

Feature location in source code: a taxonomy and survey

Bogdan Dit, Meghan Revelle, Malcom Gethers and Denys Poshyvanyk^{*,†}

Computer Science Department, The College of William & Mary, Williamsburg, VA, USA

SUMMARY

Feature location is the activity of identifying an initial location in the source code that implements functionality in a software system. Many feature location techniques have been introduced that automate some or all of this process, and a comprehensive overview of this large body of work would be beneficial to researchers and practitioners. This paper presents a systematic literature survey of feature location techniques. Eighty-nine articles from 25 venues have been reviewed and classified within the taxonomy in order to organize and structure existing work in the field of feature location. The paper also discusses open issues and defines future directions in the field of feature location. Copyright © 2011 John Wiley & Sons, Ltd.

Received 15 June 2010; Revised 27 May 2011; Accepted 24 July 2011

KEY WORDS: Feature location; concept location; program comprehension; software maintenance and evolution

1. INTRODUCTION

In software systems, a feature represents a functionality that is defined by requirements and accessible to developers and users. Software maintenance and evolution involves adding new features to programs, improving existing functionalities, and removing bugs, which is analogous to removing unwanted functionalities. Identifying an initial location in the source code that corresponds to a specific functionality is known as *feature* (or *concept*) *location* [1, 2]. It is one of the most frequent maintenance activities undertaken by developers because it is a part of the incremental change process [3]. During the incremental change process, programmers use feature location to find where in the code the first change to complete a task needs to be made. The full extent of the change is then handled by impact analysis, which starts with the source code identified by feature location and finds all the code affected by the change. Methodologically, the two activities of feature location and impact analysis are different and are treated separately in the literature and in this survey.

Feature location is one of the most important and common activities performed by programmers during software maintenance and evolution. No maintenance activity can be completed without first locating the code that is relevant to the task at hand, making feature location essential to software maintenance since it is performed in the context of incremental change. For example, Alice is a new developer on a software project, and her manager has given her the task of fixing a bug that has been recently reported. Since Alice is new to this project, she is unfamiliar with the large code base of the software system and does not know where to begin. Lacking sufficient documentation on the system and the ability to ask the code's original authors for help, the only option Alice sees is to manually search for the code relevant to her task.

^{*}Correspondence to: Denys Poshyvanyk, Computer Science Department, The College of William & Mary, Williamsburg, VA, USA.

[†]E-mail: denys@cs.wm.edu

Alice's situation is one faced by many software developers needing to understand and modify an unfamiliar codebase. However, a manual search of a large amount of source code, even with the help of tools such as pattern matchers or an integrated development environment, can be frustrating and time-consuming. Recognizing this problem, software engineering researchers have developed a number of feature location techniques (FLT) to come to aid programmers in Alice's position. The various techniques that have been introduced are all unique in terms of their input requirements, how they locate a feature's implementation, and how they present their results. Thus, even the task of choosing a suitable FLT can be challenging.

The existence of such a large body of feature location research calls for a comprehensive overview. Since there currently is no broad summary of the field of feature location, this paper provides a systematic survey and operational taxonomy of this pertinent research area. To the best of our knowledge, Wilde *et al.*'s work [4] is the only other survey, which in contrast to our survey, compares only a few FLT. Our survey includes research articles that introduce new feature location approaches; case, industrial, and user studies; and tools that can be used in support of feature location. The articles are characterized within a taxonomy that has nine dimensions, and each dimension has a set of attributes associated with it. The dimensions and attributes of the taxonomy capture key facets of typical FLT and can be useful to both software engineering researchers and practitioners [5]. Researchers can use this survey to identify what has been done in the area of feature location and what needs to be done; that is, they can use it to find related work as well as opportunities for future research. Practitioners can use this overview to determine which feature location approach is most suited to their needs.

This survey encompasses 89 articles (60 research articles and 29 tool and case study papers) from 25 venues, published between November 1992 and February 2011. These research articles were selected because they either state feature/concept location as their goal or present a technique that is essentially equivalent to feature location. The tool papers include tools developed specifically for feature location as well as program exploration tools that support feature location. The case study articles include industrial and user studies as well as studies that compare existing approaches.

There are several research areas that are closely related to feature location, such as traceability link recovery, impact analysis, and aspect mining. Traceability link recovery seeks to connect different types of software artifacts (e.g., documentation with source code), while feature location is more concerned with identifying source code associated with functionalities, not specific sections of a document. Impact analysis is the step in the incremental change process performed after feature location with the purpose of expanding on feature location's results, especially after a change is made to the source code. Feature location focuses on finding the starting point for that change. The main goal of aspect mining is to identify cross-cutting concerns and determine the source code that should be refactored into aspects, meaning the aspects themselves are not known a priori. By contrast, in the contexts in which feature location is used, the high-level descriptions of features are already known and only the code that implements them is unknown. Therefore, articles and research from these related fields are not included here as they are beyond the scope of this focused survey.

The work presented in this paper has two main contributions. The first is a systematic survey of FLT, relevant case studies, and tools. The second is the taxonomy derived from those techniques. Our online Appendix¹ lists all of the surveyed articles classified within the taxonomy. Section 2 presents the systematic review process. Section 3 introduces the dimensions of the taxonomy, and Section 4 provides brief descriptions of the surveyed approaches. Section 5 overviews the feature location tools and studies, and Section 6 provides an analysis of the taxonomy. Section 7 discusses open issues in feature location, and Section 8 concludes.

2. SYSTEMATIC REVIEW PROCESS

In this paper, we perform a systematic survey of the feature location literature in order to address the following research questions (RQ):

¹<http://www.cs.wm.edu/semeru/data/feature-location-survey/> (accessed and verified on 03/01/2011)

- RQ1: What types of analysis are used while performing feature location?
- RQ2: Has there been a change in types of analysis used to identify features in source code employed by recent FLT's?
- RQ3: Are there any limitations to current strategies for evaluating various FLT's?

In order to answer these research questions, we conducted a systematic review of the literature using the following process (see Figure 1):

- **Search:** The initial set of articles to be considered during the selection process is determined by identifying pertinent journals, conferences, and workshops.
- **Article Selection:** Using inclusion and exclusion criteria, the initial set of articles is filtered and only relevant articles are considered beyond this step.
- **Article Characterization:** Articles, which meet the selection criteria, are classified according to the set of attributes that capture important characteristics of FLT's.
- **Analysis:** Using the resulting taxonomy and systematic classification of the papers, the research questions are answered and useful insights about the state of feature location research and practice are outlined.

2.1. Search

An initial subset of papers of interest was obtained by manually evaluating articles that appear in different venues considered during our preliminary exploration. We select venues where feature location research is within their respective scope. Also, choosing such venues ensured that the selected articles meet some standard (e.g., the papers went through a rigorous peer review process).

2.2. Article selection

To adhere to the properties of systematic reviews [6], we define the following inclusion and exclusion criteria. In order to be included in the survey, a paper must introduce, evaluate, and/or complement the implementation of a source code based FLT. This includes papers that introduce novel FLT's, evaluate various existing FLT's, or present tools implementing existing or new approaches to feature location. The papers, which focused on improving the performance of underlying analysis techniques (e.g., dynamic analysis, Information Retrieval), as opposed to the feature location process, were excluded.




Figure 1. Systematic review process.

2.3. Article classification

The authors read and categorized each article according to the taxonomy and the attributes presented in Section 3. The process of classifying the articles was followed by four authors individually. Using the initial classifications produced by the authors, we identified papers that had some disagreements and further discussed those papers. The set of attributes was extracted and defined by two of the authors. Having all four authors characterize the articles allows us to verify the quality of the taxonomy, minimizing potential bias. In certain cases, disagreements served as an indication that our taxonomy and attributes or their corresponding descriptions required refinement. Through this process, we were able to improve the quality of our taxonomy and attribute set as well as improve their descriptions.

2.4. Analysis

Following the process of classifying research papers, our final step includes analysis of the results, answers to the research questions, as well as an outline of future directions for researchers and practitioners investigating FLTs. In order to complete this step, we analyzed the trends in our resulting taxonomy and observed interesting co-occurrences of various attributes across FLTs. We also investigated characteristics that rarely apply to the set of techniques considered, as well as characteristics which are currently emerging in the research literature.

3. DIMENSIONS OF THE SURVEY

The goal of this survey is to provide researchers and practitioners with a structured overview of existing research in the area of feature location. From a methodical inspection of the research literature, we extracted a number of key dimensions.² These dimensions objectively describe different techniques and offer structure to the surveyed literature. The dimensions are as follows:

- *The Type of analysis*: What underlying analyses are used to support feature location?
- *The Type of user input*: What does a developer have to provide as an input to the FLT?
- *Data sources*: What derivative artifacts have to be provided as an input for the FLT?
- *Output*: What is the type of the results and how are they provided back to the user?
- *Programming language support*: On which programming languages was this technique instantiated?
- *The Evaluation of the approach*: How was this FLT evaluated?
- *Systems evaluated*: What are the systems that were used in the evaluation?

The order in which these dimensions are presented does not imply any explicit priority or importance. Each dimension has a number of distinct attributes associated with it. For a given dimension, a FLT may be associated with multiple attributes. These dimensions and their attributes were derived by examining an initial set of articles of interest. They were then refined and generalized to succinctly characterize the properties that make FLTs unique, and can be used to evaluate and compare them. The goal of the taxonomy's dimensions and attributes is to allow researchers and practitioners to easily locate the FLTs that are most suited to their needs. The dimensions and their associated attributes that are used in the taxonomy of the surveyed articles are listed in Table I. These dimensions and attributes are discussed in the remainder of this section. The *attributes* are highlighted in italics.

3.1. Type of analysis

A main distinguishing factor of FLTs is the type or types of analyses they employ to identify the code that pertains to a feature. The most common types of analyses include dynamic, static, and textual. While these are not the only types of analysis possible, they are the ones utilized by the vast majority of FLTs, and some approaches even leverage more than one of these types of analysis. In

²Some of these dimensions were discussed at the working session on Information Retrieval Approaches in Software Evolution at 22nd IEEE International Conference on Software Maintenance (ICSM'06): <http://www.cs.wayne.edu/~amarcus/icsm2006/>

Table I. Dimensions and attributes of the feature location taxonomy.

Dimension	Attribute: Description
Type of analysis	Dynamic: Dynamic analysis is used to locate features Static: Static analysis is used to locate features Textual: Textual analysis is used to locate features Historical: Information from software repositories is used to locate features Other: Another type of analysis is used to locate features
User input	Natural Language Query: A textual description of a feature (or a bug report description) or a user-specified query is used as an input Execution Scenario: The developer uses a scenario in order to exercise a feature (or reproduce a bug) of interest in order to collect execution traces Source Code Artifact: A software artifact is used as a starting point (or seed) for the FLT
Data sources (derivative inputs)	Source code: • Compilable/executable: An executable program is used as an input to extract derivative information • Non-compilable: A source code that may or may not contain errors that prevent it to be compilable is used as an input to extract derivative information Derivative analysis data from source code: • Dependence Graph: The dependence graph can be derived directly from the source code (i.e., static) or it can be derived from execution information (i.e., dynamic) • Execution Trace: A set of methods that were executed when exercising a scenario • Historical Information: Historical information from source code repositories about the changes in the source code. This information includes log messages and actual changes • Other: Another source of information is used for feature location
Output	Source code: • File/class: The results produced by the FLT are at file/class-level granularity • Method/function: The results produced by the FLT are at method/function-level granularity • Statement: The results produced by the FLT are at statement-level granularity Non-source code artifact: The output produced by the FLT is a non-source code artifact, e.g., bug report
Programming language support	Presentation of the results: • Ranked: The results produced by the FLT are ranked based on their relevance to the user input • Visualization: The results produced by the FLT are presented using an advanced way of visualizing information
Evaluation	Java: The approach supports feature location in software written primarily in Java C/C++: The technique can find features for software systems written in C/C++ Other: Feature location in some other language is supported, e.g., Fortran, Cobol Preliminary: The evaluation is on small systems or small data sets or only preliminary evidence is given (i.e., proof of concept) Benchmark: The evaluation uses a data set that was published by other authors or the data set used in this evaluation is later used by other researchers
Systems evaluated	Human subjects: • Academic: Students or non-professional developers participated in evaluation • Professional: Professional developers participation in evaluation of the results Quantitative: The results were evaluated via comparative metrics such as precision, recall, etc. Qualitative: The paper discusses details about the characteristics of the technique/tool and/or some aspects of the results Comparison with other approaches: Comparisons of the author's approach with existing solutions Unknown/none: There is no evaluation performed, or the details are not available The software systems upon which the FLT has been applied are listed

The attributes are highlighted in bold, and their description is given after the colon separator.

Section 4, descriptions of all the surveyed articles are given, and the section is organized by the type(s) of analysis used.

Dynamic analysis refers to examining a software system's execution, and it is often used for feature location when features can be invoked and observed during runtime. Feature location using dynamic analysis generally relies on a post-mortem analysis of an execution trace. Typically, one or more feature-specific scenarios are developed that invoke only the desired feature. Then, the scenarios are run and execution traces are collected, recording information about the code that was invoked. These traces are captured either by instrumenting the system or through profiling. Once the traces are obtained, feature location can be performed in several ways. The traces can be compared with other traces in which the feature was not invoked to find code only invoked in the feature-specific traces [7, 8]. Alternatively, the frequency of execution portions of code can be analyzed to locate a feature's implementation [9–11]. Using dynamic analysis for feature location is a popular choice since most features can be mapped to execution scenarios. However, there are some limitations associated with dynamic analysis. The collection of traces can impose considerable overhead on a system's execution. Additionally, the scenarios used to collect traces may not invoke all of the code that is relevant to the feature, meaning that some of the feature's implementation may not be located. Conversely, it may be difficult to formulate a scenario that invokes only the desired feature, causing irrelevant code to be executed. Dynamic FLT's are discussed in Section 4.2.

Static analysis examines structural information such as control or data flow dependencies. In manual feature location, developers may follow program dependencies in a section of code they deem to be relevant in order to find additional useful code, and this idea is used in some approaches to feature location [12]. Other techniques analyze the topology of the structural information to point programmers to potentially relevant code [13]. While using static analysis for feature location is very close to what a developer searching for code may do, it often overestimates what is pertinent to a feature and is prone to returning many false positive results. Static approaches to feature location are summarized in Section 4.3.

Textual approaches to feature location analyze the words used in source code. The idea is that identifiers and comments encode domain knowledge, and a feature may be implemented using a similar set of words throughout a software system, making it possible to find a feature's relevant code textually. Textual analysis is performed using three main techniques: pattern matching, Information Retrieval (IR), and Natural Language Processing (NLP). Pattern matching usually involves a textual search of source code using a utility, such as `grep`.³ IR techniques, such as Latent Semantic Indexing (LSI) [14], Latent Dirichlet Allocation (LDA) [15], and Vector Space Model [16], are statistical methods used to find a feature's relevant code by analyzing and retrieving identifiers and comments that are similar to a query provided by a user. A good overview of applications of IR techniques in Software Development, Maintenance and Evolution can be found in [17, 145]. NLP approaches can also exploit a query, but they analyze the parts of speech of the words used in source code. Pattern matching is relatively robust but not very precise because of the vocabulary problem [18]; the chances of a programmer choosing query terms that match the vocabulary of unfamiliar source code are relatively low. On the other hand, NLP is more precise than pattern matching, but much more expensive. IR lies between the two. No matter the type of textual analysis used, the quality of feature location is heavily tied to the quality of the source code naming conventions and/or the user-issued query. Textual FLT's are reviewed in Section 4.4.

Feature location is not limited to just dynamic, static, or textual analysis. Many techniques draw on multiple analysis methods to find a feature's implementation, and some do not use any of these types of analyses. Existing approaches that combine two or more types of analysis do so with the goal of using one type of analysis to compensate for the limitations of another, thus achieving better results than standalone techniques. The unique ways in which multiple types of analysis are combined for feature location are described in Sections 4.5 to 4.8. Other approaches do not rely on dynamic, static, or textual analysis. For instance, two FLT's rely on historical analysis to mine version control systems in order to identify lines of code [19] or artifacts related to a feature [20]. Another

³<http://www.gnu.org/software/grep/> (accessed and verified on 03/01/2011)

technique examines the code visible to a programmer during a maintenance task and tries to infer what was important [21]. These alternative approaches are explained in Section 4.9.

3.2. Types of user input

A FLT ultimately assists developers during software maintenance tasks. This dimension describes the type of input that a developer has to provide for the FLT. The input can be a *query*, an *execution scenario*, a *source code artifact*, or a combination of these. The query can be a set of words that describe a bug or a feature that the developer is attempting to locate. The query can be generated in several ways. It can be compiled by the developer, it can be suggested by the FLT supporting automatic query expansion, or it can be extracted from the description of the feature report or some other documentation artifacts. The execution scenario is a set of steps, which the developer has to perform on an instrumented software system in order to exercise a feature of interest, with the purpose of collecting execution information about the system (i.e., execution trace). The source code artifact can be any artifact, such as a class or method, which the developer chooses as a starting point for the FLT. The rationale is that the FLT will perform some analysis starting from that artifact (e.g., a method in source code) and it will return other artifacts (e.g., source code methods) related to it. As mentioned before, based on the FLT, the developer has to provide as an input any of these types of inputs, or a combination of these.

3.3. Data sources (derivative inputs from software)

In addition to the type of user input, the FLT might require other sources of information such as the source code or derivative data from source code. Some techniques require the source code to be compilable or executable, in order to extract static dependencies or execution information, whereas other techniques that use textual information could use as input source code that contains errors and therefore may not necessarily be compilable. On the other hand, some techniques require artifacts derived from source code, such as dependence graphs, execution traces, and historical information from source code repositories concerning source code changes (i.e., change log messages and concrete source code changes). We also include the attribute *other* for the information extracted from bug repositories, forums, documentation, test cases, etc. In addition, the attribute *other* can denote the fact that the technique requires some feedback from the developer. Typically, the sources of utilized information reflect the type of analysis that can be employed. Dynamic analysis uses execution traces captured when a feature is executed. Different representations of source code, such as a program dependence graph, can be used by static analysis FLTs. Source code and documentation can be leveraged in textual analysis to find words that are relevant to a feature. Analysis applied on change history data uses version control systems, issue trackers, communication archives, etc., in order to leverage past changes in software, as a way of supporting feature location.

3.4. Output

Once a FLT identifies candidate software artifacts for a given feature, those results must be presented to the developer. The results can have different granularity levels and different presentations.

The types of granularities are *classes/files*, *methods or functions*, and *statements* (i.e., basic blocks, lines of code, variables, etc.). Throughout this survey, we refer to portions of source code at any level of granularity as **program elements**. The more fine-grained the program elements located by a technique, the more specific and expressive the FLT is. For instance, when the results are presented at the statement-level granularity, all the basic blocks or variables may be relevant, but when results are presented at class-level granularity, not all the methods from the class may pertain to the feature. Some approaches may be applicable to multiple levels of granularity, but only those program elements that are actually shown to be supported in an article are reported in this survey. On the other hand, some FLTs may produce artifacts that are not part of the source code, such as bugs, documentation, etc. We categorize these results in the *non-source code artifacts* attribute.

In this survey, we also distinguish among different ways the results can be presented to software developers. For example, one option is to present a list of candidate program elements ranked by

their relevance to the feature [9, 10, 13, 22, 23]. We use the attribute *ranked* to denote such a presentation of the results. Another way in which feature location results are presented is as an unordered set of program elements [7, 8, 24]. In other words, a set of elements is identified as being relevant to a feature, but no notion of their degree of relevance is provided. If the *ranked* attribute is not specified, we assume that a FLT presents the results in this way. Another form of presenting the results is by using *visualization* to highlight relevant program elements [25–27]. Note that the ranked and visualization attributes are not mutually exclusive. Finally, some FLTs do not automatically identify relevant program elements, but describe a process that a programmer can follow to manually search for a feature's implementation [12]. Since different FLTs present their results in different ways, comparing approaches that use dissimilar reporting styles can be challenging.

3.5. Programming language support

The programming language in which a software system is written can be a factor in the types of FLTs that can be applied to it. Textual and historical analyses are programming language agnostic at the file level, but require parsers if applied at fine granularity levels (e.g., method-level granularity). Static and dynamic analyses may be limited due to tool support for a given platform or a programming language. In this survey, all programming languages on which a technique has been applied are reported. The majority of existing feature location approaches have been exercised on Java or C/C++ systems since ample tool support is available for these languages. Other programming languages that have been supported include FORTRAN and COBOL. Knowing the languages under which an approach works can help researchers and practitioners select an appropriate technique, though the fact that an approach has not been used on a certain programming language does not imply that it is not applicable to systems implemented in that language.

3.6. Evaluation

The way in which a FLT is evaluated provides researchers and practitioners with useful information on the approach's quality, effectiveness, robustness, and practical applicability. Evaluating a FLT is difficult because defining the program elements that are relevant to a feature may be subjective at times. Despite this difficulty, researchers have devised a number of approaches to assess FLTs.

Evaluations of traditional software engineering techniques are classified as *survey*, *experiment*, *case study* and *none* [142]. However, this standard categorization does not apply to our taxonomy for two reasons. First, due to the fact that the feature location field is not as matured as other software engineering fields, there are no papers that fall into the categories survey and experiment. In addition, most of the papers evaluate their approaches using case studies or they only have a basic (i.e., preliminary) evaluation, which means that according to the standard categories, all the papers would be categorized as either case studies or none. Our second reason is that we wanted to categorize the papers in our survey using finer grained attributes, which would give more insights into the evaluation used in the papers. Thus, we choose the attributes *preliminary*, *benchmark*, *human subjects*, *quantitative*, *qualitative*, *comparison with other approaches* and *unknown/none*, which are described next. Once again, by using these attributes we wanted to distinguish as much as possible between different evaluations in different papers.

The most simplistic evaluations are *preliminary* in nature and involve small systems or a few data points, and their purpose is to provide an anecdotal evidence that the approach works (i.e., proof of concept).

More advanced evaluations use *benchmarks* that contain detailed information and datasets that could be used by other researchers in their evaluation. These datasets could contain a list of features, textual description, or documentation about the features; mappings between features or bugs and program elements that are relevant to fixing the bug or implementing the feature (referred to as *gold sets* in the literature); patches submitted to an issue tracker, etc. Benchmarks carry more weight than anecdotal evaluation, because they can be used by other researchers in other approaches and the results could be compared in a fair way. In other words, if two approaches are applied on the same dataset, their results could be easily compared to establish which technique produces more accurate results, for instance. However, if the two techniques are evaluated on different datasets (even if these datasets are part of the same software system), the comparison of the results is biased by the

difference in those datasets. In this survey, we categorize papers as using benchmarks if their evaluation is using the same datasets or a subset of these datasets, which were made available by other authors. We also categorize the evaluation of papers that initially used that dataset and made it available as a benchmark.

One of the difficulties in using benchmarks in the evaluations is that very few of them are made available up to date. One of the reasons is that constructing benchmarks requires substantial efforts, and the results are not guaranteed to be 100% correct and complete. For example, if benchmarks are constructed from source code repositories or from patches extracted from issue repositories, there is no assurance that they contain all the information required to implement the feature or fix the bug (unless this information is verified by the original designers and developers of the software system). They might contain noise, or they may only pertain to a small portion of the feature and not touch all of its program elements. Another way to evaluate a feature location approach is to have system experts or even non-experts assess the results, which is an evaluation method often used by IR-based search engines. When multiple experts or non-experts are used, the intersection of their subjective evaluations can be used to create a benchmark. However, the agreement among programmers as to what program elements are relevant to a feature has been shown to be low in some cases [28].

In this survey, we also distinguish between evaluations that use human subjects. These developers could have an *academic* background, such as undergraduate or graduate students, or they could be *professional* developers that come from industry. Other ways to differentiate the evaluations is based on their *qualitative* and/or *quantitative* results. In a quantitative evaluation, the technique is evaluated in terms of some metrics, such as precision, recall, effectiveness, etc., whereas a qualitative evaluation provides details about some particular aspect of the technique or the results. Papers which do not contain any evaluation or for which the details are not known are marked as *unknown/none*. It is important to note that these evaluation attributes are not mutually exclusive. In other words, one approach could have both a qualitative and a quantitative evaluation, whereas others could have only one of them. Or for example, one evaluation could be preliminary, but it could contain some qualitative information. One of the most important attributes of this dimension is *comparison with other approaches*. Ideally, when new FLTs are introduced, they should be directly compared with existing approaches in order to demonstrate their (expected) superior performance. Articles that include comparisons of FLTs are very useful to researchers and practitioners because they highlight the advantages and limitations of the compared approaches in certain settings. FLTs that appear frequently in comparisons are Abstract System Dependence Graphs (ASDG) [12], Dynamic Feature Traces (DFT) [10], Formal Concept Analysis-based feature location (FCA) [7], LSI-based feature location [23, 29], Probabilistic Ranking of Methods based on Execution Scenarios and Information Retrieval (PROMESIR) [30], software reconnaissance [8], and Scenario-based Probabilistic Ranking (SPR) [9]. UNIX `grep` is also another popular point of comparison because programmers often use it to search for relevant code.

3.7. Software systems used for evaluation

A wide variety of software systems have been studied in feature location research, and the size and type of systems used in a case study reflect, to a degree, the applicability of a technique. By reviewing the software systems that have previously been used for feature location, some patterns emerge. Some of the more popular systems are web browsers like Mozilla,⁴ Firefox,⁵ Mosaic, and Chimera.⁶ Other systems that have been investigated frequently are Eclipse,⁷ jEdit,⁸ and JHotDraw.⁹ For some of these systems, there are a few datasets that are repeatedly used, but the majority of papers evaluate their techniques on datasets extracted from these systems, and which are never used in other

⁴<http://www.mozilla.org/> (accessed and verified on 03/01/2011)

⁵<http://www.mozilla.org/firefox> (accessed and verified on 03/01/2011)

⁶<http://www.chimera.org/> (accessed and verified on 03/01/2011)

⁷<http://www.eclipse.org/> (accessed and verified on 03/01/2011)

⁸<http://www.jedit.org/> (accessed and verified on 03/01/2011)

⁹<http://www.jhotdraw.org/> (accessed and verified on 03/01/2011)

evaluations. In other words, there is no de facto benchmark proposed, and we argue that there should be one. Besides these popular systems, an abundance of other software systems have been studied. The systems on which a FLT has been applied are listed in the taxonomy. Having a comprehensive list of the software systems studied for feature location allows researchers to identify good candidates for systems to use in their own evaluations, or even better, to build benchmarks for these systems and make them publicly available. In addition, it allows practitioners to recognize approaches that may be successfully applied to their own software if the program they wish to apply a FLT to is similar to a system on which the approach has already been used.

3.8. Other attributes not included in the taxonomy

While classifying the papers based on this taxonomy, we initially considered including the dimensions *tool availability* and *reproducibility*, each with their two binary attributes *yes* and *no*. For example, if the paper introduces a new tool that is made publicly available for others to use, we would mark that attribute accordingly. Similarly, if the evaluation provided enough details about the design of the study so that other researchers could reproduce it, we would mark the evaluation of the paper as reproducible. However, it turned out that actual categorization was quite subjective, and there was a considerable amount of disagreement between the authors of the survey, that we decided to exclude these dimensions. In order to categorize the papers based on these dimensions, a panel of experts should evaluate these artifacts (e.g., steps toward this model have been already done for accepted research papers at ESEC/FSE 2011¹⁰ by the Artifact Evaluation Committee).

4. SURVEY OF FEATURE LOCATION TECHNIQUES

This section describes in detail the systematic process used to survey the literature as well as summarizes the 89 articles reviewed for this survey.

4.1. Performing systematic review process

In this section, we provide details on how we performed the systematic review process to institute our survey. We start this process with an initial *search* phase, which includes the identification of highly relevant venues. More specifically, we selected journal, conference, and workshop papers where research on feature location is within their respective scopes. Table II lists the abbreviations and names of the venues.

After performing our preliminary search for relevant venues, in the *article selection* phase, we manually identified literature which meets our inclusion criteria while filtering articles using our exclusion criteria (see Section 2.2). Titles, abstracts, keywords, and entire articles were cautiously inspected to determine whether or not they should be included in the survey. Two of the authors of this paper were responsible for identifying suitable articles using the selection criteria; however, the selections were later confirmed during the article classification phase by all the authors.

Figure 2 shows the distribution of articles across the venues. The height of the bars represents the number of feature location articles published. Venues at which only one surveyed paper was published are grouped together in the 'Other' bar. Filled bars represent journals, and gray bars denote conferences and workshops.

Using the taxonomy presented in the previous section, the selected articles were classified. During the *article classification* phase, all four authors separately classified each of the selected research papers. For each dimension, the authors identified whether or not the attributes of the taxonomy were applicable to the research papers. Attribute selection for a given dimension is not exclusive. That is, multiple attributes may apply to a given dimension while classifying a particular paper. Following the categorization of the papers by all the authors, an agreement was computed. Initially, there were a few disagreements between the authors' classifications; however, following a meeting in which authors discussed their reasoning, we were able to reach a consensus on the classification

¹⁰<http://2011.esec-fse.org/cfp-research-papers> (accessed and verified on 03/08/2011)

Table II. Venues that have published the articles included in this survey.

Type	Acronym	Description
Journal	JSME	Journal on Software Maintenance and Evolution: Research and Practice
	JSS	Journal on Systems and Software
	TOSEM	ACM Transactions on Software Engineering and Methodology
	TSE	IEEE Transactions on Software Engineering
Conference	AOSD	Aspect-Oriented Software Development
	APSEC	Asia Pacific Software Engineering Conference
	ASE	International Conference on Automated Software Engineering
	CSMR	European Conference on Software Maintenance and Reengineering
	ESEC/FSE	European Software Engineering Conference/ACM SIGSOFT Symposium on the Foundations of Software Engineering
	ICSE	International Conference on Software Engineering
	ICSM	International Conference on Software Maintenance
	IWPC/ICPC	International Workshop/Conference on Program Comprehension
	VISSOFT	International Workshop on Visualizing Software for Understanding and Analysis
	WCRE	Working Conference on Reverse Engineering




Figure 2. Distribution of the surveyed articles. Black bars represent journals, and gray bars denote conferences. The values above the bars represent the number of feature location papers published in that venue.

of all the papers. Additionally, the discussion helped revising some details of the taxonomy and description of the attributes by identifying ambiguous descriptions.

The remainder of Section 4 is devoted to summarizing the FLT, using the taxonomy as a means of structuring the discussion. More specifically, the dimension *type of analysis* is used to organize the remaining subsections. When summarizing the FLT and multiple articles (i.e., conference and journal versions) that describe a given approach and share identical classifications, both are cited but the summary primarily pertains to the latest version. The articles are classified by the types of analysis used for feature location, and other dimensions of the taxonomy are mentioned whenever applicable. The type of analysis employed is the most distinguishing characteristic of feature location approaches, so it is a logical choice for decomposing the survey. In the following subsections, the surveyed articles are categorized by their use of one or more types of analysis: dynamic; static; textual; dynamic and static; dynamic and textual; static and textual; dynamic, static, and textual; and other. The discussion for each paper includes a brief characterization of the approach, distinguishing features from other approaches, tool support, and some details of the evaluation. Tables III, IV, and V present the articles and their classification within the dimensions of the taxonomy. Tables III and IV present the approaches, whereas Table V presents the tools and the case studies.

4.2. Dynamic feature location

Dynamic feature location relies on collecting information from a system during runtime. Dynamic analysis has a rich research history in the area of program comprehension [31], and feature location

Table III. Classification of approaches and techniques from Section 4.2 to Section 4.9 within our taxonomy (refer to our online Appendix for references and the list of systems evaluated).

Approach	Input		Derivative						Output					Prog. Lang.			Evaluation									
	Query	Scenario	Prog. Elem.	SC compilable	SC non-comp.	Dep. graph	Exec. trace	Historical inf.	Other	File/class	Method	Statement	Non-SC artifact	Ranked	Visualization	Java	C/C++	Other	Preliminary	Benchmark	Academic	Professional	Quantitative	Qualitative	Comparison	Unknown/none
Dynamic																										
[Wilde'95]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Wilde'92]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Eisenberg'05]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Wong'99]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Eisenbarth'01b]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Eisenbarth'01c]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Safyallah'06]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Sartipi'10]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Edwards'06]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Bohnet'08b]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
Static																										
[Chen'00]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Robillard'02]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Robillard'07a]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Robillard'05a]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Robillard'08]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Saul'07]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Trifu'08]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Trifu'09]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
Textual																										
[Petrenko'08]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Wilson'10]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Marcus'04]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Poshyvanyk'07b]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Cleary'07]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Cleary'09]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Gay'09]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Grant'08]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Shepherd'06]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Hill'09]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Abebe'10]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Würsch'10]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
Dynamic + Static																										
[Eisenbarth'01a]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Eisenbarth'03]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Koschke'05]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Antoniol'05]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Antoniol'06]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Rohatgi'07]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
[Rohatgi'08]	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.</

is one subfield in which it is used. A number of dynamic approaches exist that deal with feature interactions [32–34], feature evolution [35, 36], hidden dependencies among features [37], as well as identifying a canonical set of features for a given software system [38]. These techniques are beyond the scope of this survey, which focuses only on approaches that seek to identify candidate program elements that implement a feature. This subsection summarizes articles that achieve this goal using dynamic analysis.

Software reconnaissance [8, 39] is one of the earliest FLTs, and it relies solely on dynamic information. Two sets of scenarios or test cases are defined, such that some scenarios activate the feature and the other scenarios do not, and then execution traces of all the scenarios are collected.

Table IV. Classification of approaches and techniques from Section 4.2 to Section 4.9 within our taxonomy (continued)

Approach	Input	Derivative						Output					Prog. Lang.	Evaluation											
	Query Scenario	Prog. Elem.	SC compilable	SC non-comp.	Dep. graph	Exec. trace	Historical inf.	Other	File/class	Method	Statement	Non-SC artifact	Ranked	Visualization	Java	C/C++	Other	Preliminary	Benchmark	Academic	Professional	Quantitative	Qualitative	Comparison	Unknown/none
Static + Textual																									
[Zhao'04]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Zhao'06]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Hill'07]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Shao'09]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Ratiu'06]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Ratiu'07]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Shepherd'07]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Hayashi'10b]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Dynamic + Static + Textual																									
[Eaddy'08a]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Other																									
[Chen'01a]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Cubranic'03]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Cubranic'05]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Cubranic'04]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Robillard'03a]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
[Ratanotayanon'10]	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•

For example, in a word processor, if the feature to be located is spell checking, feature-specific scenarios would activate the spell checker and the other scenarios would not. Feature location is then performed by analyzing the two sets of traces and identifying the program elements (methods) that only appear in the traces that invoked the feature, using both Deterministic and Probabilistic Formulations. This idea of comparing traces from scenarios that do and do not invoke a feature has been heavily used and extended by other researchers in the field. The software reconnaissance approach has been implemented in tools such as RECON2, and RECON3.¹¹ This approach was evaluated on small C systems, using a small number of scenarios, and the results show that the software reconnaissance is useful in indicating small portions of code where developers should look at. A study involving professional developers [8] showed evidence that software reconnaissance could be easily adopted by professional developers. In terms of potential limitations of this technique, the results produced are heavily influenced by the quality of the scenario and traces provided. In addition, the approach can be used in finding a subset of the artifacts that correspond to a particular feature, but it cannot guarantee in finding all of the related artifacts for a feature. In other words, it provides a very good starting point. Finally, this technique cannot be applied for finding features that are always present in the program, and for which there are no test cases that can and cannot exercise that feature.

An extension of the software reconnaissance approach is Dynamic Feature Traces (DFT) [10], which works as follows. The developer provides several scenarios or test cases that exercise the features in order to collect execution traces. Next, all pairs of method callers and callees are extracted from the traces, and each method is assigned a rank for the feature. The rank is based on the average of three heuristics: multiplicity, specialization, and depth. Multiplicity is the percentage of a feature's tests that exercise a method compared to the percentage of methods in each non-feature's set of tests. Specialization is the degree to which a method was only executed by a feature and no other. Depth measures how directly a set of tests exhibit a feature compared to the other test sets. This approach was implemented in a prototype tool. This paper also presents a preliminary evaluation on three Java systems where DFT and the software reconnaissance technique were

¹¹<http://www.cs.uwf.edu/~recon/> (accessed and verified on 03/01/2011)

Table V. Classification of tools and case studies.

Approach	Input	Derivative	Output	Prog. Lang.	Evaluation
	Query Scenario Prog. Elem.	SC compatible SC non-comp. Dep. graph Exec. trace Historical inf. Other	File/class Method Statement Non-SC artifact Ranked Visualization	Java C/C++ Other	Preliminary Benchmark Academic Professional Quantitative Qualitative Comparison Unknown/none
Tools for Dynamic FL					
[Lukoit'00]	. . .				
[Egyed'07]	. . .				
[Olszak'10]	. . .				
Tools for Static FL					
[Chen'01b]					
[Buckner'05]					
[Weigand-Warr'08]					
Tools for Textual FL					
[Poshyvanyk'06c]					
[Poshyvanyk'05]					
[Poshyvanyk'06b]					
[Xie'06]					
[Cleary'06]					
Other Tools for FL					
[Robillard'03b]					
[Robillard'05b]					
[Savage'10b]					
[Savage'10a]					
[Bohnet'06b]					
[Bohnet'06a]					
[Bohnet'08a]					
[Bohnet'07b]					
[Bohnet'07a]					
Case Studies					
[Wilde'96]					
[Wilde'01]					
[Wilde'03]					
[Ibrahim'03]					
[Marcus'05c]					
[Revelle'09]					
[Simmons'06]					
[Van Geet'09]					
[Revelle'05]					

compared. The results of this study revealed that software reconnaissance is not as effective as DFT in finding relevant code. Since this technique is an extension of software reconnaissance, its limitations are similar.

Wong *et al.* [40] proposed a technique based on execution slices, which is able to produce results at a finer level of granularity, such as statements, basic blocks, decisions, or variable uses, as opposed to producing results at method-level granularity. This technique requires as input a few test cases that exercise and a few test cases that do not exercise the feature of interest. The dynamic information extracted from the instrumented system, on which the test cases were run, consists of the statements that were executed, as opposed to a list of methods. Using this information, this technique is able to distinguish between code that is unique to a feature and code that is common to several features. The tool that implements this approach is called χ Vue [41], and it supports program instrumentation and collecting execution information. In addition, it provides a visualization of the statements relevant to a feature. The evaluation of this technique based on execution slices was performed on five feature of a small C system called SHARPE [42]. The results of the evaluation indicate that the output produce by the technique can be used by developers as a starting point for analyzing the feature of interest. In addition, based on a qualitative analysis of the results, the authors hypothesize that using static information might enhance the results.

Eisenbarth *et al.* [43, 44] introduced a technique for generating feature component maps utilizing dynamic information. Given execution traces resulting from various usage scenarios (covering a set of features of a system), concept analysis is applied to reveal relationships between features and components in addition to feature to feature relationships. Based on the resulting concept lattice, a feature component map is derived and used to identify components of interest given a particular feature. Additionally, a case study is performed on using Xfig as the subject software system, demonstrating the value of their approach for generating feature component maps.

Safyallah and Sartipi [11] introduced an approach that applies a data mining technique on the execution traces. This technique consists of analyzing, using sequential pattern mining, a set of execution traces, which are collected from a set of feature-specific scenarios. This technique is able to identify continuous fragments of execution traces, called execution patterns, which appear in at least a given number (i.e., *MinSupport*) among all the execution traces. These execution patterns undergo a set of refinements based on adjusting the *MinSupport* threshold. The results of this approach are a set of continuous fragments of execution traces which correspond to a particular feature. A preliminary case study was performed on a medium size C system, called Xfig,¹² and the results show that this approach is able to identify a set of core methods that are specific to the input feature. This technique has the advantage of reducing the complexity of analyzing large traces, by identifying their relevant parts. In addition, the sequence of operations from the trace is not altered, and this approach is even able to locate execution patterns specific for less visible features, such as mouse pointer handling, canvas view updating, and other. An extension of their approach [143] allows the extracted execution patterns to be distributed over a concept lattice, in order to distinguish the common group of functions from the feature-specific group of functions. The advantage of using the concept lattice is that it provides developers the chance of identifying a family of closely related features in the source code. An evaluation of this new approach on two C programs, Xfig and Pine,¹³ showed promising results for identifying feature specific functions, as well as common functions.

A FLT designed specifically for distributed systems is introduced by Edwards *et al.* [45]. This technique aims at reducing the imprecision that tends to be associated with existing dynamic FLTs applied on multithreaded or distributed systems. The imprecision of existing techniques stems from the stochastic nature of distributed systems, and from the inability of the technique to identify correctly the order and time of events. To overcome this problem, Edwards *et al.* proposes a definition of time intervals based on causal relationships among events (messages). For example, events are ordered temporally in a single process, and the events from different processes are causally ordered by message passing. The technique not only requires execution information about the system as input, but it also requires the developer to identify the first and last event associated with a feature. The technique identifies all the events that causally follow or precede a feature's starting and ending events, respectively. This sequence of events is called an interval. The output produced by the technique is a ranked list of program elements which are assigned a component relevance index, which is the proportion of executions of that element which occur during a feature's interval. A preliminary evaluation on a large-scale distributed system indicates that the technique is able to distinguish between feature-related code and not feature-related code based on the component relevance index associated with each program element. The technique is useful for any distributed system for which the program instrumentation does not significantly alter their behavior.

Bohnet *et al.* [46] proposed a technique that allows developers to visualize various characteristics of the execution information in order to gain insight into how features are implemented in the code. The developer provides as input a scenario that triggers a feature of interest in the system in order to collect an execution trace. The execution trace is used as an input for an analysis tool, and various characteristics of the execution trace are presented to the developer using advanced visualization views. These views are synchronized among each other, which mean they allow to simultaneously present the same information extracted from the trace from different perspectives. For example, the portions of the trace that are highlighted by the developer and which have a correspondence in the

¹²<http://www.xfig.org/> (accessed and verified on 03/01/2011)

¹³<http://www.washington.edu/pine/> (accessed and verified on 03/01/2011)

source code are also highlighted. A change in a view updates the information in the other views. The approach was compared against `grep` in a preliminary evaluation involving a large C/C++ software system. In this evaluation, `grep` was shown to be impractical because it returned too many results, whereas using the tool a developer was able to locate the concept of interest in less than half an hour and without having prior knowledge about the system. The advantage of this technique is that the advanced visualization allows the developer to manage the complexity associated with a large execution traces.

One of the main shortcomings of FLT based on dynamic analysis is in the overhead it imposes on a system's execution. In distributed and time-sensitive systems, the use of dynamic analysis can be prohibitive. Edwards *et al.* [47] report on their experiences using dynamic analysis to perform feature location in time-sensitive systems. Instrumenting a software system in order to collect execution traces of the program elements that are invoked affects the system's runtime performance. Edwards *et al.* developed a minimally intrusive instrumentation technique called *minist*, that reduced the number of instrumentation points while keeping test code coverage high. For an initial evaluation, Apache's `httpd`¹⁴ and several large in-house programs were used, and *minist* was compared to uninstrumented executions as well as several other tools for collecting traces. The *minist* approach increased execution time by only 1% on `httpd`, while the other tracing tools caused increases of 7% to over 2000%.

Another systematic survey of feature location approaches utilizing execution information can be found in a comprehensive survey of approaches to program comprehension via dynamic analysis [31].

4.3. Static feature location

In contrast to dynamic feature location, static feature location does not require execution information about a software system. Instead, its source code is statically analyzed and its dependencies and structure are explored manually or automatically. Some static FLTs leverage several types of control and data dependencies. Other static techniques use the structure of a software system's dependencies. In general, the static FLTs require not only a dependence graph, but also a set of software artifacts, which serve as a starting point for the analysis in order to generate program elements relevant to the initial set. The initial set of artifacts is usually specified by the developer.

Chen and Rajlich [12] introduced the concept of Abstract System Dependence Graphs (ASDG), which are based on abstracting the System Dependence Graphs (SDG). The nodes of an ASDG are either functions or global variables, and the edges between nodes represent either control dependencies between functions or data flow between variables. The static concept location technique that uses ASDGs require as input a statically built dependence graph as well as a starting node which is selected by the developer. The starting node can be a program element relevant to the feature, which is a priori known by the developer, or it can be a randomly selected node, or it can be the node corresponding to the main method. This FLT requires at each step feedback from the developer, which investigates a node and decides if it is relevant to the feature or not. The FLT keeps track of the search graph (i.e., visited components and their neighbors), and based on the developer's feedback about the relevant and irrelevant nodes, it updates and expands the search graph and tries to propose only relevant nodes. This process continues until the developer has found all the program elements related to the maintenance task. The tool that supports this static FLT based on ASDGs is called Ripples [48]. Ripples is not only able to generate the ASDG from the source code of C programs, but it also visualizes the graphs and allows the developer to input her feedback about which node is relevant. Another tool implementation of this static FLT is called JRipples [49]. This tool is an Eclipse plug-in, which supports feature location on Java systems. Unlike Ripples, JRipples does not have visualization, but it provides support for impact analysis and change propagation. The static FLT introduced by Chen and Rajlich [12] was evaluated on the Mosaic web browser. The results show that among the 984 functions in Mosaic, the developer performing concept location on a maintenance task was able to partially comprehend the system by investigating only 22 (2%) of the functions.

¹⁴<http://httpd.apache.org/> (accessed and verified on 03/01/2011)

Robillard and Murphy [50,51] developed the Concern Graphs representation, which allows the representation of a concern (or feature) in an abstract way. This abstraction allows creating and storing mappings between features and source code. A Concern Graph encapsulates a subset of program elements and a set of relations between them. These relations are based on the static dependencies between the program elements. The tool that was implemented to support Concern Graphs is called FEAT (Feature Exploration and Analysis Tool). FEAT allows developers to visualize the Concern Graphs, to explore the Concern Graph and examine the source code associated with the program elements of the graph, and to allow developers to alter the Concern Graph by adding or removing program elements and relations. In an evaluation involving developers that were required to locate a concern in a Java system, the Concern Graphs representation was shown to be appropriate for expressing the concerns during maintenance tasks, as well as for manipulation or analysis of the concerns. In addition, Concern Graphs are easy and intuitive to use even by unfamiliar developers, and they can be used even for industrial-sized systems.

Robillard [13, 52] introduced an approach that analyzes the topology of structural dependencies in a program in order to propose relevant program elements for the developer to investigate. One of the differences between this approach and ASDGs is that it requires less interaction from the developer. Robillard's approach takes as input a set I of program elements, which are marked as relevant by the developer. The approach examines the structural dependencies of the elements in I and the rest of the system and produces a suggestion set S . Both the input and output sets are fuzzy sets, which means that a program element is part of that set with only a degree of certainty. In other words, each element from the set has associated a value that signifies its relevance. The relevance values are based on two metrics, namely *specificity* and *reinforcement*. A program element's specificity measure is inversely proportional with the number of program elements related to it. In other words, a program element is more specific than others if it has fewer program elements related to it. A program element's reinforcement measure is directly proportional with the number of elements of interest related to it. A tool that supports tool the management of concerns which are scattered throughout the code is called ConcernMapper¹⁵ [53]. This Eclipse plug-in can be used by programmers to specify the initial set of interest relevant to a task. The tool that actually implements the FLT via topology analysis is called Suade¹⁶ [54]. An evaluation on a few medium-sized Java systems shows that given an initial set of program elements, the approach is able to return a ranked list of program elements, where the top entries are highly relevant to the initial program elements. In other words, the approach is able to suggest program elements worthy of investigation to developers and at the same time avoid suggesting less interesting (irrelevant) ones.

Saul *et al.* [144] introduced FRAN (Finding with RANdom walks), an approach that recommends a set of related program elements (e.g., methods) given a specific starting point as an input. FRAN only uses the structural information about a system (e.g., method call graph) and conceptually FRAN generalizes Robillard's approach [52]. This is because FRAN also takes as an input a program element e that has some interest to the developer, and builds a program dependence graph of the neighborhood of e . The main difference is that FRAN uses a larger set of related program elements and ranks them using the scores produced by applying a random walk algorithm [89] to the program dependence graph (as opposed to ranking the elements using the specificity and reinforcement metrics). An evaluation on the Apache HTTPD system was performed and FRAN was compared against Suade and FRIAR (Frequent Itemset Automated Recommender), an approach inspired from Association Rule Mining which ranks program elements based on the support values. The results showed that FRAN produced better results (in terms of returning relevant methods) than FRIAR and Suade.

While ASDGs and topology analysis use both control and data dependencies to some extent, Trifu [55] introduced an approach to feature location based only on static dataflow analysis. The input for this approach is a set of variables, called information sinks, that are selected by the developer and which are used as starting points to identify all the parts of the source code where the values from the variables propagate. This is done by tracking the dataflow dependencies in the code. Because the

¹⁵<http://www.cs.mcgill.ca/~martin/cm/> (accessed and verified on 03/01/2011)

¹⁶<http://www.cs.mcgill.ca/~swevo/suade/> (accessed and verified on 03/01/2011)

granularity of the input program elements is more fine grained (i.e., variables), the results are also more fine grained than other FLT. The tool that implements this approach is called CoDEx (Concern Discrimination and Explorer), which allows developers to mark the variables of interests and presents the parts of the source code which are related to the input variables. An initial evaluation was performed on JHotDraw. For this evaluation, the information sinks selected were all the variables with no outgoing dataflow paths, and the approach grouped the 6049 variables into 310 concerns (features). A manual inspection of a few of these concerns revealed that program elements grouped under the same concern are highly relevant to that concern (feature). Trifu [56] improved this dataflow-based concern identification approach by introducing the concept of information sources, which define boundaries for a concern.

4.4. Textual feature location

Source code comments and identifiers are embedded within textual information about features of the systems. FLT. based on textual analysis are aiming to establish a mapping between the textual description of a feature given by the developer and the parts of the source code where that feature is implemented. The approaches to establish the mapping between the description of the feature and the source code include textual search with *grep* [57], IR [23, 58–60], and NLP [61, 62].

One simple way of searching for code that is relevant to a task is by using a textual search that describes that task. For example, developers formulate a query that describes the feature they are looking for and then use a tool such as *grep* to find and investigate lines of code that match the query. Petrenko *et al.* [57] developed a FLT based on *grep* and ontology fragments. The ontology fragments store partial domain knowledge about a feature. The hypothesis of this approach is that ontology fragments help developers formulate queries and guide their investigation of the results, which would increase the effectiveness of the feature location. As programmers gain more knowledge of the system, the ontology fragments can be refined and expanded. The tool used to support the management of the ontology fragments is called Protégé.¹⁷ An exploratory study on two large systems, Eclipse and Mozilla, showed that for locating a bug in the code, ontology fragments required, on average, a few source code methods in Mozilla and Eclipse to be inspected. These results were comparable to other FLT. [22, 63] in which programmers also only had to examine about 10 methods.

Wilson [64] extended Petrenko *et al.*'s approach by introducing a systematic approach for formulating queries based on ontology fragments, which represent partial knowledge about the system. Using an ontology fragment, the developer can formulate a query based on the terms that are present in the ontology fragment, and it can provide that query as an input for *grep*. A preliminary evaluation involving four developers that were required to perform concept location on the Mozilla and Eclipse systems reveal the fact that only a small and partial knowledge about the system is sufficient for successfully locating a concept in the code.

Information Retrieval is a more advanced technique that can be used instead of the traditional *grep* pattern matching. Marcus *et al.* [23] used LSI [14] to map the feature descriptions expressed in natural language by developers to source code. LSI is an advanced IR technique that infers relations between words and passages in large bodies of text. A corpus is created from extracting all identifiers and comments from the source code of a system. The corpus is preprocessed by splitting compound identifiers based on common naming conventions. The corpus is partitioned into documents representing all terms associated with a program element. Documents can be of different granularities, such as classes or methods. The corpus is then transformed into an LSI subspace through Singular Value Decomposition (SVD). After SVD, each document in the corpus has a corresponding vector. To search for code relevant to a feature, a programmer formulates a query consisting of terms which describe the feature. The query is also transformed into a vector, and a similarity measure between the query vector and all the document vectors is used to rank documents by their relevance to the query. The similarity measure is known as the cosine similarity because it computes the cosine between the query and document vectors. The output produced by this FLT is a list of methods ranked by their textual similarity with the developer input query. The approach was evaluated on the Mosaic web

¹⁷<http://protege.stanford.edu/> (accessed and verified on 03/01/2011)

browser, using the same feature used in the study by Chen and Rajlich [48]. The new approach was compared against grep and ASDGs, and several advantages were found. LSI is as easy to use as grep, yet it produces better results. Also, LSI was able to identify some relevant program elements missed by ASDGs. Recently, LDA has been applied for bug localization [29, 65]. The proposed approach was compared with LSI and has been shown to be an effective alternative to using LSI for concept location.

Poshyvanyk and Marcus [60] added Formal Concept Analysis (FCA) to the feature location that uses LSI. FCA takes as input a matrix specifying objects and their associated attributes and then produces clusters, called concepts, of the objects based on their shared attributes. These concepts can be organized hierarchically in a concept lattice. In this case, the objects are methods and the attributes are words that appear in the source code of those methods. To combine the two types of analysis, LSI's ranked results are clustered using FCA. The top k attributes of the first n methods ranked by LSI are used to construct FCA's input matrix and create a lattice. Nodes in the lattice have associated attributes (terms) and objects (methods), and programmers can focus on the nodes with attributes similar to their query to find feature-relevant methods. The new FLT based on FCA was compared against the FLT that uses LSI alone on two maintenance tasks of Eclipse, and the results show that the new approach is able to group relevant information using concept lattices, which means that developers can locate a concept in code by analyzing fewer methods, as opposed to the case where the results are presented as a ranked list.

Cleary and Exton [58, 66] also use IR for feature location, but their solution incorporates non-source code artifacts, such as bug reports, mailing lists, external documentation, etc. Their approach, called cognitive assignment, considers indirect correspondences between query and document terms so that relevant source code can be retrieved even if it does not contain any of the query terms. Queries are expanded by analyzing term relationships from both source code and non-source code artifacts. This approach was implemented by extending the cognitive assignment Eclipse plug-in [67] to incorporate the expansion queries mechanism. A case study was conducted on Eclipse in which cognitive assignment was compared to other IR techniques, such as language modeling [68], dependency language model [69], vector space model [16], and LSI. The results show that cognitive assignment matches the performance of the other IR techniques, and in some cases, it outperforms them.

The results of any textual FLT are heavily influenced by the quality of the queries used. In other words, a textual FLT could produce more accurate results if the developer formulates more accurate queries, by refining or modifying existing ones. Gay *et al.* [59] introduced the notion of relevance feedback into textual feature location with IR. Relevance feedback incorporates user input to improve IR results. After IR returns a ranked list of program elements relevant to a query, the developer rates the top n results as relevant or irrelevant. Then, a new query that incorporates the developer feedback is automatically formulated and new results are returned, and the process repeats. A case study was performed in which a single developer was asked to use IR and relevance feedback to locate the source code associated with change requests (representing features) in Eclipse, jEdit, and Adempiere.¹⁸ Each change request had associated with it a patch that was used to implement the change request, and these patches were used to generate the gold set of methods used in the evaluation. The results indicate that relevance feedback is more effective and efficient than a pure IR-based approach.

Similarly to IR, Independent Component Analysis (ICA) [70] can examine source code text to identify features and their implementations [71]. ICA is a signal analysis technique that separates a set of input signals into statistically independent components. To apply ICA for feature location, a term by document matrix is constructed in which the rows correspond to methods, columns represent terms, and cells contain the frequency of a term in a method. ICA factors the matrix into two new matrices. The first new matrix, called the source signal matrix, stores independent signals which can be thought of as features. The second new matrix, the mixing matrix, holds information about how relevant each signal is to a method. Unlike feature location with LSI, feature location with ICA does not need a query for a specific feature because it seeks to identify multiple independent signals (features) at once. A preliminary evaluation on a medium size C program showed that ICA was able to identify a few key concepts from the code.

¹⁸<http://sourceforge.net/projects/adempiere/> (accessed and verified on 03/01/2011)

Textual FLTs are not limited to using IR only. Shepherd *et al.* [72] proposed a technique that leverages information about the use of verbs and their direct objects (nouns) in source code identifiers to create a natural language representation of the code called an Action-Oriented Identifier Graph (AOIG). In the AOIG, all the verb-direct object (verb-DO) pairs are extracted from the code, and a mapping between each verb-DO and the code is kept. An Eclipse plug-in called ViRMoVis was created to implement this approach. The developer formulates a query in the form of a verb, and ViRMoVis suggests a set of direct objects associated with that verb. The developer then selects the appropriate direct objects (i.e., refines the query), and the tool displays all the uses of the selected verb-DO pairs. An exploratory study on the Java JHotDraw system revealed that using ViRMoVis a developer was able to locate a feature by examining only a small number of classes and methods.

Hill *et al.* [61] also used NLP and the idea of query expansion and refinement in their approach to feature location based on contextual searching. Instead of focusing on verbs and direct objects, their analysis centers on three types of phrases: noun phrases, verb phrases, and prepositional phrases. The phrases are extracted from method and field names, and additional phrases are generated by also looking at a method's parameters. Once the phrases are extracted, they are grouped into a hierarchy based on partial phrase matching. The phrases are linked to the source code from which they were extracted. An Eclipse plug-in and a PHP script were created to implement this approach. A user looking for a particular feature formulates a query, and the tool searches the extracted phrases for matches. The result returned to the user is a hierarchy of phrases and the method signatures associated with them, giving some context to the results. This approach was evaluated on Rhino, on a subset of features from the benchmark created by Eaddy *et al.* [73], as well as on a subset of features used in the evaluation by Shepherd *et al.* [62]. For the evaluation, 22 developers (17 of them with 1 to 9 years of industry experience) assessed the two approaches, and the results show that contextual search has been shown to significantly outperform the verb-direct object approach both in terms of effort (number of queries needed) and effectiveness (*F*-measure).

Abebe and Tonella [74] introduced an approach that extract concepts from source code by applying NLP techniques. The identifiers from the program elements are extracted, and candidate sentences that use those identifiers are formed. Some of the sentences that do not follow certain rules are eliminated, and the remaining sentences are used as an input for creating ontologies that capture the concepts and the relations from the source code. The identifiers represent the concepts, and the semantics of the sentences establish the relations in the ontology. A preliminary evaluation on the WinMerge system revealed that enhancing the queries using information from the ontologies increases the precision of concept location by allowing developers to formulate more precise queries and by reducing the search space. The approach of Abebe and Tonella is similar to Petrenko *et al.*'s [57] approach, but the main difference is that the former approach automatically generates the ontologies, whereas for the latter approach the ontologies were generated manually by developers.

Wursch *et al.* [75] leveraged static analysis and semantic web-based technologies to provide users with a natural language guided query interface to answer program comprehension questions. More specifically, Resource Description Framework (RDF) is used to model source code entities and their properties/relationships resulting in RDF graphs. The Web Ontology Language (OWL) is used to model an ontology for source code. This allows developers to query source code with some natural language guided vocabulary within their IDE. Developers can form questions such as 'What method calls ...?' and 'What attributes have the type ...?' The authors conducted a case study with JFreeChart to demonstrate the usefulness of their technique. With that, the authors provided some first evidence that their technique is capable of answering questions in natural language form.

4.5. Combined dynamic and static feature location

The combination of dynamic and static analysis is a well-known and powerful combination in other areas of research such as testing and program analysis [76, 77]. This combination has also been applied for feature location. Dynamic analysis can be used to reduce the search space to only those program elements that were executed in a trace, and then static analysis can work on the smaller set of program elements to rank them or find additional relevant elements.

Eisenbarth *et al.* [7, 78] proposed an approach that clusters the execution information collected about the system using concept analysis. Several execution traces that exercise different features are collected from the instrumented system. The information from the traces is extracted and used as an input for concept analysis, which treats the methods as objects and the features invoked during the execution scenario as attributes. The result of concept analysis is a concept lattice which can be investigated by the developer in order to identify candidate program elements that are solely relevant to a feature or contribute to a feature but are also used by other features. The program elements located by this approach are only a subset related to a feature (i.e., a starting point), and developers seeking additional relevant code can follow an approach similar to ASDGs. In a preliminary evaluation on two C browsers, Mosaic and Chimera, the approach was able to recover a partial description of the software architecture responsible to implement the set of features used as input. The results also show that out of the large number of methods from the browsers, very few methods needed to be inspected manually. Koschke and Quante [79] adapted this approach to collect input traces at statement-level granularity, which means that the approach is able to locate features at the level of basic blocks (as opposed to method-level granularity). An evaluation conducted on two compilers, *sdcc* and *cc1*, confirmed the findings of the previous approach, and proved that this FLT could be used in practice. The advantage of combining concept analysis with dynamic information is that the concept lattice will handle the differences among the traces that arise during the invocation process.

To overcome the imprecision and noise associated with collecting dynamic data, Antoniol and Guéhéneuc [9, 80] introduced the Scenario-based Probabilistic Ranking (SPR) approach, which combines both dynamic and static data for identifying a feature's relevant program elements (methods). The idea behind SPR is to assign for each event from an execution trace a probability of that event being associated with a feature and then rank all the events. In SPR, similar to software reconnaissance, two sets of scenarios are defined, scenarios that do and do not exercise a feature, and method-level execution traces are collected for each scenario. Intervals correspond to a subsequence of contiguous events (method calls) from the traces, where I is an interval from a relevant scenario, and I' is an interval from an irrelevant scenario. Events are classified as relevant to a feature or not by determining if their frequency in interval I is greater than their frequency in interval I' . For any interval, an event's frequency is computed as the ratio of the number of times the event appears in an interval over the total number of events in the interval. Essentially, determining whether an event is relevant to a feature or not is a statistical hypothesis test. The null hypothesis is that an event's frequency in the two types of intervals is the same. A threshold, Θ , is chosen, and if an event is classified as relevant to a feature more than Θ times, the null hypothesis is rejected with a confidence level α . Events are also ranked by their relevance to a feature using a relevance index score that is computed from the number of times an event appears in relevant intervals versus the number of times it appears in irrelevant intervals. The source code is represented as an Abstract Object Language (AOL) using static analysis. This format that represents the program's architecture is used for highlighting the ranks of the elements identified using the dynamic analysis. The elements that appear in the trace that exercised or not the feature of interest will highlight the program architecture differently, which means it will be easy to compare these architecture to find out which program elements are part of the feature. SPR has been applied to a number of systems including Mozilla, Firefox, Chimera, ICEBrowser, JHotDraw, and Xfig. Case studies have compared SPR directly to feature location using *grep*, and the concept analysis-based approach [7, 78]. The results show that because SPR ranks its results, it is successful at reducing the amount of data that a programmer needs to investigate. In addition, SPR allows developers to visualize the micro-architecture of the system.

Rohatgi *et al.* [81–83] introduced an approach that locates features in source code at class-level granularity. More specifically, their technique takes as input an execution trace and a class or component dependency graph (CDG) for feature location based on impact analysis. Distinct classes are extracted from a feature-specific execution trace, and then the CDG is used to rank the classes by the impact a change to them would have on the software system. The hypothesis of this technique is that classes with the least amount of impact are most likely related to the feature. In a preliminary evaluation on two Java systems, Weka,¹⁹ a machine learning tool, and Checkstyle,²⁰ a

¹⁹<http://www.cs.waikato.ac.nz/ml/weka/> (accessed and verified on 03/01/2011)

²⁰<http://checkstyle.sourceforge.net/> (accessed and verified on 03/01/2011)

tool for formatting source code, the approach was able to identify and rank appropriately classes relevant to the feature. However, in few cases, the relevant classes were not ranked correctly.

Walkinshaw *et al.* [26] developed a FLT based on call graph slicing, which uses the concepts of landmarks and barriers. The first step of the approach is to identify landmark and barrier methods in a static call graph, where a landmark is a method that contributes to a feature and barriers are irrelevant methods. Direct paths between landmark nodes, known as hammock graphs, are found, and additional dependencies are obtained via backward slicing. Barriers and their dependencies are removed from the call graph to prevent exploration of irrelevant methods. The output of this approach is a pruned call graph. The technique was evaluated on NanoXML,²¹ Freemind,²² and JHotDraw, finding that the landmark and barrier technique substantially reduces the size of the call graph that a programmer has to investigate.

4.6. Combined dynamic and textual feature location

Dynamic and textual analyses are very synergistic when it comes to their use in feature location. Dynamic analysis generally yields good recall, whereas textual analysis has good precision. Their combination may lead to improved results over individual techniques. Both analyses can be used to rank program elements by their relevance to a feature; so a logical next step is to combine both of the rankings produced by these techniques. Another combination of dynamic and textual analyses is to use dynamic analysis to filter the program elements for textual analysis instead of ranking all the program elements in a software system.

Poshyvanyk *et al.* [30, 63] introduced the Probabilistic Ranking of Methods based on Execution Scenarios and Information Retrieval (PROMESIR) approach, which performs feature location by combining 'expert' opinions from two existing FLTs. The first is Scenario-based Probabilistic Ranking (SPR) [9, 80] and the second is IR with LSI [23]. Both approaches rank program elements according to their relevance to the feature of interest. Those rankings are combined through an affine transformation to produce PROMESIR's results. The weight given to SPR or LSI can be varied to reflect the amount of confidence that should be assigned to each of the experts. An evaluation on Eclipse and Mozilla indicates that PROMESIR outperforms the two techniques on which it is based. It is interesting to note that one of the datasets used in the evaluation was exactly the same as the one used by Antoniol *et al.* in their evaluation of SPR.

Similar to PROMESIR, the Single Trace and Information Retrieval (SITIR) [22] approach is a FLT that applies IR on the execution information collected from exercising a single scenario relevant to the feature. In other words, the technique takes as an input a developer query and a scenario relevant to the feature and produces a list of methods which are ranked based on the similarity with the query. The novelty of this approach is that it ranks only the methods that appear in the execution trace, as opposed to ranking all the methods from the system. This innovation reduces dramatically the search space and yields better results. An Eclipse plug-in that supports this approach, FLAT³ [84], was later developed. An evaluation on jEdit and Eclipse compared SITIR against the IR-based FLT, SPR, and PROMESIR. The results showed that in general, SITIR ranked the relevant methods higher than the other approaches.

Asadi *et al.* [85] proposed a FLT that identifies cohesive and decoupled fragments from execution traces which are related to concepts. The approach takes as input a scenario that exercises a feature of interest in order to collect the execution trace. The trace is preprocessed in order to remove irrelevant methods to the feature (e.g., mouse tracking methods) and to compress the trace (i.e., remove repetition of methods) for an easier analysis. The preprocessed trace is used as an input for a genetic algorithm that separates the trace into fragments that contain methods that are highly cohesive and which are highly decoupled with other fragments. The fitness function used in the genetic algorithm is based on the conceptual cohesion metric defined by Marcus and Poshyvanyk [86]. A preliminary evaluation on two Java systems, ArgoUML and JHotDraw, showed that the approach was able to locate concepts with high precision. One of the drawbacks of this technique observed in the evaluation was that precision tended to drop for different features that used similar sequences of methods.

²¹<http://devkix.com/nanoxml.php> (accessed and verified on 03/01/2011)

²²<http://freemind.sourceforge.net/> (accessed and verified on 03/01/2011)

Revelle *et al.* [87] proposed a FLT that combines textual information with the results produced by applying advanced link analysis algorithms on execution information. The approach takes as an input a query and an execution scenario that exercises a feature. A program dependence graph is generated from the execution trace by using the caller and callees methods as nodes and the relations between them as edges. Two link analysis algorithms (i.e., PageRank [88] and HITS [89]) are applied on the program dependence graph, which associate for each node in the graph (i.e., method from trace) a score, based on the nodes relative importance in the graph. Any of the PageRank scores, the HITS Authorities, or the HITS Hubs scores could be used to rank the methods from the execution trace. We refer to this novel FLT as WM. An alternative is to filter the methods that are ranked on top or bottom using the link analysis scores from the methods that appear in the execution trace, and after that to rank the remaining methods using the textual similarity between the methods and the developer's query. We refer to this novel FLT as $IR_{LSI}WM$. The evaluation of this approach was performed on two Java systems, namely Eclipse and Rhino. The Rhino data used in the evaluation represent a subset of features from the benchmark created by Eaddy *et al.* [73]. In the evaluation, the WM and $IR_{LSI}WM$ approaches were compared against each other as well as against the LSI-based FLT [23] and the SITIR [22] approach. The results showed that WM produced more accurate results than the LSI approach, but did not perform as good as the SITIR approach. On the other hand, the $IR_{LSI}WM$ outperformed all the other approaches.

Hayashi *et al.* [90] proposed iFL, an approach that combines static and dynamic analysis, along with relevance feedback to identify source code entities which comprise a feature of interest. The paper claims that the iterative approach leads to improved query formulation by end users and the evaluation of relevance during the iterative process enhances a user's understanding of features implemented in a given software system. iFL requires as input source code, a test case (used to derive dynamic dependencies), a query, and hints (relevance feedback) and returns to the user source code entities ordered by their respective evaluation scores. Evaluation of the tool is performed using five change requirements of Sched and two change requirements of JDraw systems. The results, which compare the interactive and non-interactive versions of the approach, indicate that the iterative technique is capable of reducing the understanding cost (based on improvement of lowest ranked relevant method).

4.7. Combined static and textual feature location

Several researchers have combined static and textual analyses for feature location. This combination is a natural choice because either textual analysis can be used to reduce the overestimation that static analysis is prone to produce or static analysis can be used to find additional candidate program elements given a starting set of highly relevant ones from textual analysis. Thus, combining these two types of analysis has the potential to yield better results than either static or textual analysis alone.

A static, non-interactive approach to feature location (SNIAFL) is introduced by Zhao *et al.* [91, 92]. SNIAFL uses IR in conjunction with a branch-reserving call graph (BRCG), essentially an expanded version of a call graph with branch information. An initial set of program elements (i.e., methods) specific to the feature is located using IR, and then additional relevant elements are found using the BRCG. The initial set is produced by using the vector space model to obtain and rank methods by their similarity to a query. A gap threshold technique is used to find the largest difference between the similarities of consecutive ranked methods. The methods above this gap are considered to be the initial elements specific to the feature. From the initial set, the BRCG is pruned to remove branches that are not in the initial set. Also, the relevance of branches that are included in the initial set is propagated through the graph's dependencies, essentially generating a static pseudo-execution trace. In case studies on two pieces of GNU software, SNIAFL had better precision and recall than both a pure IR approach and a purely dynamic approach, lending evidence to the fact that combining static and textual analyses is more successful than using them as standalone approaches.

Similar to the SNIAFL approach, Dora [93] combines static and textual analysis to perform feature location. Programmers formulate a query, which is used to compute a method relevance score that is based on the term frequency-inverse document frequency of words that appear in the name and body of the method. Then, starting from a seed method selected by the developer, Dora follows static caller/callee edges to identify additional relevant methods using the relevance score. Dora was

evaluated on a number of open source Java systems and compared to Suade and two naïve textual and static approaches. The datasets for these systems were created via a user study [28] in which programmers were asked to locate the implementations of several features. Dora was found to be the most successful technique in the evaluation.

In Dora and SNIAFL, one type of analysis is used to prune another. Shao and Smith [94] combine IR and static control flow information in a different manner for feature location. First, LSI is used to rank all the methods in a software system by their relevance to a query. Then, for each method in the ranked list, a call graph is constructed. A method's call graph is inspected to assign it a call graph score. The call graph score counts the number of a method's direct neighbors that also appear in LSI's ranked list. Finally, the method's cosine similarity from LSI and its call graph score are combined using an affine transformation, and a new ranked list is produced. A preliminary evaluation compared the technique against LSI on a C++ program called iVistaDesktop, which simulates Microsoft's Windows Vista operating system.

Ratiu and Deissenboeck [95,96] introduced an approach that recovers the mapping between the realworld concepts and the relevant parts of the source code. Their approach is not explicitly aimed at feature location but at interpreting programs from the point of view of the domain knowledge they implement (which could be features). They developed a framework that describes semantic defects caused by improper naming and an algorithm to recover the mappings between ontology elements and program elements. The algorithm maps concepts and program elements via graph matching. Concepts are represented in the ontology and programs are abstracted as graphs. The framework and algorithm have been applied to the Java standard library, finding actual examples of semantic defects.

Shepherd *et al.* [62] employed NLP in conjunction with static analysis for feature location. The premise of their approach involves the observation that in source code, actions are represented by verbs, and nouns correspond to objects. Their approach, which is implemented in the tool Find-Concept, has the following steps: initial query formulation, query expansion, and a search of the action-oriented identifier graph (AOIG) model. The developer creates a query consisting of a verb and a direct object. Then, Find-Concept expands the query using NLP and its knowledge of the terms used within the software's source code to recommend new queries. Once the user refines the query, the tool locates nodes in the AOIG that contain a verb and direct object from the query and returns the methods to which they are mapped. Find-Concept uses program analysis to identify any dependencies between the methods returned by the AOIG search and then presents the user with a visualization of the results as a graph. In a user study, Find-Concept's verb-direct object approach was compared to the lexical search provided by the Eclipse IDE and Google Eclipse Search [110] on a suite of open-source Java systems, and overall was found to be the most effective search technique, without requiring additional effort from the users.

Hayashi *et al.* [97] combined textual analysis and static analysis with domain ontologies to link user-specified sentences to source code fragments. Their approach returned ordered functional call-graphs (i.e., methods and their invocations of the input source code) extracted for the initial call graph of the software system under investigation. Given the source code of a software system, a sentence, and an ontology (i.e., directed relations between words), the proposed technique obtains a call graph of the system using static analysis, extracts the terms from the source code and the sentence, and subsequently traverses the call graph in search of functional call graphs. Identification of function call graphs involves using terms of the sentence to identify root nodes. Following the identification of root nodes, paths to traverse are determined by locating entities which contain terms from the user-provided sentence. The resulting functional call graphs are prioritized according to the importance determined by their name, using relations in the ontology, and the ratio of words in the input sentence. Evaluation of the technique was performed using seven features of JDraw. A comparison of precision and recall indicated that the use of an ontology provided better results than the case where the ontology was not used.

4.8. Combined dynamic, static, and textual feature location

Cerberus [24] is a FLT that utilizes three types of analysis: dynamic, static, and textual. Currently, it is the only approach that leverages all three types of analysis. At the core of Cerberus is a technique

called *prune dependency analysis* (PDA), whereby a relationship between a program element and a feature exists if the program element should be removed or modified if the feature were to be pruned from the software system. Given an initial set of relevant elements to be pruned, PDA infers additional relevant elements. Cerberus uses PROMESIR to combine rankings of program elements from execution traces with rankings from IR to produce seeds for PDA. Cerberus' authors created a large benchmark for Rhino,²³ an open source Java implementation of JavaScript, in which the code for over 400 features defined in the system's documentation were manually located. This benchmark was used to evaluate and compare Cerberus to software reconnaissance, SPR, DFT, and LSI, finding that combining the three types of analysis was the most effective approach.

4.9. Other FLTs

CVSSearch [19] is a FLT that uses textual and historical information from CVS repositories. CVS comments generally describe the change made to the lines of code which are being committed, and those comments typically hold true for many future revisions of the software. The tool maps the lines of code that were changed during CVS commits with the CVS comments associated with those commits. This means that if a line of code was changed in multiple commits, it will have associated all the CVS comments from those commits. The tool requires as input a query, and CVSSearch²⁴ returns all lines of code whose associated comments contain at least one of the query words. The textual search is done using `grep`. Each returned line also has a score indicating how well it matches the developer query. In a user study involving 74 students that were required to performed concept location on a few programs from the KDE suite, CVSSearch was compared against `grep`. The results of the evaluation showed that even though the CVS comments are a valuable source of information, the CVSSearch tool that exploits them complements the traditional `grep` technique, but does not replace it.

Hipikat [20, 98] is a feature location tool that also makes use of archival information for feature location, but instead of identifying candidate program elements, Hipikat recommends artifacts from a project's archives such as online documentation, versions, bugs, or communications. Hipikat forms a group memory from a project's history as recorded in source code repositories, issue trackers, communication channels, and web documents [99]. Links between these artifacts are inferred using IR. For example, a source code version can be linked to a bug report if the bug's id is included in a repository commit log message. This history is used to find relevant artifacts in response to a user query. The query consists of an artifact, potentially a program element, for which the user wants recommendations of related artifacts. Hipikat responds with a list of artifacts ranked by their relevance. The tool has been used in two case studies. In the first, Hipikat was validated on AVID,²⁵ and in the second, it was used to aid programmers performing a change task on Eclipse.

Robillard and Murphy [21] propose a unique approach to feature location that automatically analyzes a transcript of a program investigation session in an integrated development environment. The transcript records which program elements were visible to a developer during a maintenance task and how they were accessed: through a code browser, following a cross-reference, recalling an open window or tab, scrolling, or keyword search. For each event in the transcript, all visible program elements are determined. Then, for each visible element, a probability that it is the element the programmer was interested in is assigned to it. The probabilities are based on weights associated with each event type. Next, a correlation metric is calculated among all pairs of program elements. The correlation is based on how closely two elements were accessed in the transcript. Finally, concerns (features) are generated by clustering program elements, and the concerns can be named and saved for later retrieval.

Similarly to CVSSearch, Ratanotayanon *et al.* [100] implemented Kayley, a tool that utilized historical and textual information for the task of feature location. Ratanotayanon *et al.* introduced the concept transitive change-set, which is used to enrich the set of source code entities associated

²³<http://www.mozilla.org/rhino/> (accessed and verified on 03/01/2011)

²⁴<http://cvssearch.sourceforge.net/> (accessed and verified on 03/01/2011)

²⁵<http://people.cs.ubc.ca/~murphy/AVID/> (accessed and verified on 03/01/2011)

with a commit of a version control system. Transitive change-sets are generated by analyzing dependencies and other relations to identify other elements related to items in a given commit. Such a process allows linking of program elements at various levels of abstraction. Given a user query, related commits are identified based on the textual relationship of the user query and comments of the commits, and the associated set of program elements is augmented using the concept of transitivity, which is subsequently returned to the user.

5. FEATURE LOCATION TOOLS AND STUDIES

In addition to the many research articles that introduce FLTs, there are numerous articles describing feature location tools, case studies, industrial studies, and user studies. This section summarizes these tools and studies.

5.1. Tools

Tool support for feature location removes much of the manual burden associated with searching for a feature's program elements. In addition to providing an overview of existing FLTs, this survey also describes tools that can be used for feature location. Some of the techniques summarized in Section 4 have prototype tools that are not available; therefore, they are not listed here. Also, some tools are not directly associated with any particular approach, but they can be used for feature location, feature documentation, or program exploration, and thus they are included here.

5.1.1. Tools for dynamic feature location. TraceGraph [101] is a feature location tool that allows for the visualization of execution traces. As a software system is running, TraceGraph analyzes the execution and visualizes which program elements were invoked during a time interval. The visualization is essentially a matrix in which the rows represent program elements, the columns correspond to time intervals, and the cells indicate if the program elements were called during that time interval or not. Additionally, the first invocation of a program element is highlighted in the visualization. TraceGraph was evaluated on the Mosaic web browser as well as the Joint STARS (Joint Surveillance Target Attack Radar Subsystem), a proprietary system developed by Northrop Grumman for the United States Air Force. The tool's visualization was useful for feature location because it emphasized the first time an element was called, which often corresponded to a feature being triggered. TraceGraph was also applied in an industrial case study on feature location [102] where it was used for trace differencing and identifying code uniquely executed by a feature, and in another study on distributed simulation software [103].

Scenario-based TRAcE Detection and Analysis (STRADA) [32] is a tool developed to help programmers uncover the mappings between features and code during testing, which is based on trace analysis research ([104], [105–107]). Given a set of test cases for a feature, STRADA observes the code that is executed during testing, initially identifying all the executed code as relevant to the feature. However, since not all of the invoked code actually pertains to the feature, STRADA analyzes the traces using logical constraints to exclude irrelevant program elements. The tool visualizes its knowledge of feature-to-code mappings in a matrix. It has been evaluated on ArgoUML,²⁶ GanttProject,²⁷ and a video-on-demand player.²⁸

Olszak and Jørgensen proposed the tool Featureous [108] for locating feature implementation in legacy software. The tool is implemented as a plug-in for the NetBeans IDE and allows developers to specify a feature, a scenario and to collect execution traces that exercise the feature. The execution traces are analyzed and results are presented using advanced visualization views. Furthermore, the developer can navigate through the traceability links established through dynamic analysis. From an experience with JHotDraw, Featureous is reported to be able to support both top-down and bottom-up comprehension strategies using its visualization views.

²⁶<http://argouml.tigris.org/> (accessed and verified on 03/01/2011)

²⁷<http://www.ganttproject.biz/> (accessed and verified on 03/01/2011)

²⁸<http://peace.snu.ac.kr/dhkim/java/MPEG/> (accessed and verified on 03/01/2011)

5.1.2. Tools for static feature location. Ripples [48] is a tool that implements the ASDG approach to feature location. The tool extracts an ASDG from C code and visualizes it for the programmer who can mark relevant nodes. JRipples²⁹ [49] is an Eclipse plug-in and its functionality is similar to Ripples's tool, except that it is used on Java systems and it does not provide the visualization component. Both tools can also be used for impact analysis and change propagation by tracking and monitoring the status of program elements.

Suade [54] is an Eclipse plug-in that performs feature location using static analysis. It implements the topology analysis approach discussed in Section 4.3. Suade has been used in a case study comparing several program exploration tools [109], and it has also been directly compared to Dora [93], another static FLT.

5.1.3. Tools for textual feature location. Google Eclipse Search³⁰ (GES) [110] is an Eclipse plug-in that facilitates efficient source code searching and browsing by integrating Google Desktop Search (GDS)³¹ and Eclipse. GDS is an off-the-shelf component that uses IR. It allows users to search for files on their desktops similar to the way they would search for information on the Internet via natural language queries. By integrating GDS with Eclipse, programmers can search source code in a similar fashion. One advantage of using GDS is it unobtrusively re-indexes the search space when the source code changes. In a preliminary evaluation on the Java system Violet,³² GES was shown to produce accurate results. In addition, compared with Eclipse file search functionality, GES is considerably faster in producing the results.

IRiSS [111] and JIRiSS [112] are both tools for textual feature location. IRiSS implements IR-based feature location as an add-on to MS Visual Studio.NET, while JIRiSS is an Eclipse plug-in. Both tools work like a development environment's built-in search functionality, but instead of only displaying the lines of code that match a query, those lines' corresponding classes and methods are also listed. This allows a programmer to sort the results by different levels of granularity and to visit the classes or methods with the most matches. Also, since IR is used, the results returned from a query can be ranked by their relevance. JIRiSS is an extension to IRiSS that also includes fragment-based searches, software vocabulary extraction, query spell checking, and word suggestions to improve queries.

Xie *et al.* [27] developed a tool that supports textual analysis through visualization by combining IRiSS [111] and sv3D [113]. IRiSS performs feature location via IR and sv3D (source viewer 3D) creates 3D renderings of the results, showing poly-cylinders that represent program elements. The colors of the poly-cylinders correspond to the elements' similarity to the query following a predefined color scheme. The height of the poly-cylinder represents browsing history; so, the taller the cylinder, the more times the program element was visited in the past. The combination of these two tools allows a developer to have a visual representation of the results, as opposed to examining a ranked list of results.

Cleary and Exton implemented an Eclipse plug-in that supports the cognitive assignment technique [67]. The tool allows a developer that is unfamiliar with a system to generate and store a set of links between the problem domain concepts stored in a cognitive map and the relevant parts of the source code. The developer can select a concept (i.e., feature) which wants to investigate, and based on the textual description of the feature, the tool provides a set of results which the developer should investigate, and add them to the concept if they are relevant. The program elements suggested by the tool are the result of an analysis of links between words, using a Bayesian classifier.

5.1.4. Other tools for feature location. In this subsection, we discuss several tools that use multiple types of information in their analysis. Some of the tools use only a particular type of information, and they could be presented in different sections, but because they are highly related to one another, we present them in this section. For example, FEAT uses static analysis, and ConcernMapper uses textual analysis, but we present these tools in this subsection because these tools are highly related.

²⁹<http://jripples.sourceforge.net/> (accessed and verified on 03/01/2011)

³⁰<http://ges.sourceforge.net/> (accessed and verified on 03/01/2011)

³¹<http://desktop.google.com/> (accessed and verified on 03/01/2011)

³²<http://www.horstmann.com/violet/> (accessed and verified on 03/01/2011)

The motivation behind the following tools is that once a feature's relevant program elements have been found using feature location, they should be saved so that they could be easily accessed in the future. Features and their relevant program elements can be documented in Concern Graphs [50, 51], a model that describes which program elements pertain to a feature. An Eclipse plug-in that supports the Concern Graphs approach is FEAT³³ (Feature Exploration and Analysis Tool) [114]. ConcernMapper [53] is also an Eclipse plug-in that supports Concern Graphs, and it evolved from FEAT. One of the improvements of ConcernMapper over FEAT is that it supports a fuzzy and less rigid model for representing concerns, which is more appropriate for programmers. ConcernTagger³⁴ is another Eclipse plug-in built on top of ConcernMapper, which provides the ability to compute a number of concern-specific metrics. The Feature Location and Textual Tracing Tool³⁵ (FLAT³) [84] also extends ConcernMapper by adding support for the textual and dynamic FLT namely SITIR [22]. In each of these tools, programmers can define and name features and then associate entire or partial classes, methods, and fields with them. The tools leave feature location as a manual task and focus on documenting the features and their related elements once they are found. However, once the features and their program elements are documented, they can be saved and retrieved at a later time, thus avoiding the need to repeat searches.

Savage *et al.* [115] developed TopicXP, an Eclipse plug-in that supports developers during maintenance tasks by analyzing the unstructured information embedded in the source code identifiers and comments using an advanced IR technique, LDA [15]. LDA extracts a set of topics from the source code, which could be considered as concepts or features. The topics generated are mapped to the source code, and the relationship between the topics is determined by examining the static dependencies from the code. Using an interactive visualization view, the developer is able to navigate through these topics or to access the source code associated with them. In a preliminary study involving four graduate students, who were required to perform concept location on two Java systems, jEdit and muCommander,³⁶ using Eclipse with TopicXP or just the Eclipse IDE, the results showed that TopicXP is a functional tool, which provides comparable results to Eclipse IDE, and in some cases it even produces better results.

Bohnet and Döellner [25, 116–119] visually explore dynamically extracted information, but in this case, as a call graph. Since a call graph can be large, in order to reduce the search space for the user, the tools provide clues to identify code relevant to the feature of interest. The tools also provide a number of different types of visualizations. In one prototype, the most important view, the graph exploration view, shows other methods that pass control flow to or receive control flow from a given method. In this view, the tool only shows methods in a neighborhood if they are judged to be relevant based on execution time, while another tool has textual and 3D landscape views. These tools effectively extract dynamic call graph information and guide programmers during navigation.

5.1.5. Other tools not included in the survey. Program exploration tools support developers when performing a variety of maintenance tasks. Since feature location is central to many maintenance activities, we mention some of the program exploration tools which are somewhat out of the scope of this survey.

JQuery [120], an Eclipse plug-in, is a source code browsing tool designed to help programmers when dealing with features that have scattered implementations. The tool combines navigation based on relationships (as in a hierarchical browser) with the flexibility of query languages. Program exploration in JQuery begins with a query and a list of variables. The query determines which elements to show in the browser, and the variables establish how to organize them into a tree. The query defines the type of program element to search for given some parameters, such as its name or a type of relationship. The results of the query are returned in a hierarchical tree, and users can further explore the tree with additional queries that expand nodes into sub-trees. The tool aims to reduce the burden of program investigation on developers. It helps them remain oriented by not having to switch between multiple views, and it records their exploration path in the tree format.

³³<http://www.cs.mcgill.ca/~swevo/feat/> (accessed and verified on 03/01/2011)

³⁴<http://www.cs.columbia.edu/~eaddy/concerntagger/> (accessed and verified on 03/01/2011)

³⁵<http://www.cs.wm.edu/semeru/flat3/> (accessed and verified on 03/01/2011)

³⁶<http://www.mucommander.com/> (accessed and verified on 03/01/2011)

Ferret [121] is a tool for answering conceptual queries, which are questions about a software system a programmer may have. The Ferret model is based on the composition and integration of different sources of information into a query-able knowledge base. A source of information is known as a sphere, and examples include static, structural relationships in source code, dynamic call information from an execution trace, and revision history recorded in a software repository. Ferret supports 36 different conceptual queries such as ‘What calls this method?’, ‘What are this class’ subclasses?’, ‘What are all the fields declared by this type?’, and ‘What transactions changed this element?’ These types of queries represent questions programmers may have when investigating a software system in order to locate a feature’s implementation.

Instead of relying on dynamic information, AspectBrowser [122] is a tool that assumes that features follow the idea of information transparency [123]; design decisions that cannot be encapsulated in a single module use a common signature or terminology that can easily be exploited by search tools. The AspectBrowser tool³⁷ allows users to search a code base using pattern matching and then visualizes the results in two ways. All query matches can be highlighted in the source code, and the programmer can browse to find them. Alternatively, programmers can use a global view to see how a feature is scattered throughout the system. In the view, each line of code is represented by a row of pixels, and highlighted rows indicate lines of code that match the query. Multiple search results can be viewed at once to understand the interaction between several features.

5.2. Case studies

A number of case studies involving feature location have been performed, ranging from comparisons of existing techniques, industrial case studies, and user studies. Each type of the study is valuable to advance this research area. Comparisons evaluate few FLTs on the same systems and features, making it easier for researchers and practitioners to understand the strengths and weaknesses of each approach. Industrial case studies show the applicability of an approach in non-trivial settings. Finally, user studies provide insight into how programmers understand and search for code, and these insights can be incorporated into FLTs and tools.

5.2.1. Comparison of FLTs using case studies. RECON, RECON2, and RECON3 are tools that implement the software reconnaissance approach to feature location. Wilde and Casey [124] report on applying RECON to industrial software. In their study on using software reconnaissance for program exploration, Wilde and Casey found the tool to be very selective because it never marked more than 13 methods for a feature. They also observed that the tool seemed to find code that was near relevant program elements. In the second part of their study, they examined using software reconnaissance for traceability to build a large mapping of multiple features to code. The tool was used to run a large set of test cases that were marked with the features they exhibited, and then the collected traces were analyzed to find traceability relations that mapped features to code. With this traceability knowledge, a programmer modifying a program element is aware of the other features implemented in that program element.

When a new FLT is introduced, it is often directly compared with similar existing approaches as part of its evaluation. Some articles related to feature location focus solely on case studies comparing several techniques. Wilde *et al.* [4, 125] compare software reconnaissance, ASDG, and grep in a case study to locate two features in legacy Fortran code. The system, CONVERT3, is part of a suite of geometric modeling programs and is used to convert models to formats required by other tools. For the study, three teams each used one of the FLTs to find the code for two features of CONVERT3. The software reconnaissance and ASDG teams were able to gain sufficient understanding of the source code, but the team using grep was not. The authors concluded that grep was the least reliable approach, but it is very quick and can locate features that cannot be explicitly invoked dynamically. After grep, software reconnaissance was deemed to be the next fastest method of feature location. However, its results may not present a user with enough context to be comprehensible. The ASDG approach was the most difficult to apply, but it was the most systematic and allows for the best understanding of the relevant code.

³⁷<http://cseweb.ucsd.edu/~wgg/Software/AB/> (accessed and verified on 03/01/2011)

Ibrahim *et al.* [126] also report on their experiences applying RECON2 in the Generate Index (GI) project. Their findings echo the conclusions of the previous study. Software reconnaissance is based on test cases, but selecting appropriate scenarios to execute can be difficult. However, only a few test cases are generally needed for a feature. After the analysis, software reconnaissance is good at locating a starting point for feature location, but further investigation for additional relevant program elements should be performed.

Early FLTs were applied in the era of procedural programming paradigm. After object-oriented programming gained popularity, Marcus *et al.* [127] studied whether feature location was still needed since object-oriented code is supposed to be structured such that classes help implement well-defined problem domain concepts. They compared the performance of three static FLTs: pattern matching with grep, a depth-first dependency search, and IR using LSI. Three programmers, each assigned to a different technique, participated in a case study to locate features in Art of Illusion,³⁸ a 3D modeling studio written in Java, and in Doxygen,³⁹ a source code documentation generator written in C++. They concluded that object-orientation does not always allow for quick and easy identification of the program elements relevant to a feature. Therefore, FLTs are still needed for object-oriented systems.

Revelle and Poshyvanyk [128] performed an exploratory study evaluating several FLTs that return ranked lists of program elements (methods). The approaches they compared were IR (LSI-based feature location), IR plus dynamic analysis (SITIR), and IR plus dynamic and static analysis (similar to Cerberus). For IR, they assessed user-formulated queries as well as method seed queries in which the text of a method already known to be relevant to a feature was used as the query. For dynamic analysis, they used both full execution traces and marked traces in which only the portion of a system's execution when a feature is invoked was traced. Dynamic analysis was combined with IR by pruning unexecuted methods from the ranked list. When all three types of analysis were combined in IR + Dyn + Static, a program dependence graph was traversed starting from a seed by following dependencies only if they were executed and had textual similarities to the query that were above a given threshold. Most FLTs that return a ranked list are evaluated in terms of where the first relevant element appears on the list. This case study aimed to evaluate these approaches in terms of how well they find near-complete implementations of features, meaning how well they find as many relevant program elements as possible. Their conclusions were that none of these approaches perform particularly well in that regard because feature location is usually used to find a starting point and impact analysis tools are used to find more complete implementations. They observed that marked traces generally outperformed full traces and that the method seed queries, which can be automatically generated, performed just as well as user formulated queries.

De Alwis *et al.* [109] performed a comparative study in which programmers used three tools: JQuery; Ferret; and Suade; to plan complex maintenance tasks. Eclipse was used as a baseline for comparison. They hypothesized that programmers would find it easier to complete a task using a tool as opposed to Eclipse, need to examine less code as compared to using Eclipse, and would identify more important elements using the tool as opposed to Eclipse. The participants in the study were 18 professional programmers, and they were asked to investigate two change tasks in jEdit. In the first task, they used only Eclipse, and in the second task, they used one of the exploration tools. The order of the tasks and choice of tools was randomized. An instrumented version of Eclipse captured all events the programmers performed during their investigation. Additionally, the participants recorded the relevant elements they found in an Eclipse view built for the study. The NASA Task Load Index (TLX) was used to assess task difficulty, and distance profiles were used to gauge the degree to which the participants remained on-task. The TLX scores showed no difference in task difficulty that could be attributed to using a tool or not. Similarly, the distance profiles did not indicate that the tools had any strong effect on the tasks. Overall, the authors concluded that program exploration tools had little effect, and that the behavior of the developers seemed more impacted by the tasks performed. In addition, there was evidence that individual programmers' strategies caused them to be more or less efficient.

³⁸<http://aoi.sourceforge.net> (accessed and verified on 03/01/2011)

³⁹<http://www.stack.nl/~dimitri/doxygen/> (accessed and verified on 03/01/2011)

5.2.2. Industrial case studies. Most feature location case studies focus on open source software. However, case studies carried out on industrial software give a sense of a technique's real world applicability. Unfortunately, only a few such studies have been performed, and more are needed. As previously discussed, TraceGraph was used in an industrial setting [101], and Wilde *et al.* [4] compared a number of approaches on industrial software. In addition to these studies, Van Geet and Demeyer [129] report on their experiences of applying Eisenbarth *et al.*'s [7] formal concept analysis of execution traces FLT in an industrial setting. The context was a pre-study phase for the migration of a banking system written in COBOL. Scenarios for two features were executed using a web interface, and three separate iterations of the approach were conducted. Each iteration aimed at reducing the number of modules considered by using different combinations of scenarios that did and did not invoke the feature. A domain expert provided the modules relevant to each feature for evaluation purposes, and in two out of three iterations of the approach, all the relevant modules were in the generated concept lattice. Three additional relevant modules were also identified that had not previously been named by the domain expert.

5.2.3. User studies. Studies that focus on how programmers search and comprehend source code are important to feature location research. These types of studies provide insights on how developers find a feature's implementation or gain understanding of a system. In turn, these insights can be incorporated into feature location research in order to develop approaches that are organic and easy for programmers to use. Four relevant user studies are discussed below, and while this is not a complete list of user studies related to feature location, even more studies are necessary to advance the state of the art.

LaToza *et al.* [130] performed a user study in which 13 participants worked for three hours on understanding and improving the design of two features in jEdit. The participants' activities were recorded using think-aloud, video, and Eclipse instrumentation. The goal of the study was to answer questions about how programmers' experience affects the changes they make to code, how it affects how they work, and how they reason about design during coding tasks. LaToza *et al.* found that the more experienced programmers addressed the causes of problems, while beginners focused on the symptoms and that the experienced programmers identified relevant methods and implemented changes more quickly than the novices. They also discovered that the participants' activities centered on fact finding. The programmers sought facts relevant to their task, so they investigated certain methods and learned facts about the software system; as they learned enough facts, they were able to propose design changes. Therefore, FLTs should not only help identify relevant program elements but also aid in fact finding and program comprehension.

Robillard *et al.* [131] also conducted a study on how programmers explore source code when performing a change task. Five programmers were asked to modify jEdit so that i) users can explicitly disable the auto-save functionality and that, ii) the new version will meet five requirements. The data collected included artifacts produced or modified by the participants as well as video recordings of their screens. The programmers' success was judged in terms of time to complete the task and the quality of their change to the source code in terms of how many of the task's requirements they successfully implemented. Robillard *et al.* analyzed the behavior of each participant by transcribing the screen videos into events. Each event records the time it occurred, the method being examined at that time, how the method was accessed (scrolling, browsing, searching, etc.), and whether the method was modified. Based on their observations, they conclude that a methodical, ordered investigation of a system's source code is more effective than a systematic, opportunistic one. They found that programmers should follow a plan when exploring a program, that they should perform focused searches in the context of their plan, and that they should keep record of their findings. Based on these results, FLTs should facilitate orderly program exploration.

Revelle *et al.* [132] undertook an exploratory study on how programmers identify features and their implementations in source code. In the first study, the features of GNU sort⁴⁰ plus their relevant source code were found manually by one programmer and then compared to those of Carver and Griswold [133]. In the second study, two programmers manually located features and their implementations for a Java implementation of the Minesweeper game. Revelle *et al.* compared the actual concepts

⁴⁰<http://www.gnu.org/software/coreutils/> (accessed and verified on 03/01/2011)

recognized as features as well as the code associated with those features, looking for common trends in how developers identify and locate features. Based on their observations, they developed a set of guidelines on how to identify and recognize the existence of a feature and how to record feature's associated code in a tool called Spotlight [134]. The guidelines suggest relying on both static and textual information and flexible mappings of features to program elements of various levels of granularity.

Ko *et al.* [135, 136] performed an exploratory study to investigate developers' strategies for understanding unfamiliar code. Ten participants worked using Eclipse on five maintenance tasks associated with the Paint⁴¹ application. Screen-capture videos recorded the developers' work during the study. To simulate a more realistic working environment, the programmers were interrupted every 2.5 to 3.5 minutes and required to answer a multiplication question. Monetary incentives were offered for correctly completing the tasks, and penalties were inflicted for ignoring or incorrectly answering the multiplication questions. The study found that programmers interleave three activities when exploring source code: searching for relevant code either manually or with tools, following the dependencies of found relevant code, and collecting relevant code and information in Eclipse's interface (i.e., package explorer, tabs, and scroll bars). However, searches often failed, Eclipse's navigation tools imposed overhead when following dependencies, and developers lost track of relevant code in the interface. On average, 35% of a developer's time was spent reviewing search results and on navigation. Based on the observations of this study the authors make a number of suggestions for future tool development. First, tools need to provide better relevance cues so programmers do not miss important code or misinterpret irrelevant code. Second, dependency searches need to be more practical, such as by highlighting the dependencies of the currently selected program element. Third and finally, programmers need a better way to collect, organize, and view the relevant information they find, such as being able to see all relevant information for a given task at once. These recommendations may help programmers find task-relevant code more quickly and efficiently and were used in the design of a new debugging tool [137].

6. ANALYSIS OF THE TAXONOMY

Using the taxonomy, we are able to address the following research questions:

RQ1: What types of analysis are used when performing feature location? The results of our systematic survey point out that FLTs primarily use dynamic, static, and textual information to locate features in source code. However, analysis techniques for feature location are not limited to dynamic, static, or textual analysis, as techniques which combine various types of analysis as well as those which utilize history-based analysis are also prevalent in the literature. For a detailed discussion of the various types of analysis, see Section 3.1. Likewise, discussion classifying the various FLTs according to their respective type of analysis appears in Section 4.

RQ2: Has there been a change in types of analysis used to identify features in source code employed by recent FLTs? Our survey reveals that, in recent years, there has been an increase in the number of papers published that introduce FLTs based on textual information as well as historical data (see Section 4). This can be observed in Figure 3 (dotted line), which displays the cumulative number of FLTs that use different types of information per year. We choose this cumulative visual representation of approaches, as opposed to a direct scattering of the number of approaches per year, because the graph is much easier to understand. In addition to the emergence of new analysis techniques, our survey indicates that a growing number of researchers have demonstrated the benefits of combining multiple analysis techniques to leverage the complementary strengths of each analysis mechanism (see solid lines in Figure 3).

RQ3: Are there any limitations to current strategies for evaluating various FLTs? Two observations that can be made from our feature location survey include (1) limited comparison of proposed techniques with existing techniques and, more notably, (2) limited use of benchmarks in

⁴¹<http://www.cs.cmu.edu/~marmalade/studies.html> (accessed and verified on 03/01/2011)




Figure 3. Cumulative number of FLT's from Tables III and IV per year.

evaluations. Comparison of proposed approaches to other existing techniques should be common practice while introducing new FLT's. However, only 23 out of 60 papers (38%) that present FLT's (see Section 4 and Tables III and IV) compared their approaches with a limited number of existing approaches. In rare instances, existing techniques may not be relevant to be used during evaluation. Nevertheless, researchers ought to compare their new techniques against those that appear in the literature. The apparent lack of comparison may be attributed to the low number of techniques which provide a corresponding publicly available tool. Without tools for each of the existing techniques, researchers need to implement both the proposed technique as well as the relevant techniques which should be compared against the proposed one. Such a process is time-consuming and also introduces a threat to validity (possibility of erroneous implementation of existing FLT's). As it stands, direct comparison of approaches is limited; however, it may be increased if researchers produce publicly available tools. The second observed limitation is the lack of evaluation on benchmarks. In fact, only 3 of the 60 papers (5%) that present FLT's (see Section 4, Tables III and IV) used benchmarks in their evaluation. If benchmarks were readily available for researchers to evaluate their approaches, availability of publicly available tools would not be such a significant problem, since researchers would have been able to directly compare their results to the results, which appear on the same datasets.

7. DISCUSSION AND OPEN ISSUES

Feature location is an essential aspect of many software maintenance tasks, and because it can be challenging to perform manually, researchers have introduced many techniques to lessen the burden of searching for a feature's relevant code. Even with these numerous approaches and advancements, open issues remain in the field of feature location. One question that is unanswered is, 'What is the best way to perform feature location?' This question cannot be easily answered without an extensive comparison of analysis-specific issues and a comparison of existing approaches. Note that while addressing RQ₃, we observed that comparisons of existing approaches are simply missing from most evaluations in the feature location literature. This means that even though different approaches in their evaluations used the same systems (e.g., Eclipse, JHotDraw, jEdit, Mozilla, etc.), they did not use the same datasets (e.g., version, source code, gold set, etc.) for those systems. In addition, the approaches were not restricted to use the same evaluation metrics (e.g., precision, recall, effectiveness, etc.), which makes comparison between approaches even more difficult. Such a comparison could be facilitated by well-established benchmarks, but currently, there is no commonly accepted set of features associated with the code that implements them that could be used to compare FLT's. The issue of benchmarks was identified as we addressed RQ₃ using our systematic survey. Such a benchmark is needed in the research area. Additionally, while there are various techniques that support feature

location, not all approaches have publicly available tools, and the tools that are available do not support both locating and documenting a feature's implementation. Other open issues are usability studies of FLT's and integrating feature location into software engineering courses. The remainder of this section discusses these open issues and their associated avenues for future research. While this discussion brings to light these important topics, more panels and workshops, such as the one on the 'Identifications of concepts, features, and concerns in source code' held at the 21st IEEE International Conference on Software Maintenance (ICSM2005), are necessary to resolve many of these issues.

7.1. Comparisons

Given a wide variety of existing techniques, developers that need to perform feature location have many options, but which approach is the best for a specific situation? What parameters should be used for a certain type of analysis? Which type of analysis yields the best results, or is a combination of two or more types of analysis the best? Some case studies have been performed comparing multiple FLT's, but they only have a few data points, which impede the ability to draw statistically significant generalizations from their results. These studies are also limited in their number of examined approaches, focusing on a subset of approaches that present results in a similar fashion. An obstacle to comparing techniques is the presentation of their results. How does one evaluate one result set that ranks program elements to another that does not? Determining the best way to directly compare the performance of FLT's remains an open issue.

Not only does there need to be a comparison of techniques based on different types of analysis, but there also needs to be an evaluation of the best configuration of each type of analysis. For instance, dynamic analysis has many possible options for collecting traces. The granularity of execution traces can be classes, methods, or even lines of code. In addition, the entire execution can be logged from start up to shut down, or only select portions of the run can be captured. With static analysis, like with dynamic, granularity is also a parameter. Additionally, the type of dependencies (control or data) to take into account is another issue to consider. With textual analysis, preprocessing options, such as stemming and stop word removal, are commonly used, but their effect on feature location has not been fully studied. Also, textual analysis can be achieved through IR methods or through NLP. While the varied IR methods have been compared, the effectiveness of IR and NLP has not been compared in the context of feature location. A comparison would determine if the extra expense associated with NLP is worth the precision, or if the less expensive IR methods are sufficient. Thorough investigations comparing these different configurations of each type of analysis would reveal the most favorable settings for feature location.

There are many other open feature location issues that could potentially be resolved through comparisons. The main types of analyses are dynamic, static, and textual, although historical analysis has also been used, but not in conjunction with any other analysis. It remains to be seen if combining historical analysis with any of the others is a viable approach to feature location. Just as three types of analysis comprise the majority of existing techniques, two programming languages dominate the area of feature location. Most existing approaches have been applied to Java or C/C++. However, feature location should branch out to support more programming languages.

7.2. Benchmarks

The comparison of FLT's should be facilitated by the existence of benchmarks that could be used to consistently evaluate the approaches. Currently, there are a number of systems that have been used in the evaluation of many FLT's such as Eclipse, JHotDraw, jEdit, Mozilla, and Firefox, but sets of features used for the evaluation are not consistent. Even if two approaches are evaluated on the same system, if different features are used, comparing two techniques is difficult. Another problem with assessing feature location is in knowing the 'gold set' of program elements that implement a feature. The most commonly used method for determining the source code relevant to a feature is to mine bug tracking systems. However, the code associated with a feature may be incomplete if a bug fix only touches part of a feature's implementation. In the presence of these issues, the field of feature location research needs to establish solid standards for validation. The best solution may be

benchmarks that can facilitate easy comparison of FLT approaches. The benchmarks should include a set of features from a software system or several systems. Each feature needs to be mapped to the source code implementing that. Preferably, the benchmarks should be defined at various granularities to support different approaches that identify relevant classes, methods, or statements. Robillard *et al.* [28], Eaddy *et al.* [73], and Revelle *et al.* [87] took a step in this direction, making their datasets available in which programmers, who were not necessarily systems experts, mapped features to classes, methods, and fields in open source Java applications. Still, well-established and complete benchmarks for systems from a variety of domains and languages will make the evaluation and comparison of FLTs easier.

In this survey, we make a step toward solving this problem by making a set of benchmarks publicly available for several systems that were used frequently in case studies. The purpose of making these datasets available to researchers is twofold. First, we want to facilitate the evaluation of new approaches on larger and diverse datasets in order to make the approach more generalizable. Second, we want to ease the comparison between different FLTs. Using the same datasets, a new approach can be straightforwardly compared against an existing approach, which could provide a great indication about the value of the proposed technique. We were encouraged to make available the benchmark by observing the success of Eaddy *et al.*'s [73] benchmark. The Rhino dataset from that benchmark was used in several evaluations, and it was used by different research groups as well [24, 61, 73, 87].

The benchmarks that we make available could be downloaded from the website <http://www.cs.wm.edu/semeru/data/benchmarks/>. The website contains detailed instructions about the format of the datasets and the process used to generate them. These benchmarks include gold sets (i.e., mappings between source code and features) at method-level granularity, description of the features from bug reports, and in some cases even execution traces. In the near future, we plan to generate more benchmarks and add them to this website, but for now, we are making available the following datasets.

The first benchmark is for jEdit version 4.3, and it contains 272 features that have associated gold sets and issues from the issue tracking system, which can be used to extract the textual descriptions of the features/bugs. In addition, 150 features have associated execution traces that were generated by following the steps to reproduce enumerated in the issue description. The gold sets were generated via analysis of the SVN commits submitted between releases 4.2 and 4.3. The issue identifiers from the SVN logs were extracted and mapped to issues from the issue tracking system. On the other hand, the changes from the SVN commits were mapped to the methods from the source code that were modified by that commit.

The second benchmark we make available is for ArgoUML version 0.22, and it contains 132 features that have gold sets and issues from the issue tracking system. In addition, 91 features have associated execution traces that were generated by following the steps to reproduce described in the issue description. The process of generating the gold sets is similar to the one used for generating the benchmark for jEdit system.

The third benchmark we make available is for muCommander version 0.8.5, and it contains 92 features that have gold sets, issues from the issue tracking system, as well as execution traces. The process of gathering the data for this benchmark is similar to the ArgoUML's process.

The fourth benchmark is for JabRef⁴² version 2.6, and it contains 39 features that have associated gold sets, issues as well as execution traces. The data collection process is similar to the ArgoUML's process.

The fifth benchmark we provide is for Eclipse version 3.0, and it contains 45 features represented by bug reports submitted to Eclipse's issue tracking system. Unlike the jEdit and ArgoUML datasets, the gold sets were generated by examining the patches submitted in the bug reports, which contain information about the code that was changed to fix the bug. We also provide 45 execution traces with this dataset, which were collected by following the steps to reproduce from the description of the bugs. The Eclipse benchmark was used in the evaluation by Revelle *et al.* [87].

⁴²<http://jabref.sourceforge.net/> (accessed and verified on 03/01/2011)

7.3. Tools

Even though this survey encompasses many tools that support feature location (see Section 5.1), the majority of FLT tools do not have a publicly available tool, meaning that programmers wanting to apply such an approach may need to recreate the technique's methodology. Additionally, some tools are useful for investigating a program and locating features, while others can be used to store the mappings between features and source code; but, there are currently no tools that perform both. Combining the functionalities of finding features' implementations and being able to save them is a logical next step for tool development. Finally, de Alwis *et al.* [109] found that existing tools have little effect on the programmer's efficiency; so, further research needs to be done to improve usability of the tools.

7.4. User studies

While there have been several user studies investigating how programmers search and explore source code during maintenance, these studies are conducted with a relatively small number of developers. Further studies are needed with more programmers to be able to derive conclusive and statistically significant results. Additionally, there has been a lack of empirical studies examining usability aspects of feature location. *Do existing FLT tools reduce the amount of time and effort developers spend on maintenance? What are the practical benefits and costs of using different types of approaches?* For instance, collecting execution traces for an approach that uses dynamic analysis requires overhead in terms of the time spent to develop scenarios or test cases and capture traces. IR involves indexing the source code, which can be rather time-consuming for large software systems. Studies are needed to determine whether the overhead of collecting traces or (re)indexing a corpus yields improved feature location results and if it is worth the cost.

7.5. Feature location and education

Given that feature location is such an extensive area of research and also an important part of software maintenance, it should be taught in software engineering courses at universities and colleges. Petrenko *et al.* [138, 139] have argued for the inclusion of software maintenance and evolution in software engineering courses along with traditional development. Teaching maintenance exposes students to more realistic experiences because in industry, 70% or more of programmers' time is devoted to maintenance [140, 141]. Feature location is a significant part of the maintenance phase because before changes can be made to a system, the relevant program elements must be found. Therefore, feature location should be introduced as a topic in software engineering courses to better prepare students.

8. CONCLUSION

Through a comprehensive examination of 89 feature location articles encompassing research, tool, and case, industrial, and user studies, this survey has presented a taxonomy that classifies the literature along nine key dimensions. The taxonomy facilitates the comparison of existing FLT tools and illuminates possible areas of future research. Researchers can use the taxonomy and survey as a basis for advancing the field, while practitioners can use it to identify techniques and tools that are well-suited to their needs. This survey has also shed light on open issues in feature location, such as the need for comparisons and benchmarks. By structuring the research area of feature location, this taxonomy and survey contribute well-defined organization to the field and should aid in resolving some of the open issues. This systematic survey should serve not only academic researchers but also industrial professionals, aimed at adopting feature location tools within their organizations and development processes.

ACKNOWLEDGEMENTS

We are grateful to the anonymous JSME reviewers for their relevant and useful comments and suggestions, which helped us in significantly improving the earlier versions of this taxonomy, the survey, and the paper in general. We would also like to thank Andrian Marcus for his useful and stimulating suggestions on the

earlier drafts of this survey paper. We would also like to express our gratitude to some of the authors whose work is discussed in this survey for their insightful feedback in preparing the camera ready version of this paper. This work is supported by NSF CCF-0916260 and NSF CCF-1016868 grants. Any opinions, findings, and conclusions expressed herein are the authors' and do not necessarily reflect those of the sponsors.

REFERENCES

1. Biggerstaff TJ, Mitbender BG, Webster DE. The concept assignment problem in program understanding. Proceedings of 15th IEEE/ACM International Conference on Software Engineering (ICSE'94) May 17–21, 482–498, 1994.
2. Rajlich V, Wilde N. The role of concepts in program comprehension. Proceedings of IEEE International Workshop on Program Comprehension (IWPC'02), 2002: 271–278.
3. Rajlich V, Gosavi P. Incremental change in object-oriented programming. *IEEE Software*, 2004: 2–9.
4. Wilde N, Buckellew M, Page H, Rajlich V, Pounds L. A comparison of methods for locating features in legacy software. *Journal of Systems and Software* 2003; **65**(2): 105–114. (February 15)
5. Marcus A, Rajlich, V. Panel summary: identifications of concepts, features, and concerns in source code. Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05), Budapest, Hungary, September 25–30, 2005b: 718.
6. Kitchenham B. *Procedures for Performing Systematic Reviews*, Keele UK (ed.). Keele University: TR/SE-0401, I. 1353–I. 7776.
7. Eisenbarth T, Koschke R, Simon D. Locating features in source code. *IEEE Transactions on Software Engineering* 2003; **29**(3): 210–224. (March)
8. Wilde N, Scully M. Software reconnaissance: mapping program features to code. *Journal of Software Maintenance Research and Practice* 1995; **7**: 49–62.
9. Antoniol G, Guéhéneuc YG. Feature identification: an epidemiological metaphor. *IEEE Transactions on Software Engineering* 2006 **32**(9): 627–641.
10. Eisenberg AD, De Volder K. Dynamic Feature Traces: Finding Features in Unfamiliar Code. Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05), Budapest, Hungary, September 25–30, 2005: 337–346.
11. Safyallah H, Sartipi K. Dynamic analysis of software systems using execution pattern mining. Proceedings of International Conference on Program Comprehension (ICPC'06), 2006: 84–88.
12. Chen K, Rajlich V. Case study of feature location using dependence graph. Proceedings of 8th IEEE International Workshop on Program Comprehension (IWPC'00), Limerick, Ireland, June 2000: 241–249.
13. Robillard MP. Topology analysis of software dependencies. *ACM Transactions on Software Engineering and Methodology* 2008; **17**(4): 1–36. (August)
14. Deerwester S, Dumais ST, Furnas GW, Landauer TK, Harshman R. Indexing by latent semantic analysis. *Journal of the American Society for Information Science* 1990; **41**: 391–407.
15. Blei DM, Ng AY, Jordan MI. Latent Dirichlet allocation. *Journal of Machine Learning Research* 2003; **3**: 993–1022.
16. Salton G, McGill M. *Introduction to Modern Information Retrieval*. McGraw-Hill: New York, NY, USA. 1986.
17. Binkley D, Lawrie D. Information Retrieval Applications in Software Maintenance and Evolution. *Encyclopedia of Software Engineering*. Taylor & Francis, 2010.
18. Furnas GW, Landauer TK, Gomez LM, Dumais ST. The vocabulary problem in human–system communication. *Communications of the ACM* 1987; **30**(11): 964–971.
19. Chen A, Chou E, Wong J, Yao AY, Zhang Q, Zhang S, Michail A. CVSSearch: searching through source code using CVS comments. Proceedings of IEEE International Conference on Software Maintenance (ICSM'01), Nov., 2001a: 364–373.
20. Cubranic D, Murphy GC, Singer J, Booth KS. Hipikat: a project memory for software development. *IEEE Transactions on Software Engineering* 2005; **31**(6): 446–465. (June)
21. Robillard MP, Murphy GC. Automatically inferring concern code from program investigation activities. Proceedings of 18th International Conference on Automated Software Engineering (ASE'03), Linz, Austria, October, 2003a: 225–234.
22. Liu D, Marcus A, Poshyvanyk D, Rajlich V. Feature location via information retrieval based filtering of a single scenario execution trace. Proceedings of 22nd IEEE/ACM International Conference on Automated Software Engineering (ASE'07), Atlanta, Georgia, November 5–9, 2007: 234–243.
23. Marcus A, Sergeyev A, Rajlich V, Maletic J. An information retrieval approach to concept location in source code. Proceedings of 11th IEEE Working Conference on Reverse Engineering (WCRE'04), Delft, the Netherlands, November 9–12, 2004: 214–223.
24. Eaddy M, Aho AV, Antoniol G, Guéhéneuc YG. CERBERUS: tracing requirements to source code using information retrieval, dynamic analysis, and program analysis. Proceedings of 16th IEEE International Conference on Program Comprehension (ICPC'08), Amsterdam, the Netherlands, 2008a: 53–62.
25. Bohnet J, Döellner J. Facilitating exploration of unfamiliar source code by providing 2½D visualizations of dynamic call graphs. Proceedings of 4th IEEE International Workshop on Visualizing Software for Understanding and Analysis (VISSOFT'07), Banff, Alberta, Canada, June 25–26, 2007b: 63–66.
26. Walkinshaw N, Roper M, Wood M. feature location and extraction using landmarks and barriers. Proceedings of International Conference on Software Maintenance (ICSM'07), Paris, France, 2007: 54–63.

27. Xie X, Poshyvanyk D, Marcus A. 3D Visualization for concept location in source code. Proceedings of 28th IEEE/ACM International Conference on Software Engineering (ICSE'06), Shanghai, China, May 20–28, 2006: 839–842.
28. Robillard MP, Shepherd D, Hill E, Vijay-Shanker K, Pollock L. *An Empirical Study of the Concept Assignment Problem*. McGill University: Montreal, Quebec, Canada. 2007b.
29. Lukins S, Kraft N, Etzkorn L. Source code retrieval for bug location using latent Dirichlet allocation. Proceedings of 15th Working Conference on Reverse Engineering (WCRE'08), Antwerp, Belgium, 2008: 155–164.
30. Poshyvanyk D, Guéhéneuc, YG, Marcus A, Antoniol G, Rajlich V. Feature location using probabilistic ranking of methods based on execution scenarios and information retrieval. *IEEE Transactions on Software Engineering* 2007a; **33**(6): 420–432. (June)
31. Cornelissen B, Zaidman A, Van Deursen A, Moonen L, Koschke R. A systematic survey of program comprehension through dynamic analysis. *IEEE Trans. on Software Engineering* September–October 2010; **35**(5): 684–702.
32. Egyed A, Binder G, Grunbacher P. STRADA: a tool for scenario-based feature-to-code trace detection and analysis. Proceedings of IEEE/ACM 29th International Conference on Software Engineering (ICSE'07), 2007: 41–42.
33. Salah M, Mancoridis S. A hierarchy of dynamic software views: from object-interactions to feature-interactions. Proceedings of 20th IEEE International Conference on Software Maintenance (ICSM'04), Chicago, IL, September 11–14, 2004: 72–81.
34. Salah M, Mancoridis S, Antoniol G, Di Penta M. Scenario-driven dynamic analysis for comprehending large software systems. Proceedings of 10th IEEE European Conference on Software Maintenance and Reengineering (CSMR'06), March 22–24, 2006: 71–80.
35. Greevy O, Ducasse S, Girba T. Analyzing feature traces to incorporate the semantics of change in software evolution analysis. Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05), 2005: 347–356.
36. Greevy O, Ducasse S, Girba T. Analyzing software evolution through feature views. *Journal of Software Maintenance and Evolution: Research and Practice* 2006; **18**(6): 425–456. (November)
37. Fischer M, Pinzger M, Gall H. Analyzing and relating bug report data for feature tracking. Proceedings of IEEE Working Conference on Reverse Engineering (WCRE'03), 2003: 90–101.
38. Kothari J, Denton T, Mancoridis S, Shokoufandeh, A. On computing the canonical features of software systems. *13th IEEE Working Conference on Reverse Engineering (WCRE'06)*. Benevento, Italy, 2006.
39. Wilde N, Gomez JA, Gust T, Strasburg D. Locating user functionality in old code. Proceedings of IEEE International Conference on Software Maintenance (ICSM'92), Orlando, FL, November, 1992: 200–205.
40. Wong WE, Gokhale SS, Horgan JR, Trivedi KS. Locating program features using execution slices. Proceedings of IEEE Symposium on Application-Specific Systems and Software Engineering and Technology (ASSET'99), March 24–27, 1999: 194–203.
41. Agrawal H, Alberi JL, Horgan JR, Li JJ, London S, Wong WE, Ghosh S, Wilde N. Mining system tests to aid software maintenance. *Computer* 1998; **31**: 64–73. (July 1998)
42. Sahner RA, Trivedi KS. *SHARPE: Symbolic Hierarchical Automated Reliability and Performance Evaluator, Introduction and Guide for Users*. Duke University: Durham, NC. 1986.
43. Eisenbarth T, Koschke R, Simon D. Derivation of feature component maps by means of concept analysis. Proceedings of European Conference on Software Maintenance and Reengineering (CSMR'01), 2001b: 176–179.
44. Eisenbarth T, Koschke R, Simon D. Feature-driven program understanding using concept analysis of execution traces. Proceedings of International Workshop on Program Comprehension (IWPC'01), 2001c: 300–309.
45. Edwards D, Simmons S, Wilde N. An approach to feature location in distributed systems. *Journal of Systems and Software* 2006; **79**(1): 57–68. (Jan. 2006)
46. Bohnet J, Voigt S, Dollner J. Locating and understanding features of complex software systems by synchronizing time-, collaboration- and code-focused views on execution traces. Proceedings of 16th IEEE International Conference on Program Comprehension (ICPC'08), Amsterdam, the Netherlands, June 10–13, 2008b: 268–271.
47. Edwards D, Wilde N, Simmons S, Golden E. Instrumenting time-sensitive software for feature location. Proceedings of International Conference on Program Comprehension (ICPC'09), Vancouver, British Columbia, 2009: 130–137.
48. Chen K, Rajlich V. RIPPLES: tool for change in legacy software. Proceedings of International Conference on Software Maintenance (ICSM'01), Florence, Italy, 2001b: 230–239. (November 07–09, 2001)
49. Buckner J, Buchta J, Petrenko M, Rajlich V. JRipples: a tool for program comprehension during incremental change. Proceedings of 13th IEEE International Workshop on Program Comprehension (IWPC'05), St. Louis, Missouri, USA, May 15–16, 2005: 149–152.
50. Robillard MP, Murphy GC. Concern Graphs: finding and describing concerns using structural program dependencies. Proceedings of International conference on software engineering, 2002: 406–416.
51. Robillard MP, Murphy GC. Representing concerns in source code. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 2007a; **16**(1).
52. Robillard M. Automatic generation of suggestions for program investigation. Proceedings of Joint European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering, Lisbon, Portugal, September, 2005a: 11–20.
53. Robillard MP, Weigand-Warr F. ConcernMapper: simple view-based separation of scattered concerns. Proceedings of Eclipse Technology Exchange at OOPSLA (ETX'05), 2005b: 65–69.
54. Weigand-Warr F, Robillard MP. Suade: topology-based searches for software investigation. Proceedings of International Conference on Software Engineering, May, 2008: 780–783.

55. Trifu M. Using dataflow information for concern identification in object-oriented software systems. *Proceedings of European Conference on Software Maintenance and Reengineering (CSMR'08)*, 2008: 193–202.
56. Trifu M. Improving the dataflow-based concern identification approach. *Proceedings of European Conference on Software Maintenance and Reengineering (CSMR'09)*, 2009: 109–118.
57. Petrenko M, Rajlich V, Vanciu R. Partial domain comprehension in software evolution and maintenance. *International Conference on Program Comprehension*, 2008.
58. Cleary B, Exton C, Buckley J, English M. An empirical analysis of information retrieval based concept location techniques in software comprehension. *Empirical Software Engineering* 2009; **14**(1): 93–130.
59. Gay, G, Haiduc S, Marcus M, Menzies T. On the use of relevance feedback in IR-based concept location. *Proceedings of 25th IEEE International Conference on Software Maintenance (ICSM'09)*, Edmonton, Canada, September, 2009.
60. Poshyvanyk D, Marcus D. Combining formal concept analysis with information retrieval for concept location in source code. in *Proceedings of 15th IEEE International Conference on Program Comprehension (ICPC'07)*, Banff, Alberta, Canada, June, 2007b: 37–48.
61. Hill E, Pollock L, Vijay-Shanker K. Automatically capturing source code context of NL-queries for software maintenance and reuse. *Proceedings of 31st IEEE/ACM International Conference on Software Engineering (ICSE'09)*, May, 2009: 16–24.
62. Shepherd D, Fry Z, Gibson E, Pollock L, Vijay-Shanker K. Using natural language program analysis to locate and understand action-oriented concerns. *Proceedings of 6th International Conference on Aspect Oriented Software Development (AOSD'07)*, 2007: 212–224.
63. Poshyvanyk D, Guéhéneuc YG, Marcus A, Antoniol G, Rajlich V. Combining probabilistic ranking and latent semantic indexing for feature identification. *Proceedings of 14th IEEE International Conference on Program Comprehension (ICPC'06)*, Athens, Greece, 2006a: 137–146.
64. Wilson LA. Using ontology fragments in concept location. *Proceedings of 26th IEEE International Conference on Software Maintenance (ICSM'10)*, Timisoara, Romania, Sept 12–18, 2010: 1–2.
65. Lukins SK, Kraft NA, Etzkorn LH. Bug localization using latent Dirichlet allocation. *Information and Software Technology* 2010; **52**(9): 972–990.
66. Cleary B, Exton C. Assisting concept location in software comprehension. *Proceedings of 19th Psychology of Programming Workshop*, 2007: 42–55.
67. Cleary B, Exton C. The cognitive assignment eclipse plug-in. *Proceedings of 14th IEEE International Conference on Program Comprehension (ICPC'06)*, Athens, Greece, 2006: 241–244.(June 14–17)
68. Zhai C, Lafferty, J. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems* 2004; **22**(2): 179–214. (Apr. 2004)
69. Gao J, Nie J-Y, Wu G, Cao G. Dependence language model for information retrieval. *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval*, Sheffield, United Kingdom, 2004: 170–177.
70. Comon P. Independent component analysis, a new concept? *Signal Process.* 1994; **36**(3): 287–314. (Apr. 1994)
71. Grant S, Cordy JR, Skillicorn DB. Automated concept location using independent component analysis. *Proceedings of 15th Working Conference on Reverse Engineering (WCRE'08)*, Antwerp, Belgium, October 15–18, 2008: 138–142.
72. Shepherd D, Pollock L, Vijay-Shanker K. Towards supporting on-demand virtual remodularization using program graphs. *Proceedings of International Conference on Aspect-Oriented Software Development (AOSD'06)*, Bonn, Germany, 2006: 3–14.
73. Eaddy M, Zimmermann T, Sherwood K, Garg V, Murphy G, Nagappan N, Aho AV. Do crosscutting concerns cause defects?. *IEEE Transaction on Software Engineering* 2008b; **34**(4): 497–515. (July-August)
74. Abebe SL, Tonella P. Natural language parsing of program element names for concept extraction. *Proceedings of 18th IEEE International Conference on Program Comprehension (ICPC'10)*, Braga, Portugal, June 30 - July 2, 2010: 156–159.
75. Würsch M, Ghezzi G, Reif G, Gall HC. Supporting developers with natural language queries. *Proceedings of 32nd ACM/IEEE International Conference on Software Engineering (ICSE'10)*, Cape Town, South Africa, May 2–8, 2010: 165–174.
76. Dufour B, Ryder BG, Sevitsky G. Blended analysis for performance understanding of framework-based applications. *Proceedings of International Symposium on Software Testing and Analysis (ISSTA'07)*, London, United Kingdom, 2007: 118–128.
77. Ernst M. Static and dynamic analysis: synergy and duality. *Proceedings of ICSE Workshop on Dynamic Analysis (WODA'03)*, Portland, OR, May, 2003: 24–27.
78. Eisenbarth T, Koschke R, Simon D. Aiding program comprehension by static and dynamic feature analysis. *Proceedings of International Conference on Software Maintenance (ICSM01)*, Florence, Italy, 2001a: 602–611. (November 7–9)
79. Koschke R, Quante J. On dynamic feature location. *Proceedings of 20th IEEE/ACM International Conference on Automated Software Engineering (ASE'05)*, Long Beach, CA, USA, 2005: 86–95.
80. Antoniol G, Guéhéneuc Y. Feature identification: a novel approach and a case study. *Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05)*, Budapest, Hungary, September 25, 2005: 357–366.
81. Rohatgi A, Hamou-Lhadj A, Rilling J. Feature location based on impact analysis. *Proceedings of International Conference on Software Engineering and Applications (SEA'07)*, 2007.

82. Rohatgi A, Hamou-Lhadj A, Rilling J. An approach for mapping features to code based on static and dynamic analysis. Proceedings of 16th IEEE International Conference on Program Comprehension (ICPC'08), Amsterdam, the Netherlands, 2008: 236–241.
83. Rohatgi A, Hamou-Lhadj A, Rilling J. An approach for solving the feature location problem by measuring the component modification impact. *IET Software* 2009; **3**(4): 292–311. (August)
84. Savage T, Revelle M, Poshyvanyk D. FLAT³: feature location and textual tracing tool. Proceedings of 32nd ACM/IEEE International Conference on Software Engineering (ICSE'10), Cape Town, South Africa, 2010b: 255–258.
85. Asadi F, Di Penta M, Antoniol G, Guéhéneuc YG. A heuristic-based approach to identify concepts in execution traces. Proceedings of 14th European Conference on Software Maintenance and Reengineering (CSMR'10), Madrid, Spain, March, 2010.
86. Marcus A, Poshyvanyk D. The conceptual cohesion of classes. Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05), Budapest, Hungary, September 25–30, 2005a: 133–142.
87. Revelle M, Dit B, Poshyvanyk D. Using data fusion and web mining to support feature location in software. Proceedings of 18th IEEE International Conference on Program Comprehension (ICPC'10), Braga, Portugal, June 30 - July 2, 2010: 14–23.
88. Brin S, Page L. The anatomy of a large-scale hypertextual web search engine. Proceedings of 7th International Conference on World Wide Web, Brisbane, Australia, 1998: 107–117.
89. Kleinberg JM. Authoritative sources in a hyperlinked environment. *Journal of the ACM* 1999; **46**(5): 604–632.
90. Hayashi S, Sekine K, Saeki M. iFL: An interactive environment for understanding feature implementations. Proceedings of 26th IEEE International Conference on Software Maintenance (ICSM'10), Timisoara, Romania, Sept 12–18, 2010a: 1–5.
91. Zhao W, Zhang L, Liu Y, Sun J, Yang F. SNIAFL: towards a static non-interactive approach to feature location. Proceedings of 26th International Conference on Software Engineering (ICSE'04), May, 2004: 293–303.
92. Zhao W, Zhang L, Liu Y, Sun J, Yang F. SNIAFL: towards a static non-interactive approach to feature location. *ACM Transactions on Software Engineering and Methodologies (TOSEM)* 2006; **15**(2): 195–226.
93. Hill E, Pollock L, Vijay-Shanker K. Exploring the neighborhood with Dora to expedite software maintenance. Proceedings of 22nd IEEE/ACM International Conference on Automated Software Engineering (ASE'07), November, 2007: 14–23.
94. Shao P, Smith, RK. Feature location by IR modules and call graph. Proceedings of ACM Annual Southeast Regional Conference, Clemson, South Carolina, 2009.
95. Ratiu, D, Deissenboeck, F. How programs represent reality (and how they don't). Proceedings of 13th Working Conference on Reverse Engineering (WCRE'06), 2006: 83–92.
96. Ratiu D, Deissenboeck F. From reality to programs and (not quite) back again. Proceedings of 15th IEEE International Conference on Program Comprehension (ICPC'07), Banff, Canada, 2007: 91–102.
97. Hayashi S, Yoshikawa T, Saeki M. Sentence-to-code traceability recovery with domain ontologies. Proceedings of 17th Asia-Pacific Software Engineering Conference (APSEC'10), Sydney, Australia, Nov 30 - Dec 3, 2010b: 385–394.
98. Cubranic D, Murphy GC. Hipikat: recommending pertinent software development artifacts. Proceedings of 25th International Conference on Software Engineering (ICSE'03), Portland, OR, May, 2003: 408–418.
99. Cubranic D, Murphy GC, Singer J, Booth KS. Learning from project history: a case study for software development. Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW'04), Chicago, Illinois, USA, 2004: 82–91.
100. Ratanotayanon S, Choi HJ, Sim SE. Using transitive changesets to support feature location. Proceedings of 25th IEEE/ACM International Conference on Automated Software Engineering (ASE'10), Antwerp, Belgium, Sept 20–24, 2010: 341–344.
101. Lukoit K, Wilde N, Stowell S, Hennessey T. TraceGraph: immediate visual location of software features. Proceedings of 16th IEEE International Conference on Software Maintenance (ICSM'00), Washington DC, USA, 2000: 33–39.
102. Simmons S, Edwards D, Wilde N, Homan J, Groble M. Industrial tools for the feature location problem: an exploratory study. *Journal of Software Maintenance Research and Practice* 2006; **18**(6): 457–474.
103. Wilde N, Simmons S, Edwards D, Pounds L. But where does it DO that? Locating features in a distributed simulation. Proceedings of Fall Simulation Interoperability Workshop, Orlando, Florida, 2002.
104. Egyed A. A scenario-driven approach to trace dependency analysis. *IEEE Transactions on Software Engineering (TSE)* 2003; **29**(2): 116–132. (February)
105. Egyed A, Grünbacher P. Identifying requirements conflicts and cooperation. *IEEE Software* 2004; **21**(6): 50–58. (Nov. 2004)
106. Egyed A, Grünbacher P. Supporting software understanding with automated traceability. *International Journal of Software Engineering and Knowledge Engineering* 2005a; **15**(5): 783–810. (Oct. 2005)
107. Egyed A, Heindl M, Biffl S, Grünbacher P. Determining the cost-quality trade-off for automated software traceability. Proceedings of International Conference on Automated Software Engineering (ASE'05), Long Beach, CA, 2005b: 360–363.
108. Olszak A, Jørgensen BN. Featureous: a tool for feature-centric analysis of Java software. Proceedings of 18th IEEE International Conference on Program Comprehension (ICPC'10), Braga, Portugal, June 30 - July 2, 2010: 44–45.
109. de Alwis B, Murphy GC, Robillard M. A comparative study of three program exploration tools. Proceedings of 15th IEEE International Conference on Program Comprehension, 2007: 103–112.

110. Poshyvanyk D, Petrenko M, Marcus A, Xie X, Liu D. Source code exploration with Google. Proceedings of 22nd IEEE International Conference on Software Maintenance (ICSM'06), Philadelphia, PA, 2006c: 334–338.
111. Poshyvanyk D, Marcus A, Dong Y, Sergeyev A. IRiSS—a source code exploration tool. Proceedings of 21st IEEE International Conference on Software Maintenance (ICSM'05), Budapest, Hungary, September 25–30, 2005: 69–72.
112. Poshyvanyk D, Marcus A, Dong Y. JIRiSS—an Eclipse plug-in for source code exploration. Proceedings of 14th IEEE International Conference on Program Comprehension (ICPC'06), Athens, Greece, June 14–17, 2006b: 252–255.
113. Marcus A, Feng L, Maletic JJ. 3D representations for software visualization. Proceedings of 1st ACM Symposium on Software Visualization (SoftVis'03), San Diego, CA, June 11–13, 2003: 27–36.
114. Robillard MP, Murphy GC. FEAT: a tool for locating, describing, and analyzing concerns in source code. Proceedings of 25th International Conference on Software Engineering (ICSE03), Portland, OR, May 3–10, 2003b: 822–823.
115. Savage T, Dit B, Gethers M, Poshyvanyk D. TopicXP: exploring topics in source code using latent Dirichlet allocation. Proceedings of 26th IEEE International Conference on Software Maintenance (ICSM'10), Timișoara, Romania, September, 2010a: 12–18.
116. Bohnet J, Döellner J. *Analyzing feature implementation by visual exploration of architecturally-embedded call-graphs. Proceedings of International Workshop on Dynamic Systems Analysis (r'06)*. Shanghai: China, 2006a: 41–48.
117. Bohnet J, Döellner J. Visual exploration of function call graphs for feature location in complex software systems. *ACM Symposium on Software Visualization (SOFTVIS'06)*. Brighton, United Kingdom, 2006b: 95–104.
118. Bohnet J, Döellner J. CGA call graph analyzer—locating and understanding functionality within the Gnu compiler collection's million lines of code. Proceedings of 4th IEEE International Workshop on Visualizing Software for Understanding and Analysis (VISSOFT'07), Banff, Alberta, Canada, June 25–26, 2007a: 161–162.
119. Bohnet J, Döellner J. Analyzing dynamic call graphs enhanced with program state information for feature location and understanding. Proceedings of 30th IEEE/ACM International Conference on Software Engineering (ICSE'08), Leipzig, Germany, May 10–18, 2008a: 915–916.
120. Janzen D, Volder K. Navigating and querying code without getting lost. Proceedings of 2nd International Conference on Aspect-Oriented Software Development (AOSD'03), 2003: 178–187.
121. de Alwis B, Murphy GC. Answering conceptual queries with Ferret. Proceedings of Proceedings of the International Conference on Software Engineering (ICSE'08), Leipzig, Germany, 2008: 21–30.
122. Shonle M, Neddenriep J, Griswold W. AspectBrowser for Eclipse: a case-study in plug-in retargeting. Proceedings of OOPSLA workshop on eclipse technology eXchange, 2004: 78–82.
123. Griswold WG. Coping with crosscutting software changes using information transparency. Proceedings of 3rd International Conference on Metalevel Architectures and Separation of Crosscutting Concerns, Kyoto, Japan, September, 2001.
124. Wilde N, Casey C. Early field experience with the software reconnaissance technique for program comprehension. Proceedings of International Conference on Software Maintenance (ICSM'96), 1996: 270.
125. Wilde N, Buckellew M, Page H, Rajlich V. A case study of feature location in unstructured legacy fortran code. Proceedings of European Conference on Software Maintenance and Reengineering (CSMR'01), 2001: 68–76.
126. Ibrahim S, Idris NB, Deraman A. Case study: reconnaissance techniques to support feature location using RECON2. Proceedings of Asia-Pacific Software Engineering Conference (APSEC'03), 2003: 371–378.
127. Marcus A, Rajlich V, Buchta J, Petrenko M, Sergeyev A. Static techniques for concept location in object-oriented code. Proceedings of 13th IEEE International Workshop on Program Comprehension (IWPC'05), St. Louis, Missouri, USA, 2005c: 33–42.
128. Revelle M, Poshyvanyk D. An exploratory study on assessing feature location techniques. Proceedings of 17th IEEE International Conference on Program Comprehension (ICPC'09), Vancouver, British Columbia, Canada, May 17–19, 2009: 218–222.
129. Van Geet J, Demeyer S. Feature location in COBOL mainframe systems: an experience report. Proceedings of International Conference on Software Maintenance (ICSM'09), 2009: 361–370.
130. LaToza TD, Garlan D, Herbsleb JD, Myers BA. Program comprehension as fact finding. Proceedings of The 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The foundations of Software Engineering, 2007: 361–370.
131. Robillard MP, Coelho W, Murphy GC. How effective developers investigate source code: an exploratory study. *IEEE Transactions on Software Engineering (TSE)* 2004; **30**(12): 889–903.
132. Revelle M, Broadbent T, Coppit D. Understanding concerns in software: insights gained from two case studies. Proceedings of 13th IEEE International Workshop on Program Comprehension (IWPC'05), St. Louis, Missouri, USA, May, 2005: 23–32.
133. Carver L, Griswold WG. Sorting out concerns. Proceedings of OOPSLA'99 Workshop on Multi-Dimensional Separation of Concerns, 1999.
134. Coppit D, Painter R, Revelle M. Spotlight: a prototype tool for software plans. Proceedings of International Conference on Software Engineering (ICSE'07), 2007: 754–757.
135. Ko AJ, Aung HH, Myers BA. Eliciting design requirements for maintenance-oriented IDEs: a detailed study of corrective and perfective maintenance tasks. Proceedings of 27th IEEE/ACM International Conference on Software Engineering (ICSE'05), May 15–21, 2005: 126–135.
136. Ko AJ, Myers BA, Coblenz MJ, Aung HH. An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks. *IEEE Transactions on Software Engineering (TSE)* 2006; **32**(12): 971–987. (December)

137. Ko AJ, Meyers BA. Debugging reinvented: asking and answering why and why not questions about program behavior. *Proceedings of International Conference on Software Engineering (ICSE'08)*, Leipzig, Germany, 2008: 301–310.
138. Buchta J, Petrenko M, Poshyvanyk D, Rajlich V. Teaching evolution of open source projects in software engineering courses. *Proceedings of 22nd IEEE International Conference on Software Maintenance (ICSM'06)*, 2006: 136–144.
139. Petrenko M, Poshyvanyk D, Rajlich V, Buchta J. teaching software evolution in open source. *IEEE Computer* 2007; **40**(11): 25–31.
140. Schach S. *Object-Oriented and Classical Software Engineering* (5th edn). McGraw-Hill Higher Education: New York, 2001.
141. Sommerville I. *Software Engineering* (6th edn). Addison-Wesley: New Work, 2001.
142. Wohlin C, Runeson P, Host M, Ohlsson MC, Regnel B, Regnel B, Wesslen A. *Experimentation in Software Engineering: An Introduction*, Amsterdam, Kluwer Academic Press. 1999.
143. Sartipi K, Safyallah H. Dynamic Knowledge Extraction from Software Systems using Sequential Pattern Mining. *International Journal of Software Engineering and Knowledge Engineering (IJSEKE)* 2010; **20**(6): September, 761–782.
144. Saul MZ, Filkov V, Devanbu P, Bird C. Recommending Random Walks. *Proceedings of 11th European Software Engineering Conference held jointly with 15th ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/FSE'07)*, Dubrovnik, Croatia, 2007: 15–24.
145. Binkley D, Lawrie D. Information Retrieval Applications in Software Development. *Encyclopedia of Software Engineering*. Taylor & Francis, 2010.

AUTHORS' BIOGRAPHIES



Bogdan Dit is a Ph.D. candidate in the Computer Science Department at the College of William and Mary and he is a member of the SEMERU research group. He is advised by Dr. Denys Poshyvanyk. Bogdan obtained his B.S. from Babes-Bolyai University, Romania, and his M.S. from Wayne State University under the supervision of Dr. Andrian Marcus. His research interests include software engineering, software maintenance and evolution, mining of software repositories, information retrieval, feature location, and program comprehension. He is a student member of IEEE and ACM.



Meghan Revelle received her Ph.D. degree from the College of William and Mary in Williamsburg, Virginia, in 2010. She was advised by Dr. Denys Poshyvanyk. She received her B.S. in Computer Science from Mary Washington College in 2003, and her M.S. in Computer Science from the College of William and Mary in 2005. Her research interests include software maintenance and program comprehension.



Malcom Gethers is a Ph.D. candidate in the Computer Science Department at the College of William and Mary, where he is a member of the SEMERU research group. He is advised by Dr. Denys Poshyvanyk. Malcom obtained his B.S. from High Point University and his M.S. from the University of North Carolina at Greensboro. His research interests include software engineering, software maintenance and evolution, mining of software repositories, feature location, software measurement, and traceability link recovery and management. He is a student member of IEEE, ACM, and SIGSOFT.



Denys Poshyvanyk is an Assistant Professor at the College of William and Mary in Virginia. He received his Ph.D. degree in Computer Science from Wayne State University in 2008. He also obtained his M.S. and M.A. degrees in Computer Science from the National University of Kyiv-Mohyla Academy, Ukraine, and Wayne State University in 2003 and 2006, respectively. Since 2010, he has been serving on the steering committee of the International Conference on Program Comprehension (ICPC). He serves as a program co-chair for the 18th and 19th International Working Conference on Reverse Engineering (WCRE 2011 and WCRE 2012). He will also serve as a program co-chair for the 21st IEEE International Conference on Program Comprehension (ICPC 2013). His research interests include software engineering, software maintenance and evolution, program comprehension, reverse engineering, software repository mining, source code analysis, and metrics. He is a member of IEEE and ACM.