

HireBuild: An Automatic Approach to History-Driven Repair of Build Scripts

Anonymous Author(s)

ABSTRACT

Advancements in software build tools such as Maven reduce build management effort, but developers often still need specialized knowledge and to spend a long time to maintain build scripts and resolve build failures. More recent build tools such as Gradle give developers greater extent of customization flexibility, but can be even more difficult to maintain. According to the TravisTorrent dataset of open-source software continuous integration, 22% of code commits include changes in build script files to maintain build scripts or to resolve build failures. Automated program repair techniques have great potential to reduce cost of resolving software failures, but the existing techniques are mostly focusing on repairing the source code so that they cannot directly help resolving software build failures. To address this limitation, we propose HireBuild, the first approach to automatic patch generation for build scripts, using fix patterns automatically generated from existing build script fixes and recommending fix patterns based on build log similarity. From TravisTorrent dataset, we extracted 175 build failures and their corresponding fixes which revise Gradle build scripts. Among these 175 build failures, we used the 135 earlier build fixes for automatic fix-pattern generation and the rest later 40 build failures (fixes) for evaluation of our approach. Our experiment shows that our approach can fix 11 of 24 reproducible build failure, or 45% of the reproducible build failures within comparable time of manual fixes.

ACM Reference format:

Anonymous Author(s). 2017. HireBuild: An Automatic Approach to History-Driven Repair of Build Scripts. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference '17)*, 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Most well maintained software projects use build tools, such as Ant [36], Maven [24] and Gradle [14] to automate software building and testing process. Using these tools, developers can describe the build process of their projects with build scripts such as `build.xml` for Ant, `pom.xml` for Maven, and `build.gradle` for Gradle. With growing software size and functionality, build scripts can be complicated [23] and may need frequent maintenance. As software evolves, developers make changes to their code, test cases, system configuration, and dependencies, which may all lead to necessary changes in the build script. Adams et al. [5] found strong co-evolutionary

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference '17, July 2017, Washington, DC, USA
© 2017 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/Y/Y/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

relationship between source code and build script in their study. Hence build scripts need to be synchronized with source code and the whole build environment, and neglecting such changes in build scripts often leads build failures.

According to our statistics on TravisTorrent [9] dataset on the continuous integration of open-source software projects, 29% of code commits fail to go through a successful build on the integration server. Seo *et al.* [29] also mentioned a similar build failure proportion at Google, which is 37%. These build failures hinder projects development progress so that they need to be fixed as soon as possible. However, many developers do not have the required expertise to repair build scripts [28]. Therefore, automatic repair of build scripts can be desirable for software project managers and developers.

Automatic generation of software patches is an emerging technique, and has been addressed by multiple previous research efforts. For example, GenProg [13] and PAR [15] achieve promising result for automatic bug fixing. But these works are designed for repairing source code written in different programming languages. In contrast, repair of build scripts has its unique challenges. First, although the code similarity assumption (both GenProg and PAR are taking advantage of this assumption to fetch patch candidates from other portion of the project or other projects) still holds for build scripts, build-script repair often involves open knowledge that are not existing in the current project, such as a newly available version of a dependency or a build-tool plug-in (See Example 1). Second, unlike code bugs, build failures does not have a test suite to facilitate fault localization and to serve as the fitness function. Third, while different programming languages share similar semantics (so that code patterns / templates can be adapted and reused), the semantics of build scripts is very different from normal programs, so we need to re-develop abstract fix templates for build scripts.

Example 1 Gradle Version Dependency Change (*puniverse/quasar: 2a45c6f*)

```
task wrapper(type: Wrapper) {  
- gradleVersion = '1.11'  
+ gradleVersion = '2.0'  
}
```

On the other hand, there are also special opportunities we can take advantage of in the repair of build scripts. First, build failures often provide richer log information than normal test failures, and the build failure log can often be used to determine the reason and location of a build failure. Second, build scripts are programs in a specific domain, so it is possible to develop more specific build-fix templates (e.g., involving more domain-specific concepts such as versions, dependencies instead of general concepts like parameters, variables). Third, many build failures are related to the build tools and environments. These failures are not project-specific, so fix patterns can often be used crossing project boundary.

In this paper, we propose a novel approach, HireBuild, to generate patches for build scripts. Our insight is that, since many software projects use the same build tool (e.g., Gradle), similar build failures will result into similar build logs. Therefore, given a build failure, it is possible to use its build-failure log to locate similar build failures from a historical build-fix dataset, and adapt historical fixes for this new failures. Specifically, our technique consists of the following three phases. First, for a given build failure, based on build log similarity, we acquire a number of historically fixed build failures that have most similar build logs. We refer to these build fixes as seed fixes. Second, from build-script diff of the seed fixes, we extract a number of fix patterns based on our predefined fix-pattern templates for build scripts, and rank the patterns by their commonality among seed fixes. For build-script diff, our approach is based on an existing tool GumTree [11] which extracts changes of Java source code, XML, JavaScript code change. Third, we combine the patterns with information extracted from the build scripts and logs of the build failure to generate a ranked list of patch candidates, and the candidates can be applied to the build script until build is successful.

Although following the general generation-validation process for program repair, our technique is featured with following major differences to address the challenges and take advantage of the opportunities in build-script repair.

- **Build log analysis.** Build logs contain a lot of information about the location and reason of build failures, and sometimes even provide solutions. Our build log analysis parses build logs and extracts information relevant to build failures. Furthermore, HireBuild measures the similarity of build logs based on extracted information.
- **Build-fix-pattern templates.** There are a number of common domain-specific operations in build scripts, such as including / excluding a dependency, updating version numbers, etc. In HireBuild, we developed build-fix-pattern templates to involve these common operations specific to software build process.
- **Build validation.** In build-script repair, without test cases, we need a new measurement to validate generated patches. Specifically, we use the successful notification in the build log and the numbers of compiled source files to measure build successfulness.

In our work, we focus on repair of build scripts, so we do not consider compilation errors or unit-testing failures (although they also cause build failures) as they can be easily identified based on build logs and may be automatically repaired with existing bug repair techniques. Furthermore, we use Gradle (based on Groovy) as our targeted build tool as it is the most promising Java build tools now, and recent statistics [32] show that more than 50% of top GitHub apps have already switched to Gradle.

In our evaluation, we extracted 175 reproducible build fixes with corresponding build logs and build script changes from TraviTorrent dataset [9] on February 8, 2017, the build fixes are from 54 different projects). To evaluate HireBuild, we use the earlier 135 build fixes as our training set, and 40 later actual build failures (chronologically 135 earlier and 40 later bug fixes among the 175 regardless of which project they belong to) as our evaluations set. Among these 40 build

failures, we reproduced 24 build failure in our test environment. Empirical evaluation results show that our approach is able to generate fix for 11 of the 24 reproduced build script failures which gives same build output as developers' original fix. Overall, our work presented in the paper makes the following contributions.

- We developed a novel approach to automatic patch generation for repairing build scripts to resolve software build failures.
- We constructed a dataset of 175 build fixes which can serve as the basis and a benchmark for future research.
- We performed an empirical evaluation of our approach on the dataset of 175 real-world build fixes.
- We developed an Abstract-Syntax-Tree (AST) diff generation tool for Gradle build scripts, which potentially have more general applications.
- We generated build script fix code based on the abstract template of historical build fixes done by the developers.

The remaining part of this paper is organized as follows. After presenting a motivation example of how build-script repair is different from source-code repair in Section 2, we describe the design details of HireBuild in Section 3. Section 4 presents the evaluation of our approach, while in Section 5 presents discussion of important issues. Related works and Conclusion will be discussed in Section 6 and Section 7, respectively.

2 MOTIVATIONAL EXAMPLE

In this section, we introduce a real example from our dataset to illustrate how patch generation of build scripts is different from patch generation of source code. Example 2 shows a build failure and its corresponding patch where the upper part shows the most relevant snippet in the build-failure log and the lower part shows the code change to resolve the build failure. The project name and commit id are presented after the example title.

Example 2 A Gradle Build Failure and Patch (*puniverse/quasar: Build Failure Version:017fa18, Build Fix Version:509cd40*)

```
Could not resolve all dependencies for
  configuration ':quasar-galaxy:compile'.
> A conflict was found between the following
  modules:
- org.slf4j:slf4j-api:1.7.10
- org.slf4j:slf4j-api:1.7.7

compile ("co.paralleluniverse:galaxy:1.4") {
  exclude group: 'com.lmax', module: 'disruptor'
  exclude group: 'de.javakaffee', module: 'kryo-serializers'
  exclude group: 'com.google.guava', module: 'guava'
+ exclude group: "org.slf4j", module: '*'
}
```

In this build failure, build-failure log complains that there are two conflicting versions of `slf4j` module, and the bug fix is to add an exclusion of the module in the compilation of `Galaxy` component. Although this build fix is just a one-line simple fix, it can

illustrate differences between source-code repair and build-script repair. Specifically, we have the following observations.

First, it is possible to find from existing scripts or past fixes that we need to perform an `exclude` operation, however, since `org.slf4j` never appears in the script (it is transitively referred and will be downloaded from Gradle central dependency repository at runtime), the string “org.slf4j” can be hard to generate, and enumerating all possible strings is definitely not a feasible solution. The string can actually be generated by comparing the build-failure log and available modules in Gradle central dependency repository, but this is very different from source-code patching where all variable names to be referred to are already defined in the code (in the case when a generated fix contains a newly declared variable, the variable can have any name as long as it does not conflict with existing names in the scope).

Second, in build-script repair, we are able to, and need to consider build-specific operations. For example, we should not simply deem `exclude` as an arbitrary method name, but needs to involve its semantics into fix-pattern templates, so that we know a module name will follow the `exclude` command.

Third, the build log information is very important in that it not only provides the name of conflicting dependency, but also provide the compilation task performed when build failure happens, which can largely help patch generation tool to locate the build failure and determine where to apply the patch.

3 PROPOSED APPROACH

The overall goal of HireBuild is to generate build-script patches that can be used to resolve build failures. HireBuild achieve this goals with three steps: (1) log similarity calculation to find similar historical build fixes as seed fixes, (2) extraction of build-fix patterns from seed fixes, and (3) generation and validation of concrete patches for build scripts. In the following subsections, we first provide some preliminary knowledge for Gradle scripts, and then we describe the three steps of HireBuild with more details in the following subsections.

3.1 Gradle Build Tool

Gradle is a general purpose build management system based on Groovy and Kotlin [2]. Gradle supports the automatic download and configuration of dependencies or other libraries. It supports Maven and Ivy repositories for retrieving these dependencies. This allows reusing the artifacts of existing build systems.

A Gradle build may consist of one or more build projects. A build project corresponds to the building of the whole software project or a submodule. Each build project consists of a number of tasks. A task represents a piece of work during the building process of the build project, e.g., compile the source code or generate the Javadoc. A project using Gradle describes its build process in a `build.gradle` file. This file is typically located in the root folder of the project. In this file, a developer can use a combination of declarative and imperative statements in Groovy or Kotlin code. This build file defines a project and its tasks, and tasks can also be created and extended dynamically at runtime. Gradle is a general purpose build system hence this build file can perform any task.

3.2 Log Similarity Calculation to Find Similar Fixes

One most important characteristic of build script repair is that, a lot of software projects use the same build tools (e.g., Gradle), so that build-failure logs of different projects and versions often share the same format and output the similar error messages for similar build errors. So given a new build failure, HireBuild measures the similarity between its build-failure log and the build-failure logs of historical build failures to find its most similar build failures in history dataset.

3.2.1 Build Log Parsing. Gradle build logs typically consist of thousands of lines of text. Gradle prints these lines when performing different tasks such as downloading dependencies, compiling source files, and when facing errors during the build. Our point of interest is the error-and-exception part, which typically accounts for only a small portion of the build log. So if we use the whole build log to calculate similarity, the remaining part will bring a lot of noises to the calculation (e.g., build logs from projects that have similar dependencies may be considered similar).

Therefore, we use only the error-and-exception part of the build log to calculate similarity between build logs. An example of the error-and-exception part in Gradle build log is presented as below.

```
* What went wrong:
A problem occurred evaluating project ':android-rest'.

>
Gradle version 1.9 is required. Current version is
1.8. If using the gradle wrapper, try editing
the distributionUrl in /home/travis/build/47
deg/appsly-android-rest/gradle/wrapper/gradle-
wrapper.properties to gradle-1.9-all.zip
```

To extract the error-and-exception part, HireBuild extracts the portion of the build log after the error indicating header in Gradle (e.g., “* What went wrong”). HireBuild extracts only the last error, as the earlier ones are likely to be errors that are tolerated and are thus not likely to be the reason for the build failure. Furthermore, when there are exception stack traces in the error-and-exception part, HireBuild removes the stack traces for two reasons. First, stack traces are often very long, so they may dominate the main error message and bring noises (as mentioned above). Second, stack traces are often different from project to project so they cannot catch the commonality between build failures.

3.2.2 Text Processing. After we extracted the error-and-exception part from the build-failure log, we perform the following processing to convert the log text to standard word vector.

- **Text Normalization** breaks plain text into separate tokens and splits camel case words to multiple words.
- **Stop word Removal** removes common stop words, punctuation marks etc. For better similarity, HireBuild also removes common words for building process including ‘build’, ‘failure’, and ‘error’.
- **Stemming** is the process of reducing inflected words to their root word. As an example the word “goes” derived from word “go”. The stemming process converts “goes” to it’s root word

“go”. For stemming we applied popular Porter stemming algorithm [38].

3.2.3 Similarity Calculation. With the generated word vector from the error-and-exception part of the build-failure logs, we use the standard Term Frequency–Inverse Document Frequency (TF–IDF) [27] formula to weight all the words. Finally, we calculate cosine similarity between the log of build failure to be resolved and all build-failure logs of historical build fixes in our training set, and fetch the most similar historical fixes. We currently use the top 5 most similar historical fixes as seed fixes to generate build-fix patterns.

3.3 Generation of Build-Fix Patterns

To generate build-fix patterns, for each seed fix, we first calculate the code difference between the versions before and after the fix. The code difference consists of a list of elementary revisions including insertions, deletions and updates. Then, for each revision, we map it to one of five pre-defined build-fix pattern templates (described in Section 3.3.2), and generate a build-fix pattern. This pattern will be concretized to a set of candidates in the patch generation phase.

3.3.1 Build-Script Differencing. In this phase, for each seed fix, we extract Gradle build script commits before and after fix, and converted the script versions into AST representation. Gradle build is managed by special type of script which provides domain specific language to describe builds and this script is based on Groovy [2]. With Groovy support, AST representation of the script can be generated.

Our goal is to generate abstract representation of code changes between two commits. Having build script content represented as AST, we can apply tree difference algorithms, such as ChangeDistiller [12] or GumTree [11], to extract AST changes with sufficient abstraction. In our current implementation, we use GumTree to extract changes between two Gradle build scripts. GumTree generates a diff between two ASTs with list of actions which can be insertion, deletion, update and movement of individual AST nodes to transfer from source to destination. However, GumTree generates a list of AST revisions without node type information, so we revise GumTree to include the information. Furthermore, HireBuild also records the ancestor AST nodes of the changed AST subtree. Such ancestor AST nodes are typically the enclosing expression, statement, block, and task of the change, and they are helpful for merging different seed fixes for more general patterns, and for determining where the generated patches should be applied. As mentioned earlier, in Gradle scripts, a task is a piece of work which a build performs, and a script block is a method call with parameters as closure [3], so keeping such information helps to apply patches to certain block or task. Example 3 shows an exemplar output of HireBuild’s build-script differencing module, in which the operation, node type and ancestor nodes are extracted. Note that HireBuild extracts only one level of parent expression to avoid potential noises. As shown in the example, the task / block name can be empty if the fix is not in any tasks / blocks.

3.3.2 Hierarchical Build-Fix Patterns. In some rare cases we can directly use the concrete build-fix pattern to generate a correct patch. Example 4 provides such a patch from `Project`:

Example 3 Build Script Differencing Output (*BuildCraft/BuildCraft: 98f7196*)

```
1 <?xml version="1.0" encoding="UTF-8" standalone="no"?>
2 <patch>
3 <lineno id="30"><exp id="0">
4   <operation>Update</operation>
5   <nodetype>ConstantExpression</nodetype>
6   <nodeexp>1.7.2-10.12.1.1079</nodeexp>
7   <nodeparenttype>BinaryExpression</nodeparenttype>
8   <nodeparentexp>(version = 1.7.2-10.12.1.1079)
9 </nodeparentexp>
10  <nodeblockname>minecraft</nodeblockname>
11  <nodetaskname> </nodetaskname></exp>
12 </lineno>
13 </patch>
```

`nohana/Laevatein:a2aaca4`. There exists an exactly same build fix in the training set (from a different project).

Example 4 Training Project Fix (*journeysapps/zxing-android-embedded: 12cfa60*)

```
+ lintOptions {
+ abortOnError false
+ }
```

However, in more common scenarios, code diffs generated from seed fixes are too specific and cannot be directly applied as patches. Consider Examples 5 and 6, changes made in different project are similar, but if we consider concrete change of Example 5 as “Update 1.7.2-10.12.1.1079” then this change can hardly be applied to other scripts. Therefore we need to infer more general build-fix patterns from them.

Example 5 Gradle Build Fix (*BuildCraft/BuildCraft: 98f7196*)

```
- version = "1.7.2-10.12.1.1079"
+ version = "1.7.2-10.12.2.1121"
```

Example 6 Gradle Build Fix (*ForgeEssentials/ForgeEssentialsMain:fcbb468*)

```
-version = "1.4.0-beta7"
+version = "1.4.0-beta8"
```

Specifically, HireBuild infers a hierarchy of build-fix patterns from each seed fix by generalizing each element in the differencing output of the seed fix. The hierarchy of build-fix patterns generalized from Examples 5 and 6 are shown in Figure 1. From the figure, we can see that, HireBuild does not generalize operations and the node type of expression that are involved in the fix (i.e., `ConstantExpression`), because a change on those typically indicates a totally different fix. HireBuild also does not include the task and block information in the pattern as they are typically not a part of the fix. With the hierarchy, by choosing whether and which leaf node to be generalized, we can generate patches at different abstract levels. For example, if we generalize the parent expression from `version=1.7.2...` to `ParentExp: any`, we generate a pattern that updates a value `1.7.2...` without considering its parent. If we generalize both the

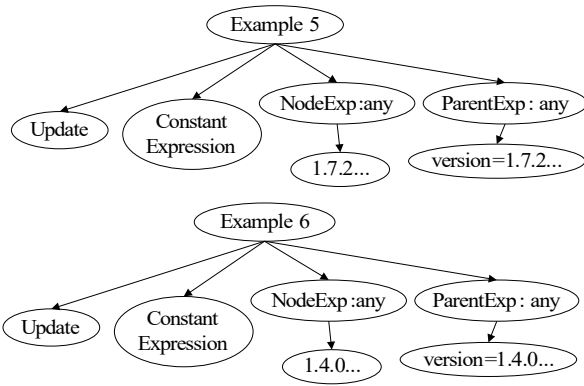


Figure 1: Hierarchies of Build-Fix Patterns




Figure 2: Merged Hierarchies

parent expression and the node expression, we generate a pattern that update any constants in the script. Note that HireBuild does not consider the cases where the node expression is generalized but the parent expression is not, as such a pattern can never match real code.

3.3.3 Merging of Build-Fix Patterns. After generating hierarchies of build-fix patterns, HireBuild first tries to merge similar hierarchies. For example, the two hierarchies in Figure 1 will be merged to a hierarchy shown in Figure 2. HireBuild merges only a pair of hierarchies with the same operation and node type (Update and Constant Expression in this case). During the merging process, HireBuild merges hierarchies recursively from their root node, and merges nodes with exactly the same value. If two nodes to be merged have different constant values, HireBuild does not merge them and their children nodes. If two nodes to be merged have different expression values, HireBuild extracts their corresponding AST tree, and merges the AST tree so that the common part of the expressions can be extracted. In Figure 2, since the expressions `version=1.7.2...` and `version=1.4.0...` share the same child nodes `version` and `=`, a node `version=*` is added. Note that more than two hierarchies can be merged in the same way if they share the same operation and node type.

3.3.4 Ranking of Build-Fix Patterns. After hierarchies are merged, HireBuild can calculate frequencies of build-fix patterns among seed fixes. If a hierarchy cannot be merged with other hierarchies, all the build-fix patterns in it have a frequency 1 among seed fixes, as they are specific to the seed fix they are from. In a merged hierarchy, the frequency of all its patterns is always the number of original hierarchies being merged. After calculating the frequencies of all build-fix patterns from hierarchies, HireBuild ranks build-fix patterns according to the frequency. For each build-fix pattern α , we counted n_{α}^t : α 's frequency among seed fixes. Then probability of α is as follows.

$$P_{\alpha} = \frac{n_{\alpha}^t}{N}$$

, where N is the total occurrences of build-fix patterns. Then, we rank the fix patterns based on the probability so that we use higher ranked build-fix patterns first to generate concrete patches. When there are ties between pattern A and B , if A is a generalization of B (A is generated by generalizing one or more leaf nodes of B), we rank B over A . The reason is that, when a build-fix pattern is generalized, it can lead to a larger number of concrete patches (e.g., update `gradleVersion` from any existing version to another existing version), so HireBuild needs to perform more build trials to exhaust all possibilities. As an example, all build-fix patterns from the hierarchy in Figure 2 have the same popularity, but the most concrete pattern: update constant expression with parent expression `version=*` will be ranked highest.

If there is no generalization relation between patterns, HireBuild ranks higher the build-fix patterns from the seed fix with higher ranking (the seed fix whose build failure log is more similar to that of the build failure to be fixed).

3.4 Generation and Validation of Concrete Patches

Before generation of concrete patches, we need to first decide which `.gradle` file to apply the fix. HireBuild uses a simple heuristic, which always choose the first `.gradle` file mentioned in the error part extracted from the build failure log. If no `.gradle` file is mentioned, HireBuild uses the `build.gradle` file in the root folder.

Given a build-fix pattern, and the buggy Gradle build script as input, to generate concrete patches, HireBuild first parses the buggy Gradle build script to AST, and then HireBuild tries to find where a patch should be applied.

For updates and deletions, HireBuild matches the build-fix patterns to nodes in the AST. For example, the build-fix pattern update constant expression with parent expression `version=*` can be mapped to an AST node of type `ConstantExpression` and its parent expression node has a value matching `version=*`. When a build-fix pattern can be mapped to multiple AST nodes (very common for general build-fix patterns), and HireBuild generates patches for all the mapped AST nodes. The only exception is when a build-fix pattern is mapped to multiple AST nodes in one block. In build scripts, within the same block, the sequence of commands typically does not matter, so HireBuild retains only the first mapped node in the block to reduce duplication.

For insertions, it is impossible to map a build-fix pattern to an existing AST node, so HireBuild matches the block and task names of the build-fix patterns to the buggy build script. When a build-fix pattern is generated from a hierarchy merging multiple seed-fixes, HireBuild considers the task and block names of all seed-fixes. If a task or block name in the buggy script is matched, HireBuild inserts the build patch at the end of the task or block.

After HireBuild determines which build-fix pattern to apply and where to apply, we finally need to concentrate on the abstract parts of the build-fix pattern and determine the possible values of the abstract nodes (e.g., determining the value to replace `*` in the pattern update constant expression with parent expression

465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522

523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580

version= *). The most commonly used values in build scripts are (1) identifiers including task names, block names, variable names, etc.; (2) names of Gradle plug-ins and third-party tools / libraries; (3) file paths within the project; and (4) version numbers. HireBuild first determines which type the value to added belongs to, base on the concrete values and AST nodes in the seed fixes leading to the build-fix pattern to be applied. HireBuild identifies version number and file paths based on regression expression matching (e.g., HireBuild can determine that `1.4.0-beta8` is a version number), and task / block / variable names by scanning the AST containing the seed fix. Other types of values including dependencies / plug-in names, and file paths are all specific to certain AST nodes so that they can be easily identified. Once the value type is determined, HireBuild generates values differently for different types as follows.

- Identifiers: HireBuild considers identifiers in the concrete seed fixes, as well as all available identifiers at the fix location.
- Names of plug-ins / libraries / tools: HireBuild considers names appearing in the concrete seed fixes, in the build failure log, and in the buggy build script.
- File paths: HireBuild considers paths appearing in the concrete seed fixes, in the build failure log, and in the buggy build script.
- Version numbers: HireBuild first locates the possible tools / libraries / plug-ins the version number is related to. This is done by searching for all occurrences of the version variable or constant in the AST of the buggy script. Once the tool / library / plug-in is determined, HireBuild searches Gradle central repository for all existing version numbers.

After the build-fix pattern, the location, and the concrete value are determined, a concrete patch is generated and put into the patch candidate list.

3.4.1 Ranking of Generated Patches. The previous steps generate a large number of patch candidates, so ranking of them is necessary to locate the actual fix as soon as possible. HireBuild ranks concrete patches with the following heuristics. Basically, we give higher priority to the patches which involve values or scopes more similar to the buggy script and the build-failure log.

- (1) Patches generated from higher ranked build-fix patterns are ranked higher than those generated from lower ranked build-fix patterns. The initial priority value of a patch is the probability value of its build-fix pattern.
- (2) If a patch p is to be applied to a location L , and p is generated from a build-fix pattern hierarchy merged from seed fixes A , B . If L resides in a task / block whose name is the same as the task / block name of A or B , HireBuild adds the priority of p by 1.0.
- (3) If a patch involves a value (any one of the four types described in Section 3.4) which appears in the build-failure log. HireBuild adds the priority of the patch by 1.0.
- (4) Rank all patches with updated priorities.

Note that, since the initial priority value is from 0 to 1, in the heuristics, we always add the priority value by 1.0 when certain condition meets, so that it go beyond all the other patches which do not satisfy the condition, no matter how high the initial priority value is.

3.4.2 Patch Application. After the ranked list of patches are generated, HireBuild applies the patches one by one until a timeout threshold is reached or the failure is fixed. HireBuild determines the failure is fixed if (1) the build process returns 0 and the build log shows build success, and (2) all source files that are compiled in the latest successfully built version are compiled if they are not deleted in between. We add the second criterion so that HireBuild can avoid trivial incorrect fixes such as changing the task to be performed from compile to clean up.

HireBuild generally focuses on one line fixes as most other software repair tool does. But it also includes a technique to generate multi-line patches if the failure is not fixed until all single line patches are applied. Multi-line patches can be viewed as a combination of single line patches, but it is impossible to exhaust the whole combination space. Example 7 shows a bug fix, which can be viewed as the combination of three one-line patches (two deletions and one insertion). To reduce the search space of patch combination, HireBuild considers only the combination that occurs in original seed fixes. Consider two one-line patches A and B , which are generated from hierarchies HA and HB . HireBuild considers the combination (A , B) only if HA and HB can be generalized from a same seed fix. After the filtering, HireBuild ranks patch combinations by the priority sum of the patches in the combination.

Example 7 Template with abstract node fix (*passy/Android-DirectoryChooser:27c194f*)

```
dependencies {
  compile files('libs/support-
    annotations-21.0.0-rc1.jar')
  - testCompile
    files('testlibs/roboelectric-2.4-
    SNAPSHOT-jar-with-dependen
    cies.jar')
  - androidTestProvided
    files('testlibs/roboelectric-2.4-
    SNAPSHOT-jar-with-dependen
    cies.jar')

  + androidTestCompile
    'org.roboelectric:roboelectric:2.3+'
  ...
}
```

4 EMPIRICAL EVALUATION

In this section, we first describe the construction of our dataset in Section 4.1. We present our experimental setting in Section 4.2, followed by research questions in Section 4.3, and experiment results in Section 4.4. Finally, we discuss the threats to validity in Section 4.5.

4.1 Dataset

We evaluate our approach to build-script repair on a dataset of build fixes extracted from the TravisTorrent dataset [9] snapshot at February 8, 2017. The tool and bug set used in our evaluation are all

Table 1: Dataset Summary

Type	Count
# Total Number of Projects	54
# Maximum Number of Fix From Single Project	25
# Minimum Number of Fix From Single Project	1
# Average Number of Fix Per Project	3.2
# Total Number of Fix	175
# Training Fix Size	135
# Testing Fix Size	40
# Reproducible Build Failure Size for Testing	24

available at our anonymized website ¹. TravisTorrent provides easy-to-use Travis CI build data to the masses through its open database. Though it provides large amount of build logs and relevant data, our point of interest is build status transition from error or fail status to pass status with changes in build scripts. From the version history of all projects in the TravisTorrent dataset, we identified as build fixes the code commits that satisfy: (1) the build status of their immediate previous version is fail / error; (2) the build status of the committed version is success; and (3) they contain only changes in gradle build scripts. Since HireBuild focuses on build script errors, we use code commits with only build-script changes so that we can filter out unit test failures and compilation failures. Our dataset may miss the more complicated build fixes that involve a combination of source-code changes and build-script changes, or a combination of build-script changes from different build tools (e.g., Gradle and Maven). HireBuild currently does not support the generation of such build fixes cross programming languages. Actually fix of such bugs are very challenging and is not supported by any existing software repair tools.

From the commit history of all projects, we extracted a dataset of 175 build fixes. More detailed information about our data set is presented in Table 1. We can see that these fixes are from 54 different projects, with maximal number of fixes in one project to be 25.

We ordered the build fixes according to the code commit time stamp, and use 135 (75%) earlier build fixes as the training set and the rest 40 build fixes (25%) as the evaluation set. **Therefore, all the build fixes in our evaluation set are chronically later than the build fixes in our training set.** Note that we combine all the projects in both training sets and evaluation sets, so our evaluation is cross-project in nature.

Among these 40 build fixes for evaluation, we successfully reproduced 24 build failures. The remaining 16 build failures can not be reproduced in our test environment for the following three reasons: (1) a missing library or build configuration file was originally missing from the central repository and caused the build failure, but they are added later; (2) an flawed third-party library or build configuration caused the build failure, but the flaws are fixed and flawed releases are no longer available on the Internet; and (3) the failure can be reproduced only with specific build commands and options which are not recorded in the repository. For case (3), we were able to reproduce some bugs by trying common build command options. We also contacted the TravisCI people about the availability of such commands / options, but they could not provide them to us.

4.2 Experiment Settings

TravisTorrent dataset provides Travis CI build analysis result as SQL dump and CSV format. We use SQL dump file for our experiment. We use a computer with 2.4 GHz Intel Core i7 CPU with 16GBs of

¹Gradle Build Script Fix Dataset and Tools: <https://sites.google.com/site/buildfix2017/>

Table 2: Project-wise Build Failure / Fix List

Project Name	#Failures	#Correctly Fixed
aol/micro-server	2	1
BuildCraft/BuildCraft	2	0
exteso/alf.io	1	1
facebook/rebound	1	1
griffon/griffon	1	0
/btrace	1	1
jMonkeyEngine/jmonkeyengine	2	0
jphp-compiler/jphp	1	0
Netflix/Hystrix	2	0
puniverse/quasar	6	2
RS485/LogisticsPipes	5	5
Total	24	11

Memory, and Ubuntu 14.10 LTS operating system. We use MySQL Server 5.7 to store build fix changes. In our evaluation, we use 600 minutes as the time out threshold for HireBuild.

4.3 Research Questions

In our research experiment, we seek to answer following research questions.

- **RQ1** How many reproducible build failures in the evaluation set can HireBuild fix?
- **RQ2** How many patches HireBuild generated and tried during the build-failure fixing?
- **RQ3** What are the amount of time HireBuild spends to fix a build failure?
- **RQ4** What are the sizes of build fixes that can be successfully fixed and that can not be fixed?
- **RQ5** What are the reasons behind unsuccessful build-script repair?

4.4 Results

RQ1: Number of successfully fixed build failures. In our evaluation, we consider a fix to be correct only if there is no build failure message in build log after applying patch, and the build result (all generated files including compiled source classes, doc files, linked third-party classes, etc.) are exactly the same as those generated by the manual fix. Among 24 reproducible build failures in the test set, we can generate the correct fix for 11 of them. Table 2 shows list of projects that are used for testing. Columns 2 and 3 represents number of build failures and number of those fixed by our approach.

Figure 3 shows the break down of successfully generated build fixes according the type of changes. With HireBuild, we can correctly generate 3 fixes about gradle option changes, 3 fixes about property changes, 2 fixes about dependency changes and external-tool option changes, respectively, and 1 fix about removing incompatible statements. Example 8 shows a build fix that is correctly generated by HireBuild falling in the category of external-tool option changes. The build failure is caused by adding a new option which is compatible only with Java 8. So the fix is to add an if condition to check the Java version. Note that this fix involves applying the combination of two insertion patches, but HireBuild still can fix it as there are a seed fix that contains both build-fix patterns.

RQ2: Patch list size. Patch list is a ranked list with all patches generated by HireBuild. Size of patch list has impact on automatic build script repair. If patch list size is too large, it will take large time span to generate fix sequence. We compare patch list size of

697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754

755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812

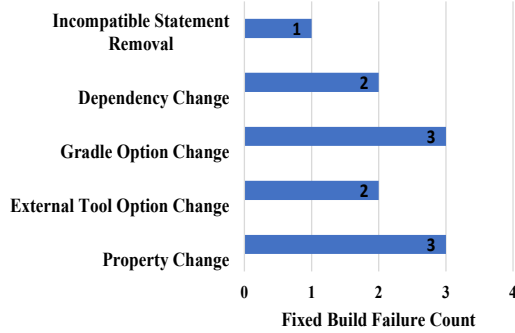


Figure 3: Breakdown of Build Fixes

Example 8 A Build Fix Correctly Generated By HireBuild (*puniverse/quasar:33bb265*)

```

+if (JavaVersion.current().isJava8Compatible()) {
+ tasks.withType(Javadoc) {
  options.addStringOption('Xdoclint:none', '-quiet')
+ }
+ }
    
```




Figure 4: Patch List Sizes

build failures we can correctly fix and patch list size of build failures we cannot correctly fix, and present the result in Figure 4. From the figure, we can see that for fixed build failures, the patch list has minimum size of 68 and maximum size of 2,245, while median is 486. For non fixed build failures, candidate list minimum, maximum and median are 8, 223 and 1,266 respectively, which are lower than fixable build failure's candidate list. The reason behind this result is that for non-fixable build failure, HireBuild cannot find similar build fixes in the training set, and thus the generated build-fix patterns cannot be easily mapped to AST nodes in buggy scripts.

RQ3: Time Spent on Fixes Time spent on fixes is very important for build failures as they need to be fixed timely. For the 11 fixable build failures, We compared in Figure 5 the time HireBuild spent on automatic fixing build failures with the manual fix time of the build failures in the commit history. We calculated the manual fix time by the time difference between the build commit where the build

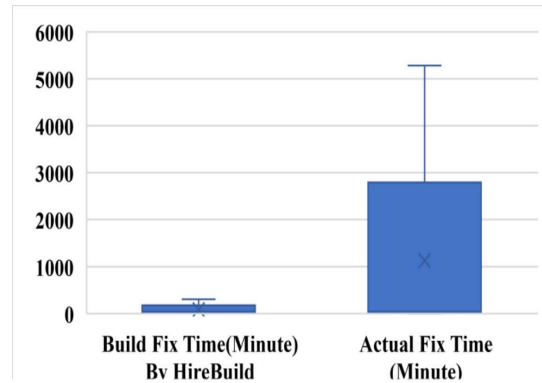


Figure 5: Amount of Time Required for Build Script Fix




Figure 6: Actual Fix Sizes

failure starts to happen and the build commit fixing the failure. From Figure 5, we can see that build script fix generated by our approach takes minimum 2 minute, maximum 305 minute and median value of 44 minute. While human fix takes minimum less than one minute, maximum 5,281 minute and median 42 minute. We can see that for the fixable build failures, HireBuild fixed them with time comparable to manual fixes.

RQ4: Actual Fix Size. Patch size has impact on automatic program repair. According to Bach et al. [17], bugs with over six lines of fix are difficult for automatic repair. In our dataset we have not performed any filtering based on actual fix size. But during result analysis, we performed statistical analysis to find out the sizes of build-fix patches, and the difference in size between the patches our approach can correctly generate and the patches our approach cannot. According to Figure 6, fixed build script failures contain minimum one, maximum two and median one. Actually 9 of the fixes contain only 1 statement change, and 2 of the fixes contain only 2 statement changes. For non fixed Gradle build script failures minimum change size is one, while maximum and median change size is 11 and 1 respectively. Therefore, our approach mainly works in the cases where the number of statement changes is small (1 or 2), which is similar to other automatic repair tools.

RQ5: Failing reasons for the rest 13 build failures. For 54.16% of evaluated build failures, our approach cannot generate build fix. So, we performed manual analysis to find out why our approach fails. We check whether the reason is related to generation of version numbers, dependency names, etc. Then we categorize these failure

Table 3: Cause of unsuccessful patch generation

Fix Type	#of Failures
Project specific change adaption	2(15%)
No matching patterns	6(46%)
Dependency resolution failures	3(23%)
Multi-location fixes	2(15%)

reasons to four major groups: (1) Project specific change adaption, (2) Non-matching patterns, (3) Dependency resolution failures, and (4) Multi-location fixes, as shown in Table 3.

Project specific change adaption indicates those changes that are dependent on project structure, file path etc. As build script manages build and its configuration, so there are project specific change issues and with our approach we can not adapt the build-fix patterns for. Example 9 shows such a build fix where it uses a specific path in build script.

Example 9 Project specific change (*Netflix/Hystrix:6600947*)

```
+ if ( dep.moduleName == 'servlet-api' ) {
+ it.artifactId[0].text() == dep.moduleName
  &&
+ asNode().dependencies[0].dependency.find {
+ ...
+ }}
+ }
```

Non-matching patterns indicates that our automatic patterns generation failed to provide required pattern that can resolve the build failure. HireBuild could not generate appropriate patterns for 6 failures which account 46% of failures. This may be due to limited size of training data and insufficient number of available build fixes that we used for template generation.

Dependency resolution failures happens for some project when HireBuild did not find actual dependency from central repository based on build log error. Even if we find dependency based on miss-compiled classes, that may not match with actual fixing. Example 10 shows such a dependency update where our approach failed to generate the dependency name.

Example 10 Dependency resolve Issue (*BuildCraft/BuildCraft:12f4f06*)

```
- mappings = 'snapshot_20160214'
+ mappings = 'stable_22'
```

Multi-Location Fixes happen when we need to apply multiple patches to fix a single build failure. HireBuild considers only limited combinations of patches as introduced in Section 3.4.2. Example 11 shows such a case where our patch generation technique generated the two “exclude” statements in two different patches. But this build failure is fixed only when we apply both “exclude” statement change simultaneously.

4.5 Threats of Validity

There are three major threats to the internal validity of our evaluation. First, there can be mistakes in our data processing and bugs in the implementation of HireBuild. To reduce this threat, we carefully

Example 11 Dependency resolve Issue (*joansmith/quasar:64e42ef*)

```
- jvmArgs '-Xbootclasspathp:
  ${System
    .getProperty('user.home')}jsr166.jar'
- testCompile 'org.testng:testng:6.9.6'
+ testCompile('org.testng:testng:6.9.6') {
+ exclude group: 'com.google.guava', module:
  '*'
+ exclude group: 'junit', module: '*'
+ }
```

double checked all the code and data in our evaluation. Second, the successful fixes generated by HireBuild may still have subtle differences with the manual fixes. Furthermore, the manual fixes that we use as the ground truth may in itself has flaws. To reduce this threat, we used a strict criterion for correct fixes. We need the automatically generated fix to generate exactly the same build results as those generated by the manual fix. Third, the manual fixing time collected in the commit history may be longer the actual fixing time as developers may choose to wait and not to fix the bug. We agree that this can happen but we believe the difference is not large as developers typically want to fix build failure as soon as possible so that it does not affect other developers.

The major threat to the external validity of our evaluation is that we use a evaluation set with limited number of reproducible bug fixes. Furthermore, our evaluation set contains only build fixes where only Gradle build scripts are changed. So it is possible that our conclusion is limited to the data set, Gradle-script fixes, or Gradle-script-only fixes. Figure 3 shows that our evaluation set already covers a large range of change types, and we plan to expand our evaluation set to more build failures and reproduce more bugs as TravisTorrent data set grows overtime. Gradle is the most widely adopted building system now, and its market share is still increasing. We also did a statistics on the number of build fixes with both Gradle-script changes and other file-changes and found 263 of them. Compared with 175 build fixes in our dataset, Gradle-script-only fixes accounts for a large portion of build script fixes for Gradle systems.

5 DISCUSSION

Source Code vs. Build Script Repair. In source code, there are lot of variation based on functionality and algorithm, on the other hand build script contains limited logic and variance. So, in that sense automatic build script repair has greater potentiality. One difference between source code repair and build script repair is the fault localization technique. In source code, we localize bugs with help of test code. But for build script we do not have such technique. Our approach of fault localization is totally dependent on fix template based on other projects.

Build Environment. Build environment defines the environment of a system that compiles source code, links module and generates assemblies. From developer’s point of view, they install all required dependencies like Java, GCC and other frameworks. But when projects are built in different environment then build problem can be generated. For example, if certain project has dependency on Java 1.8

then building the project in build environment with Java 1.7 might generate build failure. This is a challenge for build automation as well as automatic build repair. During software evaluation, developers change environment dependency based on functional requirements or efficiency. With changes version, developers build the software having those changed dependencies. But for build script repair, if we change version any dependency and keep the build environment as it was before, then fix might not resolve build failures. For Android projects environment, this issue creates greater impact as in most Gradle build script it mentions SDK version, build tool version etc. inside build script. As a result build script version dependency and build environment should be synced to avoid build breakage.

6 RELATED WORK

6.1 Automatic Program Repair

Automatic program repair is gaining research interest in the software engineering community with the focus to reduce bug fixing time and effort. Recent advancements in program analysis, synthesis, and machine learning have made automatic program repair a promising direction. Weimer et al. [13] GenProg which is one of the earliest and promising search based automatic patch generation technique based on genetic programming. Patch generated by this approach follows random mutation and use test case for the verification of the patch. Later in 2012, authors optimized their mutation operation and performed systematic evaluation 105 real bugs [19]. RSRepair [26] performs similar patch generation based on random search. S. Kim et al. [16] proposed an approach to automatic software patching by learning from common bug-fixing patterns in software version history, and later studied the usefulness of generated patches [35]. AE [37] uses deterministic search technique to generate patch. Pattern-based Automatic Program Repair (PAR) [15] uses manually generated templates learned from human written patch to prepare patch. PAR also used randomized technique to apply the fix patches. Nguyen et al. [25] proposed SemFix, which applied software synthesis to automatic program repair, by checking whether a suspicious statement can be re-written to make a failed test case pass. Le et al. [18] mines bug fix patterns for automatic template generation. Prophet [21] proposed probabilistic model learned from human written patched to generate new patch. The above mentioned approaches infer a hypothesis that new patch can be constructed based on existing source. This hypothesis also validated by Barr et al. [8] that 43 percent changes can be generated from existing code. With this hypothesis, we proposed first approach for automatic build failure patch generation. Tan and Roychoudhury proposed Relifix [34], a technique that taking advantage of version history information to repair regression faults. Smith et al. [31] reported an empirical study on the overfitting to test suites of automatically generated software patches. Most recently, Long and Rinard proposed SPR [20], which generates patching rules with condition synthesis, and searches for the valid patch in the patch candidates generated with the rules. Although our fundamental goal is same, but our approach is different than others in several aspects: 1) Our approach is applicable for build scripts, 2) We generate automatic fix template using build failure log similarity, 3) With abstract fix template matching we can generate fix candidate list with reasonable size.

6.2 Analysis of Build Configuration File

Analysis of build configuration file is growing as an important aspect for software engineering research such as dependency analysis for path expression, migration of build configuration file and empirical studies. On dependency analysis, Gunter [7] proposed a Petri-net based model to describe the dependencies in build configuration files. Adams et al. [4] proposed a framework to extract a dependency graph for build configuration files, and provide automatic tools to keep consistency during revision. Most recently, Al-Kofahi et al. [6] proposed an a fault localization approach for make files, which provides the suspiciousness scores of each statement in a make files for a building error. Wolf et al. proposed an approach [39] to predict build errors from the social relationship among developers. McIntosh et al. [22] carried out an empirical study on the efforts developers spend on the building configurations of projects. Downs et al. [10] proposed an approach to remind developers in a development team about the building status of the project. On the study of building errors, Hyunmin et al. [30] carried out an empirical study to categorize build errors at Google. Their study shows that missing types and incompatibility are the most common type of build errors, which are consistent with our findings.

The most closely related work in this category is SYMake developed by Tamrawi et al. [33]. SYMake uses a symbolic-evaluation-based technique to generate a string dependency graph for the string variables/constants in a Makefile, automatically traces these values in maintenance tasks (e.g., renaming), and detect common errors. Compared to SYMake, the proposed project plans to develop build configuration analysis for a different purpose (i.e., automatic software building). Therefore, the proposed analysis estimates run-time values of string variables with grammar-based string analysis instead of string dependency analysis, and analyzes flows of files to identify the paths to put downloaded files and source files to be involved. On migration of build configuration files, AutoConf [1] is a GNU software that automatically generates configuration scripts based on detected features of a computer system. AutoConf detects existing features (e.g., libraries, software installed) in a build environment, and configure the software based on pre-defined options.

7 CONCLUSION AND FUTURE WORK

For Source code, automatic patch generation research is already in good shape. Unfortunately, existing techniques are only concentrated to source code related bug fixing. In this work, we propose the first approach for automatic build fix candidate patch generation for Gradle build script. Our solution works on automatic build fix template generation based on build failure log similarity and historical build script fixes. For extracting build script changes, we developed GradleDiff for AST level build script change identification. Based on automated fix template we generated a ranked list of patches. In our evaluation, our approach can fix 11 out of 24 reproducible build failures.

In future, we plan to increase training and testing data size for better coverage of build failures with better evaluation. Apart from that, we are planning to apply search based technique such as genetic programming with fitness function on our patch list to better rank our generated patches and apply combination of patches.

REFERENCES

- 1161 [1] Gnu autoconf - creating automatic configuration scripts. <http://www.gnu.org/software/autoconf/manual/index.html>, accessed: 2015-10-25
- 1162 [2] The gradle build language. https://docs.gradle.org/current/userguide/writing_build_scripts.html, accessed: 2017-04-30
- 1163 [3] Gradle build script structure. <https://docs.gradle.org/3.5/dsl/>, accessed: 2017-04-30
- 1164 [4] Adams, B., Tromp, H., De Schutter, K., De Meuter, W.: Design recovery and maintenance of build systems. In: Software Maintenance, 2007. ICSM 2007. IEEE International Conference on. pp. 114–123 (Oct 2007)
- 1165 [5] Adams, B., Schutter, K.D., Tromp, H., Meuter, W.D.: The evolution of the linux build system. ECEASST 8 (2007), <http://dblp.uni-trier.de/db/journals/eceasst/eceasst8.html#AdamsSTM07>
- 1166 [6] Al-Kofahi, J., Nguyen, H.V., Nguyen, T.N.: Fault localization for build code errors in makefiles. In: Companion Proceedings of the 36th International Conference on Software Engineering. pp. 600–601. ICSE Companion 2014, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2591062.2591135>
- 1167 [7] Aoumeur, N., Saake, G.: Dynamically evolving concurrent information systems specification and validation: A component-based petri nets proposal. Data Knowl. Eng. 50(2), 117–173 (Aug 2004), <http://dx.doi.org/10.1016/j.datak.2003.10.005>
- 1168 [8] Barr, E.T., Brun, Y., Devanbu, P., Harman, M., Sarro, F.: The plastic surgery hypothesis. In: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering. pp. 306–317. ACM (2014)
- 1169 [9] Beller, M., Gousios, G., Zaidman, A.: Travisorrent: Synthesizing travis ci and github for full-stack research on continuous integration. In: Proceedings of the 14th working conference on mining software repositories (2017)
- 1170 [10] Downs, J., Plimmer, B., Hosking, J.G.: Ambient awareness of build status in collocated software teams. In: Proceedings of ICSE. pp. 507–517 (2012)
- 1171 [11] Falleri, J.R., Morandat, F., Blanc, X., Martinez, M., Monperrus, M.: Fine-grained and accurate source code differencing. In: Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering. pp. 313–324. ASE '14, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2642937.2642982>
- 1172 [12] Fluri, B., Wuersch, M., Plnzer, M., Gall, H.: Change distilling: tree differencing for fine-grained source code change extraction. IEEE Transactions on Software Engineering 33(11), 725–743 (Nov 2007)
- 1173 [13] Goues, C.L., Nguyen, T., Forrest, S., Weimer, W.: Genprog: A generic method for automatic software repair. IEEE Transactions on Software Engineering 38(1), 54–72 (Jan 2012)
- 1174 [14] Ikkink, H.K.: Gradle Dependency Management. Packt Publishing (2015)
- 1175 [15] Kim, D., Nam, J., Song, J., Kim, S.: Automatic patch generation learned from human-written patches. In: Proceedings of the 2013 International Conference on Software Engineering. pp. 802–811. ICSE '13, IEEE Press, Piscataway, NJ, USA (2013), <http://dl.acm.org/citation.cfm?id=2486788.2486893>
- 1176 [16] Kim, D., Nam, J., Song, J., Kim, S.: Automatic patch generation learned from human-written patches. In: Proceedings of the 2013 International Conference on Software Engineering. pp. 802–811 (2013)
- 1177 [17] Le, X.B.D., Lo, D., Goues, C.L.: History driven program repair. In: 2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER). vol. 1, pp. 213–224 (March 2016)
- 1178 [18] Le, X.B.D., Lo, D., Le Goues, C.: History driven program repair. In: Software Analysis, Evolution, and Reengineering (SANER), 2016 IEEE 23rd International Conference on. vol. 1, pp. 213–224. IEEE (2016)
- 1179 [19] Le Goues, C., Dewey-Vogt, M., Forrest, S., Weimer, W.: A systematic study of automated program repair: Fixing 55 out of 105 bugs for \$8 each. In: Proceedings of the 34th International Conference on Software Engineering. pp. 3–13. ICSE '12, IEEE Press, Piscataway, NJ, USA (2012), <http://dl.acm.org/citation.cfm?id=2337223.2337225>
- 1180 [20] Long, F., Rinard, M.: Staged program repair in spr. In: Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering (2015)
- 1181 [21] Long, F., Rinard, M.: Automatic patch generation by learning correct code. In: ACM SIGPLAN Notices. vol. 51, pp. 298–312. ACM (2016)
- 1182 [22] McIntosh, S., Adams, B., Nguyen, T., Kamei, Y., Hassan, A.: An empirical study of build maintenance effort. In: Software Engineering (ICSE), 2011 33rd International Conference on. pp. 141–150 (May 2011)
- 1183 [23] McIntosh, S., Adams, B., Hassan, A.E.: The evolution of java build systems. Empirical Softw. Engg. 17(4-5), 578–608 (Aug 2012), <http://dx.doi.org/10.1007/s10664-011-9169-5>
- 1184 [24] Miller, F.P., Vandome, A.F., McBrewster, J.: Apache Maven. Alpha Press (2010)
- 1185 [25] Nguyen, H.D.T., Qi, D., Roychoudhury, A., Chandra, S.: Semfix: Program repair via semantic analysis. In: Proceedings of the 2013 International Conference on Software Engineering. pp. 772–781 (2013)
- 1186 [26] Qi, Y., Mao, X., Lei, Y., Dai, Z., Wang, C.: The strength of random search on automated program repair. In: Proceedings of the 36th International Conference on Software Engineering. pp. 254–265. ICSE 2014, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2568225.2568254>
- 1187 [27] Ramos, J., et al.: Using tf-idf to determine word relevance in document queries. In: Proceedings of the first instructional conference on machine learning (2003)
- 1188 [28] Rausch, T., Hummer, W., Leitner, P., Schulte, S.: An empirical analysis of build failures in the continuous integration workflows of java-based open-source software
- 1189 [29] Seo, H., Sadowski, C., Elbaum, S., Aftandilian, E., Bowdidge, R.: Programmers' build errors: A case study (at google). In: Proceedings of the 36th International Conference on Software Engineering. pp. 724–734. ICSE 2014, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2568225.2568255>
- 1190 [30] Seo, H., Sadowski, C., Elbaum, S., Aftandilian, E., Bowdidge, R.: Programmers' build errors: A case study (at google). In: Proceedings of ICSE. pp. 724–734 (2014)
- 1191 [31] Smith, E., Barr, E., Goues, C.L., Brun, Y.: Is the cure worse than the disease? overfitting in automated program repair. In: Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering (2015)
- 1192 [32] Sulir, M., Porubán, J.: A quantitative study of java software buildability. In: Proceedings of the 7th International Workshop on Evaluation and Usability of Programming Languages and Tools. pp. 17–25. PLATEAU 2016, ACM, New York, NY, USA (2016), <http://doi.acm.org/libweb/lib.utsa.edu/10.1145/3001878.3001882>
- 1193 [33] Tamrawi, A., Nguyen, H.A., Nguyen, H.V., Nguyen, T.N.: Symake: A build code analysis and refactoring tool for makefiles. In: Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering. pp. 366–369. ASE 2012, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2351676.2351749>
- 1194 [34] Tan, S.H., Roychoudhury, A.: Relifix: Automated repair of software regressions. In: International Conference on Software Engineering (2015)
- 1195 [35] Tao, Y., Kim, J., Kim, S., Xu, C.: Automatically generated patches as debugging aids: A human study. In: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering. pp. 64–74 (2014)
- 1196 [36] Tilly, J., Burke, E.M.: Ant: The Definitive Guide. O'Reilly & Associates, Inc., Sebastopol, CA, USA, 1st edn. (2002)
- 1197 [37] Weimer, W., Fry, Z.P., Forrest, S.: Leveraging program equivalence for adaptive program repair: Models and first results. In: Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on. pp. 356–366. IEEE (2013)
- 1198 [38] Willett, P.: The porter stemming algorithm: then and now. Program 40(3), 219–223 (2006)
- 1199 [39] Wolf, T., Schroter, A., Damian, D., Nguyen, T.: Predicting build failures using social network analysis on developer communication. In: Proceedings of ICSE. pp. 1–11 (2009)
- 1200
- 1201
- 1202
- 1203
- 1204
- 1205
- 1206
- 1207
- 1208
- 1209
- 1210
- 1211
- 1212
- 1213
- 1214
- 1215
- 1216
- 1217
- 1218
- 1219
- 1220
- 1221
- 1222
- 1223
- 1224
- 1225
- 1226
- 1227
- 1228
- 1229
- 1230
- 1231
- 1232
- 1233
- 1234
- 1235
- 1236
- 1237
- 1238
- 1239
- 1240
- 1241
- 1242
- 1243
- 1244
- 1245
- 1246
- 1247
- 1248
- 1249
- 1250
- 1251
- 1252
- 1253
- 1254
- 1255
- 1256
- 1257
- 1258
- 1259
- 1260
- 1261
- 1262
- 1263
- 1264
- 1265
- 1266
- 1267
- 1268
- 1269
- 1270
- 1271
- 1272
- 1273
- 1274
- 1275
- 1276