

Reflective feature location: knowledge in mind meets information in system

Xin PENG^{1,2*}, Zhengchang XING³, Sen PAN^{1,2}, Wenyi QIAN^{1,2},
Václav RAJLICH⁴ & Wenyun ZHAO^{1,2}

¹*School of Computer Science, Fudan University, Shanghai 201203, China;*

²*Shanghai Key Laboratory of Data Science, Fudan University, Shanghai 201203, China;*

³*Research School of Computer Science, Australian National University, Acton ACT 2601, Australia;*

⁴*Department of Computer Science, Wayne State University, Detroit, MI 48202, USA*

Received August 1, 2016; accepted October 8, 2016; published online February 10, 2017

Abstract Locating code entities relevant to a feature (i.e., feature location) is an important task in software maintenance. Feature location is challenging due to the information gap between the developer's knowledge about a feature and the feature's implementation in the system. In this paper, we present a reflective feature location approach ReFLex to bridge this gap. ReFLex automatically computes a reflexion model between the developer's logical view of a feature and the feature's implemented view reverse-engineered from the code. It provides interactive reflection support for the developer to refine the logical view and the feature location results such that the logical view and the implemented view converge gradually in an iterative feature location process. We have implemented our approach as an Eclipse plugin and investigated the benefits and challenges in reflective feature location through a user study.

Keywords feature location, information seeking, reflexion model

Citation Peng X, Xing Z C, Pan S, et al. Reflective feature location: knowledge in mind meets information in system. *Sci China Inf Sci*, 2017, 60(7): 072102, doi: 10.1007/s11432-015-0860-9

1 Introduction

A software feature is defined by requirements and accessible to users [1]. Software maintenance tasks are usually specified by features, such as adding new features, extending existing features, and fixing bugs in features. To accomplish a software maintenance task, developers must identify code entities relevant to the feature. This activity is termed as feature or concept location [1–4].

Among many techniques proposed to support feature location, Information Retrieval (IR) based code search has been widely used in practice [1]. With IR-based code search, developers describe a feature using a set of keywords (i.e., a feature query). The keywords can be from feature descriptions (e.g., change request, bug report) or domain knowledge [5]. Code search engines use IR techniques to compute the relevance score between the feature query and the code entities. Relevant code entities are ranked in a list for inspection.

* Corresponding author (email: pengxin@fudan.edu.cn)

Using IR-based code search the developers are often faced with two difficulties. First, the developer may infer from the feature description that the feature should involve several correlated logical entities. However, it is difficult to express this understanding in a set of keywords. Second, the ranked list of code entities provides no ways to identify and reconcile the discrepancies between the developer's high-level understanding of the feature and the feature implementation, for example, whether the code entities have the dependencies they are or are not supposed to have.

Query reformulation techniques can help developers rewrite the query based on the user-relevance feedback on search results [6,7] or the linguistic properties of feature query and code entities [8,9]. Concept lattice [10] or multiple syntactic and semantic facets [7] can be used to facilitate the inspection of query results. Although useful, these techniques lack systematic ways to bridge the gap and reconcile the discrepancies between the developer's high-level understanding of the feature and the feature implementation in the system.

In this paper, we propose a reflective feature location approach (ReFLex) to tackle the above two difficulties. ReFLex allows developers to express their understanding and knowledge about a feature in a high-level logical view. It uses a structure-based matching algorithm to map the logical view of the feature and the implemented view of the system. Based on the mapping results, ReFLex identifies and highlights the discrepancies between the two views in a reflexion model. It provides interactive reflection support for developers to refine the logical view and/or the mapping results. As a result, the logical view and the implemented view of the feature converge gradually. Code entities relevant to the feature emerge from this reflective feature location process.

We have implemented our approach as an Eclipse plugin. To evaluate the benefits and challenges in the ReFLex process and the tool support, we conducted a user study in which two groups of nine developers performed three feature location tasks using our ReFLex tool and the standard Eclipse IDE respectively. Our results show that the ReFLex process and tool work better for coarse-grained features involving several logical entities and complex business logic. The ReFLex users demonstrated feature location behaviors that were different from those of the Eclipse users. We observed effective strategies for reflective feature location.

2 The approach

In this section, we present an overview of our ReFLex approach and then detail the underlying techniques.

2.1 Overview

Figure 1 presents an overview of our ReFLex approach. ReFLex assumes that a system expert can provide the metamodel of logical view, the metamodel of the implemented view of the system, and a metamodel mapping of types of entities and relations between the two views. It assumes that the implemented view can be reverse-engineered from the system source code according to the implemented-view metamodel.

The ReFLex process consists of 5 steps. Steps 1 and 4 involve interaction with the developer. To start a feature location process, ReFLex asks the developer to construct an initial logical view [11] (i.e., an instance of logical-view metamodel) of the feature (Step 1). This logical view captures a set of correlated logical entities of the feature. The initial logical view of the feature may not agree with the feature implementation. During feature location process, the developer can refine the logical view based on the reflexion model.

Given the logical view of the feature, and the implemented view of the system and the metamodel mapping between the two views, ReFLex uses a structure-based matching algorithm to determine the relevance of logical entities and code entities based on their lexical and structural similarities (Step 2). Based on the relevant code entities identified for each logic entity, ReFLex computes a reflexion model (Step 3) that highlights the discrepancies between the logical view of the feature and the implemented view of the system. The developer inspects the reflexion model, and refines the logical view of the feature and/or the current mapping results to reduce the discrepancies in the reflexion model (Step 4). In Step

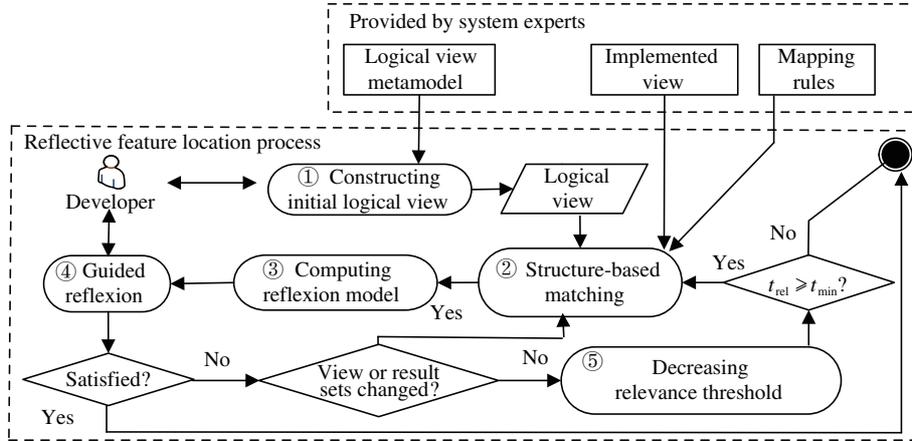


Figure 1 Reflective feature location process.

4 the developer can freeze some logical entities. ReFLex will not update the relevant-entities set of these frozen logical entities. Freezing logical entities allows the developer to focus on certain aspects of a feature by freezing other aspects. If the developer is satisfied with current feature location results (e.g., when the logical view and the implemented view converge), he can finish the process. The relevant-entities set of all the logical entities constitute the code entities relevant to the feature. Otherwise, ReFLex proceeds to the next iteration. If the logical view of the feature and/or the relevant-entities set of some logical entities have been changed in Step 4, ReFLex calls the structure-based matching algorithm to remap the logical view of the feature and the implemented view of the system, and update the reflexion model accordingly. Otherwise, ReFLex decreases the relevance threshold t_{rel} (Step 5) and then calls the structure-based matching algorithm.

2.2 Metamodels and metamodel mapping

ReFLex maps the logical view of a feature and the implemented view of the system and computes a reflexion model [12] between the two views. We define here the metamodel of the logical view, the metamodel of the implemented view, and the mapping between the two metamodels.

A software reflexion model maps a high-level model of interest and a source model from the source code based on a declarative mapping between the two models. The high-level model represents a kind of logical structure of the software system. For example, it can be a design model described in design documents or reflect the developers understanding of the logical structure of the current implementation. The reflexion model schema [12] defines HLMENTITY and SMENTITY, which define the types of high-level model entities and source model entities respectively, HLMRELATION and SMRELATION, which define the types of relations over high-level model entities and source model entities respectively, and HLMTUPLE and SMTUPLE, which define the types of tuples in the HLMRELATION and SMRELATION. The declarative mapping defines the correspondences between HLMENTITY and SMENTITY, and between HLMTUPLE and SMTUPLE. We focus on object-oriented softwares in this paper. However, our approach is not limited to object-oriented softwares because reflexion model schema is essentially an entity-relation schema that can be applied to many types of softwares.

The logical view of a feature reflects the developer’s high-level understanding and knowledge about a feature. Depending on the feature location needs, a system expert may define the metamodel of logical view as shown in Figure 2(a). A logical view of the feature may consist of four types of HLMENTITY. A Window represents a boundary object (e.g., a form or a dialog) that the user interacts with. An Operation performs certain application logic. An EntityObject represents a piece of data, while a DataStore is a subtype of EntityObject that represents a piece of persistent data. An entity has two attributes: identifier (id) and description (desc). The description stores the natural language description of the entity.

Different types of entities in the logical view have different types of usage relations (i.e., HLMTUPLE).

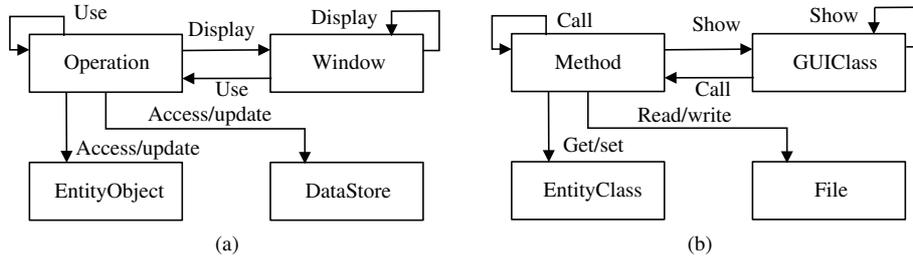


Figure 2 Examples of metamodels. (a) Logical-view metamodel; (b) implemented-view metamodel.

Table 1 Metamodel mapping

| Entity of logical view | Entity of implemented view |
|--------------------------------------|------------------------------|
| Operation | Method |
| Window | GUIClass |
| EntityObject | EntityClass |
| DataStore | File |
| Relation of logical view | Relation of implemented view |
| Operation use Operation | Method call Method |
| Operation display Window | Method show GUIClass |
| Operation access/update EntityObject | Method get/set EntityClass |
| Operation access/update DataStore | Method read/write File |
| Window use Operation | GUIClass call Method |
| Window display Window | GUIClass show GUIClass |

An Operation can use another Operation, display a Window, and access/update an EntityObject or a DataStore. A Window can use an Operation and display another Window.

The implemented view of the system describes a source model (such as a call graph or an inheritance hierarchy) from the source code. Depending on the feature location needs and the knowledge of system implementation, the system expert may define the metamodel of implemented view as shown in Figure 2(b). The implemented view of the system may consist of four types of SMENTITY. A GUIClass represents a GUI widget (e.g., a subclass of JComponent of Java Swing). A Method represents a method of a class. An EntityClass represents a class with only attributes or getter/setter methods. A File represents a piece of external data the system uses. An entity has two attributes: name and code, which are reverse-engineered from the source code.

Different types of entities in the implemented views have different types of usage dependencies (i.e., SMTUPLE). A Method can call another Method, show a GUIClass, and get/set an EntityClass or read/write a File/Database. A GUIClass can call a Method and show another GUIClass.

Given the metamodels of the logical view and the implemented view, the system expert can define the mapping between the two metamodels. Table 1 shows the mapping between the two metamodels in Figure 2(a) and (b). According to this mapping rule, ReFLex locates user interface (Window) and data object (EntityObject) implementations at class level (GUIClass and EntityClass), while it locates application logic (Operation) implementations at method level (Method).

2.3 Constructing initial logical view (Step 1)

To locate a feature’s implementation, the developer constructs an initial logical view of the feature using a graphical modeling tool. For example, the feature “Reloading From Disk” of JEdit (<http://www.jedit.org>) checks if an opened file was modified by another application when the view of the file is brought to the foreground. Even with limited knowledge of the feature implementation, a developer can identify some key abstractions for this feature by analyzing the feature’s use case description or trying out the feature with the software. For example, an initial logical view of the feature “Reloading From Disk” may consist of an operation checking whether any opened buffer was modified, a user interface (UI) object prompting

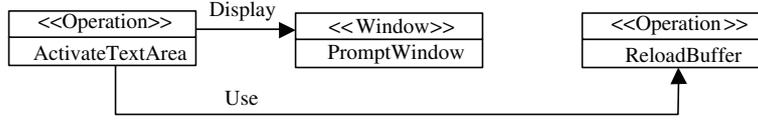


Figure 3 A logical view of “Reloading From Disk”.

whether to reload a file, and an operation reloading the file.

Figure 3 presents an initial logical view of the JEdit feature “Reloading From Disk” that one of the developers defined in our user study. Based on the feature description and the use of JEdit software, the developer believed that the implementation of the feature should include an ActivateTextArea operation to handle the activation of the text area that is brought to front; ActivateTextArea should display a PromptWindow for user confirmation, and use ReloadBuffer operation to reload the file.

The initial logical view of the feature may not agree with the system implementation, especially in early iterations of feature location. The discrepancies were later identified and fixed through reflective feature location process.

2.4 Structure-based matching (Step 2)

ReFLex’s structure-based matching algorithm is based on the computation of lexical and structural similarities.

2.4.1 Similarity metrics

Given a logical entity d and a code entity p , ReFLex encodes the description of d and the code fragment of p in the TF/IDF (Term Frequency/Inversed Document Frequency) vectors V_d and V_p , based on the vector space model computed from the code entities in the implemented view of the system. It computes the lexical similarity $S_{lex}(d, p)$ as the cosine similarity between the TF/IDF vectors V_d and V_p , i.e., $S_{lex}(d, p) = \frac{\sum_{i=1}^n V_d[i] \times V_p[i]}{\sqrt{\sum_{i=1}^n V_d[i]^2} \times \sqrt{\sum_{i=1}^n V_p[i]^2}}$ ($0 \leq S_{lex}(d, p) \leq 1$). We remark that our approach is independent of any specific computation methods of lexical similarity. Other lexical similarity computation methods, such as those based on the words extracted from class and method names as well as from class comments [13], can also be used in our approach.

Given a logical entity d and a code entity p , ReFLex computes their structural similarity $S_{stru}(d, p)$ based on the established mappings between the neighboring logical entities of d (denoted as $NGB(d)$) and the neighboring code entities of p (denoted as $NGB(p)$). Let $RS(d)$ be the relevant-entities set of d . We define $CVG(d, p)$ as the convergent neighboring logical-entity set of d , i.e., $CVG(d, p) = \{d' \in NGB(d) | \exists p' \in NGB(p) \& p' \in RS(d') \& RT(d, d') \leftrightarrow RT(p, p')\}$. $RT(d, d') \leftrightarrow RT(p, p')$ indicates the relation between d and d' and the relation between p and p' can be mapped according to the metamodel mapping. The structural similarity $S_{stru}(d, p)$ is defined as the percentage of d ’s convergent neighboring logical entities in all the d ’s neighboring logical entities, i.e., $S_{stru}(d, p) = \frac{|CVG(d, p)|}{|NGB(d)|}$ ($0 \leq S_{stru}(d, p) \leq 1$).

The overall similarity between a logical entity d and a code entity p is defined as the average of their lexical and structural similarities, i.e., $S(d, p) = (S_{stru}(d, p) + S_{lex}(d, p))/2$. In the first iteration of the ReFLex process $S(d, p)$ is set to $S_{lex}(d, p)$ as the relevant-entities sets of all the logical entities are empty.

2.4.2 Matching algorithm

The matching algorithm identifyRelevant (see Algorithm 1) takes four parameters: logical view lv , implemented view iv , candidate-entity threshold t_{cand} , and relevant-entity threshold t_{rel} . During the feature location process, each logical entity maintains a set of relevant code entities ($d.rels$) identified in previous iterations and a set of candidate code entities ($d.cands$) from which identifyRelevant identifies new relevant entities for the logical entity. The implemented view maintains a set of yet-unmapped entities ($iv.unmappeds$) which initially contains all the code entities in the implemented view.

For each logical entity $d \in lv.entities$ whose description ($d.desc$) has been changed in the previous iteration (the description of the newly introduced logical entities is considered as changed), identifyRelevant

Algorithm 1 Identify relevant code entities

```

Require: identifyRelevant(lv, iv,  $t_{\text{cand}}$ ,  $t_{\text{rel}}$ )
1: for each  $d \in \text{lv.entities} \ \& \ d.\text{descchanged}$  do
2:    $\text{iv.unmappeds} = \text{iv.unmappeds} \cup d.\text{rels}$ 
3:    $d.\text{rels} = \{\}; d.\text{cands} = \{\}$ 
4: end for
5: for each  $d \in \text{lv.entities} \ \& \ d.\text{descchanged}$  do
6:   for each  $p \in \text{iv.unmappeds}$  do
7:     if  $S_{\text{lex}}(d, p) \geq t_{\text{cand}}$  then
8:        $d.\text{cands} = d.\text{cands} \cup \{p\}$ 
9:     end if
10:  end for
11: end for
12: for each  $(d, p) : d \in \text{lv.entities} \ \& \ p \in d.\text{rels}$  do
13:    $(d, p).\text{sim} = S(d, p)$ 
14: end for
15: for each  $d \in \text{lv.entities} \ \& \ d.\text{unfrozen}$  do
16:    $d.\text{newrels} = \{\}$ 
17:   for each  $p \in d.\text{cands}$  do
18:     if  $S(d, p) \geq t_{\text{rel}}$  then
19:        $d.\text{newrels} = d.\text{newrels} \cup \{p\}$ 
20:     end if
21:   end for
22:    $\text{iv.unmappeds} = \text{iv.unmappeds} - d.\text{newrels}$ 
23:    $\text{removeCandidateEntities}(\text{lv.entities}, d.\text{newrels})$ 
24: end for
25: for each  $d \in \text{lv.entities} \ \& \ d.\text{unfrozen}$  do
26:    $d.\text{rels} = d.\text{rels} \cup d.\text{newrels}$ 
27: end for

```

invalidates the relevant-entities set $d.\text{rels}$ and the candidate-entities set $d.\text{cands}$ of d (lines 2–5). It adds code entities in $d.\text{rels}$ back to the set of yet-unmapped entities of the implemented view (iv.unmappeds), and set $d.\text{rels}$ to empty set. It also set $d.\text{cands}$ to empty set.

Then, `identifyRelevant` refinds candidate entities for the logical entity whose description has been changed (lines 6–12). If the lexical similarity of a logical entity d and a yet-unmapped code entity p (i.e., $S_{\text{lex}}(d, p)$) is above the candidate-entity threshold (t_{cand}), `identifyRelevant` adds the code entity p to $d.\text{cands}$. Candidate-entities set of different logical entities may overlap. The union of candidate entities of all the logical entities is a subset of all the yet-unmapped entities in the implement view.

Next, for each existing mapping between a logical entity and a relevant code entity (d, p) , `identifyRelevant` recomputes its overall similarity (lines 13–15). This is because the structural reflexion may increase or decrease the overall similarity of existing mappings. `identifyRelevant` does not remove the existing relevant code entities of a logical entity, even if the updated overall similarity of these code entities and the logical entity is below t_{rel} . The ReFLex tool support can highlight the similarity changes for the developer to make the decision.

Finally, `identifyRelevant` identifies new relevant entities for each unfrozen logical entity (lines 16–25). If the overall similarity between the d and a candidate code entity $p \in d.\text{cands}$ (i.e., $S(d, p)$) is above the relevant-entity threshold (t_{rel}), `identifyRelevant` adds the code entity p to a temporary relevant-entities set of d ($d.\text{newrels}$). If a code entity is identified as relevant to a logical entity, the code entity will be removed from the candidate-entities set of all the logical entities and from the yet-unmapped-entities set of the implemented view. Therefore, `identifyRelevant` produces a one-to-many mapping between the logical entities and the code entities. After processing all the unfrozen logical entities, `identifyRelevant` adds the newly identified relevant-entities of these logical entities to their relevant-entities set (lines 26–28).

2.5 Computing reflexion model (Step 3)

ReFLex computes a reflexion model between the logical view of the feature and the implemented view of the system by comparing the relations between the logical entities defined in the logical view (i.e., expected relations) and the relations between the logical entities induced from their relevant code entities (i.e., induced relations). The reflexion model consists of three kinds of edges [12], i.e., Convergences, Divergences, and Absences. Convergences identify the set of relations where the logical view agrees with the implemented view (i.e., expected relations match induced relations). Divergences identify the set of

induced relations from the relevant code entities that are missing in the logical view (i.e., no corresponding expected relations). Absences identify the set of expected relations in the logical view that are missing in the relevant code entities (i.e., no corresponding induced relations).

2.6 Guided reflection (Step 4)

The reflexion model can be visualized for the developer to refine the logical view and/or the current mapping results.

2.6.1 Structural reflexion

Structural reflexion allows the developer to reduce the structural discrepancies (i.e., divergences and absences edges) between the logical view and the implemented view. The developer can refine the logical view by adding, removing, merging or splitting logical entities. He can also add, remove, or change relations between logical entities. These refinements of the logical view can remove Divergences and Absences in the reflexion model.

The developer can select one or more Divergences edges and remove the code entities in the relevant-entities set of a logical entity that cause Divergences edges. He can also select one or more Convergences edges and remove all the code entities in the relevant-entities set of a logical entity that have no contributions to Convergences edges.

The developer can examine the source code of relevant code entities of a logical entity. If he believes a code entity is irrelevant, he can remove the code entity from the relevant-entities set of the logical entity. This helps remove Divergences. The removed code entity will be added back to the candidate-entities set of other logical entities and the yet-unmapped-entities set of the implemented view. The developer can also examine the source code of candidate code entities of a logical entity. If he believes a candidate code entity is relevant, he can add the candidate code entity to the relevant-entities set of the logical entity. This helps remove Absences. The added code entity will be removed from the candidate-entities set of other logical entities and the yet-unmapped-entities set of the implemented view.

2.6.2 Descriptive reflexion

Descriptive reflexion allows developers to refine the description of a logical entity by examining terms and topics involved in its current relevant code entities. Tag cloud and topics (mined from code using Latent Dirichlet Allocation technique [14]) allow the developer to identify important terms/topics in the relevant code entities of a logical entity. Investigating the terms/topics of a logical entity and its neighboring logical entities can also give the developer insights into the roles and relations of a set of logical entities. The developer can refine the description of a logical entity using terms/topics from its neighboring entities.

2.7 Decreasing relevance threshold (Step 5)

The identifyRelevant algorithm considers a logical entity and a code entity as relevant if their similarity is above the relevant-entity threshold t_{rel} . ReFLex uses an incrementally-decreasing-threshold strategy during reflective feature location. This strategy has been proven effective in improving feature location results [15,16]. ReFLex allows the developer to set a high initial t_{rel} to ensure the high-quality of mappings between the logical entities and the code entities in the first few iterations. The t_{rel} will be decreased by a pace in subsequent iterations to locate more code entities with lower similarities. Once the t_{rel} is lower than the minimal relevance threshold t_{min} , ReFLex ends feature location process with the current relevant-entities set of the logical entities as the final results.

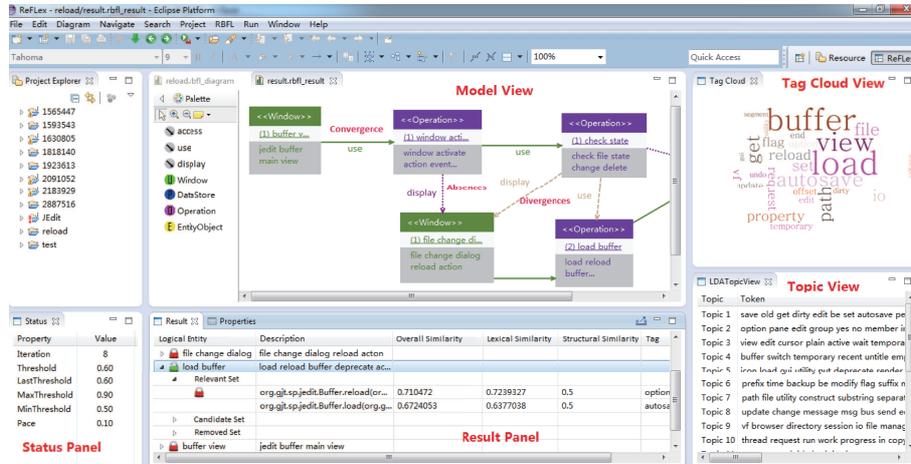


Figure 4 (Color online) The ReFLex tool.

3 Tool implementation

Figure 4 presents our ReFLex tool. We implemented the ReFLex tool using Eclipse Modeling Framework (EMF) and Graphical Editing Framework (GEF). The GUI of the ReFLex tool consists of: model view, result panel, tag cloud view, topic view, and status panel.

The model view implements a GEF-based graphical modeling tool for the developer to construct the initial logical view and refine the reflexion model. Convergences, Divergences, and Absences edges in the reflexion model are shown as solid, dashed, and dotted lines. In the model view, each logical entity is represented as a node with three compartments showing its type, identifier and description. The number before the identifier indicates the size of relevant-entities set. When a logical entity is selected, a tag cloud is generated using WordCram¹⁾ for all its current relevant code entities. A set of topics are also mined using JGibbsLDA²⁾.

The result panel shows in a tree list view of relevant-entities set, candidate-entities set, and removed-entities set of all the logical entities. The developer can expand a logical entity and inspect its result sets. The view displays similarity metrics of the relevant and candidate code entities of the logical entity. The developer can double-click a code entity to open its source code. The developer can select a logical entity and inspect the sub-reflexion-model of the selected entity and its neighboring entities in the pop-up “Filter by Interaction” window. In this window, he can remove irrelevant code entities of the selected logical entity by Convergences or Divergences edges.

During the feature location process, a set of status parameters is shown in the status panel, including the number of iterations and various threshold values.

4 Evaluation

We conducted a user study to investigate the benefits and challenges inherent in reflective feature location process, and gather user impressions and observations on our tool prototype. We investigated the following questions.

- Q1. What types of features are supported well by ReFLex, and which types are not?
- Q2. How does the ReFLex tool change the developers’ behaviors during feature location tasks?
- Q3. What are effective strategies for reflective feature location?

4.1 Experiment design

We first describe the design of our user study.

1) WordCram Website: <http://wordcram.org/>.

2) JGibbsLDA Website: <http://jgibblda.sourceforge.net/>.

4.1.1 *Participants*

Our study used matched-subjects design. We recruited 18 participants (16 master students and 2 Ph.D. students) from School of Computer Science, Fudan University. These participants were matched in pairs based on their programming capability and experience and then randomly allocated to experimental group or control group. The experimental group G_1 used the ReFLex tool to perform feature location tasks. The control group G_2 used standard Eclipse IDE to perform the same set of tasks. We chose standard Eclipse IDE as baseline tool to compare with the ReFLex tool because we are interested in behavioral changes enabled by ReFLex during feature location tasks, compared with “normal” behavior in commonly used IDEs.

4.1.2 *Tasks*

In this study, we used feature location tasks from the JEdit benchmarks [1]. The benchmark consists of 34 features location tasks for the JEdit text editor and the ground-truth code entities of these features. These 34 features were extracted from issue tracking system of the JEdit. The ground-truth was established by analyzing the commits fixing the relevant issues. We chose JEdit as the subject system because all the participants have general knowledge about text editing features, even though they do not know the implementation details of JEdit. We chose three feature-enhancement tasks from the benchmark:

Task 1. HyperSearch reports the occurrences of the string. The result of HyperSearch will be shown in a window, including the number of occurrences of the string;

Task 2. Save the editmode for recent files. In the Buffer Options window, the edit mode of current opened file could be set and saved;

Task 3. Provide a way of reloading changed buffers without prompting. Reloading settings could be changed under the General tag in the Global Options window. And when the editor pane is activated again, the opened file will be reloaded with or without prompting window according to reloading settings.

These three tasks have different characteristics. Task 1 involves only internal application logic. The ground-truth of Task 1 includes 4 classes and 8 methods. Task 2 involves both application logic and data object. The ground-truth of Task 2 includes 6 classes and 9 methods. Task 3 involves user interface, application logic and data object. The ground-truth of Task 3 includes 9 classes and 18 methods.

The description of these tasks was written in a way that a developer with limited knowledge of JEdit could understand. The relevant-entities of the features are medium-sized in scope that could be identified within limited time.

4.1.3 *Metamodels and reverse-engineering tool*

Acting as the system expert, we defined the metamodels of the logical view and the implemented view for Java Swing-based applications (see Figure 2) as well as the mapping between the two metamodels (see Table 1). We spent a half day on the discussion and definition of the metamodels and mapping rules based on our knowledge of Java Swing and understanding of the design of JEdit. We tested our metamodels and mapping rules based on some features and made several rounds of changes on the granularities of the entities. We think the efforts required to create the metamodels and mapping rules are acceptable if the system experts have sufficient understanding of the subject system, since the metamodels and mapping rules apply to most of the features of the system in a long time.

We developed an Eclipse JDT/DOM API based tool to reverse-engineer a class model from the source code of a Java Swing-based application. The tool labels JComponent subclasses as GUIClasses and classes with only attributes or getter/setter methods as EntityClasses. The tool reverse-engineers Method call Method and Method read/write Attribute relations from the source code. It aggregates the elementary relations into method-class and class-class relations. The tool extracts the name and code fragment of a code entity. It converts code fragment into a bag of words in three steps: tokenization, stop word (e.g., if, while) removal, and stemming. Given bags of words of all the code entities, the tool builds a Vector

Table 2 Performance of the ReFLex group (G_1)

| | Task 1 | | | Task 2 | | | Task 3 | | |
|------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| | P | R | F | P | R | F | P | R | F |
| P1 | 0.50 | 0.63 | 0.60 | 0.50 | 0.78 | 0.70 | 0.48 | 0.61 | 0.58 |
| P2 | 0.50 | 0.38 | 0.39 | 0.62 | 0.89 | 0.82 | 0.73 | 0.44 | 0.48 |
| P3 | 0.04 | 0.25 | 0.13 | 0.41 | 0.78 | 0.66 | 0.83 | 0.56 | 0.60 |
| P4 | 0.24 | 0.63 | 0.47 | 0.38 | 0.67 | 0.58 | 0.52 | 0.72 | 0.67 |
| P5 | 0.11 | 0.13 | 0.12 | 0.54 | 0.78 | 0.71 | 0.50 | 0.61 | 0.59 |
| P6 | 0.13 | 0.25 | 0.21 | 0.39 | 0.78 | 0.65 | 0.17 | 0.28 | 0.25 |
| P7 | 0.29 | 0.25 | 0.26 | 0.12 | 0.22 | 0.19 | 0.55 | 0.61 | 0.60 |
| P8 | 0.00 | 0.00 | 0.00 | 0.33 | 0.33 | 0.33 | 0.53 | 0.50 | 0.51 |
| P9 | 0.22 | .25 | .24 | 0.00 | 0.00 | .00 | 0.56 | 0.50 | 0.51 |
| Ave. | 0.23±0.17 | 0.31±0.20 | 0.27±0.18 | 0.36±0.19 | 0.58±0.30 | 0.52±0.26 | 0.54±0.17 | 0.54±0.12 | 0.53±0.11 |

Space Model (using Lucene API³) in which a code entity is represented as a TF/IDF vector [17].

4.1.4 Procedure

In this study, the ReFLex tool uses the following parameters: the initial relevance threshold $t_{\text{init}} = 0.9$, the pace for decreasing the relevance threshold $\text{pace} = 0.1$, the minimal relevance threshold for stopping the feature location process $t_{\text{min}} = 0.5$, the threshold for selecting candidate elements $t_{\text{cand}} = 0.2$. This set of parameters have been proven effective in the evaluation of our Iterative Context-Sensitive Feature Location algorithm [16].

Before the experiment we provided a 30-minute tutorial of the JEdit software in which we introduced and demonstrated the main features of the JEdit and its package structure. Furthermore, we provided a one-hour tutorial on the ReFLex tool to the experimental group. Because all the participants are familiar with Eclipse IDE, we did not provide tutorial and training on Eclipse IDE. We explained the concept of reflexion model and demonstrated the main features of the ReFLex tool. The tutorial did not introduce any specific style of feature location process using the ReFLex tool.

Both groups (G_1 and G_2) were given the same set of tasks and were requested to accomplish the tasks in the same order from Task 1 to Task 3. For each task, the participants were given a description of the feature. They were requested to identify as many relevant-entities as possible in 30 min. They were required to run a full-screen recorder throughout the experiment. We instrumented the ReFLex tool to collect tool usage statistics during reflective feature location process. After the experiment, we interviewed the two groups to collect their feedback on the tasks and the tool usage.

4.2 Results: performance improvement (Q1)

We evaluated the participants' performance on the three tasks by precision, recall and F-measure. Given a feature f , precision (P_f) is the percentage of correctly reported relevant code entities for the feature; recall (R_f) is the percentage of actually relevant code entities of the feature being reported. The F-measure of the feature f is computed as $F_f = (1 + b^2)/(1/P_f + b^2/R_f)$ to reflect a weighted average of precision and recall. Following the treatment in [18], we set b to 2, i.e., recall is considered four times as important as precision, because finding missing results is more difficult than removing incorrect results. Tables 2 and 3 summarize the participants' performance. In Tables 2 and 3, P, R, F denote precision, recall, and F-measure, respectively; P1 to P18 denote the 18 participants.

We introduced the following null and alternative hypotheses to evaluate how different the performance of experimental group (G_1) and control group (G_2) is.

H0. The primary null hypothesis is that there is no significant difference between the performance of G_1 and G_2 .

3) <http://lucene.apache.org/>.

Table 3 Performance of the eclipse IDE group (G_2)

| | Task 1 | | | Task 2 | | | Task 3 | | |
|------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| | P | R | F | P | R | F | P | R | F |
| P10 | 0.50 | 0.50 | 0.50 | 0.29 | 0.22 | 0.23 | 0.50 | 0.28 | 0.30 |
| P11 | 0.33 | 0.25 | 0.26 | 0.50 | 0.22 | 0.25 | 0.57 | 0.22 | 0.25 |
| P12 | 0.50 | 0.13 | 0.15 | 0.33 | 0.11 | 0.13 | 0.50 | 0.06 | 0.07 |
| P13 | 0.38 | 0.38 | 0.38 | 0.00 | 0.00 | 0.00 | 0.56 | 0.28 | 0.31 |
| P14 | 0.33 | 0.25 | 0.26 | 0.13 | 0.11 | 0.11 | 0.40 | 0.33 | 0.34 |
| P15 | 0.33 | 0.13 | 0.14 | 0.10 | 0.11 | 0.11 | 0.31 | 0.22 | 0.24 |
| P16 | 0.07 | 0.13 | 0.11 | 0.13 | 0.11 | 0.11 | 0.33 | 0.22 | 0.24 |
| P17 | 0.20 | 0.13 | 0.14 | 0.00 | 0.00 | 0.00 | 0.50 | 0.11 | 0.13 |
| P18 | 0.17 | 0.13 | 0.13 | 0.00 | 0.00 | 0.00 | 0.33 | 0.11 | 0.13 |
| Ave. | 0.31±0.14 | 0.22±0.13 | 0.23±0.13 | 0.16±0.16 | 0.10±0.08 | 0.11±0.09 | 0.44±0.10 | 0.20±0.09 | 0.22±0.09 |

Table 4 Results of Wilcoxon tests of the null hypothesis H_0 , for precision, recall, and F-measure. The following measurements are reported: minimum value, maximum value, median, means (μ), variance (σ^2), degrees of freedom (DF), pearson correlation coefficient (PC), Z statistics, and statistical significance (p).

| Task | Var | Approach | Samples | Min | Max | Median | μ | σ^2 | DF | PC | Z | p | Decision |
|--------|-----|----------|---------|------|------|--------|-------|------------|----|-------|--------|-------|----------|
| Task 1 | P | ReFLex | 9 | 0.00 | 0.50 | 0.22 | 0.23 | 0.03 | 8 | 0.148 | -1.120 | 0.156 | Accept |
| | | Eclipse | 9 | 0.07 | 0.50 | 0.33 | 0.31 | 0.02 | 8 | | | | |
| | R | ReFLex | 9 | 0.00 | 0.63 | 0.25 | 0.31 | 0.04 | 8 | 0.816 | -1.732 | 0.074 | Accept |
| | | Eclipse | 9 | 0.13 | 0.50 | 0.13 | 0.22 | 0.02 | 8 | | | | |
| | F | ReFLex | 9 | 0.00 | 0.60 | 0.24 | 0.27 | 0.04 | 8 | 0.811 | -0.770 | 0.248 | Accept |
| | | Eclipse | 9 | 0.11 | 0.50 | 0.15 | 0.23 | 0.02 | 8 | | | | |
| Task 2 | P | ReFLex | 9 | 0.00 | 0.62 | 0.39 | 0.36 | 0.04 | 8 | 0.630 | -2.380 | 0.008 | Reject |
| | | Eclipse | 9 | 0.00 | 0.50 | 0.13 | 0.16 | 0.03 | 8 | | | | |
| | R | ReFLex | 9 | 0.00 | 0.89 | 0.78 | 0.58 | 0.10 | 8 | 0.637 | -2.585 | 0.004 | Reject |
| | | Eclipse | 9 | 0.00 | 0.22 | 0.11 | 0.10 | 0.01 | 8 | | | | |
| | F | ReFLex | 9 | 0.00 | 0.81 | 0.65 | 0.52 | 0.08 | 8 | 0.655 | -2.521 | 0.004 | Reject |
| | | Eclipse | 9 | 0.00 | 0.25 | 0.11 | 0.11 | 0.01 | 8 | | | | |
| Task 3 | P | ReFLex | 9 | 0.17 | 0.83 | 0.53 | 0.54 | 0.03 | 8 | 0.557 | -1.599 | 0.064 | Accept |
| | | Eclipse | 9 | 0.31 | 0.57 | 0.50 | 0.44 | 0.01 | 8 | | | | |
| | R | ReFLex | 9 | 0.28 | 0.72 | 0.56 | 0.54 | 0.02 | 8 | 0.296 | -2.675 | 0.002 | Reject |
| | | Eclipse | 9 | 0.06 | 0.33 | 0.22 | 0.20 | 0.01 | 8 | | | | |
| | F | ReFLex | 9 | 0.25 | 0.67 | 0.58 | 0.53 | 0.02 | 8 | 0.138 | -2.666 | 0.002 | Reject |
| | | Eclipse | 9 | 0.07 | 0.35 | 0.24 | 0.22 | 0.01 | 8 | | | | |

H1. An alternative hypothesis to H_0 is that there is significant difference between the performance of G_1 and G_2 .

We used Wilcoxon’s matched-pairs signed-ranked tests to evaluate the null hypothesis H_0 in terms of precision, recall and F-measure for the three tasks with $p = 0.05$. Table 4 presents the results of these tests. Based on the results, we accept the null hypothesis H_0 for the precision, recall and F-measure of Task 1. That is, there is no significant difference between the performance of G_1 and G_2 in Task 1. Tables 2 and 3 show that the ReFLex users achieved better recall but worse precision in Task 1 than the Eclipse IDE users. The performance of the two groups is very close in terms of F-measure. We reject the null hypothesis H_0 for the precision, recall and F-measure of Task 2 and for the recall and F-measure of Task 3. Therefore, we accept the alternative hypothesis H_1 , i.e., there is significant difference between the performance of G_1 and G_2 for the precision, recall and F-measure of Task 2 and the recall and F-measure of Task 3. Tables 2 and 3 show that the ReFLex users achieved better precision, recall, and F-measure in Tasks 2 and 3 than the Eclipse IDE users. Thus, the ReFLex users achieved significantly better performance than the Eclipse IDE users in Tasks 2 and 3, except that the precision improvement in Task 3 is not significant.

The results show that ReFLex can help the developers improve the recall of feature location results. In both Tasks 2 and 3 the ReFLex users archived significantly better recall than the Eclipse users. In Task 1 the ReFLex users archived better (but not significantly) recall than the Eclipse users. The

results show that the impact of ReFLex on the precision is mixed. In Task 2 the ReFLex users archived significantly better precision than the Eclipse users. In Task 3 the ReFLex users archived better (but not significantly) precision than the Eclipse users. In Task 1 the ReFLex users archived worse (but not significantly) precision than the Eclipse users.

Considering the granularity and complexity of the three tasks, the results suggest that ReFLex may work well on coarse-grained features (such as Tasks 2 and 3) that involve two or more types of logical entities and complex business logic. However, ReFLex may not work well on fine-grained features (such as Task 1) that involve mainly fine-grained application logic. Such fine-grained features often cannot be well broken down into a set of correlated logical entities. For example, the feature of the Task 1 was implemented in 8 methods. The logical boundary among these methods is not clear. Therefore, it becomes challenging in constructing proper logical view for the feature. This challenge is evident in the statistics of the ReFLex tool usage (e.g., the number of restarts) and the quality of the final reflexion model in the Task 1.

4.3 Results: behavioral changes (Q2)

We qualitatively compared behavioral differences between the experimental group and the control group by watching the participants' task-completion videos and by analyzing the field notes of interviews. We focused on behavioral differences in how a developer starts feature location, what he does during feature location, and when he finishes feature location.

We observed that the Eclipse users usually started feature location by identifying keywords from feature descriptions (i.e., IR-based search pattern in [19]). They then tried to find entry points by examining source code of code entities in the search results and/or exploring related entities from the search results. In contrast, the ReFLex users had to think about the logical entities and their relations involved in the feature. Some ReFLex users mentioned in the post-experiment interview that ReFLex requires them to consider not only entry points but also other related aspects required for the conceptual integrity of the feature.

Some ReFLex users mentioned that sometimes it was hard to construct the initial logical view due to the lack of understanding of the feature. In such cases, they used the ReFLex tool to learn the logical structure of the feature in a trial-and-error manner from the reflexion model. Once they had deeper understanding of the feature, they discarded the initial logical view and restarted with a clearer logical view. The ReFLex's usage statistics support this qualitative observation. The performance metrics and the quality of final reflexion models suggest that the difficulty in constructing the proper logical view seems to affect the ReFLex users' feature location performance.

During feature location process, the Eclipse users refined queries based on hints collected from the result list or reading the code. They often selected some methods and examined other methods declared in the same class. They also explored code entities by following type hierarchy or call relations, but they usually can explore only two or three layers before they were lost.

In contrast, the ReFLex users interpreted the reflexion model and performed the structural and descriptive reflexion. They focused on macro issues such as Convergences and Absences edges in the reflexion model, the number of relevant code entities identified, and the terms and topics of relevant code entities. The ReFLex users usually only examined source code to confirm the relevant code entities or remove the irrelevant ones when the relevant-entities set of the logical entity is small and the edges of this logical entity in the reflexion model become Convergences.

The Eclipse users usually ended their feature location when they felt they cannot find any more relevant code entities. They may have confidence in the precision of their feature location results, but they often had no idea about the recall. Because they lacked an overall understanding of the feature. In contrast, 63% of ReFLex users ended their feature location tasks when the reflexion model became convergent and the relevant-entities set of the logical entities had a reasonable size. The ReFLex users often had high confidence in their results due to the overall understanding of the feature and its implementation obtained from the reflexion model. They usually had confidence in the recall of their feature location

results when the reflexion model became convergent, as they knew that all the logical constituents of the feature had their corresponding implementation entities found in source code and these code entities implement all the expected relations between the logical constituents.

4.4 Results: reflexion strategy (Q3)

We also observed a few effective reflexion strategies in the three tasks. First, in the first few iterations of reflective feature location process the relevant-entities set of some logical entities can be large (e.g., ≥ 50 elements in our study). Such a number of relevant entities can make the analysis of Divergences or Absences edges meaningless, as there may be hundreds of edges. In such situations, an effective strategy is to refine the description of these logical entities by descriptive reflexion in order to reduce the size of their relevant-entities sets.

Second, when a logical entity has a medium-sized relevant-entities set (e.g., 20–30 elements), the developer can focus on reducing Divergences edges in the reflexion model by removing irrelevant code entities of the logical entities. This is much easier than reducing Absences edges, as the developer only needs to examine each of the Divergences edges and consider to remove code entities that are related to the edge.

Third, when a logical entity has a small relevant-entities set (e.g., ≤ 10 elements), the developer should try to reduce Absences edges in the reflexion model by refining the structure of the logical view or adding candidate code entities to the relevant-entities set. This is usually harder than reducing Divergences edges, as the developer may need to examine a lot of candidate code entities to select one from them to reduce a Divergences edge. The developer can also proceed to the next iteration to identify more relevant code entities with lower relevant-entity threshold. Furthermore, the developer can examine the relevant code entities in details (e.g., by reading its code) to confirm their relevance to the logical entity. He can freeze the logical entity if he has high confidence in its relevant-entities set and does not want it to be changed.

4.5 Threats to validity

A major threat to the internal validity of our study is the “equivalence” of the experimental group and the control group. To address this threat, we carefully compared participants’ experience and capabilities and randomly allocated comparable participants into the study groups. Another threat to the internal validity is “ground-truth” feature location results provided by the JEdit benchmark. As the benchmark has been used for evaluating several feature location approaches [1], the quality of the benchmark should be high. Furthermore, we manually examined and confirmed “ground-truth” results of the features used in our study. The third threat to the internal validity is the insufficient training of the tools (ReFLex and Eclipse IDE). This factor has greater influence on the experimental group using ReFLex, since the participants use Eclipse IDE regularly while ReFLex is a completely new tool for the participants.

A major threat to the external validity of our study lies in the fact that we only conducted one experimental study. The study was limited by the following factors: a relatively small Java system (i.e., JEdit), three features, graduate students as participants, and familiar application domain. Another threat to the external validity is that we only compared ReFLex with Eclipse. Other IDEs with more advanced search mechanisms such as IntelliJ IDEA or other advanced feature location tools may weaken the advantage of ReFLex. Therefore, our findings may not be applicable to other systems with millions lines of code or developed in other programming languages. And the assumed metamodels and mapping rules in Subsection 2.2 may not be applicable for other kinds of features with different characteristics. Further studies are required to generalize our findings in different kinds of systems, with different kinds of features, with professional developers, with unfamiliar application domain.

5 Case study

We illustrate the process of applying ReFLex for feature location with a case from our user study. In this case the user is told to complete Task 1, i.e., locating relevant classes and methods for the feature “HyperSearch reports the occurrences of the string” of JEdit.

To start feature location with ReFLex, the user first runs JEdit and experiences the feature “HyperSearch reports the occurrences of the string”. Based on his understanding of the feature, he constructs an initial logical view with three logical entities: a HyperSearch operation (with the description “hyper, search”) invokes a CountLineNumber operation (with the description “count, line, number”) and shows a HyperSearchWindow window (with the description “hyper, search, result”).

Iteration 1. Based on this initial logical view, ReFLex runs an automatic mapping process and returns an initial mapping result: 53 code entities mapped for HyperSearch, 9 mapped for CountLineNumber, and 2 mapped for HyperSearchWindow. The user then inspects the mapping results of the three logical entities using the “Filter by Interaction window. As there are many code entities mapped for HyperSearch and CountLineNumber, he does not filter out their results. For HyperSearchWindow, he confirms that a candidate code entity HyperSearchResults class matches the logical entity well while the other one is not relevant, so he removes the irrelevant one and freezes the results of HyperSearchWindow. Moreover, the user decides to revise the description of CountLineNumber to “count, line, file, node”, as he sees the terms “file” and “node” when inspecting the candidate results and thinks that hyper search is implemented by accessing file nodes.

Iteration 2. Based on the filtered result set and updated description, ReFLex returns an updated mapping result: 53 code entities mapped for HyperSearch, 2 mapped for CountLineNumber, and 1 mapped for HyperSearchWindow. The user inspects the candidate results of the three logical entities and revises the descriptions of HyperSearch and CountLineNumber to “do, hyper, search” and “count file node” respectively.

Iteration 3. Based on the updated description, ReFLex returns an updated mapping result: 53 code entities mapped for HyperSearch, 19 mapped for CountLineNumber, and 1 mapped for HyperSearchWindow. The resulted reflexion model is shown in Figure 5. As the numbers of candidate results for HyperSearch and CountLineNumber are large, the user decides to filter the candidate results by interaction. He chooses all the candidate results of CountLineNumber that match the interaction from HyperSearch to CountLineNumber (a Convergences edge) by clicking the arrow and removes all the other candidate results. He inspects the rest candidate results of CountLineNumber by examining their similarity and source code. Based on the examination, he further filters out some candidates and freezes the results of CountLineNumber. For HyperSearch, he decides to keep all the candidate results of it that match the interaction from HyperSearch to CountLineNumber (a Convergences edge). On the other hand, there are a Divergences edge from HyperSearchWindow to HyperSearch and an Absences edge from HyperSearch to HyperSearchWindow. Based on the implementation knowledge revealed by the reflexion model, the user thinks that maybe the edge from HyperSearchWindow to HyperSearch is correct, so he decides to keep all the candidate results of HyperSearch that match this edge. All the other candidate results of HyperSearch are then removed. Moreover, the user removes the Absences edge from HyperSearch to HyperSearchWindow in the logical view.

Iteration 4. Based on the updated result sets and logical view, ReFLex returns an updated mapping result: 17 code entities mapped for HyperSearch, 5 mapped for CountLineNumber, and 1 mapped for HyperSearchWindow. The user thinks that the only problem now is further filtering the result set of HyperSearch. To this end, he examines the overall similarity and source code of the candidate results of HyperSearch and manually removes some of them. On the other hand, he adds an invoke edge from HyperSearchWindow to HyperSearch in the logical view.

Iteration 5. Based on the updated result sets and logical view, ReFLex returns an updated mapping result: 5 code entities mapped for HyperSearch, 5 mapped for CountLineNumber, and 1 mapped for HyperSearchWindow. And the reflexion model now contains two Convergences edges: HyperSearchWindow invokes HyperSearch, HyperSearch invokes CountLineNumber. Based on the reflexion model, the user

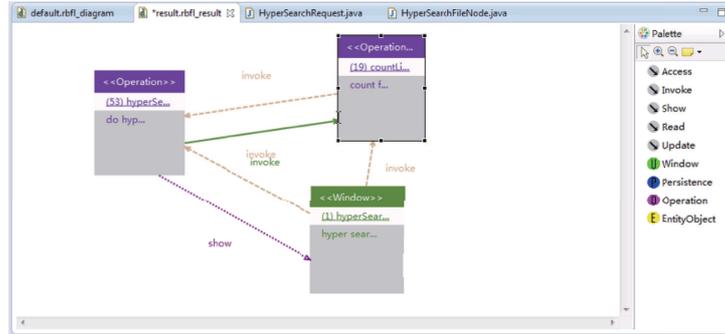


Figure 5 (Color online) The reflexion model of the third iteration.

thinks he has done with task. The final results make a precision of 50% and a recall 62.5%.

6 Discussion

The effectiveness of ReFLex highly relies on the two predefined metamodels of logical view and implemented view and the mapping rules between them. Different metamodels may lead to different feature location results of different quality. For example, if the developers of a system have the knowledge that a feature is implemented by message-based communication, their knowledge cannot be used for feature location based on the metamodels shown in Figure 2. The implemented-view metamodel in Figure 2 cannot reveal the implementation of message-based communication in source code. In this case more complete and fine-grained metamodels such as the one defined by Spoon [20] should be used to enable the definition of message-based communication in source code. The metamodels usually can be defined by technical experts considering the characteristics of specific business domains and technical platforms and the convention of the development team. For example, the metamodels shown in Figure 2 reflect the characteristics of Java Swing-based applications. The definition of logical view metamodel should allow the developers to fully express their knowledge of the implementation of features, i.e., the logical constituents of a feature and their interactions. The definition of implemented view metamodel, on the other hand, should consider the following two aspects. First, the implementation entities and relationships can be recovered from source code. Second, the implementation entities and relationships can be well mapped to the entities and relationships in logical view according to carefully defined mapping rules.

Therefore, the problem usually lies in the lack of clear and consistent mappings between high-level model entities in the logical view and source model entities in the implementation view. One possibility is that the implementation lacks modularity, for example, different kinds of responsibilities (user interface, business logic, technical details) are mixed in the same methods or classes. The other possibility is that some relations in the implemented view cannot be identified by static analysis, for example some functional dependencies implemented by asynchronous messaging. We remark that, although these problems may influence the accuracy of feature location for some features, our approach can generally improve the current feature location practice. Our approach prompts the developers to consider the logical constituents of a feature, which essentially reduces the gap between features and code entities by introducing an intermediate layer. On the other hand, we can ensure the positive influence of relation mapping by only involving high-confident relation mappings. In this way, the feature location results of some features with correct relation mappings can be improved, while those of the other features will not be influenced by noises introduced by relation mapping.

ReFLex asks the developers to construct a logical view of the feature and refine the reflexion model. This seems to be an additional effort for feature location, compared with IR-based code search. However, this seemingly additional effort may bring in benefits to not only feature location in particular but also software maintenance in general. ReFLex provides to the developers a systematic way to evaluate the feature location results and refine their search. This can improve the developers' performance (espe-

cially recall) in feature location tasks (especially for coarse-grained features). Furthermore, the resulting convergent reflexion model emerged from reflective feature location process can increase the developers' confidence in their feature location results and their overall understanding of the feature and its implementation. Feature location has been recognized as the prerequisite for many software maintenance tasks [1, 4]. The overall understanding of the feature and its implementation is necessary to accomplish a software maintenance task. ReFLex makes the development of this understanding an integral part of reflective feature location process. Therefore, ReFLex may have longer-term impact on the maintenance of a software system, because it puts feature location back into the context of software maintenance tasks in which it is required.

7 Related work

IR-based code search can be enhanced by query reformulation techniques [8, 9], better presentation of search results [7, 10], and interactive exploration tools [7, 21, 22]. Beck et al. [23] propose a feature location approach with improved user interface to better support developers during browsing and inspecting the retrieved code entities. Feature location can be facilitated by grouping source artifacts that use similar vocabulary by using topic mining techniques such as the semantic clustering technique proposed by Kuhn et al. [24]. It can also be facilitated by extracting conceptual descriptions from source code by using source code summarization techniques [25]. The user studies by Wilson et al. [5] and Fritz et al. [26] suggest that developers form high-level concept map or mental model of the feature to assist their feature location and code navigation tasks. However, these approaches do not provide systematic ways that allows developers to express their high-level view of the feature and to reconcile the discrepancies between the developers' high-level understanding of the feature and the feature's implementation in the system.

Techniques have been proposed to exploit structural information of software system to improve feature location [27–29]. For example, Robillard's approach [28] identifies relevant code entities based on structural relevance of candidate entities to previously identified relevant entities. Hill et al. [27] explore lexical and structural information to suggest code entities based on a set of entrance points. Portfolio [30] uses network analysis of code structure to improve the ranking of the search results. None of these techniques explore the structural information at both logical and code levels in an iterative feature location process.

Our study on how developers perform feature location tasks suggest that feature location is a human-centered, iterative information seeking process [19]. Incremental feature location approaches allow developers to interactively investigate candidate results, provide feedbacks, and update feature queries. iFL [22] iteratively improves relevance scores of candidate elements based on a relevance feedback mechanism that examines method invocations of code entities. Lucia et al. [21] propose an incremental process for IR-based traceability recovery, which allows developers to tune the threshold in an iterative process and control the number of validated correct links. In [15], Lucia et al. further compare the performances of “one-shot” approach and incremental approach with an empirical study. Their results indicated that incremental process can improve the performances of traceability link recovery and reduce required effort. Poshyvanyk et al. [10] propose to organize query results in a labeled concept lattice to facilitate the exploration of relevant categories in the lattice. Our previous work [7] supports multi-faceted interactive feature location. It allows developers to interactively group, sort, and filter feature location results with multiple syntactic and semantic facets automatically mined from candidate code entities. These approaches do not capture and use logical design structure of a feature in the exploration process.

Our Iterative Context-aware Feature Location algorithm (ICFL) [16] compares both the lexical description and structural context of features and code entities to determine their relevance. The ReFLex's structure-based matching algorithm adopts structure matching concept similar to that of ICFL. However, ICFL takes as input a feature model (i.e., multiple features and their dependencies), while the ReFLex's matching algorithm takes as input a logical view of one feature. More importantly, ICFL is just an algorithm, while ReFLex supports reflective feature location process. Structure-based matching is only one step of this process.

Hill et al. [31] empirically investigate how feature role information, i.e., the role a code element plays in the larger feature, can be utilized to strengthen evaluation studies of feature location techniques. Their study provides empirical support for the assumption of our approach, i.e., feature location can be improved by considering the logical constituents (roles) of features.

In contrast, our ReFLex approach explicitly supports developers to express their understanding of the feature in a high-level logical view. It computes the reflexion model [12] to help developers to reconcile the discrepancies between the high-level logical view of the feature and the implemented view of the system through a reflexion process. The usefulness of software reflexion model has been demonstrated in various software maintenance tasks such as architecture compliance checking [32], change impact analysis [33], and reengineering [34]. Our approach is the first attempt to adopt reflexion model for feature location.

8 Conclusion and future work

In this paper we presented our reflective feature location approach and the tool support. Our approach provides a systematic way for the developers to express their high-level understanding of the feature and to identify and reconcile the discrepancies between their high-level understanding of the feature and the feature implementation. Our user study yielded several promising observations of the reflective feature location. First, developers can express their high-level understanding of the feature in a logical view through reflective feature location process. Second, reflexion on the discrepancies between the logical view of the feature and the system implementation can improve feature location results, especially for coarse-grained features. Third, reflective feature location may lead to longer-term impact on software maintenance as it can help developers gain a deeper understanding of the feature and its implementation. In the future, we will investigate more reflexion mechanisms to better support developers to refine logical views and/or feature location results.

Acknowledgements This work was supported by National Basic Research Program of China (973 Program) (Grant No. 2016YFB1000801) and National Natural Science Foundation of China (Grant No. 61370079).

Conflict of interest The authors declare that they have no conflict of interest.

References

- 1 Dit B, Revelle M, Gethers M, et al. Feature location in source code: a taxonomy and survey. *J Softw Evol Proc*, 2013, 25: 53–95
- 2 Biggerstaff T J, Mitbender B G, Webster D E. Program understanding and the concept assignment problem. *Commun ACM*, 1994, 37: 72–82
- 3 Rajlich V. *Software Engineering: the Current Practice*. 1st ed. BocaRaton: CRC Press, 2012
- 4 Rajlich V, Wilde N. The role of concepts in program comprehension. In: *Proceedings of the 10th International Workshop on Program Comprehension*, Paris, 2002. 271–278
- 5 Wilson L A, Petrenko M, Rajlich V. Using concept maps to assist program comprehension and concept location: an empirical study. *J Inf Knowl Manag*, 2012, 11: 1250018
- 6 Gay G, Haiduc S, Marcus A, et al. On the use of relevance feedback in ir-based concept location. In: *Proceedings of the 25th IEEE International Conference on Software Maintenance*, Edmonton, 2009. 351–360
- 7 Wang J, Peng X, Xing Z, et al. Improving feature location practice with multi-faceted interactive exploration. In: *Proceedings of the 35th International Conference on Software Engineering*, San Francisco, 2013. 762–771
- 8 Haiduc S, Bavota G, Marcus A, et al. Automatic query reformulations for text retrieval in software engineering. In: *Proceedings of the 35th International Conference on Software Engineering*, San Francisco, 2013. 842–851
- 9 Haiduc S, Bavota G, Oliveto R, et al. Automatic query performance assessment during the retrieval of software artifacts. In: *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*, Essen, 2012. 90–99
- 10 Poshyvanyk D, Marcus A. Combining formal concept analysis with information retrieval for concept location in source code. In: *Proceedings of the 15th International Conference on Program Comprehension*, Banff, 2007. 37–48
- 11 Kruchten P. The 4+ 1 view model of architecture. *IEEE Softw*, 1995, 12: 42–50
- 12 Murphy G C, Notkin D, Sullivan K J. Software reflexion models: bridging the gap between design and implementation. *IEEE Trans Softw Eng*, 2001, 27: 364–380
- 13 Lucia A D, Penta M D, Oliveto R, et al. Labeling source code with information retrieval methods: an empirical study. *Empir Softw Eng*, 2014, 19: 1383–1420

- 14 Blei D M, Ng A Y, Jordan M I. Latent dirichlet allocation. *J Mach Learn Res*, 2003, 3: 993–1022
- 15 Lucia A D, Oliveto R, Tortora G. IR-based traceability recovery processes: an empirical comparison of “one-shot” and incremental processes. In: *Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering*. Washington: IEEE Computer Society, 2008. 39–48
- 16 Peng X, Xing Z, Tan X, et al. Improving feature location using structural similarity and iterative graph mapping. *J Syst Softw*, 2013, 86: 664–676
- 17 Baeza-Yates R, Ribeiro-Neto B. *Modern Information Retrieval*. 1st ed. New Jersey: Addison Wesley, 1999
- 18 Hayes J H, Dekhtyar A, Sundaram S K. Advancing candidate link generation for requirements tracing: the study of methods. *IEEE Trans Software Eng*, 2006, 32: 4–19
- 19 Wang J, Peng X, Xing Z, et al. An exploratory study of feature location process: distinct phases, recurring patterns, and elementary actions. In: *Proceedings of the IEEE 27th International Conference on Software Maintenance, Williamsburg*, 2011. 213–222
- 20 Pawlak R, Monperrus M, Petitprez N, et al. Spoon: a library for implementing analyses and transformations of java source code. *Softw Pract Exper*, 2015, 46: 1155–1179
- 21 Lucia A D, Fasano F, Oliveto R, et al. Recovering traceability links in software artifact management systems using information retrieval methods. *ACM Trans Softw Eng Meth*, 2007, 16: 13
- 22 Hayashi S, Sekine K, Saeki M. iFL: an interactive environment for understanding feature implementations. In: *Proceedings of the 26th IEEE International Conference on Software Maintenance, Timisoara*, 2010. 1–5
- 23 Beck F, Dit B, Velasco-Madden J, et al. Rethinking user interfaces for feature location. In: *Proceedings of the IEEE 23rd International Conference on Program Comprehension, Florence*, 2015. 151–162
- 24 Kuhn A, Ducasse S, Girba T. Semantic clustering: identifying topics in source code. *Inform Softw Tech*, 2007, 49: 230–243
- 25 Haiduc S, Aponte J, Marcus A. Supporting program comprehension with source code summarization. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering, Cape Town*, 2010. 223–226
- 26 Fritz T, Shepherd D C, Kevic K, et al. Developers’ code context models for change tasks. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, Hong Kong*, 2014. 7–18
- 27 Hill E, Pollock L L, Vijay-Shanker K. Exploring the neighborhood with dora to expedite software maintenance. In: *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering, Atlanta*, 2007. 14–23
- 28 Robillard M P. Topology analysis of software dependencies. *ACM Trans Softw Eng Meth*, 2008, 17: 18
- 29 Wang X, Lo D, Cheng J, et al. Matching dependence-related queries in the system dependence graph. In: *Proceedings of the 25th IEEE/ACM International Conference on Automated Software Engineering, Antwerp*, 2010. 457–466
- 30 McMillan C, Grechanik M, Poshyvanyk D, et al. Portfolio: finding relevant functions and their usage. In: *Proceedings of the 33rd International Conference on Software Engineering, Waikiki*, 2011. 111–120
- 31 Hill E, Shepherd D, Pollock L L. Exploring the use of concern element role information in feature location evaluation. In: *Proceedings of the IEEE 23rd International Conference on Program Comprehension, Florence*, 2015. 140–150
- 32 Knodel J, Popescu D. A comparison of static architecture compliance checking approaches. In: *Proceedings of the 6th Working IEEE / IFIP Conference on Software Architecture*. Washington: IEEE Computer Society, 2007. 12
- 33 Kim T, Kim K, Kim W. An interactive change impact analysis based on an architectural reflexion model approach. In: *Proceedings of the 34th Annual IEEE International Computer Software and Applications Conference*. Washington: IEEE Computer Society, 2010. 297–302
- 34 Murphy G C, Notkin D. Reengineering with reflexion models: a case study. *IEEE Comput*, 1997, 30: 29–36