appendix A.

Several publications report empirical results from multiple evaluations. Consequently, our mapping includes 132 unique empirical contributions, i.e. the mapping comprises results from 132 unique combinations of an applied IR model and its corresponding evaluation on a dataset. As described in Section 4.1, we denote such a unit of empirical evidence a ‘study’, to distinguish from ‘publications’.




Fig. 5 IR-based trace recovery publication trend. The curve shows the number of publications, while the bars display empirical studies in these publications.

5.1 IR models applied to trace recovery (RQ1)

In Figure 6, reported studies in the primary publications are mapped according to the (one or several) IR models applied, as defined in Section 2. The most frequently reported IR models are the algebraic models, VSM and LSI. For LSI, the dimensionality reductions applied in previous studies is reported in Appendix A. Various probabilistic models have been applied in 29 of the 132 evaluations, including 14 applications of statistical LMs. Five of the applied approaches do not fit in the taxonomy; examples include utilizing swarm techniques (Sultanov and Huffman Hayes, 2010*) and B-splines (Capobianco et al., 2009*[b]). As shown in Figure 6, VSM is the most applied model 2008-2011, however repeatedly as a benchmark to compare new IR models against. An apparent trend is that trace recovery based on LMs has received an increasing research interest during the last years.

Only 47 (72%) of the 65 primary publications with technical foci report which preprocessing operations were applied to NL text. Also, in several publications one might suspect that the complete preprocessing was not reported (e.g. (Chen, 2010*) and (Kong et al., 2009*)), possibly due to page restriction. As a result, a reliable report of feature selection for IR-based trace recovery is not possible. Furthermore, several papers do not report any differences regarding preprocessing of NL text and source code (on the other hand some papers make a clear distinction, e.g. (Wang, Lai, and Liu, 2009*)). Among the publications reporting preprocessing, 32 report conducting stop word removal and stemming, making it the most common combination. The remaining publications report other combinations of stop word removal, stemming and ID splitting. Also, two publications report applying Google Translate as a preprocessing step to translate NL text to English (Li et al., 2008*; Huffman Hayes et al., 2011*). Figure 7 presents in how many primary publications different preprocessing steps are explicitly mentioned, both for NL text and source code.

Regarding NL text, most primary publications select all terms that remain after preprocessing as features. However, two publications select only nouns and verbs (Zhao et al., 2003*; Zhou and Yu, 2007*), and one selects only nouns (Capobianco et al., 2009*[b]). Also, Capobianco et al. (2009*[a]) have explicitly explored the semantic role of nouns. For the purposes of the mapping of primary publications dealing with source code, a majority unfortunately does not clearly report about the feature selection (i.e.

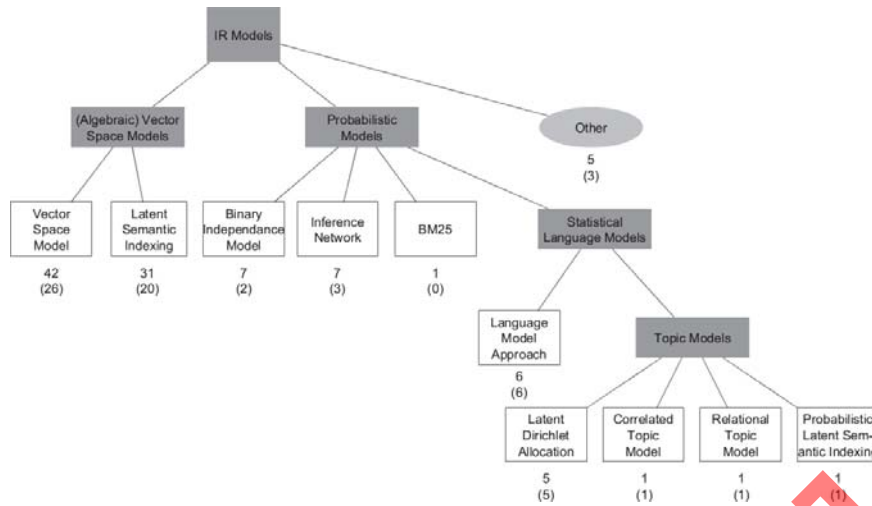


Fig. 6 Taxonomy of IR models in trace recovery. The numbers show in how many of the primary publications a specific model has been applied, the numbers in parentheses show IR models applied since 2008.

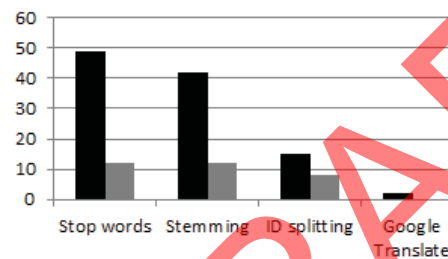


Fig. 7 Preprocessing operations used in IR-based trace recovery. The figure shows the number of times a specific operation has been reported in the primary publications. Black bars refer to preprocessing of NL text, gray bars show preprocessing of text extracted from source code.

selecting which subset of terms to extract to represent the artifact). Seven publications report that only IDs were selected, while four publications selected both IDs and comments. Three other publications report more advanced feature selection, including function arguments, return types and commit comments (Canfora and Cerulo, 2006*; Abadi, Nisenson, and Simionovici, 2008*; Ali, Guéhéneuc, and Antoniol, 2011*[b]).

Among the primary publications, the weighting scheme applied to selected features is reported in 58 articles. Although arguably more tangible for algebraic retrieval models, feature weighting is also important in probabilistic retrieval. Moreover, most weighing schemes are actually families of configuration variants (Salton and Buckley, 1988), but since this level of detail often is omitted in publications on IR-based trace recovery, as also noted by Oliveto (2008), we were not able to investigate this further. Figure 8 shows how many times, in the primary publications, various types of feature weighting schemes have been applied. Furthermore, one publication reports upweight-

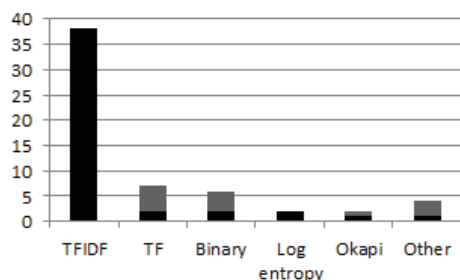


Fig. 8 Feature weighting schemes in IR-based trace recovery. Bars depict how many times a specific weighting scheme has been reported in the primary publications. Black color shows reported weighting in publications applying algebraic IR models.

ing of verbs in the TFIDF weighting scheme, motivated by verbs' nature of describing the functionality of software (Mahmoud and Niu, 2010*).

Several enhancement strategies to improve the performance of IR-based trace recovery tools are proposed, as presented in Figure 9. The figure shows how many times different enhancement strategies have been applied in the primary publications. Most enhancements aim at improving the precision and recall of the tool output, however also a computation performance enhancement is reported (Jiang et al., 2008*). The most frequently applied enhancement strategy is relevance feedback, applied by e.g. De Lucia, Oliveto, and Sgueglia (2006*) and Huffman Hayes et al. (2007*), giving the human a chance to judge partial search results, followed by re-executing an improved search query. The following most frequently applied strategies, further described in Section 2.1, are applying a thesaurus to deal with synonyms, (e.g. proposed by Huffman Hayes, Dekhtyar, and Sundaram (2006*) and Leuser and Ott (2010*)), clustering results based on for instance document structure to improve presentation or ranking of recovered trace links, (explored by e.g. Duan and Cleland-Huang (2007*) and Zhou and Yu (2007*)), and phrasing, i.e. going beyond the BoW model by considering sequences of words, e.g. as described by Zou, Settini, and Cleland-Huang (2006*) and Chen and Grundy (2011*). Other enhancement strategies repeatedly applied include: up-weighting terms considered important by applying a project glossary, e.g. (Zou, Settini, and Cleland-Huang, 2008*), machine learning approaches to improve results based on for example the existing trace link structure, e.g. (Di and Zhang, 2009*), and combining the results from different retrieval models in voting systems, e.g. (Gethers et al., 2011*). Yet another set of enhancements have only been proposed in single primary publications, such as query expansion (Gibiec, Czauderna, and Cleland-Huang, 2010*), analyses of call graphs (Zhao et al., 2003*), regular expressions (Chen and Grundy, 2011*), and smoothing filters (De Lucia et al., 2011*).

5.2 Types of software artifacts linked (RQ2)

Figure 10 maps onto the classical software development V-model the various software artifact types that are used in IR-based trace recovery evaluations. Requirements, the left part of the model, include all artifacts that specify expectations on a system, e.g. market requirements, system requirements, functional requirements, use cases, and




Fig. 9 Enhancement strategies in IR-based trace recovery. Bars depict how many times a specific strategy has been reported in the primary publications. Black color represents enhancements reported in publications using algebraic IR models.

design specifications. The distinction between these are not always possible to derive from the publications, and hence we have grouped them together under the broad label ‘requirements’. The right part of the model represents all artifacts related to verification activities, e.g. test case descriptions and test scripts. Source code artifacts constitute the bottom part of the model. Note however, that our inclusion/exclusion criteria, excluding duplication analyses and studies where neither source nor target artifacts are dominated by NL text, results in fewer links between requirements-requirements, code-code, code-test, test-test and defect-defect than would have been the case if we had studied the entire field of IR applications within software engineering.

The most common type of links that has been studied was found to be between requirements (37 evaluations), either of the same type or of different levels of abstraction. The second most studied artifact linking is between requirements and source code (32 evaluations). Then, in decreasing order, mixed links in an information space of requirements, source code and tests (10 evaluations), links between requirements and tests (9 evaluations) and links between source code and manuals (6 evaluations). Less frequently studied trace links include links between source code and defects/change requests (e.g. (Gethers et al., 2011)) and links between tests (Lormans, Van Deursen, and Gross, 2008*). In three primary publications, the types of artifacts traced are unclear, either not specified at all or merely described as ‘documents’ (e.g. (Chen, Hosking, and Grundy, 2011*)).

5.3 Strength of evidence (RQ3)

An overview of the datasets used for evaluations in the primary publications is shown in Figure 11. In total we identified 132 evaluations; in 42 (32%) cases proprietary artifacts were studied, either originating from development projects in private companies or the US agency NASA. Nineteen (14%) evaluations using artifacts collected from open source projects have been published and 65 (49%) employing artifacts originating from a university environment. Among the datasets from university environments, 34 consist

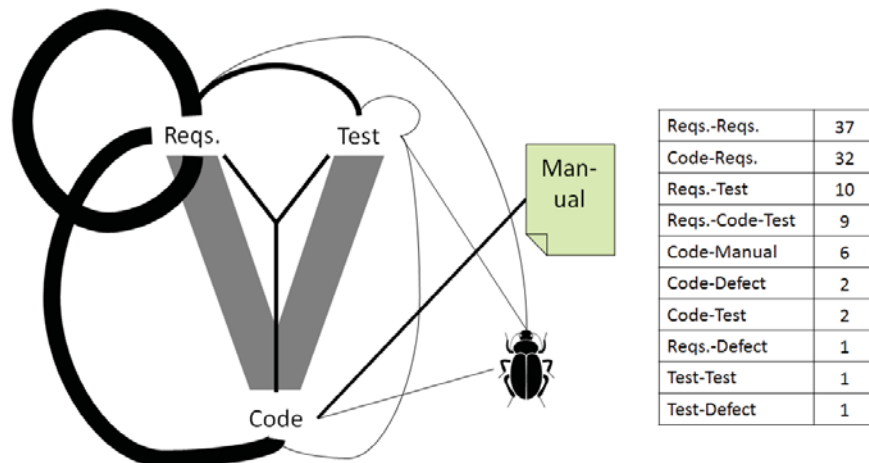


Fig. 10 Types of links recovered in IR-based trace recovery. The table shows the number of times a specific type of link is the recovery target in the primary publications, also represented by the weight of the edges in the figure.

of artifacts developed by students. In six primary publications, the origin of the artifacts is mixed or unclear (e.g. (Park et al., 2000*; Li et al., 2008*; Parvathy, Vasudevan, and Balakrishnan, 2008*). Figure 11 also depicts the sizes of the datasets used in the evaluations, wrt. the number of artifacts. The majority of the evaluations in the primary publications were conducted using an information space of less than 500 artifacts. In 38 of the evaluations, less than 100 artifacts were used as input. The primary publications with the by far highest number of artifacts, evaluated links between 3,779 business requirements and 8,334 market requirements at Baan (Natt och Dag et al., 2004*) (now owned by Infor Global Solutions), and trace links between 9 defect reports and 13,380 test cases at Research in Motion (Kaushik, Tahvildari, and Moore, 2011*).

Table 6 presents the six datasets that have been most frequently used in evaluations of IR-based trace recovery, sorted by the number of primary studies in which they were used. CM-1, MODIS, and EasyClinic are publicly available from the CoEST web page⁵. Note that most publicly available datasets except EasyClinic are *bipartite*, i.e. the dataset contains only links between two disjunct subsets of artifacts.

All primary publications report some form of empirical evaluations, a majority (80%) conducting “studies of methods” (Huffman Hayes, Dekhtyar, and Sundaram, 2006*). Fourteen publications (18%) report results regarding the human analyst, two primary publications study both methods and human analysts (Antoniol et al., 2002*; De Lucia et al., 2006*). Figure 12 shows the primary publications mapped to the four levels of the context taxonomy described in Section 3.3. Note that a number of publications cover more than one environment, due to either mixed artifacts or multiple studies. Also, two publications did not report the environment, and could not be mapped. A majority of the publications (50), exclusively conducted evaluations taking place in the innermost retrieval context, the so-called “cave of IR evaluation” (Ingwersen and Järvelin, 2005). As mentioned in Section 2, evaluations in the cave display

⁵ coest.org




Fig. 11 Datasets used in studies on IR-based trace recovery. Bars show number and origin of artifacts.

#	Dataset	Artifacts	Links	Origin	Development characteristics	Size ^a	Lang.
17	CM-1	Requirements specifying system requirements and detailed design	Bipartite dataset, many-to-many links	NASA	Embedded software development in governmental agency	455	English
16	EasyClinic	Use cases, sequence diagrams, source code, test case descriptions	Many-to-many links	Univ. of Salerno	Student project	150	Italian
8	MODIS	Requirements specifying system requirements and detailed design	Bipartite dataset, many-to-many links.	NASA	Embedded software development in governmental agency	68	English
7	Ice-Breaker System (IBS)	Functional requirements and source code	Not publicly available in full detail	Robertson and Robertson (1999)	Textbook on requirements engineering	185	English
6	LEDA	Source code and user documentation	Bipartite dataset, many-to-one links	Max Planck Inst. for Informatics Saarbrücken	Scientific computing	296	English
5	Event-Based Traceability (EBT)	Functional requirements and source code	Not publicly available	DePaul Univ.	Tool from research project	138	English

^a Size is presented as the total number of artifacts.

Table 6 Summary of the datasets most frequently used for evaluations.

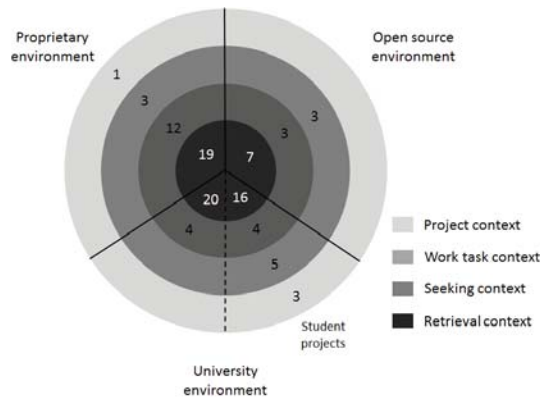


Fig. 12 Contexts of evaluations of IR-based trace recovery, along with study environments. Numbers show the number of primary publications that target each combination.

an inconsistent use of terminology. Nineteen (38%) of the primary publications refer to their evaluations in the retrieval context as experiments, 22 (44%) call them case studies, and in nine (18%) publications they are merely referred to as studies.

Since secondary measures were applied, fourteen publications (18%) are considered to have been conducted in the seeking context. Eleven primary publications conducted evaluations in the work context, mostly through controlled experiments with student subjects. Only three evaluations are reported in the outermost context of IR evaluation, the project context, i.e. evaluating the usefulness of trace recovery in an actual end user environment. Among these, only a single publication reports an evaluation from a non-student development project (Li et al., 2008*).

6 Discussion

This section discusses the results reported in the previous section and concludes on the research questions. Along with the discussions, we conclude every question with concrete suggestions on how to advance research on IR-based trace recovery. Finally, in Section 6.4, we map our recommendations to the traceability challenges articulated by CoEST (Gotel et al., 2012).

6.1 IR models applied to trace recovery (RQ1)

During the last decade, a wide variety of IR models have been applied to recover trace links between artifacts. Our study shows that the most frequently applied models have been algebraic, i.e. Salton’s classic VSM from the 60s (Salton, Wong, and Yang, 1975) and LSI, the extension developed by Deerswester in the 90s (Deerwester et al., 1990). Also, we show that VSM has been implemented more frequently than LSI, in contrast to what was reported by Binkley and Lawrie (2010). The interest in algebraic models might have been caused by the straightforwardness of the techniques; they have concrete geometrical interpretations, and are rather easy to understand also for non-IR experts. Moreover, several open source implementations are available. Consequently,

1 the algebraic models are highly applicable to trace recovery studies, and they constitute
2 feasible benchmarks when developing new methods. However, in line with the devel-
3 opment in the general IR field (Zhai, 2007), LMs (Ponte and Croft, 1998) have been
4 getting more attention in the last years. Regarding enhancements strategies, relevance
5 feedback, introduction of a thesaurus and clustering of results are the most frequently
6 applied.

7 While implementing an IR model, the developers inevitably have to make a vari-
8 ety of design decisions. Consequently, this applies also to IR-based trace recovery
9 tools. As a result, tools implementing the same IR model can produce rather differ-
10 ent output (Borg, Runeson, and Brodén, 2012). Thus, omitting details in the report-
11 ing obstructs replications and the possibility to advance the field of trace recovery
12 through secondary studies and evidence-based software engineering techniques (Jedl-
13 itschka, Ciolkowski, and Pfahl, 2008). Unfortunately, even fundamental information
14 about the implementation of IR is commonly left out in trace recovery publications.
15 Concrete examples include feature selection and weighting (particularly neglected for
16 publications indexing source code) and the number of dimensions of the LSI sub-
17 space. Furthermore, the heterogeneous use of terminology is an unnecessary difficulty
18 in IR-based trace recovery publications. Concerning general traceability terminology,
19 improvements can be expected as Cleland-Huang, Gotel, and Zisman (2012) dedicated
20 an entire chapter of their recent book to this issue. However, we hope that Section 2.1
21 of this paper is a step toward aligning also the IR terminology in the community.

22 To support future replications and secondary studies on IR-based trace recovery, we
23 suggest that:

- 24 – Studies on IR-based trace recovery should use IR terminology consistently, e.g.
25 as presented in Table 1 and Figure 6, and use general traceability terminology as
26 proposed by Cleland-Huang, Gotel, and Zisman (2012).
- 27 – Authors of articles on IR-based trace recovery should carefully report the imple-
28 mented IR model, including the features considered, to enable aggregating empirical
29 evidence.
- 30 – Technology-oriented experiments on IR-based trace recovery should adhere to rigor-
31 ous methodologies such as the evaluation framework by Huffman Hayes and Dekht-
32 yar (2005a).

33 6.2 Types of software artifacts linked (RQ2)

34 Most published evaluations on IR-based trace recovery aim at establishing trace links
35 between requirements in a wide sense, or between requirements and source code. Ap-
36 parently, the artifacts of the V&V side of the V-model are not as frequently in focus
37 of researchers working on IR-based trace recovery. One can think of several reasons for
38 this unbalance. First, researchers might consider that the structure of the document
39 subspace of the requirement side of the V-model is more important to study, as it is
40 considered the “starting point” of development. Second, the early public availability
41 of a few datasets containing requirements of various kinds, might have paved the way
42 for a series of studies by various researchers. Third, publicly available artifacts from
43 the open source community might contain more requirements artifacts than V&V ar-
44 tifacts. Nevertheless, research on trace recovery would benefit from studies on a more
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 diverse mix of artifacts. For instance, the gap between requirements artifacts and V&V
2 artifacts is an important industrial challenge (Sabaliauskaite et al., 2010). Hence, ex-
3 ploring whether IR-based trace recovery could be a way to align “the two ends of
4 software development” is worth an effort.

5 Apart from the finding that requirement-centric studies on IR-based trace recov-
6 ery are over-represented, we found that too few studies go beyond trace recovery in
7 bipartite traceability graphs. Such simplified datasets hardly represent the diverse in-
8 formation landscapes of large-scale software development projects. Exceptions include
9 studies by De Lucia et al., who repeatedly have evaluated IR-based trace recovery
10 among use cases, functional requirements, source code and test cases (De Lucia et al.,
11 2004*; De Lucia et al., 2006*; De Lucia, Oliveto, and Sgueglia, 2006*; De Lucia, Oliveto,
12 and Tortora, 2008*; De Lucia, Oliveto, and Tortora, 2009*[a]; De Lucia, Oliveto, and
13 Tortora, 2009*[b]; De Lucia et al., 2011*), however originating from student projects,
14 which reduces the industrial relevance.
15

16 To further advance the research of IR-based trace recovery, we suggest that:

- 17 – Studies should be conducted on diverse datasets containing a higher number of
18 artifacts, to explore recovery of different types of trace links.
- 19 – Studies should go beyond bipartite datasets to better represent the heterogeneous
20 information landscape of software engineering, thus enabling studies on several
21 types of links within the same datasets.
22

23 6.3 Strength of evidence (RQ3)

24 Most evaluations on IR-based trace recovery were conducted on bipartite datasets con-
25 taining fewer than 500 artifacts. Obviously, as pointed out by several researchers, any
26 software development project involves much larger information landscapes, that also
27 consist of heterogeneous artifacts. A majority of the evaluations of datasets containing
28 more than 1,000 artifacts were conducted using open source artifacts, an environment
29 in which fewer types of artifacts are typically maintained (Scacchi, 2002; Canfora and
30 Cerulo, 2006*), thus links to or from source code are more likely to be studied. Even
31 though small datasets might be reasonable to study, only two primary publications
32 report from evaluations containing more than 10,000 artifacts (Natt och Dag et al.,
33 2004*; Kaushik, Tahvildari, and Moore, 2011*). As a result, the question of whether
34 the state-of-the-art IR-based trace recovery scales to larger document spaces or not
35 remains unanswered. In the empirical NLP community, Banko and Brill (2001) showed
36 that some conclusions (related to machine learning techniques for NL disambiguation)
37 drawn on small datasets may not carry over to very large datasets. Researchers on
38 IR-based trace recovery appear to be aware of the scalability issue however, as it is
39 commonly mentioned as a threat to external validity and suggested as future work in
40 the primary publications (Huffman Hayes et al., 2004*; De Lucia et al., 2007*; Leuser,
41 2009*; Wang, Lai, and Liu, 2009*; Gibiec, Czauderna, and Cleland-Huang, 2010*; Mah-
42 moud and Niu, 2011*). On the other hand, one reason for the many studies on small
43 datasets is the challenge involved in obtaining the complete set of correct trace links,
44 i.e. a gold standard or ground truth, required for evaluations. In certain domains, e.g.
45 development of safety-critical systems, such information might already be available. If
46 such information is missing however, a traceability researcher first needs to establish
47 the gold standard, which requires much work for a large dataset.
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 Regarding the validity of datasets used in evaluations, a majority used artifacts origi-
2 nating from university environments as input. Furthermore, most studies on propri-
3 etary artifacts used only the CM-1 or MODIS datasets collected from NASA projects,
4 resulting in their roles as de-facto benchmarks from an industrial context. Clearly,
5 again the external validity of state-of-the-art trace recovery must be questioned. On
6 one hand, benchmarking can be a way to advance IR tool development, as TREC
7 have demonstrated in the general IR research (Smeaton and Harman, 1997), but on
8 the other hand it can also lead the research community to over-engineering tools on
9 specific datasets (Borg, Runeson, and Brodén, 2012). Thus, the community needs to
10 consider the risk of optimization against those specific benchmarks, which may make
11 the final result less suitable in the general case, if the benchmarks are not representative
12 enough. The benchmark discussion has been very active in the traceability community
13 the last years (Dekhtyar and Huffman Hayes, 2006; Dekhtyar, Huffman Hayes, and
14 Antoniol, 2007; Cleland-Huang et al., 2011; Ben Charrada et al., 2011*; Gotel et al.,
15 2012).

16 A related problem, in particular for proprietary datasets that cannot be disclosed,
17 is that datasets often are poorly described (Borg, Wnuk, and Pfahl, 2012). In some
18 particular publications, NL artifacts in datasets are only described as ‘documents’.
19 Thus, as already discussed related to RQ1 in Section 6.1, inadequate reporting obstructs
20 replications and secondary studies. Moreover, providing information about the datasets
21 and their contexts is also important for interpreting results and their validity, in line
22 with previous work by Ali, Guéhéneuc, and Antoniol (2012) and Borg, Wnuk, and
23 Pfahl (2012). For example, researchers should report as much of the industrial context
24 from which the dataset arose as encouraged by Kitchenham et al. (2002). As a starting
25 point, researchers could use the preliminary framework for describing industrial context
26 by Petersen and Wohlin (2009).

27 As discussed in Section 3.4, P-R values can be reported from IR-based trace re-
28 covery evaluations in different ways. Unfortunately, the reported values are not always
29 properly explained in the primary publications. In the evaluation report, it is central to
30 state whether a query-based or matrix-based evaluation style has been used, as well as
31 which cut-off strategies were applied. Furthermore, for query-based evaluations (closer
32 resembling traditional IR), we agree with the opinion of Spärck Jones, Walker, and
33 Robertson (2000), that reporting only precision at standard recall levels is opaque.
34 The figures obscure the actual numbers of retrieved documents needed to get beyond
35 low recall, and should be complemented by P-R values from a constant cut-point cut-
36 off strategy. Moreover, regarding both query-based and matrix-based evaluation styles,
37 reporting also secondary measures (such as MAP and DCG) is a step toward more
38 mature evaluations.

39 Most empirical evaluations of IR-based trace recovery were conducted in the inner-
40 most of IR contexts, i.e. a clear majority of the research was conducted “in the cave”
41 or just outside (Ingwersen and Järvelin, 2005). For some datasets, the output accuracy
42 of IR models has been well-studied during the last decade. However, more studies on
43 how humans interact with the tools are required; similar to what has been explored
44 by Huffman Hayes et al. (Huffman Hayes et al., 2004*; Huffman Hayes and Dekhtyar,
45 2005b; Dekhtyar, Huffman Hayes, and Larsen, 2007*; Cuddeback, Dekhtyar, and Huff-
46 man Hayes, 2010*) and De Lucia et al. (De Lucia, Oliveto, and Sgueglia, 2006*; De
47 Lucia, Oliveto, and Tortora, 2008*; De Lucia, Oliveto, and Tortora, 2009*[a]). Thus,
48 more evaluations in a work task context or a project context are needed. Regarding
49 the outermost IR context, only one industrial in-vivo evaluation (Li et al., 2008*) and
50

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

three evaluations in student projects (De Lucia et al., 2005*; De Lucia et al., 2006*; De Lucia et al., 2007*) have been reported. Finally, regarding the innermost IR contexts, the discrepancy of methodological terminology should be harmonized in future studies.

To further advance evaluations of IR-based trace recovery, we suggest that:

- The community should continue its struggle to acquire a set of more representative benchmarks.
- Researchers should better characterize both the context and the datasets used in evaluations, in particular when they cannot be disclosed for confidentiality reasons.
- P-R values should be complemented by secondary measures such as MAP and DCG, and it should be made clear whether a query-based or matrix-based evaluation style was used.
- Focus on tool enhancements “in the cave” should be shifted towards evaluations in the work task or project context.

6.4 In the light of the CoEST research agenda

Gotel et al. (2012) recently published a framework of challenges in traceability research, a CoEST community effort based on a draft from 2006 (Cleland-Huang, Huffman Hayes, and Dekhtyar, 2006). The intention of the framework is to provide a structure to direct future research on traceability. CoEST defines eight research themes, addressing challenges that are envisioned to be solved in 2035, as presented in Table 7. Our work mainly contributes to three of the research themes, *purposed traceability*, *trusted traceability*, and *scalable traceability*. Below, we discuss the three research themes in relation to IR-based trace recovery, based on our empirical findings.

The research theme *purposed traceability* charts the development of a classification scheme for traceability contexts, and a collection of possible stakeholder requirements on traceability. Also, a “Traceability Book of Knowledge” is planned, including terminology, methods, practices and the like. Furthermore, the research agenda calls for additional empirical studies. Our contribution intensifies CoEST’s call for additional industrial case studies, by showing that a majority of IR-based trace recovery studies have been conducted in the “cave of IR evaluation”. To guide future empirical studies, we propose an adapted version of the model of IR evaluation contexts by Ingwersen and Järvelin (2005), tailored for IR-based trace recovery. Also, we confirm the need for a “Traceability Book of Knowledge” and an aligned terminology in the traceability community, as our secondary study was obstructed by language discrepancies.

Trusted traceability comprises research to gain improvements in the quality of creation and maintenance of automatic trace links. Also, the research theme calls for empirical evidence as to the quality of traceability methods and tools with respect to the quality of the trace links. Our work, founded in evidence-based software engineering approaches, aggregated the empirical evidence of IR-based trace recovery until December 2011. Based on this, we provide several advice on how to advance future evaluations.

Finally, the research theme *scalable traceability* calls for the traceability community to obtain and publish large industrial datasets from various domains to enable researchers to investigate scalability of traceability methods. Also this call for research is intensified by our work, as we empirically show that alarmingly few evaluations of

Research theme	Goal to reach by 2035
Purposed traceability	to define and instrument prototypical traceability profiles and patterns
Cost-effective traceability	to perform systematic quality assessment and assurance of the traceability
Configurable traceability	to provide for levels of abstraction and granularity in traceability techniques, methods and tools, facilitated by improved trace visualizations, to handle very large datasets and the longevity of these data
Trusted traceability	to develop cost-benefit models for analyzing stakeholder requirements for traceability and associated solution options at a fine-grained level of detail
Scalable traceability	to use dynamic, heterogeneous and semantically rich traceability information models to guide the definition and provision of traceability
Portable traceability	to agree upon universal policies, standards, and a unified representation or language for expressing traceability concepts
Valued traceability	to raise awareness of the value of traceability, to gain buy-in to education and training, and to get commitment to implementation
Ubiquitous traceability	to provide automation such that traceability is encompassed within broader software and systems engineering processes, and is integral to all tool support

Table 7 Traceability research themes defined by CoEST (Gotel et al., 2012). Ubiquitous traceability is referred to as “the grand challenge of traceability”, since it requires significant progress in the other research themes

IR-based trace recovery have been conducted on industrial datasets of representative sizes.

7 Summary and Future Work

Our review of IR-based trace recovery compares 79 publications containing 132 empirical studies, systematically derived according to established procedures (Kitchenham and Charters, 2007). Our study constitutes the most extensive summary of publications of IR-based trace recovery yet published.

More than 10 IR models have been applied to trace recovery (RQ1). More studies have evaluated algebraic IR models (i.e. VSM and LSI) than probabilistic models (e.g. BIM, PIN, LM, LDA). A visible trend is, in line with development in the general field of IR, that the probabilistic subset of statistical language models have received increased attention in recent years. While extracting data from the primary publications, it became clear that the inconsistent use of IR terminology is an issue in the field. In an attempt to homogenize the language, we present structure in the form of a hierarchy of IR models (Figure 6) and a collection of IR terminology (Table 1).

In the 132 mapped empirical studies, artifacts from the entire development process have been linked (RQ2). The dominant artifact type is requirements at various levels of abstraction, followed by source code. Substantially fewer studies have been conducted on test artifacts, and only single publications have targeted user manuals and defect reports. Furthermore, a majority of the evaluations of IR-based trace recovery have

1 been made on bipartite datasets, i.e. only trace links between two disjoint sets of
2 artifacts were recovered.

3
4 Among the 79 primary publications mapped in our study, we conclude that the het-
5 erogeneity of reporting detail obstructs the aggregation of empirical evidence (RQ3).
6 Also, most evaluations have been conducted on small bipartite datasets containing
7 fewer than 500 artifacts, which is a severe threat to external validity. Furthermore, a
8 majority of evaluations have been using artifacts originating from a university envi-
9 ronment, or a dataset of proprietary artifacts from NASA. As a result, the two small
10 datasets EasyClinic and CM-1 constitute the de-facto benchmark in IR-based trace
11 recovery. Another validity threat to the applicability of IR-based trace recovery is that
12 a clear majority of the evaluations have been conducted in “the cave of IR evaluation”
13 as reported in Figure 12. Instead, the strongest empirical evidence in favor of IR-based
14 trace recovery tools comes from a set of controlled experiments on student subjects,
15 reporting that tool-supported subjects outperform manual control groups. Thus, we
16 argue that industrial in-vivo evaluations are needed to motivate the feasibility of the
17 approach and further studies on the topic, in which IR-based trace recovery should be
18 studied within the full complexity of an industrial setting. As such, our empirical find-
19 ings intensify the recent call for additional empirical studies by CoEST (Gotel et al.,
20 2012).

21 In several primary publications it is not made clear whether a query-based or
22 matrix-based evaluation style has been used. Also, the different reporting styles of P-R
23 values make secondary studies on candidate trace link accuracies challenging. We argue
24 that both the standard measures precision at fixed recall levels and P-R at specific
25 document cut-offs should be reported when applicable, complemented by secondary
26 measures such as MAP and DCG.

27 As a continuation of this literature study, we intend to publish the extracted data
28 to allow for collaborative editing, and for interested readers to review the details. A
29 possible future study would be to conduct a deeper analysis of the enhancement strate-
30 gies that have been reported as successful in the primary publications, to investigate
31 patterns concerning in which contexts they have been successfully applied. Another
32 option for the future is to aggregate results from the innermost evaluation context,
33 as P-R values repeatedly have been reported in the primary studies. However, such a
34 secondary study must be carefully designed to allow a valid synthesis across different
35 studies. Finally, future work could include other mapping dimensions, such as catego-
36 rizing the primary publications according to other frameworks, e.g positioning them
37 related to the CoEST research themes.

38 39 40 **Acknowledgement**

41
42 This work was funded by the Industrial Excellence Center EASE – Embedded Appli-
43 cations Software Engineering⁶. Thanks go to our librarian Mats Berglund for working
44 on the search strings, and Lorand Dali for excellent comments on IR details.

45 ⁶ <http://ease.cs.lth.se>

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Primary Publications

- 1
2
3
4 Abadi, A., M. Nisenson, and Y. Simionovici (2008*). “A traceability technique for
5 specifications”. In: *Proceedings of the 16th International Conference on Program*
6 *Comprehension*, pp. 103–112.
- 7 Ali, N., Y-G. Guéhéneuc, and G. Antoniol (2011*[a]). “Requirements traceability for
8 object oriented systems by partitioning source code”. In: *Proceedings of the 18th*
9 *Working Conference on Reverse Engineering*, pp. 45–54.
- 10 — (2011*[b]). “Trust-Based Requirements Traceability”. In: *Proceedings of the 19th*
11 *International Conference on Program Comprehension*, pp. 111–120.
- 12 Antoniol, G. et al. (1999*). “Recovering code to documentation links in OO systems”.
13 In: *Proceedings of the 6th Working Conference on Reverse Engineering*, pp. 136–
14 144.
- 15 Antoniol, G. et al. (2000*). “Tracing object-oriented code into functional require-
16 ments”. In: *Proceedings of the 8th International Workshop on Program Compre-*
17 *hension*, pp. 79–86.
- 18 — (2002*). “Recovering traceability links between code and documentation”. In: *Trans-*
19 *actions on Software Engineering*. Vol. 28, pp. 970–983.
- 20 Asuncion, H., A. Asuncion, and R. Taylor (2010*). “Software traceability with topic
21 modeling”. In: *Proceedings of the International Conference on Software Engineer-*
22 *ing*, pp. 95–104.
- 23 Ben Charrada, E. et al. (2011*). “Towards a benchmark for traceability”. In: *Pro-*
24 *ceedings of the 12th International Workshop on Principles on Software Evolution*,
25 pp. 21–30.
- 26 Borg, M. and D. Pfahl (2011*). “Do better IR tools improve the accuracy of engineers’
27 traceability recovery?” In: *Proceedings of the International Workshop on Machine*
28 *Learning Technologies in Software Engineering*, pp. 27–34.
- 29 Canfora, G. and L. Cerulo (2006*). “Fine grained indexing of software repositories to
30 support impact analysis”. In: *Proceedings of the International Workshop on Mining*
31 *Software Repositories*, pp. 105–111.
- 32 Capobianco, G. et al. (2009*[a]). “On the Role of the Nouns in IR-based Traceabil-
33 ity Recovery”. In: *Proceedings of the 17th International Conference on Program*
34 *Comprehension*, pp. 148–157.
- 35 — (2009*[b]). “Traceability Recovery using Numerical Analysis”. In: *Proceedings of*
36 *the 16th Working Conference on Reverse Engineering*, pp. 195–204.
- 37 Chen, X. (2010*). “Extraction and visualization of traceability relationships between
38 documents and source code”. In: *Proceedings of the International Conference on*
39 *Automated Software Engineering*, pp. 505–509.
- 40 Chen, X. and J. Grundy (2011*). “Improving automated documentation to code trace-
41 ability by combining retrieval techniques”. In: *Proceedings of the 26th International*
42 *Conference on Automated Software Engineering*, pp. 223–232.
- 43 Chen, X., J. Hosking, and J. Grundy (2011*). “A combination approach for enhancing
44 automated traceability”. In: *Proceeding of the 33rd International Conference on*
45 *Software Engineering, (NIER track)*, pp. 912–915.
- 46 Cleland-Huang, J. et al. (2005*). “Utilizing supporting evidence to improve dynamic
47 requirements traceability”. In: *Proceedings of the 13th International Conference on*
48 *Requirements Engineering*, pp. 135–144.
- 49 Cleland-Huang, J. et al. (2007*). “Best Practices for Automated Traceability”. In:
50 *Computer* 40.6, pp. 27–35.
- 51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Cleland-Huang, J. et al. (2010*). “A machine learning approach for tracing regulatory
2 codes to product specific requirements”. In: *Proceedings International Conference*
3 *on Software Engineering*, pp. 155–164.
- 4 Cuddeback, D., A. Dekhtyar, and J. Huffman Hayes (2010*). “Automated requirements
5 traceability: The study of human analysts”. In: *Proceedings of the 18th Interna-*
6 *tional Requirements Engineering Conference*, pp. 231–240.
- 7 Czauderna, A. et al. (2011*). “Traceability challenge 2011: using TraceLab to evaluate
8 the impact of local versus global idf on trace retrieval”. In: *Proceeding of the 6th In-*
9 *ternational Workshop on Traceability in Emerging Forms of Software Engineering*,
10 pp. 75–78.
- 11 De Lucia, A., R. Oliveto, and P. Sgueglia (2006*). “Incremental approach and user
12 feedbacks: A silver bullet for traceability recovery?” In: *Proceedings of the Interna-*
13 *tional Conference on Software Maintenance*, pp. 299–308.
- 14 De Lucia, A., R. Oliveto, and G. Tortora (2008*). “IR-based traceability recovery pro-
15 cesses: An empirical comparison of ”one-shot” and incremental processes”. In: *Pro-*
16 *ceedings of the 23rd International Conference on Automated Software Engineering*,
17 pp. 39–48.
- 18 — (2009*[a]). “Assessing IR-based traceability recovery tools through controlled ex-
19 periments”. In: *Empirical Software Engineering* 14.1, pp. 57–92.
- 20 — (2009*[b]). “The Role of the Coverage Analysis during IR-based Traceability Recov-
21 ery: a Controlled Experiment”. In: *Proceedings of the 25th International Conference*
22 *on Software Maintenance*, pp. 371–380.
- 23 De Lucia, A. et al. (2004*). “Enhancing an artefact management system with trace-
24 ability recovery features”. In: *Proceedings of the 20th International Conference on*
25 *Software Maintenance*, pp. 306–315.
- 26 — (2005*). “ADAMS re-trace: A traceability recovery tool”. In: *Proceedings of the 9th*
27 *European Conference on Software Maintenance and Reengineering*, pp. 32–41.
- 28 — (2006*). “Can Information Retrieval techniques effectively support traceability link
29 recovery?” In: *Proceedings of the 14th International Conference on Program Com-*
30 *prehension*, pp. 307–316.
- 31 — (2007*). “Recovering traceability links in software artifact management systems
32 using information retrieval methods”. In: *Transactions on Software Engineering*
33 *and Methodology* 16.4.
- 34 De Lucia, A. et al. (2011*). “Improving IR-based Traceability Recovery Using Smooth-
35 ing Filters”. In: *Proceedings of the 19th International Conference on Program Com-*
36 *prehension*, pp. 21–30.
- 37 Dekhtyar, A., J. Huffman Hayes, and J. Larsen (2007*). “Make the most of your time:
38 How should the analyst work with automated traceability tools?” In: *Proceedings*
39 *of the 3rd International Workshop on Predictor Models in Software Engineering*.
- 40 Dekhtyar, A. et al. (2007*). “Technique integration for requirements assessment”.
41 In: *Proceedings of the 15th International Requirements Engineering Conference*,
42 pp. 141–152.
- 43 Dekhtyar, A. et al. (2011*). “On human analyst performance in assisted requirements
44 tracing: Statistical analysis”. In: *Proceedings of the 19th International Requirements*
45 *Engineering Conference*, pp. 111–120.
- 46 Di, F. and M. Zhang (2009*). “An Improving Approach for Recovering Requirements-
47 to-Design Traceability Links”. In: *Proceedings of the International Conference on*
48 *Computational Intelligence and Software Engineering*, pp. 1–6.
- 49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Di Penta, M., S. Gradara, and G. Antoniol (2002*). “Traceability recovery in RAD
2 software systems”. In: *Proceedings of the 10th International Workshop on Program*
3 *Comprehension*, pp. 207–216.
- 4 Duan, C. and J. Cleland-Huang (2007*). “Clustering support for automated tracing”.
5 In: *Proceedings of the International Conference on Automated Software Engineer-*
6 *ing*, pp. 244–253.
- 7 Gethers, M. et al. (2011). “An adaptive approach to impact analysis from change
8 requests to source code”. In: *Proceedings of the 26th International Conference on*
9 *Automated Software Engineering*, pp. 540–543.
- 10 Gethers, M. et al. (2011*). “On integrating orthogonal information retrieval methods to
11 improve traceability recovery”. In: *Proceedings of the 27th International Conference*
12 *on Software Maintenance*, pp. 133–142.
- 13 Gibiec, M., A. Czauderna, and J. Cleland-Huang (2010*). “Towards mining replace-
14 ment queries for hard-to-retrieve traces”. In: *Proceedings of the International Con-*
15 *ference on Automated Software Engineering*, pp. 245–254.
- 16 Huffman Hayes, J., A. Dekhtyar, and J. Osborne (2003*). “Improving requirements
17 tracing via information retrieval”. In: *Proceedings of the 11th International Require-*
18 *ments Engineering Conference*, pp. 138–147.
- 19 Huffman Hayes, J., A. Dekhtyar, and S. Sundaram (2005*). “Text mining for software
20 engineering: how analyst feedback impacts final results”. In: *Proceedings of the*
21 *International Workshop on Mining Software Repositories*, pp. 1–5.
- 22 — (2006*). “Advancing candidate link generation for requirements tracing: the study
23 of methods”. In: *Transactions on Software Engineering* 32.1, pp. 4–19.
- 24 Huffman Hayes, J. et al. (2004*). “Helping analysts trace requirements: An objective
25 look”. In: *Proceedings of the 12th International Conference on Requirements Engi-*
26 *neering*, pp. 249–259.
- 27 Huffman Hayes, J. et al. (2007*). “REquirements TRacing On target (RETRO): Im-
28 proving software maintenance through traceability recovery”. In: *Innovations in*
29 *Systems and Software Engineering* 3.3, pp. 193–202.
- 30 Huffman Hayes, J. et al. (2011*). “Software verification and validation research labora-
31 tory (SVVRL) of the University of Kentucky: traceability challenge 2011: language
32 translation”. In: *Proceeding of the 6th International Workshop on Traceability in*
33 *Emerging Forms of Software Engineering*. ACM, pp. 50–53.
- 34 Jiang, H. et al. (2008*). “Incremental latent semantic indexing for automatic traceabil-
35 ity link evolution management”. In: *Proceedings of the 23rd International Confer-*
36 *ence on Automated Software Engineering*, pp. 59–68.
- 37 Kaushik, N., L. Tahvildari, and M. Moore (2011*). “Reconstructing traceability be-
38 tween bugs and test cases: An experimental study”. In: *Proceedings of the 18th*
39 *Working Conference on Reverse Engineering*, pp. 411–414.
- 40 Klock, S. et al. (2011*). “Traceclipse: an eclipse plug-in for traceability link recovery
41 and management”. In: *Proceedings of the 6th International Workshop on Traceabil-*
42 *ity in Emerging Forms of Software Engineering*, pp. 24–30.
- 43 Kong, L. et al. (2009*). “A requirement traceability refinement method based on rele-
44 vance feedback”. In: *Proceedings of the 21st International Conference on Software*
45 *Engineering and Knowledge Engineering*.
- 46 Kong, W. and J. Huffman Hayes (2011*). “Proximity-based traceability: An empirical
47 validation using ranked retrieval and set-based measures”. In: *Proceedings of the*
48 *1st International Workshop on Empirical Requirements Engineering*, pp. 45–52.
- 49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Kong, W. et al. (2011*). “How do we trace requirements: an initial study of analyst be-
2 havior in trace validation tasks”. In: *Proceeding of the 4th International Workshop*
3 *on Cooperative and Human Aspects of Software Engineering*, pp. 32–39.
- 4 Leuser, J. (2009*). “Challenges for semi-automatic trace recovery in the automotive
5 domain”. In: *Proceedings of the International Workshop on Traceability in Emerging*
6 *Forms of Software Engineering*, pp. 31–35.
- 7 Leuser, J. and D. Ott (2010*). “Tackling semi-automatic trace recovery for large speci-
8 fications”. In: *Requirements Engineering: Foundation for Software Quality*, pp. 203–
9 217.
- 10 Li, Y. et al. (2008*). “Requirement-centric traceability for change impact analysis: A
11 case study”. In: *International Conference on Software Process*, pp. 100–111.
- 12 Lormans, M. and A. van Deursen (2006*). “Can LSI help reconstructing requirements
13 traceability in design and test?” In: *Proceedings of the 10th European Conference*
14 *on Software Maintenance and Reengineering*, pp. 45–54.
- 15 Lormans, M., A. Van Deursen, and H-G. Gross (2008*). “An industrial case study
16 in reconstructing requirements views”. In: *Empirical Software Engineering* 13.6,
17 pp. 727–760.
- 18 Lormans, M. et al. (2006*). “Monitoring requirements coverage using reconstructed
19 views: An industrial case study”. In: *Proceedings of the 13th Working Conference*
20 *on Reverse Engineering*, pp. 275–284.
- 21 Mahmoud, A. and N. Niu (2010*). “Using semantics-enabled information retrieval in
22 requirements tracing: An ongoing experimental investigation”. In: *Proceedings of*
23 *the International Computer Software and Applications Conference*, pp. 246–247.
- 24 — (2011*). “Source code indexing for automated tracing”. In: *Proceeding of the 6th*
25 *International Workshop on Traceability in Emerging forms of Software Engineering*,
26 pp. 3–9.
- 27 Marcus, A., J. Maletic, and A. Sergeyev (2005*). “Recovery of traceability links be-
28 tween software documentation and source code”. In: *International Journal of Soft-*
29 *ware Engineering and Knowledge Engineering* 15.5, pp. 811–836.
- 30 McMillan, C., D. Poshyvanyk, and M. Revelle (2009*). “Combining Textual and Struc-
31 tural Analysis of Software Artifacts for Traceability Link Recovery”. In: *Proceedings*
32 *of the International Workshop on Traceability in Emerging Forms of Software En-*
33 *gineering*, pp. 41–48.
- 34 Natt och Dag, J., T. Thelin, and B. Regnell (2006*). “An experiment on linguistic tool
35 support for consolidation of requirements from multiple sources in market-driven
36 product development”. In: *Empirical Software Engineering* 11.2, pp. 303–329.
- 37 Natt och Dag, J. et al. (2002*). “A Feasibility Study of Automated Natural Language
38 Requirements Analysis in Market-Driven Development”. In: *Requirements Engi-*
39 *neering* 7.1, pp. 20–33.
- 40 Natt och Dag, J. et al. (2004*). “Speeding up Requirements Management in a Product
41 Software Company: Linking Customer Wishes to Product Requirements through
42 Linguistic Engineering”. In: *Proceedings of the 12th International Requirements*
43 *Engineering Conference*, pp. 283–294.
- 44 Oliveto, R. et al. (2010*). “On the equivalence of information retrieval methods for
45 automated traceability link recovery”. In: *Proceedings of the 18th International*
46 *Conference on Program Comprehension*, pp. 68–71.
- 47 Park, S. et al. (2000*). “Implementation of an Efficient Requirements Analysis Support-
48 ing System Using Similarity Measure Techniques”. In: *Information and Software*
49 *Technology* 42.6, pp. 429–438.
- 50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Parvathy, A. G, B. G Vasudevan, and R. Balakrishnan (2008*). “A comparative study
2 of document correlation techniques for traceability analysis”. In: *Proceedings of*
3 *the 10th International Conference on Enterprise Information Systems, Information*
4 *Systems Analysis and Specification*, pp. 64–69.
- 5 Port, D. et al. (2011*). “Experiences with text mining large collections of unstructured
6 systems development artifacts at JPL”. In: *Proceedings of the 33rd International*
7 *Conference on Software Engineering*, pp. 701–710.
- 8 Settimi, R. et al. (2004*). “Supporting software evolution through dynamically retriev-
9 ing traces to UML artifacts”. In: *Proceedings of the 7th International Workshop on*
10 *Principles of Software Evolution*, pp. 49–54.
- 11 Sultanov, H. and J. Huffman Hayes (2010*). “Application of swarm techniques to
12 requirements engineering: Requirements tracing”. In: *Proceedings of the 18th In-*
13 *ternational Requirements Engineering Conference*, pp. 211–220.
- 14 Sundaram, S., J. Huffman Hayes, and A. Dekhtyar (2005*). “Baselines in requirements
15 tracing”. In: *Proceedings of the Workshop on Predictor Models in Software Engi-*
16 *neering*, pp. 1–6.
- 17 Sundaram, S. et al. (2010*). “Assessing traceability of software engineering artifacts”.
18 In: *Requirements Engineering* 15.3, pp. 313–335.
- 19 Wang, X., G. Lai, and C. Liu (2009*). “Recovering Relationships between Documenta-
20 tion and Source Code based on the Characteristics of Software Engineering”. In:
21 *Electronic Notes in Theoretical Computer Science* 243, pp. 121–137.
- 22 Winkler, S. (2009*). “Trace retrieval for evolving artifacts”. In: *Proceedings of the In-*
23 *ternational Workshop on Traceability in Emerging Forms of Software Engineering*,
24 pp. 49–56.
- 25 Yadla, S., J. Huffman Hayes, and A. Dekhtyar (2005*). “Tracing requirements to de-
26 fect reports: an application of information retrieval techniques”. In: *Innovations in*
27 *Systems and Software Engineering* 1, pp. 116–124.
- 28 Zhao, W. et al. (2003*). “Understanding how the requirements are implemented in
29 source code”. In: *Proceedings of the 10th Asia-Pacific Software Engineering Con-*
30 *ference*, pp. 68–77.
- 31 Zhou, X. and H. Yu (2007*). “A clustering-based approach for tracing object-oriented
32 design to requirement”. In: *Proceedings of the 10th International Conference on*
33 *Fundamental Approaches to Software Engineering*, pp. 412–422.
- 34 Zou, X., R. Settimi, and J. Cleland-Huang (2006*). “Phrasing in dynamic requirements
35 trace retrieval”. In: *Proceedings of the 30th International Computer Software and*
36 *Applications Conference*, pp. 265–272.
- 37 — (2008*). “Evaluating the Use of Project Glossaries in Automated Trace Retrieval”.
38 In: *Proceedings of the International Conference on Software Engineering Research*
39 *and Practice*, pp. 157–163.
- 40 — (2010*). “Improving automated requirements trace retrieval: A study of term-based
41 enhancement methods”. In: *Empirical Software Engineering* 15.2, pp. 119–146.
- 42
43
44

45 References

- 46 Aitchison, J., D. Bawden, and A. Gilchrist (2000). *Thesaurus Construction and Use:*
47 *A Practical Manual*. 4th ed. Routledge.
- 48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Ali, N., Y-G. Guéhéneuc, and G. Antoniol (2012). “Factors Impacting the Inputs of
2 Traceability Recovery Approaches”. In: *Software and Systems Traceability*. Ed. by
3 J. Cleland-Huang, O. Gotel, and A. Zisman. Springer.
- 4 Antoniol, G. et al. (1999). “Evolving object oriented design to improve code traceabil-
5 ity”. In: *Proceedings of the 7th International Workshop on Program Comprehension*,
6 pp. 151–160.
- 7 Antoniol, G. et al. (2000). “Information retrieval models for recovering traceability
8 links between code and documentation”. In: *Conference on Software Maintenance*,
9 pp. 40–49.
- 10 Assawamekin, N., T. Sunetnanta, and C. Pluempitiwiriwaj (2010). “Ontology-based
11 multiperspective requirements traceability framework”. In: *Knowledge and Infor-*
12 *mation Systems* 25.3, pp. 493–522.
- 13 Ayari, K. et al. (2007). “Threats on building models from CVS and Bugzilla reposi-
14 tories: The Mozilla case study”. In: *Proceedings of the Conference of the Center for*
15 *Advanced Studies on Collaborative Research*, pp. 215–228.
- 16 Bacchelli, A., M. Lanza, and R. Robbes (2010). “Linking e-mails and source code
17 artifacts”. In: *Proceedings of the 32nd International Conference on Software Engi-*
18 *neering*, pp. 375–384.
- 19 Baeza-Yates, R. and B. Ribeiro-Neto (2011). *Modern Information Retrieval: The Con-*
20 *cepts and Technology behind Search*. Addison-Wesley.
- 21 Banko, M. and E. Brill (2001). “Scaling to very very large corpora for natural language
22 disambiguation”. In: *Proceedings of the 39th Annual Meeting on Association for*
23 *Computational Linguistics*, pp. 26–33.
- 24 Bianchi, A., A. Fasolino, and G. Visaggio (2000). “An exploratory case study of the
25 maintenance effectiveness of traceability models”. In: *Proceedings of the 8th Inter-*
26 *national Workshop on Program Comprehension*, pp. 149–158.
- 27 Binkley, D. and D. Lawrie (2010). “Information retrieval applications in software
28 maintenance and evolution”. In: *Encyclopedia of software engineering*. Ed. by J.
29 Marciniak. 2nd ed. Taylor & Francis.
- 30 Blei, D. and J. Lafferty (2007). “A correlated topic model of science”. In: *Annals of*
31 *Applied Statistics* 1.1, pp. 17–35.
- 32 Blei, D., A. Ng, and M. Jordan (2003). “Latent dirichlet allocation”. In: *The Journal*
33 *of Machine Learning Research* 3.4-5, pp. 993–1022.
- 34 Borg, M., P. Runeson, and L. Brodén (2012). “Evaluation of Traceability Recovery in
35 Context: A Taxonomy for Information Retrieval Tools”. In: *Proceedings of the 16th*
36 *International Conference on Evaluation & Assessment in Software Engineering*.
- 37 Borg, M., K. Wnuk, and D. Pfahl (2012). “Industrial comparability of student artifacts
38 in traceability recovery research - An exploratory survey”. In: *Proceedings of the*
39 *16th European Conference on Software Maintenance and Reengineering*.
- 40 Borillo, M. et al. (1992). “Applying linguistic engineering to spatial software engineer-
41 ing: the traceability problem”. In: *Proceedings of the 10th European Conference on*
42 *Artificial Intelligence*, pp. 593–595.
- 43 Bras, M. and Y. Toussaint (1993). “Artificial intelligence tools for software engineer-
44 ing: Processing natural language requirements”. In: *Applications of Artificial Intel-*
45 *ligence in Engineering*, pp. 275–290.
- 46 Brereton, P. et al. (2007). “Lessons from applying the systematic literature review pro-
47 cess within the software engineering domain”. In: *Journal of Systems and Software*
48 80.4, pp. 571–583.
- 49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Carnegie Mellon Software Engineering Institute (2010). *CMMI for Development, Ver-*
2 *sion 1.3*.
- 3 Castell, N. et al. (1994). “Quality control of software specifications written in natural
4 language”. In: *Proceedings of the 7th International Conference on Industrial and*
5 *Engineering Applications of Artificial Intelligence and Expert Systems*, pp. 37–44.
- 6 Chang, J. and D. Blei (2010). “Hierarchical relational models for document networks”.
7 In: *The Annals of Applied Statistics* 4.1, pp. 124–150.
- 8 Charikar, M. et al. (1997). “Incremental clustering and dynamic information retrieval”.
9 In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*,
10 pp. 626–635.
- 11 Cleland-Huang, J., C. K Chang, and M. Christensen (2003). “Event-based traceability
12 for managing evolutionary change”. In: *Transactions on Software Engineering* 29.9,
13 pp. 796–810.
- 14 Cleland-Huang, J., O. Gotel, and A. Zisman, eds. (2012). *Software and Systems Trace-*
15 *ability*. Springer.
- 16 Cleland-Huang, J., J. Huffman Hayes, and A. Dekhtyar (2006). *Center of Excellence for*
17 *Traceability: Problem statement and grand challenges in traceability (v0.1)*. Tech-
18 nical Report COET-GCT-06-01-0.9.
- 19 Cleland-Huang, J., W. Marrero, and B. Berenbach (2008). “Goal-Centric Traceability:
20 Using Virtual Plumblines to Maintain Critical Systemic Qualities”. In: *Transactions*
21 *on Software Engineering* 34.5, pp. 685–699.
- 22 Cleland-Huang, J. et al. (2011). “Grand Challenges, Benchmarks, and TraceLab: De-
23 veloping Infrastructure for the Software Traceability Research Community”. In:
24 *Proceedings of the 6th International Workshop on Traceability in Emerging Forms*
25 *of Software Engineering*.
- 26 Cleverdon, C. (1991). “The significance of the Cranfield tests on index languages”. In:
27 *Proceedings of the 14th Annual International SIGIR Conference on Research and*
28 *Development in Information Retrieval*, pp. 3–12.
- 29 Croft, B., H. Turtle, and D. Lewis (1991). “The use of phrases and structured queries
30 in information retrieval”. In: *Proceedings of the 14th Annual International ACM*
31 *SIGIR Conference on Research and Development in Information Retrieval*, pp. 32–
32 45.
- 33 De Lucia, A., F. Fasano, and R. Oliveto (2008). “Traceability management for impact
34 analysis”. In: *Frontiers of Software Maintenance*, pp. 21–30.
- 35 De Lucia, A. et al. (2006a). “COCONUT: COde COmprehension Nurturant Using
36 Traceability”. In: *Proceedings of the 22nd International Conference on Software*
37 *Maintenance*, pp. 274–275.
- 38 — (2006b). “Improving comprehensibility of source code via traceability information:
39 A controlled experiment”. In: *Proceedings of the 14th International Conference on*
40 *Program Comprehension*, pp. 317–326.
- 41 De Lucia, A. et al. (2012). “Information Retrieval Methods for Automated Traceability
42 Recovery”. In: *Software and Systems Traceability*. Ed. by J. Cleland-Huang, O.
43 Gotel, and A. Zisman. Springer.
- 44 Deerwester, S. et al. (1990). “Indexing by latent semantic analysis”. In: *Journal of the*
45 *American Society for Information Science* 41.6, pp. 391–407.
- 46 Dekhtyar, A. and J. Huffman Hayes (2006). “Good Benchmarks are Hard To Find:
47 Toward the Benchmark for Information Retrieval Applications in Software Engi-
48 neering”. In: *Proceedings of the 22th International Conference on Software Main-*
49 *tenance*.
- 50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Dekhtyar, A., J. Huffman Hayes, and G. Antoniol (2007). “Benchmarks for Traceability?” In: *Proceedings of the International Symposium on Grand Challenges in Traceability*.
- 2
3
4
5 Dit, B. et al. (2011). “Feature location in source code: a taxonomy and survey”. In: *Journal of Software Maintenance and Evolution: Research and Practice*.
- 6
7 Dömges, R. and K. Pohl (1998). “Adapting traceability environments to project-specific needs”. In: *Communications of the ACM* 41.12, pp. 54–62.
- 8
9
10 Egved, A. and P. Grunbacher (2002). “Automating requirements traceability: Beyond the record replay paradigm”. In: *Proceedings of the 17th International Conference on Automated Software Engineering*, pp. 163–171.
- 11
12 Eisenbarth, T., R. Koschke, and D. Simon (2003). “Locating features in source code”. In: *Transactions on Software Engineering* 29.3, pp. 210–224.
- 13
14 Falessi, D., G. Cantone, and G. Canfora (2010). “A comprehensive characterization of NLP techniques for identifying equivalent requirements”. In: *Proceedings of the 4th International Symposium on Empirical Software Engineering and Measurement*.
- 15
16 Felizardo, K. R et al. (2011). “Using Visual Text Mining to Support the Study Selection Activity in Systematic Literature Reviews”. In: *Proceedings of the 5th International Symposium on Empirical Software Engineering and Measurement*, pp. 77–86.
- 17
18
19 Fiutem, R. and G. Antoniol (1998). “Identifying design-code inconsistencies in object-oriented software: A case study”. In: *Proceedings of the International Conference on Software Maintenance*, pp. 94–102.
- 20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
- Gay, G. et al. (2009). “On the use of relevance feedback in IR-based concept location”. In: *Proceedings of the 25th International Conference on Software Maintenance*, pp. 351–360.
- Gotel, O. and C. Finkelstein (1994). “An analysis of the requirements traceability problem”. In: *Proceedings of the First International Conference on Requirements Engineering*, pp. 94–101.
- Gotel, O. et al. (2012). “The Grand Challenge of Traceability (v1.0)”. In: *Software and Systems Traceability*. Ed. by J. Cleland-Huang, O. Gotel, and A. Zisman. Springer.
- Heindl, M. and S. Biffl (2005). “A case study on value-based requirements tracing”. In: *Proceedings of the 10th European Software Engineering Conference held jointly with the 13th SIGSOFT International Symposium on Foundations of Software Engineering*, pp. 60–69.
- Hofman, T. (2001). “Unsupervised Learning by Probabilistic Latent Semantic Analysis”. In: *Machine Learning* 42.1-2, pp. 177–196.
- Huffman Hayes, J., G. Antoniol, and Y-G. Guéhéneuc (2008). “PREREQIR: Recovering Pre-Requirements via Cluster Analysis”. In: *Proceedings of the 15th Working Conference on Reverse Engineering*, pp. 165–174.
- Huffman Hayes, J. and A. Dekhtyar (2005a). “A framework for comparing requirements tracing experiments”. In: *International Journal of Software Engineering and Knowledge Engineering* 15.5, pp. 751–781.
- (2005b). “Humans in the traceability loop: can’t live with ’em, can’t live without ’em”. In: *Proceedings of the 3rd International Workshop on Traceability in Emerging Forms of Software Engineering*, pp. 20–23.
- Ingwersen, P. and K. Järvelin (2005). *The Turn: Integration of Information Seeking and Retrieval in Context*. Springer.
- International Electrotechnical Commission (2003). *IEC 61511-1 ed 1.0, Safety instrumented systems for the process industry sector*.

- 1 International Organization for Standardization (2011). *ISO 26262-1:2011 Road vehicles*
2 *– Functional safety –*.
- 3 Järvelin, K. and J. Kekäläinen (2000). “IR evaluation methods for retrieving highly
4 relevant documents”. In: *Proceedings of the 23rd Annual International ACM SIGIR*
5 *Conference on Research and Development in Information Retrieval*, pp. 41–48.
- 6 Jedlitschka, A., M. Ciolkowski, and D. Pfahl (2008). “Reporting Experiments in Soft-
7 ware Engineering”. In: *Guide to Advanced Empirical Software Engineering*. Ed. by
8 F. Shull, J. Singer, and D. Sjöberg. London: Springer, pp. 201–228.
- 9 Katta, V. and T. Stålhane (2011). “A Conceptual Model of Traceability for Safety Sys-
10 tems”. In: *Proceedings of the Complex Systems Design & Management Conference*.
- 11 Kekäläinen, J. and K. Järvelin (2002). “Evaluating information retrieval systems under
12 the challenges of interaction and multidimensional dynamic relevance”. In: *Proceed-*
13 *ings of the COLIS 4 Conference*, pp. 253–270.
- 14 Kitchenham, B., D. Budgen, and P. Brereton (2011). “Using mapping studies as the
15 basis for further research A participant-observer case study”. In: *Information and*
16 *Software Technology* 53.6, pp. 638–651.
- 17 Kitchenham, B. and S. Charters (2007). “Guidelines for performing Systematic Liter-
18 ature Reviews in Software Engineering”. In: *EBSE Technical Report*.
- 19 Kitchenham, B. et al. (2002). “Preliminary guidelines for empirical research in soft-
20 ware engineering”. In: *Transactions on Software Engineering and Methodology* 28.8,
21 pp. 721–734.
- 22 Kruchten, P. (2004). *The Rational Unified Process: an Introduction*. Addison-Wesley
23 Professional.
- 24 Lewis, D. (1998). “Naive (Bayes) at forty: The independence assumption in information
25 retrieval”. In: *Machine Learning: ECML-98*. Vol. 1398. Springer, pp. 4–15.
- 26 Liddy, E. (2001). *Natural Language Processing*. 2nd. Encyclopedia of Library and In-
27 formation Science. Marcel Decker.
- 28 Lin, J. et al. (2006). “Poirot: A Distributed Tool Supporting Enterprise-Wide Au-
29 tomated Traceability”. In: *Proceedings of the 14th International Conference on*
30 *Requirements Engineering*, pp. 363–364.
- 31 Lindvall, M. et al. (2009). “Searching for relevant software change artifacts using seman-
32 tic networks”. In: *Proceedings of the Symposium on Applied Computing*, pp. 496–
33 500.
- 34 Manning, C., P. Raghavan, and H. Schütze (2008). *Introduction to Information Re-*
35 *trieval*. Cambridge University Press.
- 36 Marcus, A. and J. Maletic (2003). “Recovering documentation-to-source-code traceabil-
37 ity links using latent semantic indexing”. In: *Proceedings of the 25th International*
38 *Conference on Software Engineering*, pp. 125–135.
- 39 Marcus, A. et al. (2004). “An information retrieval approach to concept location in
40 source code”. In: *Proceedings of the 11th Working Conference on Reverse Engi-*
41 *neering*, pp. 214–223.
- 42 Maron, M. and J. Kuhns (1960). “On Relevance, Probabilistic Indexing and Informa-
43 tion Retrieval”. In: *Journal of the ACM* 7.3, pp. 216–244.
- 44 Oliveto, R. (2008). “Traceability Management meets Information Retrieval Methods:
45 Strengths and Limitations”. PhD Thesis. University of Salerno.
- 46 Olsson, T. (2002). “Software Information Management in Requirements and Test Doc-
47 umentation”. Licentiate Thesis. Lund University.
- 48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 Petersen, K. and C. Wohlin (2009). "Context in industrial software engineering re-
2 search". In: *Proceedings of the 3rd International Symposium on Empirical Software*
3 *Engineering and Measurement*, pp. 401–404.
- 4 Petersen, K. et al. (2008). "Systematic Mapping Studies in Software Engineering". In:
5 *Proceedings of the 12th International Conference on Evaluation and Assessment in*
6 *Software Engineering*, pp. 71–80.
- 7 Pohl, K., G. Bockle, and F. van der Linden (2005). *Software product line engineering:*
8 *foundations, principles, and techniques*. Birkhuser.
- 9 Ponte, J. and B. Croft (1998). "A language modeling approach to information re-
10 trieval". In: *Proceedings of the 21st Annual International SIGIR Conference on*
11 *Research and Development in Information Retrieval*, pp. 275–281.
- 12 Randolph, J. (2005). "Free-Marginal Multirater Kappa (multirater K[free]): An Al-
13 ternative to Fleiss' Fixed-Marginal Multirater Kappa". In: *Joensuu Learning and*
14 *Instruction Symposium*.
- 15 Robertson, S. (1977). "The probability ranking principle in IR". In: *Journal of Docu-*
16 *mentation* 33.4, pp. 294–304.
- 17 Robertson, S. and J. Robertson (1999). *Mastering the Requirements Process*. Addison-
18 Wesley Professional.
- 19 Robertson, S. and H. Zaragoza (2009). "The probabilistic relevance framework: BM25
20 and Beyond". In: *Foundation and Trends in Information Retrieval* 3.4, pp. 333–
21 389.
- 22 Robertson, S. E. and S. Jones (1976). "Relevance weighting of search terms". In: *Jour-*
23 *nal of the American Society for Information Science* 27.3, pp. 129–146.
- 24 Rocchio, J. (1971). "Relevance feedback in information retrieval". In: *The SMART Re-*
25 *trieval System: Experiments in Automatic Document Processing*. Ed. by G. Salton.
26 Prentice-Hall, pp. 313–323.
- 27 Runeson, P., M. Alexandersson, and O. Nyholm (2007). "Detection of Duplicate Defect
28 Reports Using Natural Language Processing". In: *Proceedings of the 29th Interna-*
29 *tional Conference on Software Engineering*, pp. 499–510.
- 30 Runeson, P. et al. (2012). *Case Study Research in Software Engineering. Guidelines*
31 *and Examples*. Wiley.
- 32 Sabaliauskaite, G. et al. (2010). "Challenges in Aligning Requirements Engineering and
33 Verification in a Large-Scale Industrial Context". In: *Requirements Engineering:*
34 *Foundation for Software Quality*, pp. 128–142.
- 35 Salton, G. and C. Buckley (1988). "Term-weighting approaches in automatic text re-
36 trieval". In: *Information Processing and Management* 24.5, pp. 513–523.
- 37 Salton, G., A. Wong, and C. Yang (1975). "A vector space model for automatic index-
38 ing". In: *Communications of the ACM* 18.11, pp. 613–620.
- 39 Scacchi, W. (2002). "Understanding the requirements for developing open source soft-
40 ware systems". In: *IEEE Software* 149.1, pp. 24–39.
- 41 Shull, F. et al. (2008). "The role of replications in Empirical Software Engineering".
42 In: *Empirical Software Engineering* 13.2, pp. 211–218.
- 43 Singhal, A. (2001). "Modern information retrieval: a brief overview". In: *Data Engi-*
44 *neering Bulletin* 24.2, pp. 1–9.
- 45 Smeaton, A. and D. Harman (1997). "The TREC experiments and their impact on
46 Europe". In: *Journal of Information Science* 23.2, pp. 169–174.
- 47 Spanoudakis, G., A. d'Avila-Garcez, and A. Zisman (2003). "Revising Rules to Capture
48 Requirements Traceability Relations: A Machine Learning Approach". In: *Proceed-*
49
- 50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 *ings of the 15th International Conference in Software Engineering and Knowledge*
2 *Engineering.*
- 3 Spanoudakis, G. et al. (2004). “Rule-based generation of requirements traceability re-
4 lations”. In: *Journal of Systems and Software* 72.2, pp. 105–127.
- 5 Spärck Jones, K., S. Walker, and S. E. Robertson (2000). “A probabilistic model of
6 information retrieval: development and comparative experiments”. In: *Information*
7 *Processing and Management* 36.6, pp. 779–808.
- 8 Stone, A. and P. Sawyer (2006). “Using pre-requirements tracing to investigate require-
9 ments based on tacit knowledge”. In: *Proceedings of the 1st International Confer-*
10 *ence on Software and Data Technologies*, pp. 139–144.
- 11 Torchiano, M. and F. Ricca (2010). “Impact analysis by means of unstructured knowl-
12 edge in the context of bug repositories”. In: *Proceedings of the 4th International*
13 *Symposium on Empirical Software Engineering and Measurement*, 47:1–47:4.
- 14 Turtle, H. and B. Croft (1991). “Evaluation of an inference network-based retrieval
15 model”. In: *Transactions on Information Systems* 9.3, pp. 187–222.
- 16 Van Rompaey, B. and S. Demeyer (2009). “Establishing Traceability Links between
17 Unit Test Cases and Units under Test”. In: *Proceedings of the 13th European Con-*
18 *ference on Software Maintenance and Reengineering*, pp. 209–218.
- 19 Voorhees, E. (2005). *TREC: Experiment and Evaluation in Information Retrieval*. MIT
20 Press.
- 21 Winkler, S. and J. Pilgrim (2010). “A survey of traceability in requirements engineering
22 and model-driven development”. In: *Software & Systems Modeling* 9.4, pp. 529–565.
- 23 Wohlin, C. et al. (2012). *Experimentation in Software Engineering: A Practical Guide*.
24 Springer.
- 25 Zhai, C. (2007). *A Brief Review of Information Retrieval Models*. Technical Report.
26 University of Illinois at Urbana-Champaign.
- 27 — (2008). “Statistical Language Models for Information Retrieval A Critical Review”.
28 In: *Foundations and Trends in Information Retrieval* 2.3, pp. 137–213.
- 29 Zhai, C. and J. Lafferty (2001). “Model-based feedback in the language modeling ap-
30 proach to information retrieval”. In: *Proceedings of the 10th International Confer-*
31 *ence on Information and Knowledge Management*, pp. 403–410.
- 32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

A Classification of Primary Publications

Tables 8-12 present our classification of the primary publications, sorted by number of citations according to Google Scholar (July 1, 2012). Note that the well-cited works by Marcus and Maletic (2003) (354 citations) and Antoniol et al. (2000) (85 citations) are not listed. Applied IR models are reported in the fourth column. For LSI, the number of dimensions (k) in the reduced term-document space is reported in parenthesis, divided per dataset when possible. The number of dimensions is reported either as a fixed number of dimensions, an interval of dimensions, a dimensionality reduction in percent, or ‘N/A’ when the information is not available. A bold number represents that the best choice, as concluded by the original authors. Regarding LDA, the number of topics (t) is reported. Datasets are classified according to origin: proprietary (Ind), open source (OS), university (Univ), student (Stud), not clearly reported (Unclear), and mixed origin (Mixed). Numbers in parentheses show the number of artifacts studied, i.e. the total number of artifacts in the dataset, ‘N/A’ is used when it is not reported. Unless the full dataset name is presented, the following abbreviations are used: IBS (Ice Breaker System), EBT (Event-Based Traceability), LC (Light Control system), TM (Transient Meter). Evaluation, the rightmost column, maps primary publications to the context taxonomy described in Section 3 (Level 1-4 = retrieval context, seeking context, work task context, project context). Finally, Table 13 shows the distinctly most productive authors and affiliations, based upon our primary publications.

Cit.	Title	Authors	IR mod.	Dataset	Evaluation
486	Recovering Traceability Links between Code and Documentation	Antoniol, Canfora, De Lucia, Merlo	BIM, VSM	Univ: LEDA (296), Stud: Albergate (116)	Level 1, Level 3 (8 subj.)
205	Advancing Candidate Link Tracing: Generation for Requirements The Study of Methods	Huffman Hayes, Dekhtyar, Sundaram	VSM, LSI (k=10 (MODIS), 100 (CM-1))	Ind: MODIS (68), CM-1 (455)	Level 2
169	Improving Requirements Tracing via Information Retrieval	Huffman Hayes, Dekhtyar, Osborne	VSM	Ind: MODIS (68)	Level 1
140	Recovering Traceability Links in Software Artifact Management Systems Using Information Retrieval Methods	De Lucia, Fasano, Oliveto, Tortora	LSI (k=30-100%)	Stud: (Multiple projects)	Level 4 (150 subj.)
99	Utilizing Supporting Evidence to Improve Dynamic Requirements Traceability	Cleland-Huang, Settimi, Duan, Zou	PIN	Univ: IBS (252), EBT (114), LC (61)	Level 1
79	Best Practices for Automated Traceability	Cleland-Huang, Berenbach, Clark, Settimi, Romanova	PIN	Ind: Siemens Logistics and Automation (N/A), Univ: IBT (255), EBT (114)	Level 1
74	Helping Analysts Trace Requirements: An Objective Look	Huffman Hayes, Dekhtyar, Sundaram, Howard	VSM	Ind: MODIS (68)	Level 2
70	Can LSI help Reconstructing Requirements Traceability in Design and Test?	Lormans, van Deursen	LSI (k=20%)	Ind: Philips (359), Stud: PacMan (46), Callisto (N/A)	Level 1
68	Supporting Software Evolution through Dynamically Retrieving Traces to UML Artifacts	Settimi, Cleland-Huang, Khadra, Mody, Lukasik, DePalma	VSM	Univ: EBT (138)	Level 1
64	Enhancing an Artefact Management System with Traceability Recovery Features	De Lucia, Fasano, Oliveto, Tortora	LSI (k=10-50%)	Stud: EasyClinic (150)	Level 1

Table 8 Classification of primary publications, part I.

Cit.	Title	Authors	IR mod.	Dataset	Evaluation
58	Recovery of Traceability Links Between Software Documentation and Source Code	Marcus, Maletic, Sergeev	LSI (N/A)	Univ: LEDA (228-803), Stud: Albergate (73)	Level 1
44	Recovering Code to Documentation Links in OO Systems	Antoniol, Canfora, De Lucia, Marlo	BIM	Univ: LEDA (296)	Level 1
40	Fine grained indexing of software repositories to support impact analysis	Canfora, Cerulo	BM25	OS: Gedit (233), ArgoUML (2208), Firefox (680)	Level 1
38	ADAMS Re-Trace: A Traceability Recovery Tool	De Lucia, Fasano, Oliveto, Tortora	LSI (N/A)	Stud: (48, 50, 54, 55, 73, 74, 111)	Level 4 (7 proj.)
36	On the Equivalence of Information Retrieval Methods for Automated Traceability Link Recovery	Oliveto, Gethers, Poshyvanyk, De Lucia	VSM, LSI (N/A), LM, LDA (t=50-300)	Stud: EasyClinic (77), eTour (174)	Level 1
33	Incremental Approach and User Feedbacks: a Silver Bullet for Traceability Recovery	De Lucia, Oliveto, Sgueglia	VSM, LSI (k=10, 19, (MODIS), 60 (EasyClinic))	Ind: MODIS (68), Stud: EasyClinic (150)	Level 1
30	A machine learning approach for tracing regulatory codes to product specific requirements	Cleland-Huang, Czauderna, Gibiec, Emenecker	PIN	Mixed: (254)	Level 2
30	Assessing IR-based traceability recovery tools through controlled experiments	De Lucia, Oliveto, Tortora	LSI (N/A)	Stud: EasyClinic (150)	Level 3 (20, 12 subj.)
29	A Traceability Technique for Specifications	Abadi, Nisenson, Simionovici	VSM, LSI (k=5-100, 16 (SCA), 96 (CORBA)), PLSI (k=5-128), SDR (k=5-128), LM	OS: SCA (1311), CORBA (3340)	Level 2
29	Can Information Retrieval Techniques Effectively Support Traceability Link Recovery?	De Lucia, Fasano, Oliveto, Tortora	LSI (k=20%)	Stud: EasyClinic (150), Univ: ADAMS (309), LEDA (803)	Level 1, Level 4 (150 subj.)
29	Software traceability with topic modeling	Asuncion, Asuncion, Taylor	LSI (k=10), LDA (t=10, 20,30)	Univ: ArchStudio (N/A), Stud: EasyClinic (160)	Level 1
29	Speeding up Requirements to Management in a Product Software Company: Linking Customer Wishes Product Requirements through Linguistic Engineering	Natt och Dag, Gervasi, Brinkkemper, Regnell	VSM	Ind: Baan (12083)	Level 2
29	Tracing Object-Oriented Code into Functional Requirements	Antoniol, Canfora, De Lucia, Casazza, Merlo	BIM	Stud: Albergate (76)	Level 1
28	Clustering support for automated tracing	Duan, Cleland-Huang	PIN	Univ: IBS (185)	Level 1
27	Text mining for software engineering: how analyst feedback impacts final results	Huffman Hayes, Dekhtyar, Sundaram	N/A	Ind: MODIS (68)	Level 3 (3 subj.)
26	A feasibility study of automated natural language requirements analysis in market-driven development	Natt och Dag, Regnell, Carlshamre, Andersson, Karlsson	VSM	Ind: Telelogic (1891, 1089)	Level 1
26	Implementation of an Efficient Requirements Analysis Supporting System Using Similarity Measure Techniques	Park, Kim, Ko, Seo	Sliding window, syntactic parser	Ind: Unclear (33)	Level 1
25	Traceability Recovery in RAD Software Systems	Di Penta, Gradara, Antoniol	BIM	Univ: TM (49)	Level 1
23	REquirements TRacing On target (RETRO): improving software maintenance through traceability recovery	Huffman Hayes, Dekhtyar, Sundaram, Holbrook, Vadlamudi, April	VSM	Ind: CM-1 (74)	Level 3 (30 subj.)
22	Phrasing in Dynamic Requirements Trace Retrieval	Zou, Settimi, Cleland-Huang	PIN	Univ: IBS (235), LC (59), EBT (93)	Level 1

Table 9 Classification of primary publications, part II.

Cit.	Title	Authors	IR mod.	Dataset	Evaluation
21	Combining Textual and Structural Analysis of Software Artifacts for Traceability Link Recovery	McMillan, Poshyvanyk, Revelle	LSI (k=15, 25, 50, 75)	Univ: CoffeeMaker (143)	Level 1
20	Tracing requirements to defect reports: an application of information retrieval techniques	Yadla, Huffman Hayes, Dekhtyar	VSM	Ind: CM-1 (68,118)	Level 2
18	Automated Requirements Traceability: the Study of Human Analysts	Cuddeback, Dekhtyar, Huffman Hayes	VSM	OS: BlueJ Plugin (49)	Level 3 (26 subj.)
18	Incremental Latent Semantic Indexing for Automatic Traceability Link Evolution Management	Jiang, Nguyen, Chen, Jaygarl, Chang	LSI (k=10%)	Univ: LEDA (634)	Level 1
18	Understanding how the requirements are implemented in source code	Zhao, Zhang, Liu, Juo, Sun	VSM	OS: Desktop Calculator (123)	Level 1
17	Improving Automated Requirements Trace Retrieval: A Study of Term-Based Enhancement Methods	Zou, Settimi, Cleland-Huang	PIN	Ind: CM-1 (455), Univ: IBS (235), EBT (93), LC (89), Stud: SE450 (521)	Level 2
17	IR-Based Traceability Recovery Processes: An Empirical Comparison of "One-Shot" and Incremental Processes	De Lucia, Oliveto, Tortora	LSI (N/A)	Stud: EasyClinic (150)	Level 3 (30 subj.)
17	Make the Most of Your Time: How Should the Analyst Work with Automated Traceability Tools?	Dekhtyar, Huffman Hayes, Larsen	VSM	Ind: CM-1 (455)	Level 2
16	Baselines in requirements tracing	Sundaram, Huffman Hayes, Dekhtyar	VSM, LSI (k=10,19,29 (MODIS), 100,200 (CM-1))	Ind: CM-1 (455), MODIS (68)	Level 2
11	Challenges for semi-automatic trace recovery in the automotive domain	Leuser	VSM, LSI (N/A)	Ind: Daimler AG (1500)	Level 1
11	Monitoring Requirements Coverage Using Reconstructed Views: An Industrial Case Study	Lormans, Gross, van Deursen, Stehouwer, van Solingen	LSI (N/A)	Ind: LogicaCMG (219)	Level 1
11	On the role of the nouns in IR-based traceability recovery	Capobianco, De Lucia, Oliveto, Panichella, Panichella	LSI (N/A), LM	Stud: EasyClinic (150)	Level 1
10	An experiment on linguistic tool support for consolidation of requirements from multiple sources in market-driven product development	Natt och Dag, Thelin, Regnell	VSM	Stud: PUSS (299)	Level 3 (23 subj.)
9	An Industrial Case Study in Reconstructing Requirements Views	Lormans, van Deursen, Gross	LSI (k=40%)	Ind: LogicaCMG (293)	Level 1
9	Towards Mining Replacement Queries for Hard-to-Retrieve Traces	Gibiec, Czauderna, Cleland-Huang	VSM	Mixed: (254)	Level 2
8	Recovering Relationships between Documentation and Source Code based on the Characteristics of Software Engineering	Wang, Lai, Liu	LSI (N/A), BIM	Univ: LEDA (597), Univ: IBS (270)	Level 1
8	Trace retrieval for evolving artifacts	Winkler	LSI (k=15%)	Ind: Robert Bosch GmbH (500), MODIS (68)	Level 1
8	Traceability Recovery using Numerical Analysis	Capobianco, De Lucia, Oliveto, Panichella, Panichella	VSM, LSI (N/A), LM, B-splines	Stud: EasyClinic (150)	Level 1
7	Assessing Traceability of Software Engineering Artifacts	Sundaram, Huffman Hayes, Dekhtyar, Holbrook	VSM, LSI (k=10,25, 30,40,60 (MODIS), 10,25,100,200, 400 (CM-1), 5,10,15,25,40 (Waterloo))	Ind: MODIS (68), CM-1 (455), Stud: 22* Waterloo (65)	Level 2
7	Requirement-centric traceability for change impact analysis: A case study	Li, Li, Yang, Li	VSM	Unclear: Requirements Management System (501)	Level 4 (5 subj.)

Table 10 Classification of primary publications, part III.

Cit.	Title	Authors	IR mod.	Dataset	Evaluation
6	How do we trace requirements: an initial study of analyst behavior in trace validation tasks	Kong, Huffman Hayes, Dekhtyar, Holden	N/A	OS: BlueJ plugin (49)	Level 3 (13 subj.)
6	Technique Integration for Requirements Assessment	Dekhtyar, Huffman Hayes, Sundaram, Holbrook, Dekhtyar	VSM, LSI (N/A), BIM LDA (N/A), Chi2 key extr.	Ind: CM-1 (455)	Level 1
4	Application of Swarm Techniques for Requirements Engineering: Requirements Tracing	Sultanov, Huffman Hayes	VSM, Swarm	Ind: CM-1 (455), Univ: PINE (182)	Level 1
4	On Integrating Orthogonal Information Retrieval Methods to Improve Traceability Recovery	Gethers, Oliveto, Posyvanyk, De Lucia	VSM, LM, RTM	Stud: eAnsi (194), eAnsi (67), EasyClinic (57) EasyClinic (100), eTour (232), SMOS (167)	Level 1
3	A clustering-based approach for tracing object-oriented design to requirement	Zhou, Yu	VSM	Univ: Resource Management Software (33)	Level 1
3	Evaluating the Use of Project Glossaries in Automated Trace Retrieval	Zou, Settini, Cleland-Huang	PIN	Ind: CM-1 (455), Univ: IBS (235), Stud: SE450 (61)	Level 1
3	On Human Analyst Performance in Assisted Requirements Tracing: Statistical Analysis	Dekhtyar, Dekhtyar, Holden, Huffman Hayes, Cuddeback, Kong	VSM	OS: BlueJ (49)	Level 3 (84 subj.)
3	Tackling Semi-automatic Trace Recovery for Large Specifications	Leuser, Ott	VSM	Ind: Daimler (2095, 944)	Level 1
2	Extraction and visualization of traceability relationships between documents and source code	Chen	Unclear	OS: JDK1.5 (N/A), uDig 1.1.1 (N/A)	Level 1
2	Source code indexing for automated tracing	Mahmoud, Niu	VSM	Stud: eTour (174), iTrust (264)	Level 1
2	Traceability challenge 2011: using tracelab to evaluate the impact of local versus global idf on trace retrieval	Czauderna, Gibiec, Leach, Li, Shin, Keenan, Cleland- Huang	VSM	Ind: CM-1 (75), WV-CCHIT (1180)	Level 2
2	Trust-Based Requirements Traceability	Ali, Guéhéneuc, Antoniol	VSM	OS: Pooka (388), SIP (1853)	Level 1
1	An Adaptive Approach to Impact Analysis from Change Requests to Source Code	Gethers, Kagdi, Dit, Posyvanyk	LSI (N/A)	OS: ArgoUML (qualitative analysis)	Level 2
1	Do Better IR Tools Improve the Accuracy of Engineers' Traceability Recovery?	Borg, Pfahl	VSM	Ind: CM-1 (455)	Level 3 (8 subj.)
1	Experiences with text mining large collections of unstructured systems development artifacts at JPL	Port, Nikora, Hihn, Huang	LSI (N/A)	Unclear	Level 3
1	Improving Automated Documentation to Code Traceability by Combining Retrieval Techniques	Chen, Grundy	VSM	OS: JDK (431)	Level 1
1	Improving IR-based Traceability Recovery Using Smoothing Filters	De Lucia, Di Penta, Oliveto, Panichella, Panichella	VSM, LSI (N/A)	Univ: PINE (131), Stud: EasyClinic (150)	Level 1
1	Using semantics-enabled information retrieval in requirements tracing: An ongoing experimental investigation	Mahmoud, Niu	VSM	Ind: CM-1 (455)	Level 1
1	Traceclipse: an eclipse plug-in for traceability link recovery and management	Klock, Gethers, Dit, Posyvanyk	Unclear	Ind: CM-1 (455), Stud: EasyClinic (150)	Level 1
0	A combination approach for enhancing automated traceability: (NIER track)	Chen, Hosking, Grundy	VSM	OS: JDK 1.5 (N/A)	Level 1

Table 11 Classification of primary publications, part IV.

Cit.	Title	Authors	IR mod.	Dataset	Evaluation
0	A Comparative Study of Document Correlation Techniques for Traceability Analysis	Parvathy, Vasudevan, Balakrishnan	VSM, LSI (k=10), LDA (t=21), CTM	Unclear: (43), (261)	Level 1
0	A requirement traceability refinement method based on relevance feedback	Kong, Li, Li, Yang, Wang	VSM, LM	Ind: Web app (511)	Level 1
0	An Improving Approach for Recovering Requirements-to-Design Traceability Links	Di, Zhang	BIM	Ind: CM-1 (455), MODIS (68)	Level 1
0	Proximity-based traceability: An empirical validation using ranked retrieval and set-based measures	Kong, Huffman Hayes	VSM	Ind: CM-1 (75), OS: Pine (182), Univ: StyleChecker (49), Stud: EasyClinic (77)	Level 2
0	Reconstructing Traceability between Bugs and Test Cases: An Experimental Study	Kaushik, Tahvildari, Moore	LSI (k=50-500, 150-200)	Ind: RIM (13389)	Level 1
0	Requirements Traceability for Object Oriented Systems by Partitioning Source Code	Ali, Guéhéneuc, Antoniol	VSM	OS: Pooka (388), SIP (1853), Univ: iTrust (526)	Level 1
0	Software verification and validation research laboratory (SVVRL) of the University of Kentucky: traceability challenge 2011: language translation	Huffman Hayes, Sultanov, Kong, Li	VSM	Stud: EasyClinic (150), eTour (174)	Level 2
0	The role of the coverage analysis during IR-based traceability recovery: A controlled experiment	De Lucia, Oliveto, Tortora	LSI (N/A)	Stud: EasyClinic (150)	Level 3 (30 subj.)
0	Towards a Benchmark for Traceability	Ben Charrada, Casper, Jeanneret, Glinz	VSM	Univ: AquaLush (793)	Level 1

Table 12 Classification of primary publications, part V.

Author	Publications
Andrea De Lucia	16 (9)
Jane Huffman Hayes	16 (6)
Alexander Dekhtyar	15 (3)
Rocco Oliveto	13 (1)
Jane Cleland-Huang	10 (3)
Affiliation	Publications
University of Kentucky, United States	13
University of Salerno, Italy	11
DePaul University, United States	10
University of Sannio, Italy	5

Table 13 Most productive authors and affiliations. For authors, the first number is the total number of primary publications, while the number in parenthesis is first-authored primary publications. For affiliations, the numbers show the number of primary publications first-authored by an affiliated researcher.