

Unravel Programming Sessions with THRESHER: Identifying Coherent and Complete Sets of Fine-granular Source Code Changes

Marcel Taeumel, Stephanie Platz, Bastian Steinert,
Robert Hirschfeld, and Hidehiko Masuhara

Development teams benefit from version control systems, which manage shared access to code repositories and persist entire project histories for analysis or recovery. Such systems will be efficient if developers commit coherent and complete change sets. These best practices, however, are difficult to follow because multiple activities often interleave without notice and existing tools impede unraveling changes before committing them. We propose an interactive, graphical tool, called THRESHER, that employs adaptable scripts to support developers to group and commit changes—especially for fine-granular change tracking where numerous changes are logged even in short programming sessions. We implemented our tool in Squeak/Smalltalk and derived a foundation of scripts from five refactoring sessions. We evaluated those scripts’ precision and recall, which indicate a reduced manual effort because developers can focus on project-specific adjustments. Having such an interactive approach, they can easily intervene to accurately reconstruct activities and thus follow best practices.

1 Introduction

Software developers benefit from version control systems (vcss), which manage collaborative work and persist whole project histories. In systems such as Git^{†1} and Mercurial,^{‡2} developers can *commit* changed source code with descriptive messages into distinct branches to separate features or bug fixes. When committing, branching, or merging, new versions of the project arise and a particular development trace emerges in the vcs. Such historical information can then be used to ease code compre-

hension [37] or to reason about software evolution patterns [3].

Software evolution induces very high costs in software development [2] and hence developers are advised to follow best practices to maximize the benefits of vcss. One important practice is *continuous integration* [8]: commit and test code on a regular basis. When committing, only store *small, coherent, complete* change sets with a brief yet descriptive message to improve comprehensibility [35]. A well-organized project history can also improve code review processes [1] and recommender systems [4]. —Robbes et al. [26] observed that developers frequently commit once per working day on average. Further research [16][18][26], however, shows that developers tend to commit *incoherent* change sets. Buse and Weimar [5] also found that only two-thirds of commits get assigned with a descriptive message.

We argue that one challenge for developers is that multiple activities interleave [21]—often without being noticed even in short sessions as illustrated in Figure 1. Manual refactorings, for example, typically include changes to affected code that

完全かつ一貫した細粒度コード変更集合の検出によりプログラミングセッションを解きほぐす THRESHER

Marcel Taeumel, Stephanie Platz, Bastian Steinert, and Robert Hirschfeld, Hasso Plattner Institute, University of Potsdam, Germany.

増原英彦, 東京工業大学情報理工学院, Department of Mathematical and Computing Science, Tokyo Institute of Technology.

コンピュータソフトウェア, Vol.34, No.1 (2017), pp.103–118. [研究論文] 2016年3月24日受付.

†1 <http://git-scm.com>

‡2 <http://mercurial.selenic.com>

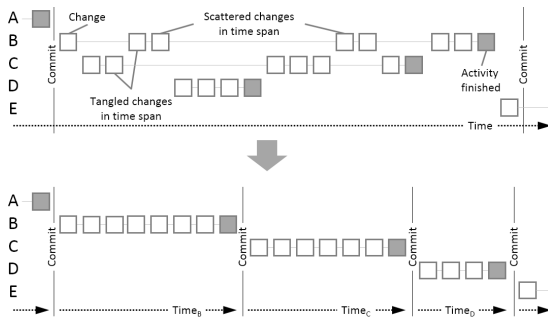


Fig. 1 Multiple activities (A to E) often interleave without notice (top) and best practices suggest to untangle and assemble changes before committing them (bottom). This is often challenging with existing tools.

get accidentally postponed because developers simply forgot. Existing tool support typically poses high effort on following best practices, especially for fine-granular change models that track numerous changes even in short sessions [33]. Developers tend to explain all tangled and scattered changes from one session in a *single* commit message—if any. Consequently, understanding project histories will be impeded. Based on these observations, we address the following research question:

How can we support developers to identify coherent and complete sets of fine-granular changes when committing code to a repository?

We define a change set as *coherent* if it contains only changes that origin from a single activity; it is *complete* if it contains all changes from that (completed) activity. We focus on fine-granular change models where every modification to the code base is tracked, even two consecutive changes that rescind each other. We think that semi-automatic tool support should favor *coherent sets over complete ones*, that is, avoid false-positives but accept false-negatives. Then, developers do only have to merge change sets to make them complete. In fact, we believe that there cannot be one fully automatic approach but the developer has to be part of the process to benefit from tacit project knowledge.

We propose a scriptable, interactive, graphical tool that supports developers to (1) identify consecutive changes that indicate continuity of an activity and (2) collect scattered runs of changes that

belong to the same activity. We make use of a simple scripting language as a flexible extension mechanism, because we think that software projects have domain-specific characteristics to be specified by the who are developers actually working on the project. Our scripts are basically just Smalltalk code:

```
[[:grouper |
  "Add all moved methods in a group."
  grouper addGroup: grouper
  moveMethodChanges].
```

Based on a lab study [33] including 5 programming sessions with up to 250 fine-granular changes each, we derived a foundation of scripts that can be combined to semi-automatically reveal activity-specific change sets. When evaluated, a script assigns changes into *groups* and adds *recommendations*^{†3} if changes are already part of distinct groups. After this scriptable analysis, developers can explore all proposed groups, adjust them manually, or resolve recommendations. We implemented the graphical tool THRESHER in Squeak/Smalltalk,^{†4} which visualizes intermediate results concisely and reduces the effort for manual adjustments. The underlying rules of identifying coherent and complete sets of changes can easily be accessed and modified.

In this paper, we make the following contributions:

- A design for an interactive, scriptable tool we call THRESHER, which supports developers to assemble coherent and complete change sets for fine-granular source code changes
- A description of an exemplary scripting interface, which supports the Squeak/Smalltalk change model, to emphasize the simplicity of extending THRESHER for domain-specific needs
- A brief description of scripts we extracted from a lab study to support Smalltalk-based refactoring sessions—including an evaluation in terms of precision (almost 100%) and recall (33% to 57%) of the proposed groups, which illustrates coherence over completeness

The next Section 2 summarizes related work.

^{†3} Note that we use the term “recommendation” unrelated to the domain of recommendation systems.

^{†4} The Squeak/Smalltalk programming environment: www.squeak.org

Section 3 describes existing challenges in state-of-the-art commit tool support. In Section 4, we present our tool THRESHER and give an introduction to its scripting capabilities; we also present a foundation of scripts for developers to get started. We evaluate those scripts in Section 5 and discuss limitations and applicability to other environments. Section 6 sketches hypotheses and next steps for this kind of commit tool support. Finally, we conclude our thoughts in Section 7.

2 Related Work

There are not many projects that try to proactively involve developers in the process of committing changes. Although, there are many projects that investigate how to automate the identification of coherent change sets—only after commits have been taken place. Some projects derive rules by mining project histories, other projects create rules for changes that have detailed context information. Most of them deal with domain-specific characteristics, which emphasizes the need of adaptability and manual intervention when using such tools.

Note that there is always the possibility to ask programmers to explicitly set a scope *before* making any changes, hopefully producing a clean, untangled history record. For example, the Mylar/Myllyn plugin for Eclipse provides that level of support [19][25]. We argue that such an approach is not prepared for many software projects out there because it expects too much cautiousness while programmers get distracted by side-tasks too easily. At the end of the day, untangling existing changes becomes an issue and tools have to support that.

The tool Historef [14] takes an approach similar to THRESHER. There, programmers can untangle edits in Java source code files to form task-level commits or do selective undo. — While Historef accounts for the concept of a custom *edit recorder*, THRESHER was designed to work with several kinds of change models, even coarse-grained ones. Both tools do not provide automatic measures to prevent bugs introduced by reordering history events. THRESHER operates on an object-oriented level and abstracts from files by means of packages, classes, and methods. Additionally, THRESHER incorporates the idea of adjusting the automated grouping

rules to accommodate domain-, project-, or task-specific requirements.

When mining repositories, existing changes are often modeled and analyzed at a fixed level of detail. Herzig and Zeller [15][16] applied machine learning to adjust project-specific weights or thresholds for data dependencies, lexical distance measures, or test impact couplings. They analyzed coherent change sets that were already assigned to particular bug fixes. Kawrykow and Robillard [18] looked for behavior-preserving thus non-essential changes, which include rename-refactorings. They claim that such changes affect coherence negatively. Kim and Notkin [20] derived logical rules from change sets, which also identify anomalies and hence tangling changes such as “All methods X in classes that implement interface Y got deleted, *except class Z.*” This should also help describe commit messages. — We derived a foundation of scripts from 5 programming sessions, which arguably reflect many important characteristics of object-oriented projects in Smalltalk. One can easily extend this approach with mining some Squeak/Smalltalk repositories to identify more corner cases. As scripts represent concise and accessible descriptions of a particular change model, we argue that developers can easily adapt those scripts to accommodate domain or project-specific needs.

When approaches are open to extensions in the change model, identification and classification methods can benefit from any kind of context information such as developer activity tracking. For example, Coman and Sillitti [6] tracked tool interactions and created the notion of *time intervals of intensive access* to help detect activity switches; Zou and Godfrey [41] verified this approach in an industrial setting. Robbes et al. [27][28] also track timestamps and use time-coupling to improve recall. Besides that, they track tool window usage, code authors, and annotations to indicate automated refactorings. If actually noticed, the developer can express the beginning of a new activity explicitly. Yoon et al. [39] log all low-level events of the Eclipse code editor in an XML format to be processed by other tools. — THRESHER promotes fine-granular change tracking without making assumptions about the level of detail that is tracked. For example in our Squeak/Smalltalk environment, there were logs for tool-controlled refactoring activ-

ities, which could be used to simplify script code. Other extensions to the change model could further improve both scripts, results, and hence the overall utility.

Working with numerous fine-granular changes raises challenges for graphical tools from an information visualization perspective. Many approaches apply a timeline metaphor to provide overview and details on demand as pluggable visualizations in programming environments. Examples include Azurite [40] and CodeTimeline [23]. — We argue that common list, table, or tree views are capable of presenting huge amounts of data in a supportive way. Problems will arise primarily if developers have no simple means to configure the views' level of detail. THRESHER employs those means by giving developers that required configuration support in terms of adaptable scripts with direct feedback after script modifications.

3 Challenges in Existing Tools

In this section, we highlight important factors that substantiate the need for better tool support when committing source code to a vcs. We cover the nature of interleaving programming activities, state-of-the-art tool support, and specifics of the Squeak/Smalltalk change model—as we built THRESHER in that programming environment.

3.1 Interleaving Programming Activities

The more programming activities interleave during a session, the more important it is to untangle them to commit only coherent and complete change sets. The number of actually traceable activities will be influenced by the developer's task and the level of change tracking editing tools provide. As for the task, a simple bugfix may spawn only one activity but “drive-by corrections” and other activities interleave quite often [21][24]—those should be documented in separate commits. As for the change tracking, the more editing events get tracked, the higher the chance that a developer's activity drifts get documented for the long term—all the more if there are many modifications that rescind each other during a session.

File-based environments (such as Java + Eclipse + Git) compare the local working copy with the latest version in the vcs to determine changes at

the level of text lines, which typically match program statements. We consider such an approach as *coarse-grained* because overwritten changes are not tracked as coherent program entities but only at the level of character ranges. Object-based environments (such as Smalltalk + Squeak + CoEXIST [32]) log any modification to the code base by means of method and class changes. This allows for rewinding or replaying single decisions and whole programming sessions. Such a change model is more *fine-grained* and persists more details of a project's history.

There is a programming approach called *explorative/exploratory programming* [9], where developers are encouraged to make any change as they go without having a specific task in mind. They do not know upfront whether they will fix some bugs or add some neat features or clean-up some code. If software projects encourage emerging specifications, a programming environment that supports explorative programming will help programmers write those specifications [31][30] while the software system is growing. This approach frees programmers from being too specific too early in highly uncertain scenarios like many research projects face. At the end of an explorative session, developers can have numerous changes to untangle. We argue that there is a need for better commit tool support to even promote explorative programming and still ease comprehension of and learning from project histories.

3.2 Interactive Commit Tool Support

Figure 2 gives an impression of how existing commit tools make use of graphics and interaction. First and foremost, those tools present changes in scrollable lists with support for manual selection only. “Advanced” filtering covers presets such as *all*, *none*, and *invert selection*. It doesn't matter whether changes are file, method, or line-based: Such interfaces do not support untangling many changes very well. Furthermore, multiple activities have to be considered one after another.

However, we do consider lists or tables as appropriate views because they can visualize chronological order in a compact fashion. What existing tools miss is a way to filter and group changes according to some rules, which may not be expressed with a single button-click but rather simple scripts.

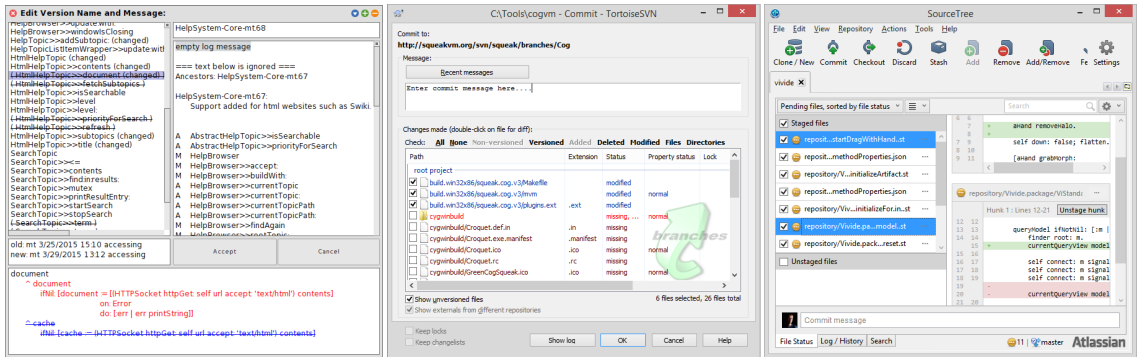


Fig. 2 Usually, commit tools show changes in long, scrollable lists, which developers have to untangle by tedious manual selection for each single commit. LTR: Monticello for Squeak (<http://www.squeak.org>), TortoiseSVN (<http://tortoisesvn.net>), SourceTree (<http://www.sourcetreeapp.com>)

Additionally when seeing a complete and coherent change set, developers may be better suited to provide a descriptive commit message.

3.3 The Squeak Change Model

Squeak has a direct notion of changes, change files, and change sets^{†5}—but all in a local sense and not shared like common VCSS such as Git. Tracked changes include any source code modification affecting classes, methods, and their categories. In Squeak, change sets represent means to manually organize changes to easily file-out and share them, for example, via mailing lists. There is a tool called *Change Sorter*, which supports moving changes between sets in case the developer forgot to switch before starting her activity. Overall, Squeak’s notion of changes provides a local history of everything that happens in the programming environment. However, not all changes are reversible per se but, at the least, they can be replayed.

The Smalltalk community has a VCS called *SqueakSource*^{†6} and an in-image tool called *Monticello* as counterpart (left in Figure 2). When committing changes, Monticello only sees the latest versions of classes and methods. Cherry-picking is supported at the level of methods—not lines or ranges of text.

^{†5} Except for this section, we do not use the term *change set* in a Squeak-specific way but in a broader sense throughout this paper.

^{†6} <http://www.squeaksource.com>

We make use of a recent research project called CoEXIST [32], which extends Squeak’s change model by storing data to revert changes. It also adds the capability to Monticello to store fine-granular changes in SqueakSource such as two consecutive modifications of the same method. In that sense, CoEXIST elevates the fine-granular but local change tracking of Squeak to the level of Monticello/SqueakSource, which is now comparable to other VCSS.

Building on top of many general concepts, we created THRESHER for Squeak/Smalltalk but are arguably able to transfer our findings to other programming languages, tools, and environments.

3.4 The VIVIDE Tool Building Framework

We implemented our tool THRESHER with the VIVIDE^{†7} tool building framework [34]. Vivide provides a direct mapping between all graphical parts of the user interface and the internal tool logic. It is implemented in Squeak/Smalltalk and builds on top of the Morphic framework, which supports direct manipulation of all graphical objects. With employing this directness, the developer can easily find responsible data transformation scripts starting from a visual impression and express modifications in the script source code. Due to this simple yet powerful abstraction, the framework can update all running tools consistently. Thus, VIVIDE is both a programming environment and a tool build-

^{†7} <https://github.com/hpi-swa/vivide>




Fig. 3 THRESHER works bottom-up: groups of changes will be merged (1-3) according to scriptable rules. There are no partial merges (4) but recommendations (5a) to be resolved manually by the developer.

ing framework. Building tools means composing widgets and writing script code.

The underlying concept of VIVIDE, that is describing tools as data processing pipelines, aligns with our idea about analyzing code changes to recover programming activities. Developers should open it after finishing a programming session to assemble recent changes into groups, add descriptive messages, and eventually commit them to the vcs with only a few clicks.

4 Thresher

In this section, we explain structure and use of our tool as well as details about how changes are processed with scripts. The name “thresher” derives from the fact that we want to identify all changes that belong to a certain activity and ignore the rest (for now)—just like the agricultural machine does with *winnowing*.

4.1 Changes, Groups, and Recommendations

The overall goal is to make developers commit coherent and complete change sets with a descriptive message. We classify a change set as *coherent* if it contains only changes that origin from a single activity; it is *complete* if it contains all changes from that (completed) activity. In THRESHER, scripts support to detect activity continuations and thus grouping changes coherently. Completeness might often not be achieved on a scripting level because

some developer knowledge might be hard to express or even ambiguous. Then, manual grouping via THRESHER’s graphical interface will take place.

There are multiple scripts that modularize the change analysis process in a sequence of stages. At the first stage, each change has its own group assuming that each change belongs to its own programming activity. Intermediate stages either expand those groups by merging them or propose recommendations to be resolved manually. In the last stage, each resulting group is meant to be committed to the vcs. The concept is illustrated in Figure 3.

In script code, programmers can iterate over all changes, check for specific properties, and assign matches to new groups. Those checks depend on the information available; a stage may also add new information to changes by integrating external data sources. For example, results of static or dynamic code analysis can be embedded this way.

4.2 How to Use and When to Use

THRESHER’s graphical user interface is shown in Figure 4. It consists of four main views:

- ① **Sources View** shows a chronologically ordered list with all recent changes. The order can be adapted.
- ② **Result View** shows resulting groups and recommendations. Developers can resolve recommendations with button clicks, change groups via drag-and-drop, and add descriptive messages.
- ③ **Diff View** shows details about all selected groups to support developers to revise their decisions in detail.
- ④ **Stages View** reveals intermediate results of the involved scripts, supports adapting (add/remove/reorder) them as well as reviewing (or debugging) their effects.

There are several tasks that the developer accomplishes when working with THRESHER. Those describe a seamless workflow and should encourage to follow best practices by committing only coherent, complete change sets as well as adding descriptive messages.

At first, the freshly opened THRESHER window analyzes recent changes with the existing scripts to form groups and add recommendations. The source view ① lists all changes chronologically with

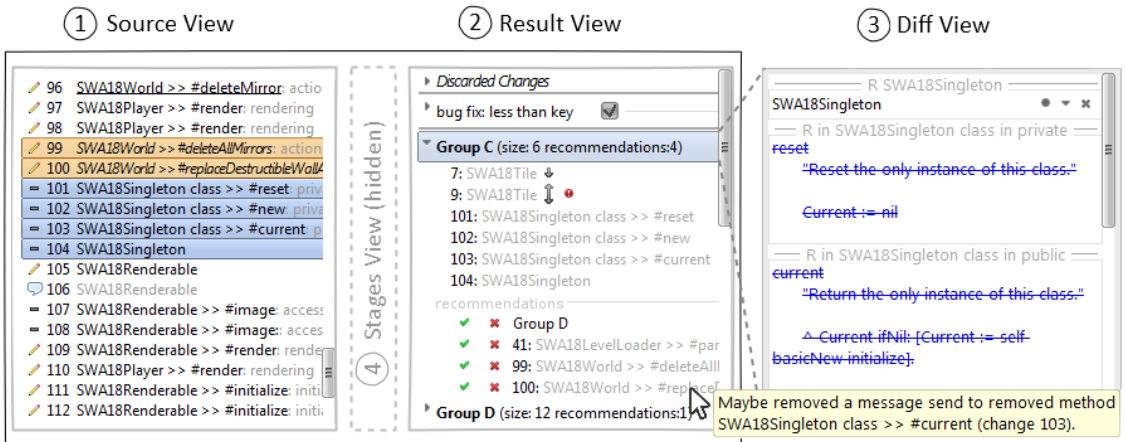


Fig. 4 The main views of THRESHER: ① local, uncommitted changes, ② computed change sets with recommendations, and ③ detailed source code diffs.

a distinct label. The result view ② shows all resulting groups, which contain these changes and maybe several ungrouped^{†8} ones. At the top of that view, *discarded changes* are listed such as recognized debugging code. All groups get single-letter labels for easier recognition. At this point, the diff view ③ is not used because no group is selected.

After that, the developer reviews all group's contents to verify coherence and completeness. Selecting a group in the result view ② will reveal more details in the diff view and highlight the contents, which are automatically selected in the source view. Any group can be marked as *ready-to-commit* by clicking the checkbox near the group name. This will separate that group visually by moving it to the top as shown in Figure 4 for the group “bug fix: less than key.”

Group C contains 6 changes and 4 recommendations, which are highlighted in blue and orange. To see the content of single changes, the developer can hover over one of them in the source view. To discard the changes, the developer selects them and drag-drops them onto the “Discarded Changes” group.

Recommendations can either be single changes or whole groups as proposed by previous stages. When hovering over a recommendation, the reason is shown in a tooltip (Figure 4, bottom right). The

developer can accept recommendations by clicking the tick-icon next to it; the cross-icon rejects it.

After resolving all recommendations, the developer can merge groups if she discovers a more appropriate intent. This happens by simply drag-and-drop groups on top of each other. It is also possible to move selected changes (or groups) between groups.

Finally after reviewing a group, the developer can propose a descriptive name, which is displayed as the group's name to help identify complete ones. When she decides to commit, name and description are used as *commit message*. Besides manual naming, a short description of structural changes is appended automatically such as “added class `SWA18World`” and “modified method `Player >> render:`.”

Developers can invoke THRESHER whenever they want. They do not have to actually do the commit after they organized their changes, but they can continue programming. However, we designed THRESHER to handle hundreds of fine-granular changes at once. Hence we do not expect developers to continuously use it *within* a programming session. We encourage developers to focus on their tasks and not to worry about tangled history records.

^{†8} With “ungrouped” or “single changes” we mean single-change groups.

4.3 Scripting to Accommodate the Project's Domain

Developers can accommodate domain or project-specific needs by adding, removing, or reordering stages. Each stage has a current script, which developers can access, modify, and re-execute to update the results. Once visible, the *stages view* (Figure 5 resp. ④ in Figure 4) shows intermediate grouping results. Scripts access the change model and create groups with the help of a *grouper* like this:

```
[:grouper | |newGroup|
  newGroup := CustomGroup new.
  grouper allChanges do: [:change | "See below."].
  grouper addGroup: newGroup.
].
```

In our implementation, the scripting language is Smalltalk and scripts are just blocks evaluated with the *grouper* instance automatically. Each script is like an anonymous function that will be called with a *grouper* as argument. That function has no return value but can trigger side effects via the *grouper*'s interface, that is, creating and adding groups of changes.

If there are already existing groups, developers can inspect their contents and consider them. At the beginning, there are as many groups as there are changes. For each match in a script, that change's current group is fetched and the nested structure of groups (Figure 3) is extended. The following snippet extends the example above:

```
"..." [:change |
  change property = 'SomeFilter' ifTrue: [
    group add: (
      (grouper groupOf: change) -> 'Reason_for_
      addition.'
    )]] . "..."
```

To write effective scripts, developers need to know about possible change properties and the scripting language to write scripts. Here, filters have the conciseness and expressiveness of the Smalltalk language and are only limited by the environment's loaded libraries. In particular, external data sources may also be queried and their answers used in such expressions. For example, one might think of a library that supports connecting email accounts with source code and hence reveal further information about the original author of the

changed piece of code.

THRESHER's scripting capabilities can be summarized like this: **Scripts** are high-level descriptions of change characteristics and create or merge appropriate *groups* while using a *grouper* for navigation. Different **group** classes represent distinct change characteristics with an optional *pivot change*, which is a dedicated change used for identification. For example, a "class X renamed" change can be the pivot change for subsequent code updates. Initially, each change is in a single-change group. Arbitrary changes can be put into a custom group. Groups are units of reuse in *scripts*. A **grouper** supports navigating changes in the current group hierarchy and manages top-level *groups* across all *scripts*. It is not meant to be extended. **Recommendations** are created automatically if changes are associated with two distinct groups.

To present a more elaborate example, the following script analyzes changes for method additions/removals and combines them with all other method changes that update corresponding message sends:

```
[:grouper |
  grouper allChanges
  select: [:change | change
    isMethodAddOrRemove]
  thenDo: [:pivot | | group |
    group := MethodAddOrRemoveGroup new.
    group add: (
      (grouper groupOf: pivot) -> pivot name).
    grouper methodChanges do: [:change |
      "1) Check for removed methods."
      pivot isRemoval &
      (change isSendRemoved: pivot) ifTrue: [
        group add: (
          (grouper groupOf: change)
          -> ('Send_removed:', pivot
            selector)]].
      "2) Check for added methods."
      pivot isAddition &
      (change isSendAdded: pivot) ifTrue: [
        group add: (
          (grouper groupOf: change)
          -> ('Send_added:', pivot
            selector)]]].
    "Add new group to hierarchy."
    grouper addGroup: group]].
```

Here, *associations*, which attach a textual description as *reason* to each change, are used to fill groups with changes. Developers do not have to

provide reasons but debugging unexpected results can benefit from such contextual explanations.

With only 19 lines of Smalltalk code, the developer can automate such an untangling of changes. If existing changes are ungrouped, the grouper will just create new groups. If they were already assigned to groups in preceding stages, THRESHER will handle conflict resolution and may add recommendations because changes can only have a single group. Developers do not have to deal with conflict handling when writing scripts.

Scripts can also modify existing groups that were created by previous stages. By default, scripts do not modify previous decisions but *add*, *merge*, or *recommend* new changes from emerging groups. For example, two groups of the same kind such as `MethodAddOrRemoveGroup` are usually merged and not added as sub-groups. Eventually, top-level groups are important for the developer because those groups are meant to be committed.

Developers can add, remove, or edit THRESHER's scripts whenever they want—even during a regular programming session. They can try out ideas, keep the good ones, discard the useless ones. We think that once the project domain and source code conventions become clear developers do rarely have to modify the existing scripts. Hence, scripting is no ongoing activity in projects that make use of THRESHER. We expect the additional costs to learn and use the scripting language to amortize quickly.

4.4 Towards a Foundation of Scripts

We implemented a set of complementary scripts that aim for detecting most characteristics as extracted from several programming sessions [33]. In most cases, they use structural/code change information but they may also consider chronological affinity. These assumptions encode the Smalltalk language characteristics besides Squeak's change model and hence serve as a valid baseline for other projects in the Squeak/Smalltalk environment. Due to space constraints we cannot print all scripts' details here but will only summarize their intents:

1. **Consecutive grouping** All consecutive changes on the same source code artifact^{†9} are

grouped.

2. **Refactoring grouping** All changes that were triggered by tool-driven refactorings are grouped. Depends in respective hints being present in the change data.
3. **Renaming grouping** Detects a manual or tool-driven renaming of a source code artifact and groups it with all consecutive changes of the same kind. Affects also renamed variables.
4. **Organizing grouping** All consecutive changes that reorganize software artifacts are grouped. This includes class categories and method protocols.
5. **Added/removed-method grouping** The addition or removal of a method is grouped with all changes that update the respective message sends in the source code. Applies static analysis to locate all sends.
6. **Added/removed-variable grouping** The addition or removal of an instance/class variable is grouped with all changes that update the respective references in the source code.
7. **Added/removed-class grouping** The addition or removal of a class is grouped with all changes that concern this class' methods, variables, and references.

The last three stages add recommendations:

8. **Close-to-artifact recommendation** Single changes are *recommended* to groups whose changes affect the same artifact. Considers chronological affinity.
9. **Close-to-polymorphism recommendation** Single changes are *recommended* to groups that concern method additions/removals but are ambiguous due to polymorphism. Considers chronological affinity.
10. **Pretty-print recommendation** Single changes of cosmetic nature are *recommended* to groups whose changes affect the same artifact if chronological affinity is high. If there is no such group, those changes are merged into a new group. Such changes include modification of whitespace or comments, renaming of temporal variables, or cascading message sends.

With this focus on rules to combine rather than split changes, these scripts do not detect activity *switches* but rather describe *continuations*. From one stage to another, scripts merge existing groups of changes into bigger ones. Therefore, the or-

^{†9} Source code artifacts can be packages, classes, or methods.

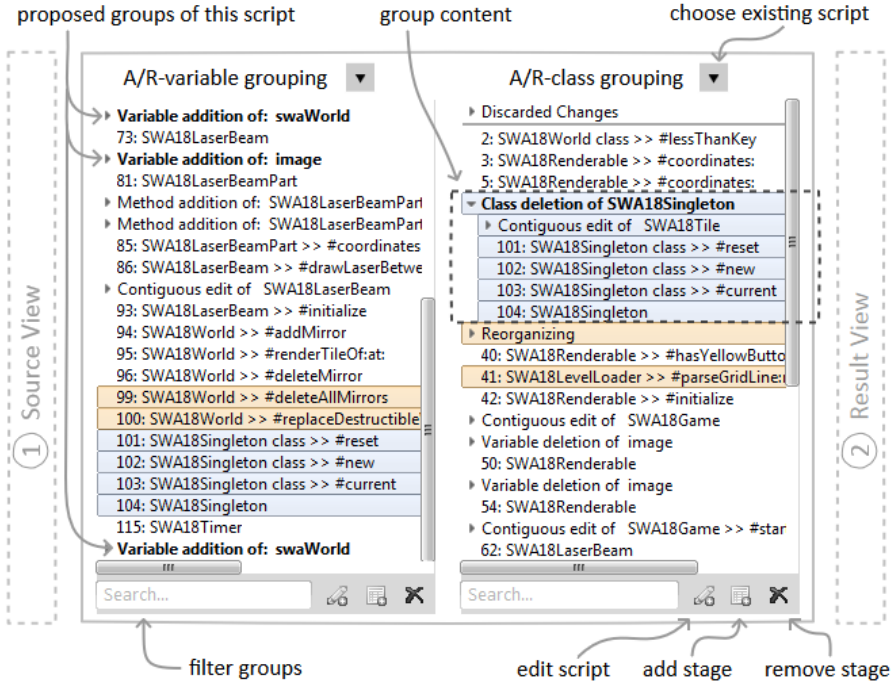


Fig. 5 The stages view in THRESHER supports exploring intermediate results and adapting the scripts.

der of scripts can influence the results. For now, THRESHER should stop processing if further merging might decrease coherence of groups. Splitting up groups in scripts and hence describe activity switches would be a valuable addition to our concept and is considered future work.

5 Evaluation of Scripts

In this section, we evaluate the accuracy of the basic scripts in THRESHER, which we derived from a programming study [33]. We then estimate the manual effort that remains to complete the mapping of all changes to their activities—by interacting with the tool’s graphical interface.

We have not yet evaluated to overall utility of THRESHER. To some extent, we rely on the positive effects of the VIVIDE programming environment [34], in which THRESHER runs. This, however, remains subject to further research.

5.1 Analyzing Programming Sessions

Our proposed scripts are based on sample data, which Steinert et al. collected during a lab

study [33] with 22 developers that had the task to improve source code of a game in 2-hour sessions. The number of changes range from 40 to 250 per session. We manually mapped those changes to activities, which included debugging, refactoring, and code clean-up. From this laborious activity, we derived several scripts that should automate this mapping process in terms of *coherent groups of consecutive changes* and which can be combined to *assemble scattered change runs*. Reconsider Figure 1 for the distinction of tangled and scattered changes.

In particular, we chose 5 out of the 22 sessions as representatives because their changes contained many refactorings or frequently interrupted activities with much tangling and many revisions; some of them contained at least one larger refactoring. We reviewed each fine-grained change *manually* and added it to coherent, complete sets to get kind of a *gold standard* for our analysis. Some developers created their own commits with Monticello, which we only partially took into account because at that time, it forced them to commit all changes at once. We did not interview the developers but were fa-

miliar with the domain of the game and its source code. Eventually, we identified recurring characteristics that indicate strong relationships or obvious activity switches as summarized in Table 1. We used the 6th data set to “simulate” a programming session whose contents were not explicitly used to derive those scripts.

5.2 Precision, Recall, and Manual Effort

We quantify the accuracy of THRESHER’s basic scripts in terms of *precision* and *recall*, as known from the domain of information retrieval [36]. The result’s precision describes the percentage of true positives, the result’s recall accounts for false negatives. We describe the accuracy with *sets of changes* to reflect the way THRESHER works. In addition to precision and recall, we quantify the *manual effort* that is necessary to get an ideal result.

For the accuracy, we compare the expected set of related changes with the calculated one. Let the expected groups be $A_E = \{3, 5, 6\}$ and $B_E = \{1, 2, 4\}$ where digits indicate a change’s running number. Let the calculated groups be $A_C = \{3, 5\}$, $B_C = \{6\}$, and $C_C = \{1, 2, 4\}$. The expected set of relations would be

$$\mathbb{E} = \{\{1, 2\}, \{1, 4\}, \{2, 4\}, \{3, 5\}, \{3, 6\}, \{5, 6\}\}$$

and the calculated one would be

$$\mathbb{C} = \{\{1, 2\}, \{1, 4\}, \{2, 4\}, \{3, 5\}\}$$

by expanding the groups into pair sets. The *precision* is now defined as $\mathbb{P} = |\mathbb{E} \cap \mathbb{C}| / |\mathbb{C}|$ and *recall* as $\mathbb{R} = |\mathbb{E} \cap \mathbb{C}| / |\mathbb{E}|$. In this example, precision is 1.0 (100%) and recall is 0.67 (67%), which indicates that *all* calculated groups are coherent but manual reviewing and merging is required for completion.

For the manual effort, we consider two extreme cases: developers are either directly able to assign changes to activities after reviewing them once ($\min = 2 \cdot n_{\text{changes}}$), or they have to look through all changes for *each* activity again ($\max = n_{\text{activities}} \cdot n_{\text{changes}}$). When using THRESHER, this effort is reduced because it depends on the number of proposed groups instead of single changes. The developer clicks to accept/reject single recommendations or all at once. Drag-and-drop interactions can merge groups or single changes.

5.3 Results

We achieved a precision of 96% to 100% for our five data sets as summarized in Table 2. Thus, we met our goal of avoiding false-positives in groups by favoring coherence over completeness in the scripts.

The upper part in the table shows information about the session size by means of *number of changes* and *number of expected groups*. The middle part shows the accuracy as described above. It includes the number of calculated groups that *exactly match* the expected commits and the number of recommendations that should help reduce manual effort. The lower part approximates the manual effort including the minimal amount of click/drag-and-drop *actions* in the user interface.

The recall indicates that 33% to 57% of the relations, which express activity continuation, were detected by our scripts. For the data set V, we simply adapted the script to increase recall; we expect developers to do the same when facing many small groups. Regarding the number of groups compared to the expected commits, an expected commit is roughly spread over two to three groups. Since some commits were completely detected (see “# Exact Matches”), other commits are spread over more than three groups.

We explicitly counted the number of manual actions; this number can be approximated as two times the number of expected commits because this corresponds to the distribution of commits over proposed groups. For some data sets this is even less, which means that the given recommendations help reduce manual effort. Notice that the developer could also accept or reject *all* recommendations for a particular group with a single click.

Programming session VI was analyzed to give a first impression of the quality and reusability of THRESHER’s base scripts. We extracted patterns from sessions I to V, which had to be adapted to consider several corner cases in session V. A precision of 100% and a recall of 25% for session VI, however, indicate a satisfying result without further adaptation needed. The developer does only have to assemble the proposed groups with the help of recommendations and manual drag-and-drop interactions.

Table 1 Characteristics that rendered data sets interesting for deriving a foundation of scripts in THRESHER.

Characteristics	Data sets
Programming breaks between different refactorings (activity switch)	I II IV
Changed statements for debugging output or breakpoints (to be ignored)	I V
Consecutive changes in the same artifact	I II III IV V
Consecutive renamings of artifacts including updates of their uses in the code	IV
Consecutive reorganization of methods' protocols	I V VI
Changes of methods in different classes with shared protocol	I III V VI
Similar patterns in identifiers of artifacts	II III IV VI
Added/removed a method and update their uses	I II III IV V VI
Added/removed an instance/class variable and update their uses	I IV V VI
Added a class and initial methods/variables as well as initial uses	I VI
Removed a class and update their uses	I II III V VI

**Table 2 Evaluation results comparing manual analysis (expected groups) with the scripts of THRESHER (calculated groups).
[* Scripts adapted for V due to low recall]**

Data set	I	II	III	IV	V	V*	VI
Number of Changes	121	57	39	74	157	157	211
Number of Expected Groups	8	5	6	7	4	4	4
Accuracy							
Number of Calculated Groups	20	10	7	12	58	34	40
Number of Exact Matches	1	3	4	1	1	1	0
Number of Recommendations	20	4	2	21	32	31	22
Precision (%)	96	100	100	100	100	100	100
Recall (%)	35	43	57	35	1	33	25
Effort							
Minimum Interactions	242	114	78	148	314	422	
with THRESHER	40	20	14	24	116	48	80
Maximum Interactions	968	285	234	518	628	-	844
with THRESHER	160	50	42	84	232	96	160
Ratio (%)	17	18	18	16	37	15	19
Number of Manual Merge Actions	16	6	4	14	-	16	31

5.4 Discussion and Threats to Validity

After the manual analysis of all changes, we confirmed that THRESHER's scripts cannot run fully automated but developers have to intervene. One could have expected even better results because we designed the scripts to detect the groups as found manually in the first place. Still, our evaluation approach confirmed that even a simple set of scripts can support developers in the commit process. We argue that there can hardly be a fully-automated approach that works without any human intervention. On the one hand, several decisions rely on the developer's insights in the partic-

ular domain and project, which were *not materialized* in the change data. On the other hand, our scripts may not correctly reflect the characteristics found. Although we cross-checked several characteristics among us, there was not always consensus. For example, there were several method extractions across the code base and one of us argued to combine them into a *refactoring group* but another one would rather consider more *domain-specific information* and split them apart. We found similar conclusions in [6][11][29]. THRESHER is well suited to interactively help developers to decide which particular characteristics to pursue and when to merge

proposed groups.

We derived the scripts and tool requirements from only 5 programming sessions. Although we found an arguably representative spectrum of change characteristics for the Squeak/Smalltalk programming environment, other sessions may contain more special cases. As different kinds of activities influence the number and structure of changes, we should analyze more sessions that do not mainly deal with refactorings and bug fixing but also with feature implementations like classified by Hattori et al. [12]. Those could reveal further characteristics [27][38] such as more additions and less modifications of methods.

The concept of recommendations accounts for scripts' interdependence; it helps developers adjust the proposed groups before committing them. However, misleading recommendations may negatively affect the coherence of a change set. THRESHER can miss to recommend the actually correct group; we have no means to prevent that.

Our scripts detect activity continuations rather than switches and thus related changes rather than checking for "unrelatedness"; which is undecidable [15]. This may be contrary to how developers would manually select coherent change sets from a chronologically ordered list of recent changes. Additionally, our scripts do not cope with *successive activities* but would merge them into a single one. For example, the developer may rename methods as a refactoring (A) and then use those to implement a new feature (B succeeding A).

We reduced the implementation and computation effort of the scripts by prepending a *static analysis* of the methods' source code to add useful information such as changed references, variables, and message sends. When implementing such a tool in a different environment, the benefits of such caching depends on the existing change model.

Our scripts only detect low-level relations between changes. Future work should include more abstract characteristics such as design patterns [10] as previously investigated by [17][22]. Dynamic analysis could also reveal interesting relations by, for example, collecting concrete type information [13] to be used in scripts.

Applicability to other object-oriented languages and environments is straightforward because the concept of classes and methods is reflected in many

of our scripts. Languages with fundamentally different concepts, such as logic programming, may reveal quite different characteristics. The scripting language may also differ from the one used in the programming environment, which was not the case in Squeak/Smalltalk. We think that it is a good idea to provide THRESHER's level of support in other environments as well. Further research has to reveal whether there are any drawbacks if one has to design a custom scripting language for programmers to make the idea of THRESHER feasible in more static or closed environments such as Eclipse or Visual Studio.

We could also apply THRESHER directly to other programming languages. For example, we can modify the scripts of THRESHER to provide a level of support comparable to Historef [14], which targets files with Java source code. In general, there are ways to provide language-agnostic tool support in a Smalltalk environment because programming languages can be integrated as needed [7].

6 Hypotheses and Future Work

Based on our experience with implementing THRESHER in Squeak/Smalltalk and providing a foundation of scripts based on the lab study [33], we argue that this can only be the first step towards a tool that supports developers in any kind of software project. Our foundation of scripts as presented in Section 4.4 has to be stress-tested with additional data. At the time of writing, we derive the following hypotheses:

- Given a foundation of scripts that considers a particular programming language and its change model, developers will rarely have to adapt those scripts except for project-specific needs.
- If developers are knowledgeable about the scripting capabilities in such a commit tool, they are more likely to improve their personal workflow by taking advantage of that.
- If developers adapt such scripts for project-specific needs, they will make fewer mistakes and save time during the commit process.

Besides running experiments to test those hypotheses, we are eager to improve the general idea of how fine-granular changes might be combined into coherent and complete groups. In addition

to describing the *continuity of activities* between changes, we think that it is important to also split up bigger groups to indicate *activity switches* and thus separate commits. Finally, expressing rules of *dependency relationships* in scripts is valuable so that the correct commit order of change sets can be ensured, too.

Furthermore, we want to try out visualizations based on structures other than lists, tables, or trees. For example, recommendations across group boundaries might be better displayed with a graph similar to Figure 3. Finally, we have to generalize our scripting approach to other programming systems and their change models. File-based models, for example, may reveal different traits than object-based ones.

For all these improvements, we see a need for data from programming sessions that go beyond simple refactoring tasks such as adding features or fixing bugs. Only then, we can establish a useful scripting interface to support programmers to accommodate project-specific requirements.

7 Conclusion

We presented THRESHER, a scriptable, interactive, graphical tool that developers can use after programming sessions to identify coherent and complete sets of fine-granular changes to be committed to a VCS. A foundation of scripts automatically detects related changes that indicate activity continuations; developers then manually assemble those scattered groups into complete activity descriptions. We argue that it is very important to involve the actual developer into this process because there is much tacit knowledge about the project and its domain that cannot be manifested in code or change structures.

We showed that our foundation of scripts propose groups with a precision of almost 100% indicating that developers can rely on their coherence and focus on adjusting their completeness with respect to an activity. We argue that such tool support promotes the usage of fine-granular changes, which is beneficial for explorative programming strategies—omitting explicit checkpoints—and detailed tracking of project histories.

Acknowledgments

We wish to thank Jens Lincke, Lena Herscheid, Lysann Schlegel, Mark Rooney, Marko Röder, and Philipp Tessenow for fruitful discussions and valuable feedback. We gratefully acknowledge the financial support of HPI's Research School^{†10} and the Hasso Plattner Design Thinking Research Program.^{†11}

References

- [1] Barnett, M., Bird, C., Brunet, J. and Lahiri, S. K.: Helping developers help themselves: Automatic decomposition of code review changesets, in *Proceeding of the 37rd International Conference on Software Engineering (ICSE)*, ACM/IEEE, 2015.
- [2] Boehm, B. W.: The high cost of software, *Practical Strategies for Developing Large Software Systems*, (1975), pp. 3–15.
- [3] Bohner, S. A. and Arnold, R. S.: *Software Change Impact Analysis*, John Wiley & Sons, 1996.
- [4] Borg, M. and Runeson, P.: Changes, evolution, and bugs: Recommendation systems for issue management, in *Recommendation systems in software engineering*, Springer, 2014, pp. 477–509.
- [5] Buse, R. P. L. and Weimer, W. R.: Automatically documenting program changes, in *Proceedings of the 25th International Conference on Automated Software Engineering (ASE)*, IEEE/ACM, 2010, pp. 33–42.
- [6] Coman, I. D. and Sillitti, A.: Automated identification of tasks in development sessions, in *Proceedings of the 16th IEEE International Conference on Program Comprehension (ICPC)*, IEEE, 2008, pp. 212–217.
- [7] Stéphane, D., Gîrba, T., Lanza, M. and Demeyer Moose, S.: : A collaborative and extensible reengineering environment, in *Tools for Software Maintenance and Reengineering (RCOST)*, Software Technology Series, 2005.
- [8] Duvall, P. M., Matyas, S. and Glover, A.: *Continuous integration: Improving software quality and reducing risk*, Pearson Education, 2007.
- [9] Gabriel, R. P.: I Throw Itching Powder at Tulips, in *Proceedings of the ACM Symposium for New Ideas, New Paradigms, and Reflections on Everything to do with Programming and Software (Onward!)*, ACM, 2014, pp. 301–319.
- [10] Gamma, E., Helm, R., Johnson, R. and Vlissides, J.: *Design patterns: Elements of reusable object-oriented software*, Pearson Education, 1994.
- [11] Gold, N. and Mohan, A.: A framework for un-

^{†10} www.hpi.uni-potsdam.de/research_school

^{†11} www.hpi.de/en/research/design-thinking-research-program

- derstanding conceptual changes in evolving source code, in *Proceedings of the 19th International Conference on Software Maintenance (ICSM)*, IEEE, 2003, pp. 431–439.
- [12] Hattori, L. P. and Lanza, M.: On the nature of commits, in *Proceedings of the 23rd International Conference on Automated Software Engineering Workshops*, 2008, pp. 63–71.
- [13] Haupt, M., Perscheid, M. and Hirschfeld, R.: Type harvesting: A practical approach to obtaining typing information in dynamic programming languages, in *Proceedings of the 2011 Symposium on Applied Computing (SAC)*, ACM, 2011, pp. 1282–1289.
- [14] Hayashi, S., Hoshino, D., Matsuda, J., Saeki, M., Omori, T. and Maruyama, K.: Historef: A tool for edit history refactoring, in *Proceedings of the 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, IEEE, 2015, pp. 469–473.
- [15] Herzig, K. and Zeller, A.: Untangling changes. *Unpublished manuscript, September*, 2011. <http://www.st.cs.uni-saarland.de/publications/files/herzig-tmp-2011.pdf>, accessed: July 1, 2014.
- [16] Herzig, K. and Zeller, A.: The impact of tangled code changes, in *Proceedings of the 10th International Workshop on Mining Software Repositories (MSR)*, IEEE, 2013, pp. 121–130.
- [17] Heuzeroth, D., Holl, T., Hogstrom, G. and Lowe, W.: Automatic design pattern detection, in *Proceedings of the 11th International Workshop on Program Comprehension (IWPC)*, IEEE, 2003, pp. 94–103.
- [18] Kawrykow, D. and Robillard, M. P.: Non-essential changes in version histories, in *Proceedings of the 33rd International Conference on Software Engineering (ICSE)*, ACM/IEEE, 2011, p. 351.
- [19] Kersten, M. and Murphy, G. C.: Mylar: A Degree-of-Interest Model for IDEs, in *Proceedings of the 4th International Conference on Aspect-oriented Software Development (AOSD)*, ACM, 2005, pp. 159–168.
- [20] Kim, M. and Notkin, D.: Discovering and representing systematic code changes, in *Proceedings of the 31st International Conference on Software Engineering (ICSE)*, IEEE, 2009, pp. 309–319.
- [21] Ko, A. J., DeLine, R. and Venolia, G.: Information Needs in Collocated Software Development Teams, in *Proceedings of the 29th International Conference on Software Engineering (ICSE)*, ACM/IEEE, 2007, pp. 344–353.
- [22] Kramer, C. and Prechelt, L.: Design recovery by automated search for structural design patterns in object-oriented software, in *Proceedings of the 3rd Working Conference on Reverse Engineering (WCRE)*, IEEE, 1996, pp. 208–215.
- [23] Kuhn, A. and Stocker, M.: CodeTimeline: Storytelling with Versioning Data, in *International Conference on Software Engineering (ICSE)*, IEEE/ACM, 2012, pp. 1333–1336.
- [24] Meyer, A., Fritz, T., Murphy, G. C. and Zimmermann, T.: Software Developers’ Perceptions of Productivity, *SIGSOFT FSE*, ACM, 2014.
- [25] Murphy, G. C., Kersten, M., Robillard, M. and Cubranic, D.: the emergent structure of development tasks, in *Proceedings of the European Conference on Object-oriented Programming (ECOOP)*, Springer, 2005, pp. 734–749.
- [26] Robbes, R. and Lanza, M.: A change-based approach to software evolution, *Electronic Notes in Theoretical Computer Science*, Vol. 166 (2007), pp. 93–109.
- [27] Robbes, R. and Lanza, M.: Characterizing and understanding development sessions, in *Proceedings of the 15th International Conference on Program Comprehension (ICPC)*, IEEE, 2007.
- [28] Robbes, R., Pollet, D. and Lanza, M.: Logical coupling based on fine-grained change information, in *Proceedings of the 15th Working Conference on Reverse Engineering (WCRE)*, IEEE, 2008, pp. 42–46.
- [29] Robillard, M. P., Coelho, W. and Murphy, G. C.: How effective developers investigate source code: An exploratory study, *IEEE Transactions on Software Engineering*, Vol. 30, No. 12 (2004), pp. 889–903.
- [30] Sandberg, D. W.: Smalltalk and exploratory programming, *ACM SIGPLAN Notices*, Vol. 23, No. 10 (1988), pp. 85–92.
- [31] Sheil, B.: Environments for exploratory programming, *Datamation*, Vol. 29, No. 7 (1983), pp. 131–144.
- [32] Steinert, B., Cassou, D. and Hirschfeld, R.: CoExist: Overcoming aversion to change - Preserving immediate access to source code and run-time information of previous development states, in *Proceedings of the 8th Symposium on Dynamic Languages (DLS)*, ACM, 2012, pp. 107–118.
- [33] Steinert, B. and Hirschfeld, R.: How to compare performance in program design activities: Towards an empirical evaluation of CoExist, in *Design Thinking Research: Understanding Innovation*, Springer, 2014, pp. 219–238.
- [34] Taeumel, M., Perscheid, M., Steinert, B., Lincke, J. and Hirschfeld, R.: Interleaving of Modification and Use in Data-driven Tool Development, in *Proceedings of the ACM Symposium for New Ideas, New Paradigms, and Reflections on Everything to do with Programming and Software (Onward!)*, ACM, 2014, pp. 185–200.
- [35] Tao, Y., Dang, Y., Xie, T., Zhang, D. and Kim, S.: How do software engineers understand code changes?: An exploratory study in industry, in *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, ACM, 2012, p. 51.

- [36] van Rijsbergen, C. J.: *Information retrieval (second edition)*, Butterworths, 1979.
- [37] von Mayrhauser, A. and Vans, A. M.: Program comprehension during software maintenance and evolution. *Computer (IEEE)*, Vol. 28, No. 8 (1995), pp. 44–55.
- [38] Williams, B. J. and Carver, J. C.: Characterizing software architecture changes: A systematic review, *Information and Software Technology*, Vol. 52, No. 1 (2010), pp. 31–51.
- [39] Yoon, Y. S. and Myers, B. A.: Capturing and Analyzing Low-Level Events from the Code Editor, in *International Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU)*, 2011, pp. 25–30.
- [40] Yoon, Y. S., Myers, B. A. and Koo, S.: Visualization of Fine-grained Code Change History, in *Proceedings of the 2013 IEEE Symposium on Visual Languages and Human-centric Computing (VL/HCC)*, IEEE, 2013, pp. 119–126.
- [41] Zou, L. and Godfrey, M. W.: An industrial case study of Coman’s automated task detection algorithm: What worked, what didn’t, and why, in *Proceedings of the 28th International Conference on Software Maintenance (ICSM)*, IEEE, 2012, pp. 6–14.



Marcel Taeumel

Marcel Taeumel (marcel.taeumel@hpi.de) is a Research Assistant in the Software Architecture Group at the Hasso-Plattner-Institute (HPI).

He received a master’s degree in Software Engineering from the Hasso-Plattner-Institute, University of Potsdam, Germany. See also <http://marcel.taeumel.eu>.



Stephanie Platz

Stephanie Platz (stephanie.platz@student.hpi.de) is a Software Engineer at MarleySpoon GmbH, Germany. She received a master’s degree in Software Engineering from

the Hasso-Plattner-Institute, University of Potsdam, Germany.



Bastian Steinert

Bastian Steinert (bastian.steinert@gmail.com) is a Product Manager at Signavio GmbH, Germany. He received a Ph.D. in Computer Science from the Hasso-Plattner-

Institute, University of Potsdam, Germany. See also <http://www.bastiansteinert.org>.



Robert Hirschfeld

Robert Hirschfeld (robert.hirschfeld@hpi.de) is a Professor of Computer Science at the Hasso-Plattner-Institute (HPI) at the University of Potsdam, Germany. He received

a Ph.D. in Computer Science from the Technical University of Ilmenau, Germany. See also <http://www.hpi.uni-potsdam.de/swa>.



Hidehiko Masuhara

Hidehiko Masuhara (masuhara@acm.org) is a professor at the Department of Mathematical and Computing Science, Tokyo Institute of Technology. He received his

B.Sc., M.Sc., and Ph.D degrees from Department of Information Science, University of Tokyo in 1992, 1994, and 1999, respectively. His research interest is in programming languages and programming environments, especially advanced modularization mechanisms, optimization techniques, code recommendations, and debuggers. Before joining Tokyo Institute of Technology, he worked as an assistant professor, a lecturer, and an associate professor at Graduate School of Arts and Sciences, the University of Tokyo.