

Current Challenges in Automatic Software Repair

Claire Le Goues¹ and Stephanie Forrest² and Westley Weimer¹

¹ University of Virginia, Charlottesville, VA 22904, USA,
{legoues,weimer}@cs.virginia.edu

² University of New Mexico, Albuquerque, NM 87131, USA,
forrest@cs.unm.edu

Abstract. The abundance of defects in existing software systems is unsustainable. Addressing them is a dominant cost of software maintenance, which in turn dominates the lifecycle cost of a system. Recent research has made significant progress on the problem of automatic program repair, using techniques such as evolutionary computation, instrumentation and run-time monitoring, and sound synthesis with respect to a specification. This article serves three purposes. First, we review current work on evolutionary computation approaches, focusing on GenProg, which uses genetic programming to evolve a patch to a particular bug. We summarize algorithmic improvements and recent experimental results. Second, we review related work in the rapidly growing subfield of automatic program repair. Finally, we outline important open research challenges that we believe should guide future research in the area.

Keywords: automatic program repair, software engineering, evolutionary computation

1 Introduction

Program evolution and repair are major components of software maintenance, which consumes a daunting fraction of the total cost of software production [1]. Although there are many tools available to help with bug triage (e.g., [2]), localization (e.g., [3,4]), validation (e.g., [5]) and even confirmation (e.g., [6]), generating repairs remains a predominantly manual, and thus expensive, process. The trend is clear: There is a pressing need for automatic techniques to supplement manual software development with inexpensive tools.

Research in *automated program repair* has focused on reducing repair costs by enabling continued program execution in the face of runtime errors (e.g., Juzi [7], ClearView [8], or Demsky *et al.* [9]); using code contracts or formal specifications to synthesize repairs (e.g., AutoFix-E [10], Axis [11], AFix [12], or Gopinath *et al.* [13]); or using evolutionary computation (EC) (e.g., by co-evolving test cases and repairs [14,15], or via language-specific operators and representations [16,17]). In this latter category, we introduced *GenProg* [18–20], which uses genetic programming (GP) to repair a wide range of defect types in

legacy software (e.g., infinite loops, buffer overruns, segfaults, integer overflows, format string vulnerabilities, and general incorrect output) without requiring *a priori* knowledge or formal specifications.

The breadth and depth of recent activity in this area is exciting. Automatic repair work has been evaluated by DARPA red teams [8] and won awards for human-competitive results produced by genetic and evolutionary computation [21]. Harman sums up the challenge succinctly: “If finding bugs is technically demanding and yet economically vital, how much more difficult yet valuable would it be to automatically fix bugs?” [22]

This article provides a high-level overview of the state of current research and existing challenges in automatic program repair, making several contributions. We begin with an update on the GenProg tool (in Section 2) and provide an overview and summary of recent experimental results (in Section 3). In conjunction with an overview of related work (Section 4), we use our experience to motivate a discussion of open research problems (Section 5), which we outline as challenges to the field. We conclude in Section 6.

2 GenProg

Over the past several years, we have described and evaluated several versions of *GenProg* [18–20] an automated method that uses *genetic programming (GP)* [23, 24] to search for a source-level patch³ that causes an off-the-shelf program to avoid a known defect while retaining key functionality. GP is a search technique based on the principles of biological evolution. As applied to program repair, GP maintains and evolves populations of program patches, seeking a patch that repairs the buggy behavior. In this section, we describe the current state of the algorithm, summarizing previously published work and highlighting recent improvements that enable GenProg to efficiently repair real bugs in large real-world programs.

2.1 Illustrative example

For the purposes of clarifying insights underlying our approach, we begin by presenting a running example adapted from a publicly available vulnerability report. Consider the pseudocode shown in Figure 1(a), adapted from a remote-exploitable heap buffer overflow vulnerability in the `nullhttpd v0.5.0` webserver. Function `ProcessRequest` processes an incoming request based on data copied from the request header. Note that on line 13, the call to `calloc` trusts the content length provided by a POST request, copied from the header on line 7. A malicious attacker can provide a negative value for `Content-Length` and a malicious payload in the request body to overflow the heap and kill or remotely gain control of the running server [25].

³ GenProg can also effect repairs in assembly code, binary files, and (recently) the llvm intermediate representation.

<pre> 1 char* ProcessRequest() { 2 ... 3 while(l=sgets(l,sock)) { 4 if(l=="Request:") 5 strcpy(req_type,l+12) 6 if(l=="Content-Length:") 7 len=atoi(l+16); 8 } 9 if(req_type=="GET") 10 buff=DoGETReq(sock,len); 11 if(req_type=="POST") { 12 sz=sizeof(char); 13 buff=calloc(len,sz); 14 rc=recv(sock,buff,len) 15 buff[len]='\0'; 16 } 17 return buff; 18 } </pre>	<pre> 3 ... 4 if(l=="Request:") 5 strcpy(req_type,l+12) 6 if(l=="Content-Length:") 7 len=atoi(l+16); 8 } 9 if(req_type=="GET") 10 buff=DoGETReq(sock,len); 11 if(req_type=="POST") { 12 + if (len <= 0) 13 + return null; 14 sz=sizeof(char); 15 buff=calloc(len,sz); 16 rc=recv(sock,buff,len) 17 buff[len]='\0'; 18 } 19 return buff; 20 } </pre>
(a) Webserver code snippet.	(b) Patched webserver.

Fig. 1: Pseudocode of a webserver that contains a bug (a), and a repaired version of the same program (b).

This buffer overflow vulnerability can be repaired fairly simply by adding a check on the content length before using it in the call to `calloc`. This candidate patch to the original program is shown in Figure 1(b).

At a high level, the goal of GenProg is to get from Figure 1(a) to Figure 1(b) automatically. In subsequent sections, we periodically refer back to this example to illustrate the algorithmic presentation and its underlying design insights. GenProg can also successfully repair the defect on which this example is based [18].

2.2 Overview

High-level pseudocode for GenProg’s main GP loop is shown in Figure 2. GenProg takes as input a program and a set of test cases that encode the bug (referred to as negative test cases) as well as required behavior that cannot be changed (the positive test cases). GenProg uses the test cases to localize the fault and compute context-sensitive information to guide the repair search (Section 2.6). Each individual, or variant, is represented as a *patch*, or a sequence of edit operations to the original program (Section 2.3). The goal is to produce a patch that causes the original program to pass all test cases.

Line 1 of Figure 2 initializes the population by creating a number of random initial patches. Lines 2–6 correspond to one iteration, or *generation*, of the algorithm. On Line 3, *tournament selection* [26] selects high-fitness (Section 2.4) individuals to be *parents* in the next generation. In general, fitness evaluation

Input: Full fitness predicate $\text{FullFitness} : \text{Patch} \rightarrow \mathbb{B}$
Input: Sampled fitness $\text{SampleFit} : \text{Patch} \rightarrow \mathbb{R}$
Input: Mutation operator $\text{mutate} : \text{Patch} \rightarrow \text{Patch}$
Input: Crossover operator $\text{crossover} : \text{Patch}^2 \rightarrow \text{Patch}^2$
Input: Parameter PopSize
Output: Patch that passes FullFitness
1: **let** $\text{Pop} \leftarrow \text{map } \text{mutate}$ over PopSize copies of $\langle \rangle$
2: **repeat**
3: **let** $\text{parents} \leftarrow \text{tournSelect}(\text{Pop}, \text{Popsize}, \text{SampleFit})$
4: **let** $\text{offspr} \leftarrow \text{map } \text{crossover}$ over parents , pairwise
5: $\text{Pop} \leftarrow \text{map } \text{mutate}$ over $\text{parents} \cup \text{offspr}$
6: **until** $\exists \text{candidate} \in \text{Pop}. \text{FullFitness}(\text{candidate})$
7: **return** candidate

Fig. 2: High-level pseudocode for the main GenProg loop; figure adapted from [19]. Typically, multiple *trials* (instances of the main repair loop) will be run in parallel for a given bug, with each trial initialized with a different random seed. Each trial is run either until a patch is found (line 6), or until a resource limit is reached (e.g., a certain number of iterations of the loop). The resource limit is externally checked and does not appear in the pseudocode.

dominates GenProg run time. The “parents” are selected pairwise at random to undergo *crossover*, in which a single point is chosen randomly and the subsequences up to the point in each parent are swapped, creating two new “offspring” variants. Each parent and each offspring is mutated once (*mutate*), and the result forms the incoming population for the next generation (Section 2.5). The GP loop terminates if a patch is found that causes the input program to pass all of its test cases, or when resources are exhausted (i.e., a predetermined time limit is exceeded). If GenProg succeeds in producing a repair, the resulting patch can be minimized in a deterministic post-processing step that combines tree-structured differencing [27] with delta-debugging [28]. Multiple executions of the algorithm are typically run in parallel for a given bug, each with a distinct random seed. Each execution, or run, is referred to as a *trial*.

In theory, the “best known” patch could be returned if no repair is found within the given time limit. Previous work has shown that developers address bug reports associated with a candidate patch more quickly than when no suggested patch accompanies the bug report, even if the proposed patch is incorrect [29]. A “partial solution,” or the best known patch found at a certain point, might serve as a useful guide for developers faced with repairing a bug by hand when the automated process fails. In an alternative use case, human advice or input might be solicited by the automated process when it is struggling, perhaps based on the “best known” patch to date. We have not yet investigated these scenarios in detail, but we speculate that they might provide alternative use cases to improve fault localization and debugging processes in practice.

2.3 Representation

The current version of GenProg represents each variant as a *patch*, or a sequence of edit operations with respect to the input program. In earlier work, and in keeping with a considerable proportion of the GP literature, GenProg represented an individual by its entire abstract syntax tree (AST), combined with a novel weighted execution path [30]. We subsequently found that the full AST representation limits scalability. For example, for at least 36 of the 105 defects in our largest dataset of real, historical defects in open-source programs [19], a population of 40–80 full ASTs did not fit in 1.7 GB of main memory. Over the same dataset, however, half of patches were 25 lines or less. Thus, two unrelated variants are likely to differ by at most 2×25 lines, with all other AST nodes in common. As a result, we now favor representing individuals as patches to avoid storing redundant copies of untouched code [19]. This design choice allows each individual in the population to be stored more compactly, and it scales sublinearly with the size of the code to which GenProg is being applied, a clear efficiency advantage.

Returning to our running example, one random individual in the population might correspond to “Delete the statement on line 10.” We index statements by assigning them unique integer values when the program is initially parsed, and thus the candidate patch can be represented as “Delete(N)”, where N is a unique identifying integer. This consumes much less storage than an entire secondary copy of the code, with the code from line 10, `buff=DoGETReq(sock,len);`, replaced by an empty block. To evaluate each candidate, the edits are applied to the input program in order to produce a new AST, whose fitness is measured as described in the next subsection.

2.4 Fitness

The *fitness* function guides a GP search. The fitness of an individual in a program repair task should assess how well the patch causes the program to avoid the bug while still retaining all other required functionality. We use test cases to measure fitness by applying a candidate patch to the original program and then rerunning the test suite on the result.

We typically take these test suites from the regression tests associated with many open-source projects. Regardless of its provenance, the input test suite should contain at least one case that initially fails, encoding the bug under repair, as well as at least one (but typically several) that initially pass, encoding required functionality that should be maintained post-patch.

For the running example, we write a test case that demonstrates the bug by sending a POST request with a negative content-length and a malicious payload to the webserver in order to try to crash it, and then we check whether the webserver is still running. Unmodified `nullhttpd` fails this test case.

However, defining desired program behavior exclusively by what we want `nullhttpd` to *not* do may lead to undesirable results. Consider the following variant of `nullhttpd`, created by a patch that replaces the body of the function with `return null`:

```
char* ProcessRequest() { return null; }
```

This version of `ProcessRequest` does not crash on the bug-encoding test case, but it also fails to process any requests at all. The repaired program should pass the error-encoding test case, but it must also retain core functionality before it can be considered acceptable. Such functionality can also be expressed with test cases, such as a regression test case that obtains `index.html` and compares the retrieved copy against the expected output.⁴

Running test cases typically dominates GenProg’s runtime, so we use several strategies to reduce the time to evaluate candidate patches. First, test cases can often be evaluated in parallel [30]. Second, our problem, like many GP problems, is tolerant of noisy fitness functions [31], which allows us to evaluate candidates on subsamples of the test suite. The function `SampleFit` evaluates a candidate patch on a random sample of the positive tests and on all of the negative test cases. For efficiency, only variants that maximize `SampleFit` are fully tested on the entire test suite (using `FullFitness`). The final fitness of a variant is the weighted sum of the number of tests it passes, where negative tests are typically weighted more heavily than the positive ones. This biases the search towards patches that repair the defect [20]. Programs that do not compile are assigned fitness zero.

We have experimented with several test suite sampling strategies and found that a random approach works well: The benefits gained by more precise sampling are outweighed by the additional computation time to select the samples. Sampling introduces *noise* into fitness evaluation, in that the value produced by `SampleFit` may differ from the value produced by `FullFitness` for the same individual. Too much noise could lead to more fitness evaluations over the course of the search. Although we have not been able to characterize the amount of noise `SampleFit` introduces across bug or benchmark type, our experiments show that it can vary from around 15% to as high as 70%. Overall, we observe the additional cost of increased sampling (more fitness evaluations required to find a successful repair) is strongly outweighed by the much smaller cost per evaluation achieved through sampling.

2.5 Mutation and Crossover

Mutation operates on AST nodes corresponding to C statements (e.g., excluding expressions or declarations), which limits the size of the search space. In each mutation, a destination statement d is chosen from the set of permitted statements according to a probability distribution (Section 2.6). In GenProg there are three distinct types of mutation, and the algorithm chooses randomly which one to apply. We have experimented with different mutation operators, but recent versions of GenProg use **delete**, **insert**, or **replace**. If **insert** or **replace** are selected, a second statement s is selected randomly from elsewhere in the same program. Statement d is then either replaced with s or with a new statement

⁴ In practice, we use several test cases to express program requirements. We describe only one here for brevity.

consisting of d followed by an inserted s . These changes are appended to the variant’s current list of edits.

Crossover selects two variants and exchanges subsequences between the two list of edits. The motivation for this operator is that valid partial solutions might be discovered by different variants, and crossover can combine them efficiently, helping to avoid local optima in the search. GenProg currently uses one-point crossover [32] as follows: Given parent individuals p and q , *crossover* selects crossover points p_n and q_m . The first portion of p is appended to the second portion of q , and vice versa, creating two offspring, both of which are evaluated by `SampleFit`.

2.6 Search space

Because the space of all possible edits to a program is so large, GenProg restricts the search to a smaller space that is likely to contain a repair. Consider again the bug in `nullhttpd` (Figure 1(a)). This code snippet represents only a small portion of the 5575-line program. Displaying all 5575 lines is unnecessary, however, because *not all program locations are equally likely to be good choices for changes to fix the bug*. Fault localization reduces the number of destination statements d that can be selected as locations for mutation.

Once a location for the mutation has been chosen, GenProg next selects the source statement s to be used as insertion or replacement code. We observe that *a program that makes a mistake in one location often handles a similar situation correctly in another* [33]. As a result, GenProg selects source statements s from code found elsewhere in the same program. This approach applies to `nullhttpd`. Although the POST request handling in `ProcessRequest` does not perform a bounds check on the user-specified content length, the `cgi_main` function, implemented elsewhere, does:

```
502 if (length <= 0) return null;
```

This code can be copied and inserted into the buggy region, as shown in the repaired version of the program (Figure 1(b)).

The search space for program repair is therefore defined by the locations that can be changed, the mutations that can be applied at each location, and the statements that can serve as sources of the repair. We parameterize these components of the search along two key dimensions:

Fault space. GenProg mutates only statements that are associated with incorrect behavior, and the statements are weighted to influence mutation probability. The input program is instrumented and executed to identify which statements are executed by which test cases. We initially believed that mutation should be biased heavily towards statements visited exclusively by the negative test cases [30]. However, we subsequently found that this intuition does not hold on our largest dataset: A uniform weighting, or one in which statements executed by both positive and negative test cases are weighted more heavily, was found to be preferable [20], although we do not consider this issue completely resolved.

Fix space. We use the term *fix localization* (or *fix space*) to refer to the *source* of insertion or replacement code. Candidate fixes are restricted to those within the original program, and they are currently restricted to statements visited by at least one test case (because we hypothesize that common behavior is more likely to be correct). In addition, GenProg rules out insertions that include variables that would be out-of-scope at the destination (to avoid type checking errors). Such localization improves search efficiency because it greatly reduces the proportion of generated variants that do not compile [34].

3 Evaluation

We have evaluated GenProg along several dimensions. We established *generality* by showing that GenProg can repair many different types of bugs in real-world programs, and we demonstrated *scalability* by showing that GenProg can repair programs containing millions of lines of code, without requiring special coding practices or annotations. We have characterized and improved the algorithm in both of these dimensions. Performing these evaluations has highlighted the challenge of developing benchmarks for automated defect repair, a problem we have approached from multiple angles. In this section, we summarize several of our recent evaluation efforts, focusing on high-level goals, results, and challenges. We elide some details in the interest of space, and direct the reader to associated publications where relevant.

3.1 Benchmarks

One of the greatest challenges we have faced in evaluating GenProg has been finding a good set of benchmark bugs and programs. Good benchmarks are critical to high-quality empirical science: “Since benchmarks drive computer science research and industry product development, which ones we use and how we evaluate them are key questions for the community.” [36] A good benchmark defect set should be *indicative* and *generalizable*, and it should therefore be drawn from a variety of programs representative of real-world systems. The defects should illustrate real bugs that human developers would consider important, and be easy to reproduce. Existing benchmark suites such as SPEC or Siemens [37] do not fulfill these requirements. The SPEC programs were designed for performance benchmarking and do not contain intentional semantic defects that are required for the automated repair problem. The Siemens suite does provide programs with test suites and faults. However, it was designed for controlled testing of software testing techniques, and therefore, the test suites maximize statement coverage, the faults are almost exclusively seeded, and the programs are fairly small.

A number of studies of automatic bug finding, localization, and fixing techniques have used bugs “in the wild,” found through case studies, careful search through bug databases, industrial partnerships, and word-of-mouth (e.g., [6, 8]). We have also taken this approach, identifying as broad a range of defects in as

Program	LOC Description	Fault
<code>gcd</code>	22 example	infinite loop
<code>zune</code>	28 example [35]	infinite loop†
<code>uniq utx</code>	1146 duplicate text processing	segmentation fault
<code>look utx</code>	1169 dictionary lookup	segmentation fault
<code>look svr</code>	1363 dictionary lookup	infinite loop
<code>units svr</code>	1504 metric conversion	segmentation fault
<code>deroff utx</code>	2236 document processing	segmentation fault
<code>nullhttpd</code>	5575 webserver	remote heap buffer overflow (code)†
<code>opendap</code>	292598 directory protocol	non-overflow denial of service†
<code>ccrypt</code>	7515 encryption utility	segmentation fault†
<code>indent</code>	9906 source code processing	infinite loop
<code>lighttpd</code>	51895 webserver	remote heap buffer overflow (vars)†
<code>flex</code>	18775 lexical analyzer generator	segmentation fault
<code>atris</code>	21553 graphical tetris game	local stack buffer exploit†
<code>pphp</code>	764489 scripting language	integer overflow†
<code>wu-ftpd</code>	67029 FTP server	format string vulnerability†
total	1246803	

Table 1: A set of benchmark programs used in experiments to evaluate GenProg’s generality, with size of the program measured in lines of code (LOC). The dataset contains bugs spanning 8 different fault types. A † indicates an openly available exploit. See <http://genprog.cs.virginia.edu/> for all benchmarks and source code used in our evaluations. Table adapted from [18].

many different types of programs as possible to substantiate our claim that GenProg is *general* (Table 1). for the benchmarks used in many of our studies). The programs total 1.25M lines of C code and the bugs in the dataset cover 8 different fault types; a number are taken from public vulnerability reports (indicated with a † in the table).

To enable large-scale evaluation of GenProg’s scalability and real-world utility, we recently developed a larger benchmark defect set, leveraging source control and regression tests suites of open-source C programs in a systematic way [19]. Given a set of popular programs from open-source repositories, we searched systematically through each program’s source history, looking for revisions that caused the program to pass test cases that failed in a previous revision. Such a scenario corresponds to a human-written repair for the bug defined by the failing test case. Table 2 summarizes the programs and defects in this dataset, which allows, to the best of our knowledge, the largest evaluation of automatic program repair to date. In total, it comprises 8 open-source C programs and 105 defects, with at least 2 defects per program.

We used this larger set to evaluate GenProg’s real-world utility [19] (i.e. what proportion of real bugs can be repaired automatically and the cost of a repair on publicly available cloud compute resources) and to conduct in-depth studies

Program	Description	LOC	Tests	Bugs
<code>fb</code>	legacy compiler for Basic	97,000	773	3
<code>gmp</code>	precision math library	145,000	146	2
<code>gzip</code>	data compression utility	491,000	12	5
<code>libtiff</code>	image manipulation library	77,000	78	24
<code>lighttpd</code>	lightweight web server	62,000	295	9
<code>php</code>	web programming language interpreter	1,046,000	8,471	44
<code>python</code>	general programming language interpreter	407,000	355	11
<code>wireshark</code>	network packet analyzer	2,814,000	63	7
total		5,139,000	10,193	105

Table 2: A benchmark set of subject C programs, test suites and historical defects, designed to allow large-scale, indicative, and systematic evaluation of automatic program repair techniques. Tests were taken from the most recent version available in May, 2011; Defects are defined as test case failures that were repaired by developers in previous versions. See <http://genprog.cs.virginia.edu/> for all benchmarks and source code used in the evaluations, including virtual machine images and pre-packaged bug scenarios that can be used to reproduce these defects. Table adapted from [19].

of critical algorithmic choices. [20] These latter studies allow us to ask questions about the nature of the search, how and why it works, why it does not always work, and how we may improve it, and are ongoing.

3.2 Generality

GenProg repairs all of the bugs in Table 1 in 356.5 seconds, on average, using relatively small sets of regression test cases (automatically or human-generated or taken from the existing test suites) on a machine with 2 GB of RAM and a 2.4 GHz dual-core CPU. These bugs cover a variety of defect types, including one of the earliest reported format string vulnerabilities (`wu-ftpd`). Of the sixteen patches, seven insert code, seven delete code, and two both insert and delete code. We note that patches that delete code do not necessarily degrade functionality, because the code may have been included erroneously, or the patch may compensate for the deletion with another insertion. Similarly, it is also possible to insert code without negatively affecting functionality, because the inserted code can be guarded so it applies only to relevant inputs (i.e., zero-valued arguments or tricky leap years).

Although a comprehensive code review is beyond the scope of this article, manual inspection (and quantitative evaluation, results not shown [18]) suggests that the patches are acceptable, in that they appear to address the underlying defect without introducing new vectors of attack. In our experiments and experience with patches that GenProg has produced, we observe that lost functionality in response to inadequate positive test cases appears more likely than

the introduction of new vulnerabilities. Overall, GenProg patches are typically highly localized in their effects.

Using commodity cloud resources, and limiting all repair runs to a maximum of 12 hours (simulating an overnight repair run), GenProg repaired approximately half of the bugs in Table 2, including at least one per program. Recall that this dataset was intended to evaluate real-world cost and expressive power. Modifying certain parameter values or changing various selection probabilities in the algorithm can influence GenProg’s ability to find a repair, especially for the more “difficult” repair scenarios (that is, those on which GenProg’s random success rate is lower on average). For example, altering the probability distribution used to select the mutation type, changing the crossover algorithm, and changing the fault and fix space weightings allowed GenProg to repair 5 new bugs when compared against a default baseline [20]. Similarly, running the repair algorithm for longer [38] causes GenProg to repair at least another 6 of the 105 scenarios, as compared to the 12-hour scenario.

We have investigated several of the important parameter, operator, and representation choices (see especially [20]), including two representations and four versions of the crossover operator. We also investigated the mutation operators and their selection probability as well as fault and fix space modifications and probability distributions. These investigations leave open the possibility of additional parameter sweeps in future work. Our results suggest additional avenues of future inquiry. For example, the patch representation (Section 2.3) appears to be more effective than the original abstract syntax tree/weighted path representation [20], but the mechanism behind this remains unknown. While the two representation choices encode the same types of changes, we hypothesize that differences in the way they are applied to the AST result in slightly different search space traversals in each case. Regardless, the GP algorithm that we use in GenProg is quite different from that typically used in many GP applications, a fact that motivates careful consideration of the various operator and parameter choices that underlie its implementation.

Overall, we view the successful and efficient repair of at least 8 different defect types in 16 programs and half of 105 systematically identified defects from programs totaling 5.1 million lines of code as a strong result, indicating the potential for generic automated repair algorithms such as GenProg.

3.3 Scalability and success

GenProg has repaired bugs in large programs with large test suites, as shown both in Table 1 and Table 2. We have found that the time to execute test cases dominates repair time (comprising 64% of the time to repair the benchmarks in Table 1, for example), which motivated our efforts to find ways to reduce the time necessary for fitness evaluation [31].

We are still working to characterize the conditions that influence GenProg’s success rate and time to repair. However, we have investigated a number of potential relationships. We consistently find a weak but statistically significant power law relationship between fault localization size and both time to repair and

probability of success [39, 18, 19]. As fault space size increases, the probability of repair success decreases, and the number of fitness evaluations required to find a repair (an algorithmic measure of search time) increases. We have also found a negative correlation between the fix space size and repair time. We speculate that larger fix spaces include more candidate repair options, thus reducing the time to find any given one.

We have also analyzed the relationship between repair success and external metrics such as human repair time and size, and defect severity. The only significant correlation we have identified using such metrics is between the number of files touched by a human-generated patch and repair success: The more files the humans changed to address the defect, the less likely GenProg was to find a repair. We have found no significant correlation between “bug report severity” and “GenProg’s ability to repair,” which we consider encouraging.

3.4 Example patch

In this subsection, we describe one patch produced by GenProg for a `php` bug from Table 2, and compare it to the one produced by humans for the same defect. We adapted this description from [19] for the purposes of illustrating the types of patches that GenProg can produce.

The `php` interpreter uses reference counting to determine when dynamic objects should be freed. User programs written in `php` may overload internal accessor functions to specify behavior when undefined class fields are accessed. Version 5.2.17 of `php` had a bug related to a combination of these features. At a high level, the “read property” function, which handles accessors, always calls a deep reference count decrement on one of its arguments, potentially freeing both that reference and the memory it points to. This is the correct behavior *unless* that argument points to `$this` when `$this` references a global variable—a situation that arises if the user program overrides the internal accessor to return `$this`. In such circumstances, the global variable has its reference count decremented to zero and its memory is mistakenly freed while it is still globally reachable.

The human-written patch replaces a line that always calls the deep decrement with a simple if-then-else: in the normal case (i.e., the argument is not a class object) calling the deep decrement as before, otherwise calling a separate shallow decrement function. The shallow decrement function will free the pointer, but not the object to which it points.

The GenProg patch adapts code from a nearby “unset property” function. The deep decrement is unchanged, but additional code is inserted to check for the abnormal case. In the abnormal case, the reference count is deeply incremented (through machinations involving a new variable) and then the same shallow decrement is called.

Thus, at a very high level, the human patch changes the call to `deep_Decr()` to:

```
1 if (normal) deep_Decr(); else shallow_Decr();
```

while the GP-generated patch changes it to:

```

1 deep_Decr();
2 if (abnormal) { deep_Incr(); shallow_Decr(); }

```

The logical effect is the same but the command ordering is not, and both patches are of comparable length. The human patch is perhaps more natural: It avoids the deep decrement rather than performing it and then undoing it.

4 Related Work

Automatic program repair and related problems have received considerable attention in recent years, including work on debugging and debugging assistance; error preemption, recovery, and repair; and evolutionary search, GP, and search-based software engineering.

Debugging. Work on debugging and debugging assistance focuses on identifying defects or narrowing the cause of a defect to a small number of lines of code. Recent debugging advances include replay debugging [40], cooperative statistical bug isolation [41], and statically debugging fully-specified programs [42]. Other techniques mine program history and related artifacts to suggest bug repairs or otherwise provide debugging support [43, 44]. Trace localization [45], minimization [46], and explanation [47] projects also aim to elucidate faults in the context of static defect detection.

Such work is best viewed as complementary to automated repair: A defect found or localized automatically could also be explained and repaired automatically. However, a common underlying assumption of such work is that unannotated programs must be repaired manually, albeit with additional information or flexibility presented to the developer. We propose several ways that these approaches might be extended or improved to improve automated repair, and related challenges, in Section 5.1.

Automated error preemption and defect repair. One class of approaches to automatic error handling uses source code instrumentation and runtime monitoring to prevent harmful effects from certain types of errors. Programs with monitoring instrumentation can detect data structure inconsistency or memory over- or under-flows. Various strategies are then used to enable continued execution [7, 9], generate trace logs, attack signatures and candidate patches for the system administrators [48, 49], or dispatch to custom error handlers [50, 51]. Jolt [52] assists in the dynamic detection of and recovery from infinite loops.

Other research efforts (including our own) focus directly on patch generation. At the binary level, ClearView [8] uses monitors and instrumentation to flag erroneous executions and generate candidate patches. AutoFix-E [10] uses contracts present in Eiffel code to propose semantically sound fixes. Gopinath *et al.* also use formal specifications to transform buggy programs [13]. Axis [11] and AFix [12] soundly patch single-variable atomicity violations, while Bradbury *et al.* propose to use GP to address concurrency errors [53]; PACHIKA [54] infers object behavior models to propose candidate fixes.

Many of these techniques are designed for particular types of defects, making use of pre-enumerated repair strategies or templates. Buffer overruns are particularly well-handled in the previous work, but overall generality remains a dominant research concern. Additionally, concurrency bugs remain a significant challenge. AFix and Axis are two of the only techniques to address them explicitly, although several of the other techniques can repair deterministic bugs in multi-threaded programs. However, nondeterministic bugs remain very difficult to test, and addressing that challenge is largely independent of integrating any solution into an automated repair framework that depends on testing. Some techniques, such as AutoFix-E, require specifications or annotations. While this enables semantic soundness guarantees, formal specifications are rare in practice [55]. We discuss several potentially fruitful research directions suggested by the current state of the art in program repair in the next section.

One major challenge in comparing these techniques for expressive power, generality, or generality utility is the relative dearth of agreed upon benchmark defects and experimental frameworks. Many researchers, ourselves included, identify bugs on which to test by combing through bug database histories, borrowing from other previous work, or following word-of-mouth. It can be difficult to reproduce the datasets from other work for direct comparison. We have begun to propose a process for identifying candidate benchmark defects (as described in Section 3.1 and [19]), but community consensus would go a long way towards enabling comprehensive comparative studies of both previously and newly proposed techniques.

SBSE and evolutionary computation. Search-Based Software Engineering (SBSE) [56] traditionally uses evolutionary and related methods for software testing, e.g., to develop test suites [57, 58]. These techniques typically focus on automatically generating high-coverage test suites [59], often with multiple competing search objectives (e.g., high-coverage tests, optimizing for a smaller test suite [60]). SBSE also uses evolutionary methods to improve software project management and effort estimation [61], find safety violations [62], and in some cases re-factor or re-engineer large software bases [63].

Arcuri [64] proposed to use GP to co-evolve defect repairs and unit test cases; the idea has since been significantly expanded [14, 15]. These approaches use competitive coevolution: The test cases evolve to find more bugs in the program, and the program evolves in response [65]. Techniques along these lines tend to rely at least in part on formal program specifications to define program correctness, and thus the fitness function [64, 66]. More work remains to increase the scalability, usability, and applicability of specification and verification-based approaches. We discuss potential extensions along these lines in the next section.

There has been considerable recent interest in and independent validations of the potential of GP [53] or random mutation [67] for program repair. It has been applied to new languages [16, 17, 68] and domains [69, 70], and has improved non-functional program properties, particularly execution time [71].

5 Open Research Challenges

The research results summarized in the previous sections, both our own and those of others, suggest that the prospects for automated software repair are promising. Transforming research results into practicality, however, raises important challenges, which we believe should guide future work. In this section we summarize the near-term challenges that we believe are most important to automated program repair. At a high level, these challenges fall into two broad categories:

- **Real-world practicality.** How can we transform automatic software repair research into widely used real-world software maintenance techniques? Challenges include scalability (how quickly they find repairs, how many lines of code they can handle) and generality (what sorts of programs and bugs they can address). An additional challenge is establishing the credibility of automated repairs in terms of programmers’ confidence in them and their understandability.
- **Theory of unsound repair methods.** How and why do current unsound repair techniques work? The success of existing approaches, particularly those that are unsound or stochastic (e.g., ClearView [8] or GenProg [19]) has been demonstrated empirically, but we lack a theoretical underpinning. The empirical results raise questions about the nature of extant software, its robustness in the face of random modifications, and how robustness and evolvability can be leveraged and enhanced.

The following subsections outline a set of specific challenges, each of which falls into one of these two general categories.

5.1 Adapting fault localization for automated repair applications

There exists considerable research on the problem of *fault localization* which automatically identifies (potentially) faulty code regions. In the context of automatic repair, these techniques identify likely locations for making a code modification, i.e., where the code is likely broken. We have consistently observed that well-localized bugs are more likely to be repaired and take less time to repair than poorly localized bugs. As a result, there is a concern that data-only bugs such as SQL injection vulnerabilities will be less suitable candidates for automated repair. Because the fault localization methods we have used to date are quite simple, there is more sophisticated methods (e.g., [72, 73]) could be incorporated to allow automated program repair of data-only bugs.

Contrary to our initial intuition, statements that are executed exclusively by the bug-inducing input may not be the best locations for a repair. This suggests the need to revisit our ideas about fault localization, because GenProg and related techniques may not benefit from increasingly precise identification of buggy statements. Rather, we need fault localization techniques that identify good candidates for code change to affect the buggy execution in a positive way,

without breaking non-buggy execution. This reconception of fault localization is a subtle but potentially important shift from its traditional purposes. It may be necessary to modify, extend, or reimagine existing fault localization techniques, designed for human consumption to help developers identify regions of code associated with a bug, to this new task of identifying regions of code that are good locations to change for repairing a particular bug.

Other extensions of the repair technique may require entirely novel innovations in fault localization. For example, Schulte *et al.* extended the GenProg approach to repair programs at the assembly level [74]. They proposed a novel stochastic, sampled fault localization technique to smoothly identify good candidate mutation locations along the assembly-level execution path. Applying the algorithm at other, different levels of abstraction (at the component or software architectural level, for example) will almost certainly demand similar types of innovation.

5.2 Automatically finding or generating code that is likely to repair software defects

Beyond the challenge of identifying good repair locations, it also is desirable to understand, formalize, and automatically predict how best to make a repair, addressing the *fix localization* problem. Progress in this area could increase the applicability and scalability of automated repair techniques, and it might improve our understanding of bug repair in general.

For example, GenProg currently restricts inserted code to statements that appear elsewhere in the same program. New programming paradigms or APIs, or novel bugs or vulnerability types, could stymie this paradigm. Considerable research attention has been devoted to mining information from source code or repositories [75]; this work could suggest augmented mutation operators including repair templates, even for novel bugs. We envision a technique that suggest candidate repairs by mining other related projects or modules, perhaps adapting methods from specification mining [76] or guided by software quality metrics [77, 78]. Mined templates could also potentially introduce expression- or declaration-level modifications in a domain-intelligent way, without necessarily expanding the size of the search space prohibitively. These types of modifications may enable the repair of more complex defects than GenProg and related techniques have addressed to date.

Even within the same project, identifying particularly promising code to serve as sources for automated repair remains a largely open problem. We took an initial step in this direction by restricting the genetic operators such that they do not move variables out of scope. Going beyond this general approach will likely involve special knowledge about the nature of both the bug and the program under repair. For example, in the security domain, specialized transformations have been proposed to repair or preclude buffer overruns [49, 50, 79]. If GenProg had reason to believe that the bug in question is a buffer overrun (because of the results of a pre-processing static analysis step, for example), it could apply such transformations with higher probability than it might otherwise.

The fix localization challenge affects all automated repair techniques, not just EC- or GP-based stochastic search methods, including those that use pre-specified templates and those that make semantically guaranteed mutations. Which repair template is best for a given defect type? This challenge is motivating other researchers in the area, and we expect it will continue to do so [?,80]. By accurately refining the search space for repairs and identifying more semantically expressive fixes, fix localization could enable more efficient searches for repairs and the repair of more complex bugs.

5.3 Formalizing patch quality and maintainability

Patch quality is an important impediment to the practical adoption of automated repair techniques. Human developers, who may never fully “leave the loop,” must be confident that a given patch correctly repairs a defect without violating other system requirements. This problem arises in all techniques that do not provide soundness guarantees. In the absence of fully automated programming, it is desirable for repair techniques, whether sound or unsound, to produce patches that humans can understand and later maintain. Formal models of repair quality would allow tools to present only the best repairs to humans or provide confidence estimates of the repair’s quality. They could provide useful insight about patches from any source, both human- and tool-generated.

Quantitative measures of quality could include existing techniques such as held-out test suites and black-box fuzzing or new techniques which have not yet been discovered. We anecdotally observe that GenProg is more likely to reduce program functionality in the face of inadequate test cases than it is to introduce a new malicious and exploitable vulnerability. Additional research in measuring, predicting, or ensuring functional patch quality might profitably begin by focusing on generating high-coverage tests to validate the behavior impacted by a patch, rather than on designing new operators to mitigate the probability that new vulnerabilities are introduced.

A formal model of quality, with related metrics, could also be used to analyze and improve existing techniques with repair quality in mind. Such a model would likely require human studies to understand and quantify what factors affect programmer understanding of a patch. Many programmers appear to have an instinct about whether a not a patch is “good.” The research challenge is in formalizing, quantifying, and ultimately predicting this judgment. We have previously published results that studied the impact of patches on certain measures of maintainability [81]. Much more work needs to be done in this general area, both to develop additional measures, study other repair methods, and improve existing repair search techniques. We believe that a full model that predicts, measures, and controls for the effect of a patch on higher-level system concerns will likely incorporate both syntactic and deeper semantic features of a patch.

Extending such an approach so it can *predict* the impact of a given change may allow automated repair techniques to optimize for both functional and non-functional patch properties. Tools like GenProg and ClearView can often generate multiple patches for the same bug, and a notion of patch quality could be

used to select from multiple options, or as an additional parameter in the fitness function.

Patch explanation or documentation techniques [82] could increase developer confidence in the output of automated tools, improving usability and adoption. Such work could also apply to human-generated patches. While existing change documentation techniques produces what may be referred to as “what” explanations, or explanations of the externally visible behavior the patch impacts and under what conditions, it may be possible to develop more helpful descriptions of what a patch is doing and why. For example, dataflow or slicing analyses may be able to describe the input and output variables affected by a proposed patch, or provide an execution “frontier” beyond which a patch will no longer have an effect. Alternatively, a patch could come with an automatically generated “justification” for each change that explains what happens (which test cases fail, for example) if it is removed.

5.4 Automatic generation of full test cases: inputs and *oracles*

Repair techniques such as GenProg use testing to guide a search or measure acceptability. Test suites, found much more commonly in practice than formal specifications, serve as a proxy for complete specifications and are used by many program analysis techniques (e.g. [3]). However, despite their prevalence in industrial practice, few programs are comprehensively tested, constraining the applicability of automated repair techniques that depend on test suites.

Although automated test case generation is a popular topic, most techniques intended for large programs produce test inputs, not full test cases (e.g., [83, 84]). These methods typically generate successive inputs that help maximize code coverage, for example, but are not full test cases in the sense of the oracle-comparator model [85]. Unless the program crashes, it is often difficult to determine whether the program has “passed” when presented with the input. There is promising recent work in this arena [86, 87], but we think there is much more to be done, both for automated testing and automatic program repair.

Another natural way to lift the test case assumption is to integrate GenProg with existing bug-finding tools, particularly those that identify faulty inputs or execution paths associated with buggy behavior [88]. Such approaches could mitigate the need for humans to identify the failing test cases that serve as observers for the bug in question. Reconstructing inputs or even execution traces from static bug finding tools is a known challenge [89]. The problem of creating high quality test oracles applies here as well. Previous work has suggested that program-specific temporal safety specifications are more useful for identifying bugs than library- or API-level specifications (such as “lock should always eventually be followed by unlock”) [90]. This suggests that advances in specification mining for extant software may prove useful in this domain [91, 76].

5.5 Combining unsound repair algorithms and formal methods

Today, few production programs come with a complete formal specification, and most commercial software is validated through test suites. Formal methods, however, have made enormous advances in the past decade, and there are several intriguing possibilities for integrating the two approaches. In GenProg, for example, the fitness function is sensitive only to whether or not a test case was passed, a binary signal. Measuring correctness only through test cases is a crude and possibly misleading approach. For example, consider a bug that requires the addition of the synchronization primitives `lock` and `unlock` around the use of a shared variable. A candidate patch that inserts only a `lock` will almost certainly lead to code that performs worse than inserting nothing at all, since it will likely lead to deadlock.

It is therefore desirable to have a finer-grained fitness function that can report how close a program is to passing a test case, especially the negative test cases. This might be accomplished by incorporating program analysis information, such as learned invariants [92–94], into the fitness function. In other work, we performed an initial study of such an approach [31]. We found that, on a particular case study, it was possible to use behavior as measured by such invariants to construct a fitness function that is more accurate than test cases alone [95]. However, the model we constructed was fairly opaque, and in subsequent efforts we were unable to generalize it to different bugs or programs. This is not to say that such a general model of intermediate patch correctness as profiled by dynamic invariants is impossible. Rather, using observed behavior on dynamic program predicates to characterize intermediate program behavior en route to a repair shows promise as an approach, but it will require additional insight before it can achieve general utility.

Other possibilities for integrating sound program analysis techniques into unsound approaches for program repair include automatically generating proofs for repairs in the context of proof carrying code [96], or automatically mutating code in a semantics-preserving way with a goal of simplifying automatic verification (because equisatisfiable verification conditions are not always equally easy to discharge) [97]. These potential applications could lead to higher-quality repairs as well as repairs that are easier for humans to validate and trust.

5.6 Robustness and evolvability as first-order concerns

GenProg and other evolutionary search techniques modify code using guided random mutations. Techniques such as ClearView use template-based repair systems that modify the code to avoid faulty behavior, without necessarily knowing *a priori* what the specified correct behavior should be (beyond universal “don’t crash” policies). Given that these techniques often succeed, in the sense that the software passes all tests [19], is just as maintainable by humans [81], and defeats red teams [8], they raise important questions about software robustness in the face of external perturbations (random or not). In this context, software is considered *robust* if it continues to function correctly in the face of perturbations to

its environment, operating conditions, or program code. In our own experience, software appears surprisingly robust to random statement-level changes, at both the abstract-syntax tree and assembly language levels, regardless of the coverage of its test suite [98]. Why is it that seemingly random mutations can improve software with respect to a set of test cases?

The success these unsound techniques challenges common definitions of *acceptability*: Continued execution of an adapted or evolved system may often be a better outcome than complete failure. This is the general approach taken by Rinard *et al.* in their failure-oblivious computing models, where programs are dynamically modified to continue executing in the face of a defect such as a buffer overrun or infinite loop [8, 52, 99]. Additional exploration of such ideas is warranted, both to characterize why current repair techniques can succeed at all and to gain better insight into the nature of software.

These observations apply to program repair as well as to software systems in general. Software today is deployed in highly dynamic environments. Nearly every aspect of a computational system is likely to change during its normal life cycle [100]. New users interact with software in novel ways, often finding new uses for old code; the specification of what the system is supposed to do changes over time; the owner and maintainers of the system are replaced by new developers; the system software and libraries on which the computation depends are upgraded or replaced; and the hardware on which the system is deployed is continuously modified.

A formal notion of software change and evolvability would open new avenues for understanding, prescribing, and predicting software behavior. It also suggests extensions to more traditional lines of software engineering research. One example might be new methods for developing and maintaining evolving test suites that continue to encode correct behavior for a given program as it evolves naturally throughout the software engineering process.

5.7 Fully automated software development

Beyond the goal of automated software repair lies the challenge of fully automating the entire software development process, including synthesizing complex programs from scratch using unsound techniques such as genetic programming. GenProg starts with a program that is almost correct and has a single identified defect. What would it take for automated repair to start with a partially debugged program and iteratively improve the program until all of the defects were eliminated, perhaps in the style of extreme programming [101]? Or more ambitiously, what technical impediments are there to using a technique like GenProg to synthesize full-length programs, perhaps beginning with a small set of functional test cases and then co-evolving the test suite with the program? We do not know the answers to these questions, and we consider this a long-term goal for automated software engineering rather than a near-term challenge. We note, however, that maintaining large systems often consists of iteratively debugging small sections of code with a goal of ever-increasing program functionality.

It is not beyond the realm of possibility to imagine a future GenProg, or another unsound method, developing software from scratch using a “programming as iterated debugging” development paradigm.

6 Conclusions

Software quality is an expensive problem, and the need for automated techniques for defect repair is pressing. Existing work in the area is promising, receiving attention from DARPA red teams and other external measures of success. In this article, we have described the current state of *GenProg*, a technique that uses genetic programming to repair unannotated legacy programs. It does so by evolving a set of changes that causes a given buggy program to pass a given set of test cases. In recent evaluations, we have established that GenProg is expressive and scalable, evaluated against a varied set of bugs as well as a large, systematically generated set of defects.

As a result of our initial experience with automated program repair, we have identified a number of design decisions that are critical to scalability and repair success, e.g., evolving patches rather than programs, focusing the search on commonly visited code areas, etc. In addition, we have identified a number of challenges that we believe lie at the heart of future progress in automated program repair: locating possible fixes, evaluating the quality of repairs, operating without full test suites or formal specifications, understanding change, and ultimately taking new steps towards automated software development.

We have released the source code for GenProg as well as our benchmark programs and defects at <http://genprog.cs.virginia.edu>. We hope other researchers focusing on bugs in legacy software will find utility in a relatively large set of objects to study, and that encourage others to download, extend, and compare against our tool. Overall, we hope that this work will motivate researchers to tackle this important and promising area, which is still in its infancy.

References

1. Seacord, R.C., Plakosh, D., Lewis, G.A.: Modernizing Legacy Systems: Software Technologies, Engineering Process and Business Practices. Addison-Wesley Longman Publishing Co., Inc. (2003)
2. Anvik, J., Hiew, L., Murphy, G.C.: Who should fix this bug? In: International Conference on Software Engineering. (2006) 361–370
3. Jones, J.A., Harrold, M.J.: Empirical evaluation of the Tarantula automatic fault-localization technique. In: Automated Software Engineering. (2005) 273–282
4. Saha, D., Nanda, M.G., Dhoolia, P., Nandivada, V.K., Sinha, V., Chandra, S.: Fault localization for data-centric programs. In: Foundations of Software Engineering. (2011)
5. Yin, Z., Yuan, D., Zhou, Y., Pasupathy, S., Bairavasundaram, L.N.: How do fixes become bugs? In: Foundations of Software Engineering. (2011) 26–36

6. Liblit, B., Naik, M., Zheng, A.X., Aiken, A., Jordan, M.I.: Scalable statistical bug isolation. In: *Programming Language Design and Implementation*. (2005) 15–26
7. Elkarablieh, B., Khurshid, S.: Juzi: A tool for repairing complex data structures. In: *International Conference on Software Engineering*. (2008) 855–858
8. Perkins, J.H., Kim, S., Larsen, S., Amarasinghe, S., Bachrach, J., Carbin, M., Pacheco, C., Sherwood, F., Sidiroglou, S., Sullivan, G., Wong, W.F., Zibin, Y., Ernst, M.D., Rinard, M.: Automatically patching errors in deployed software. In: *Symposium on Operating Systems Principles*. (2009)
9. Demsky, B., Ernst, M.D., Guo, P.J., McCamant, S., Perkins, J.H., Rinard, M.C.: Inference and enforcement of data structure consistency specifications. In: *International Symposium on Software Testing and Analysis*. (2006)
10. Wei, Y., Pei, Y., Furia, C.A., Silva, L.S., Buchholz, S., Meyer, B., Zeller, A.: Automated fixing of programs with contracts. In: *International Symposium on Software Testing and Analysis*. (2010) 61–72
11. Liu, P., Zhang, C.: Axis: Automatically fixing atomicity violations through solving control constraints. In: *International Conference on Software Engineering*. (2012) 299–309
12. Jin, G., Song, L., Zhang, W., Lu, S., Liblit, B.: Automated atomicity-violation fixing. In: *Programming Language Design and Implementation*. (2011)
13. Gopinath, D., Malik, M.Z., Khurshid, S.: Specification-based program repair using sat. In: *Tools and Algorithms for the Construction and Analysis of Systems*. Volume 6605 of *Lecture Notes in Computer Science*, Springer (2011) 173–188
14. Arcuri, A.: Evolutionary repair of faulty software. *Applied Soft Computing* **11**(4) (June 2011) 3494–3514
15. Wilkerson, J.L., Tauritz, D.R., Bridges, J.M.: Multi-objective coevolutionary automated software correction. In: *Genetic and Evolutionary Computation Conference*. (2012) 1229–1236
16. Orlov, M., Sipper, M.: Flight of the FINCH through the Java wilderness. *Transactions on Evolutionary Computation* **15**(2) (2011) 166–192
17. Schulte, E., Forrest, S., Weimer, W.: Automatic program repair through the evolution of assembly code. In: *Automated Software Engineering*. (2010) 33–36
18. Le Goues, C., Nguyen, T., Forrest, S., Weimer, W.: GenProg: A generic method for automated software repair. *Transactions on Software Engineering* **38**(1) (2012) 54–72
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: *International Conference on Software Engineering*. (2012) 3–13
20. Le Goues, C., Forrest, S., Weimer, W.: Representations and operators for improving evolutionary software repair. In: *Genetic and Evolutionary Computation Conference*. (2012) 959–966
21. Koza, J.R.: Awards for human-competitive results produced by genetic and evolutionary computation. <http://www.genetic-programming.org/hc2009/cfe2009.html> (2009)
22. Harman, M.: Automated patching techniques: the fix is in (technical perspective). *Communications of the ACM* **53**(5) (May 2010) 108
23. Koza, J.R.: *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press (1992)
24. Forrest, S.: Genetic algorithms: Principles of natural selection applied to computation. *Science* **261** (Aug 1993) 872–878

25. Nullhttpd: Bug: <http://nvd.nist.gov/nvd.cfm?cvename=CVE-2002-1496>. Exploit: <http://www.mail-archive.com/bugtraq@securityfocus.com/msg09178.html> (2002)
26. Miller, B.L., Goldberg, D.E.: Genetic algorithms, selection schemes, and the varying effects of noise. *Evolutionary Computing* **4**(2) (1996) 113–131
27. Al-Ekram, R., Adma, A., Baysal, O.: diffX: an algorithm to detect changes in multi-version XML documents. In: Conference of the Centre for Advanced Studies on Collaborative research, IBM Press (2005) 1–11
28. Zeller, A.: Yesterday, my program worked. Today, it does not. Why? In: Foundations of Software Engineering. (1999)
29. Weimer, W.: Patches as better bug reports. In: Generative Programming and Component Engineering. (2006) 181–190
30. Weimer, W., Nguyen, T., Le Goues, C., Forrest, S.: Automatically finding patches using genetic programming. In: International Conference on Software Engineering. (2009) 364–367
31. Fast, E., Le Goues, C., Forrest, S., Weimer, W.: Designing better fitness functions for automated program repair. In: Genetic and Evolutionary Computation Conference. (2010) 965–972
32. Rowe, J.E., McPhree, N.F.: The effects of crossover and mutation operators on variable length linear structures. In: Genetic and Evolutionary Computation Conference. (2001) 535–542
33. Engler, D.R., Chen, D.Y., Chou, A.: Bugs as inconsistent behavior: A general approach to inferring errors in systems code. In: Symposium on Operating Systems Principles. (2001)
34. Orlov, M., Sipper, M.: Genetic programming in the wild: Evolving unrestricted bytecode. In: Genetic and Evolutionary Computation Conference. (2009) 1043–1050
35. BBC News: Microsoft Zune affected by ‘bug’. In: <http://news.bbc.co.uk/2/hi/technology/7806683.stm>. (December 2008)
36. Blackburn, S.M., Garner, R., Hoffman, C., Khan, A.M., McKinley, K.S., Bentzur, R., Diwan, A., Feinberg, D., Frampton, D., Guyer, S.Z., Hirzel, M., Hosking, A., Jump, M., Lee, H., Moss, J.E.B., Phansalkar, A., Stefanović, D., VanDrunen, T., von Dincklage, D., Wiedermann, B.: The DaCapo benchmarks: Java benchmarking development and analysis. In: Object-Oriented Programming, Systems, Languages, and Applications. (2006) 169–190
37. Hutchins, M., Foster, H., Goradia, T., Ostrand, T.: Experiments of the effectiveness of dataflow-and control flow-based test adequacy criteria. In: International Conference on Software Engineering. (1994) 191–200
38. von Laszewski, G., Fox, G., Wang, F., Younge, A., Kulshrestha, A., Pike, G., Smith, W., Vckler, J., Figueiredo, R., Fortes, J., Keahey, K.: Design of the future-grid experiment management framework. In: Gateway Computing Environments Workshop. (2010) 1–10
39. Forrest, S., Weimer, W., Nguyen, T., Le Goues, C.: A genetic programming approach to automated software repair. In: Genetic and Evolutionary Computation Conference. (2009) 947–954
40. Albertsson, L., Magnusson, P.S.: Using complete system simulation for temporal debugging of general purpose operating systems and workload. In: International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems. (2000) 191

41. Liblit, B., Aiken, A., Zheng, A.X., Jordan, M.I.: Bug isolation via remote program sampling. In: *Programming Language Design and Implementation*. (2003) 141–154
42. He, H., Gupta, N.: Automated debugging using path-based weakest preconditions. In: *Fundamental Approaches to Software Engineering*. (2004) 267–280
43. Jeffrey, D., Feng, M., Gupta, N., Gupta, R.: BugFix: A learning-based tool to assist developers in fixing bugs. In: *International Conference on Program Comprehension*. (2009)
44. Ashok, B., Joy, J., Liang, H., Rajamani, S.K., Srinivasa, G., Vangala, V.: DebugAdvisor: A recommender system for debugging. In: *Foundations of Software Engineering*. (2009) 373–382
45. Ball, T., Naik, M., Rajamani, S.K.: From symptom to cause: Localizing errors in counterexample traces. *SIGPLAN Notices* **38**(1) (2003) 97–105
46. Groce, A., Kroening, D.: Making the most of BMC counterexamples. *Electronic Notes in Theoretical Computer Science* **119**(2) (2005) 67–81
47. Chaki, S., Groce, A., Strichman, O.: Explaining abstract counterexamples. In: *Foundations of Software Engineering*. (2004) 73–82
48. Smirnov, A., Chiueh, T.C.: Dira: Automatic detection, identification and repair of control-hijacking attacks. In: *Network and Distributed System Security Symposium*. (2005)
49. Smirnov, A., Lin, R., Chiueh, T.C.: PASAN: Automatic patch and signature generation for buffer overflow attacks. In: *Systems and Information Security*. (2006) 165–170
50. Sidiroglou, S., Giovanidis, G., Keromytis, A.D.: A dynamic mechanism for recovering from buffer overflow attacks. In: *Information Security*. (2005) 1–15
51. Sidiroglou, S., Keromytis, A.D.: Countering network worms through automatic patch generation. *IEEE Security and Privacy* **3**(6) (2005) 41–49
52. Carbin, M., Misailovic, S., Kling, M., Rinard, M.C.: Detecting and escaping infinite loops with Jolt. In: *European Conference on Object Oriented Programming*. (2011)
53. Bradbury, J.S., Jalbert, K.: Automatic repair of concurrency bugs. In: *International Symposium on Search Based Software Engineering - Fast Abstracts*. (Sept. 2010) 1–2
54. Dallmeier, V., Zeller, A., Meyer, B.: Generating fixes from object behavior anomalies. In: *Automated Software Engineering*. (2009) 550–554
55. Palshikar, G.: Applying formal specifications to real-world software development. *IEEE Software* **18**(5) (2001) 89–97
56. Harman, M.: The current state and future of search based software engineering. In: *International Conference on Software Engineering*. (2007) 342–357
57. Wappler, S., Wegener, J.: Evolutionary unit testing of object-oriented software using strongly-typed genetic programming. In: *Genetic and Evolutionary Computation Conference*. (2006) 1925–1932
58. Michael, C.C., McGraw, G., Schatz, M.A.: Generating software test data by evolution. *IEEE Transactions on Software Engineering* **27**(12) (2001) 1085–1110
59. Jia, Y., Harman, M.: An analysis and survey of the development of mutation testing. *Transactions on Software Engineering* **37**(5) (2011) 649–678
60. Lakhotia, K., Harman, M., McMinn, P.: A multi-objective approach to search-based test data generation. In: *Genetic and Evolutionary Computation Conference*. (2007) 1098–1105

61. Barreto, A., de O. Barros, M., Werner, C.M.: Staffing a software project: a constraint satisfaction and optimization-based approach. *Computers and Operations Research* **35**(10) (2008) 3073–3089
62. Alba, E., Chicano, F.: Finding safety errors with ACO. In: *Genetic and Evolutionary Computation Conference*. (2007) 1066–1073
63. Seng, O., Stammel, J., Burkhart, D.: Search-based determination of refactorings for improving the class structure of object-oriented systems. In: *Genetic and Evolutionary Computation Conference*. (2006) 1909–1916
64. Arcuri, A., Yao, X.: A novel co-evolutionary approach to automatic software bug fixing. In: *Congress on Evolutionary Computation*. (2008) 162–168
65. Adamopoulos, K., Harman, M., Hierons, R.M.: How to overcome the equivalent mutant problem and achieve tailored selective mutation using co-evolution. In: *Genetic and Evolutionary Computation Conference*. (2004) 1338–1349
66. Wilkerson, J.L., Tauritz, D.R.: A guide for fitness function design. In: *Genetic and Evolutionary Computation Conference*. (2011) 123–124
67. Debroy, V., Wong, W.E.: Using mutation to automatically suggest fixes for faulty programs. In: *International Conference on Software Testing, Verification, and Validation*. (2010) 65–74
68. Ackling, T., Alexander, B., Grunert, I.: Evolving patches for software repair. In: *Genetic and Evolutionary Computation*. (2011) 1427–1434
69. Sitthi-Amorn, P., Modly, N., Weimer, W., Lawrence, J.: Genetic programming for shader simplification. *ACM Transactions on Graphics* **30**(5) (2011)
70. Langdon, W.B., Harman, M.: Evolving a CUDA kernel from an nVidia template. In: *Congress on Evolutionary Computation*. (2010) 1–8
71. White, D.R., Arcuri, A., Clark, J.A.: Evolutionary improvement of programs. *Transactions on Evolutionary Computation* **15**(4) (2011) 515–538
72. Chen, M.Y., Kiciman, E., Fratkin, E., Fox, A., Brewer, E.: Pinpoint: Problem determination in large, dynamic Internet services. In: *International Conference on Dependable Systems and Networks*. (2002) 595–604
73. Abreu, R., Zoetevej, P., van Gemund, A.J.C.: An evaluation of similarity coefficients for software fault localization. In: *Pacific Rim International Symposium on Dependable Computing*, IEEE Computer Society (2006) 39–46
74. Schulte, E., DiLorenzo, J., Forrest, S., Weimer, W.: Automated repair of binary and assembly programs for cooperating embedded devices. In: *Architectural Support for Programming Languages and Operating Systems*. (2013)
75. Lanza, M., Penta, M.D., Xi, T., eds.: *IEEE Working Conference on Mining Software Repositories*. In Lanza, M., Penta, M.D., Xi, T., eds.: *MSR, IEEE* (2012)
76. Robillard, M.P., Bodden, E., Kawrykow, D., Mezini, M., Ratchford, T.: Automated api property inference techniques. *Transactions on Software Engineering* **99**(PP) (2012)
77. Buse, R.P.L., Weimer, W.: A metric for software readability. In: *International Symposium on Software Testing and Analysis*. (2008) 121–130
78. McCabe, T.J.: A complexity measure. *IEEE Trans. Software Eng.* **2**(4) (1976) 308–320
79. Barrantes, E.G., Ackley, D.H., Palmer, T.S., Stefanovic, D., Zovi, D.D.: Randomized instruction set emulation to disrupt binary code injection attacks. In: *Computer and Communications Security*. (2003) 281–289
80. Kim, D., Nam, J., Song, J., Kim, S.: Automatic patch generation learned from human-written patches. In: *International Conference on Software Engineering*, San Francisco, CA (2013 (To appear))

81. Fry, Z.P., Landau, B., Weimer, W.: A human study of patch maintainability. In: Heimdahl, M.P.E., Su, Z., eds.: *International Symposium on Software Testing and Analysis*. (2012) 177–187
82. Buse, R.P.L., Weimer, W.: Automatically documenting program changes. In: *Automated Software Engineering*. (2010) 33–42
83. Godefroid, P., Klarlund, N., Sen, K.: Dart: directed automated random testing. In: *Programming Language Design and Implementation*. (2005) 213–223
84. Sen, K.: Concolic testing. In: *Automated Software Engineering*. (2007) 571–572
85. Binder, R.V.: *Testing object-oriented systems: models, patterns, and tools*. Addison-Wesley Longman Publishing Co., Inc. (1999)
86. Fraser, G., Zeller, A.: Mutation-driven generation of unit tests and oracles. *Transactions on Software Engineering* **38**(2) (2012) 278–292
87. Fraser, G., Zeller, A.: Generating parameterized unit tests. In: *International Symposium on Software Testing and Analysis*. (2011) 364–374
88. Cadar, C., Ganesh, V., Pawlowski, P.M., Dill, D.L., Engler, D.R.: EXE: automatically generating inputs of death. In: *Computer and communications security*. (2006) 322–335
89. Jhala, R., Majumdar, R.: Path slicing. In: *Programming Language Design and Implementation*, New York, NY, USA, ACM Press (2005) 38–47
90. Weimer, W., Necula, G.C.: Mining temporal specifications for error detection. In: *Tools and Algorithms for the Construction and Analysis of Systems*. (January 2005) 461–476
91. Le Goues, C., Weimer, W.: Measuring code quality to improve specification mining. *IEEE Transactions on Software Engineering* **38**(1) (2012) 175–190
92. Ernst, M.D., Perkins, J.H., Guo, P.J., McCamant, S., Pacheco, C., Tschantz, M.S., Xiao, C.: The Daikon system for dynamic detection of likely invariants. *Sci. Comput. Program.* **69**(1–3) (2007) 35–45
93. Nguyen, T., Kapur, D., Weimer, W., Forrest, S.: Using dynamic analysis to discover polynomial and array invariants. In: *International Conference on Software Engineering*. (2012) 683–693
94. Gabel, M., Su, Z.: Testing mined specifications. In: *Foundations of Software Engineering*, ACM (2012) 1–11
95. Jones, T., Forrest, S.: Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In: *International Conference on Genetic Algorithms*. (1995) 184–192
96. Necula, G.C.: Proof-carrying code. In: *Principles of Programming Languages*, New York, NY, USA, ACM (1997) 106–119
97. Yin, X., Knight, J.C., Weimer, W.: Exploiting refactoring in formal verification. In: *International Conference on Dependable Systems and Networks*. (2009) 53–62
98. Schulte, E., Fry, Z.P., Fast, E., Forrest, S., Weimer, W.: Software mutational robustness: Bridging the gap between mutation testing and evolutionary biology. *CoRR* **abs/1204.4224** (2012)
99. Rinard, M.C., Cadar, C., Dumitran, D., Roy, D.M., Leu, T., Beebe, W.S.: Enhancing server availability and security through failure-oblivious computing. In: *Operating Systems Design and Implementation*. (2004) 303–316
100. Le Goues, C., Forrest, S., Weimer, W.: The case for software evolution. In: *Workshop on the Future of Software Engineering Research*. (2010) 205–210
101. Beck, K.: *Extreme Programming Explained: Embrace Change*. Addison-Wesley (2000)