

Defect Prediction & Fault Localisation

CS454 AI-Based Software Engineering

Shin Yoo

Defect Prediction

Defect Prediction

- “I have made one or more changes to my project. Which change is the most likely to break some tests?”
- Ideal: you want to know about the breakage as early as possible, as cheaply as possible.
- If you try to answer this question by changing the order of test executions, then we have test prioritisation problem (i.e., heuristics that try to execute the failing test case the first).
- If you try to answer this question **by discriminating changes**, you have defect prediction problem (i.e., which change is the most problematic?)

Traditional Formulation

(non-agile)

- You are about to enter the testing phase of the development cycle.
- Your project contains various changes.
- Which file/module is most likely to contain faults? If you can predict this correctly, your testing can be more effective, i.e., you will direct your testing efforts more towards these files and modules.
- But how do you know? What are the signs?

Predictive Modelling for Defectiveness

- There is no way of knowing about faults in advance, so the signals have to be *circumstantial* and *correlational*. What would you suspect?
- A file is more likely to be faulty if it _____.
 - “is very complex”
 - “has gone through a LOT of changes”
 - “is very long”
 - “is written by Shin”

Features for Predictive Modelling

- Ideally features should be:
 - Cheap to collect (i.e., static)
 - Credible source of correlation to defects
- Churn: rates of change
- Age: too old or too young can both contribute
- Change: the relative degree of change, or historical records

Table 4. List of Change metrics used in the study.

Metric name	Definition
REVISIONS	Number of revisions of a file
REFACTORINGS	Number of times a file has been refactored ¹
BUGFIXES	Number of times a file was involved in bug-fixing ²
AUTHORS	Number of distinct authors that checked a file into the repository
LOC_ADDED	Sum over all revisions of the lines of code added to a file
MAX_LOC_ADDED	Maximum number of lines of code added for all revisions
AVE_LOC_ADDED	Average lines of code added per revision
LOC_DELETED	Sum over all revisions of the lines of code deleted from a file
MAX_LOC_DELETED	Maximum number of lines of code deleted for all revisions
AVE_LOC_DELETED	Average lines of code deleted per revision
CODECHURN	Sum of (added lines of code – deleted lines of code) over all revisions
MAX_CODECHURN	Maximum CODECHURN for all revisions
AVE_CODECHURN	Average CODECHURN per revision
MAX_CHANGESET	Maximum number of files committed together to the repository
AVE_CHANGESET	Average number of files committed together to the repository
AGE	Age of a file in weeks (counting backwards from a specific release)
WEIGHTED_AGE	See equation (1)

Modern Formulation: Just-in-Time DP

Under Agile, we test all changes all the time, so...

- We now predict whether a specific commit is likely to break any tests.
- Features are still circumstantial; the prediction results can influence both code reviewing and testing efforts.
- The agile workflow as well as version control systems provide additional information.

Feature Sets in Just-in-Time DP

Bludau and Pretschner, PROMISE 2022 (<https://doi.org/10.1145/3558489.3559068>)

Table 1: State-of-the-art features for JIT defect prediction

	Feature	Description
Size	LA	Number of lines added by the change
	LD	Number of lines deleted by the change
	LT	Number of lines in the files before the change
Diffusion	NS	Number of modified subsystems
	ND	Number of modified directories
	NF	Number of modified files
	ENT	Spread of modified lines across modified files
History	NUC	Number of unique changes to all modified files
	NDEV	Number of past developers modifying the changed files
	AGE	Time between last and current change to modified files
Experience	EXP	Number of prior developer/reviewer changes
	REXP	Number of prior developer changes weighted by the age of the respective change
	SEXP	Number of prior changes to modified subsystems that an developer has participated in
	AWARE	Proportion of prior changes that an developer has participated in
Review	ITER	Number of reviews that were performed
	NREV	Number of persons, reviewing the changes
	NCOM	Number of review comments
	AGEREV	Time between the review request and the final approval

Table 2: Software development workflow-based features for JIT defect prediction

	Feature	Description
Size	DUR	Feature development time until the current change
	NTH	Position of this change within the feature
	C-LA	Number of added lines in the feature
	C-LD	Number of deleted lines in the feature
	SHARE	Share of current changes compared to whole feature development
Diffusion	C-NS	Number of modified subsystems in the feature
	C-ND	Number of modified directories in the feature
	C-NF	Number of modified files in the feature
	C-ENT	Spread of modified lines across modified files in the feature
Continuity	REL-NTH	Number of commits since last release
	REL-DUR	Passed time since last release
	DEV-NTH	Number of commits since last developer change
	DEV-DUR	Passed time since last developer change
Comment	PRCOM	Number of comments
	PRCOM-R	Number of comment reactions

Table 3: AST-based features for JIT defect prediction

	Feature	Description
Methods	FUN	Number of functions in all files before the change
	FUNT	Size of all functions before the change
	FUNDIFF	Difference to FUNT after the change
	FUNA	Number of added functions
	FUND	Number of removed function
	FUNU	Number of changed functions
Nodes	ASTA	Number of added nodes to the AST
	ASTD	Number of removed nodes from the source AST
	ASTU	Number of updated nodes between both AST versions
	SASTA	Number of added special nodes to the AST
	SASTD	Number of removed special nodes from the AST
	CASTA	Number of added comment nodes to the AST
	CASTD	Number of removed comment nodes from the AST
	PASTA	Number of added primitive nodes to the AST
	PASTD	Number of removed primitive nodes from the AST
Depth	DEPTHA	Maximal depth of added nodes in the AST
	DEPTHD	Maximal depth of removed nodes in the AST
	DEPTHU	Maximal depth of updated nodes in the AST

What else?

We need a predictive model, and the supervision.

- For the features discussed so far, any classical classifier model would work.
- This is a supervised learning, i.e., you try to predict the label (“this change actually broke tests / this change was okay”) based on these features.
- To adopt DP, you need a good issue tracking: unless you have commit-test breakage data, you cannot train any model.
- From this, another popular theme in DP research emerges: **cross-project DP**, i.e., you try to train the model with project A and then apply it to project B. This is to mitigate the **cold-start** problem.

Deep Learning for DP

We hope to work out better features.

- First wave: use a similar feature set, but with DNNs, so that meaningful internal features can be extracted.
- Second wave: use embedding to directly represent the unit (either commits, or sometimes even lines), and build classifier on top of the embedding vectors.
- Important questions remain, despite success of varying degrees.
 - The cost of modelling vs. the cost saving of DP
 - Circumstantial or Semantic? Which is more feasible?

Input Reduction

A Motivating Example

```

        <td align=left valign=top>
            <SELECT NAME="op sys" MULTIPLE SIZE=7>
<OPTION VALUE="All">All<OPTION VALUE="Windows 3.1">Windows 3.1<OPTION VALUE="Windows 95">Windows 95<OPTION VALUE="Windows
  98">Windows 98<OPTION VALUE="Windows ME">Windows ME<OPTION VALUE="Windows 2000">Windows 2000<OPTION VALUE="Windows
  NT">Windows NT<OPTION VALUE="Mac System 7">Mac System 7<OPTION VALUE="Mac System 7.5">Mac System 7.5<OPTION VALUE="Mac
  System 7.6.1">Mac System 7.6.1<OPTION VALUE="Mac System 8.0">Mac System 8.0<OPTION VALUE="Mac System 8.5">Mac System
  8.5<OPTION VALUE="Mac System 8.6">Mac System 8.6<OPTION VALUE="Mac System 9.x">Mac System 9.x<OPTION VALUE="MacOS X">MacOS
  X<OPTION VALUE="Linux">Linux<OPTION VALUE="BSDI">BSDI<OPTION VALUE="FreeBSD">FreeBSD<OPTION VALUE="NetBSD">NetBSD<OPTION
    VALUE="OpenBSD">OpenBSD<OPTION VALUE="AIX">AIX<OPTION
      VALUE="BeOS">BeOS<OPTION VALUE="HP-UX">HP-UX<OPTION VALUE="IRIX">IRIX<OPTION VALUE="Neutrino">Neutrino<OPTION
        VALUE="OpenVMS">OpenVMS<OPTION VALUE="OS/2">OS/2<OPTION VALUE="OSF/1">OSF/1<OPTION VALUE="Solaris">Solaris<OPTION
          VALUE="SunOS">SunOS<OPTION VALUE="other">other</SELECT></td> <td align=left valign=top>
            <SELECT NAME="priority" MULTIPLE SIZE=7>
<OPTION VALUE="--">--<OPTION VALUE="P1">P1<OPTION VALUE="P2">P2<OPTION VALUE="P3">P3<OPTION VALUE="P4">P4<OPTION
          VALUE="P5">P5</SELECT>
        </td>
        <td align=left valign=top>
            <SELECT NAME="bug severity" MULTIPLE SIZE=7>
            <OPTION VALUE="blocker">blocker<OPTION VALUE="critical">critical<OPTION VALUE="major">major<OPTION
              VALUE="normal">normal<OPTION VALUE="minor">minor<OPTION VALUE="trivial">trivial<OPTION VALUE="enhancement">enhancement</SELECT>
        </tr>
    </table>
```


This HTML code caused Mozilla to crash: what is the actual cause?

Delta Debugging

- Delta debugging is a technique designed to minimise failure inducing inputs.
 - Given a long input that causes a program to fail, delta debugging returns the shortest input that still causes the failure.
 - It is based on a simple idea: binary search.

A. Zeller and R. Hildebrandt. Simplifying and isolating failure-inducing input. IEEE Transactions on Software Engineering, 28(2):183–200, Feb. 2002.

Delta Debugging




Delta Debugging

- The basic binary search version is simple:
 - Cut and throw away half of the input, and see if the output is still wrong
 - If so, go back to the first step
 - If not, go back to the previous state, and throw away the other half

Issues to Consider

- Given N parts, delta debugging requires $O(\log N)$ steps to find the root cause (binary search). Depending on problems, this may or may not be acceptable.
- It requires fully automated oracle.
- What if the root cause of the failure is actually composed of two separate parts?
- What if the failure inducing input cannot be cut in half?

DD: What if the root cause of the failure is actually composed of two separate parts?




This is a not uncommon case. To deal with such disjoint causes, the actual Delta Debugging algorithm is more than the basic binary search.

Recursive Delta Debugging

- Line 1: if there is a single input left, return it.
- Line 2-7: see if program still fails with halves of the given input - if it does, continue halving recursively
- Line 8-10: otherwise, make two recursive calls
 - First: keep the first half of the given input, and apply DD to the second half
 - Second: keep the second half of the given input, and apply DD to the first half

```
DD( $P, \{i_1, \dots, i_n\}$ )
(1)   if  $n == 1$  then return  $i_1$ 
(2)    $P_1 = (P + \{i_1, \dots, i_{\frac{n}{2}}\})$ 
(3)    $P_2 = (P + \{i_{\frac{n}{2} + 1}, \dots, i_n\})$ 
(4)   if  $P_1$  fails
(5)       return DD ( $P, \{i_1, \dots, i_{\frac{n}{2}}\}$ )
(6)   else if  $P_2$  fails
(7)       return DD ( $P, \{i_{\frac{n}{2} + 1}, \dots, i_n\}$ )
(8)   else
(9)       return DD ( $P_2, \{i_1, \dots, i_{\frac{n}{2}}\}$ ) +
(10)      DD ( $P_1, \{i_{\frac{n}{2} + 1}, \dots, i_n\}$ )
```



DD: What if the failure inducing input cannot be cut in half?

- Suppose we are testing a C compiler. The following code actually made GCC 2.95.2 crash.
- What if we try to cut this faulty source code in two halves?
- Neither is a valid C code anymore, not to mention a failure inducing input.

```
double mult( double z[], int n ) {
    int i;
    int j;
    for (j= 0; j< n; j++) {
        i= i+j+1;
        z[i]=z[i]*(z[0]+0);
    }
    return z[n];
}

int copy(double to[], double from[], int count)
{
    int n= (count+7)/8;
    switch (count%8) do {
        case 0: *to++ = *from++;
        case 7: *to++ = *from++;
        case 6: *to++ = *from++;

        case 5: *to++ = *from++;
        case 4: *to++ = *from++;
        case 3: *to++ = *from++;
        case 2: *to++ = *from++;
        case 1: *to++ = *from++;
    } while (--n > 0);
    return (int)mult(to,2);
}

int main( int argc, char *argv[] ) {
    double x[20], y[20];
    double *px= x;

    while (px < x + 20)
        *px++ = (px-x)*(20+1.0);
    return copy(y,x,20);
}
```

Hierarchical Delta Debugging

- When input is highly structured, simply cutting it in half is likely to invalidate the input structure, damaging the efficiency of DD.
- Cutting C source code by halves will eventually work: DD will recurse down to the single character level, and find the structural boundaries that will work.
- But this will take a significant amount of test execution that is just invalid.

Hierarchical Delta Debugging

- Hierarchical Delta Debugging (HDD) does not cut at the raw data level, but handles structural elements.
 - If the input is source code, HDD deletes Abstract Syntax Tree nodes, instead of half of the lines.
 - If the input is HTML, HDD deletes XML nodes.
- Instead of recursive binary search based deletions, HDD recursively goes into deeper nested structures.
 - For example, HDD first tries to delete an entire function from source code; if this fails, HDD then tries to delete half of the basic blocks inside the function, etc.
- HDD always produce syntactically correct output, and preserves atomic elements such as identifiers.

Delta Debugging (ddmin)

```
t(double z[], int)
{
    int i;
    int j;
    for(;;j++) {
        i=i+j+1;
        z[i]=z[i]*(z[0]+0);
    }
}
```

HDD

```
mult(double *z, int n)
{
    int i;
    int j;
    for (;;) {
        i=i+j+1;
        z[i]=z[i]*(z[0]+0);
    }
}
```

DD is sensitive to whitespaces and deletion boundaries.

Consider the `int`, or `j++`: in both cases, ddmin can delete them only when its binary search deletion has the deletion size of 3 and matches these characters exactly.

Note that ddmin also reduced `mult` to `t`.

Delta Debugging Tools

- DD.py (<https://www.st.cs.uni-saarland.de/dd/DD.py>) and tutorial (<https://www.st.cs.uni-saarland.de/dd/ddusage.php3>)
- delta (<http://delta.tigris.org>) which implements HDD
- Lithium (<http://www.squarefree.com/lithium/using.html>), a python implementation of DD
- Git Bisect (<https://git-scm.com/docs/git-bisect>), which performs DD to a series of git commits

Variations of Delta Debugging

- In some sense, delta-debugging is an iterative refinement (optimisation?) of test inputs: as long as the failure signal is retained, we can keep getting rid of parts of the input.
- This is a very broadly applicable, strong idea (mostly due to its simplicity). Let us examine two related topics.
 - Program Slicing / Program Reduction
 - Agentic Failure Attribution

Program Slicing / Reduction

- Suppose I am only interested in parts of the computation done by my program. For example, I only want to know prod, not sum.
- How do I automatically identify only the relevant parts?

```
def prodsum(n):  
    sum = 0  
    prod = 1  
    for i in range(1, n+1):  
        sum += i  
        prod *= i  
    print("sum:", sum)  
    print("product:", prod)
```

Dependency Analysis

- You mark both control- and data-flow dependencies.
- Take the backward closure from the variable you are interested (e.g., `prod` in the line 9)
- This can be done on top of static dependence analysis, called static program slicing.

```
1 def prodsum():
2     n = int(input())
3     sum = 0
4     prod = 1
5     for i in range(1, n+1):
6         sum += i
7         prod *= i
8     print("sum:", sum)
9     print("product:", prod)
```

The diagram illustrates dependencies in the provided code. Red arrows represent data dependencies, showing the flow of data from the assignment of `prod` on line 4 to its use in the loop on line 7, and from the loop on line 5 to the final print statement on line 9. A blue arrow represents a control dependency, showing the flow from the `for` loop condition on line 5 to the loop body on line 6.

→ Data Dependency

→ Control Dependency

Slicing, DD way

- Suppose you do Hierarchical DD to program prodsum itself.
- It is not an input to compiler.
- Instead of retaining failure (DD), we want to retain the value of prod. That is, as long as line 9 prints the same value for the given input n, continue with HDD.

```
1 def prodsum():
2     n = int(input())
3     sum = 0
4     prod = 1
5     for i in range(1, n+1):
6         sum += i
7         prod *= i
8     print("sum:", sum)
9     print("product:", prod)
```

Observational Slicing

Binkley et al., FSE 2014 (<https://doi.org/10.1145/2635868.2635893>)

- Instead of doing bisection/deletion, apply a deletion window of size k:
 - Try deleting each line; if deleting a single line does not work, increase the number of liens, up to k.
- Repeat multiple passes; stop when the last pass did not delete anything.

```
pass 2
1 def prodsum():
2     n = int(input())
3
4     prod = 1
5     for i in range(1, n+1):
6
7         prod *= i
8
9     print("product:", prod)
```

Why do it this way?

Some dependencies are harder to capture statically.

```
1 get_command(int *command, int *prio, float* ratio)
2 {
3     char buf[CMDSIZE];
4
5     if(fgets(buf, CMDSIZE, stdin))
6     {
7         *prio = *command = -1;
8         sscanf(buf, "%d", command);
9
10        while(buf[strlen(buf)-1] != '\n'
11              && fgets(buf, CMDSIZE, stdin))
12        {
13        }
14    }
15 }
```

Fig. 9. When slicing this code, the static slice incorrectly omits the while loop because it does not account for the *implicit* dependencies caused by the input stream. This is one example of an unsafe static slice highlighted by ORBS.

Indirect dependency
via i/o stream

```
class checker {
    public static void main(String[] args) {
        int dots = 0;
        int chars = 0;
        for (int i = 0; i < args.length; i++) {
            if (args[i].charAt(0) == '.') {
                ++dots;
            } else if ((args[i].charAt(0) == '$') || (args[i].charAt(0) == '%')) {
                ++chars;
            }
        }
        System.out.println("Dots: " + dots);
        System.out.println("Chars: " + chars);
    }
}
```

```
#include <stdlib.h>
#include <stdio.h>
#include <locale.h>

int main(int argc, char** argv) {
    setlocale(LC_ALL, "");
    struct lconv *cur;
    if (atoi(argv[1]) < 0) {
        printf("%s\n", "Invalid argument");
    } else {
        printf("%s\n", "Valid argument");
    }
    return 0;
}
```

```
# Glue reader and checker together.
import commands
import sys

use_locale = True
currency = "?"
decimal = ","

if use_locale:
    currency = commands.getoutput('./reader 0')
    decimal = commands.getoutput('./reader 1')

cmd = ('java checker ' + currency
      + sys.argv[1] + decimal + sys.argv[2])
print commands.getoutput(cmd)
```

Dependencies across
multiple languages

What does this compile into? :)

```

\begin{tikzpicture}[xscale=-1,
  ray/.style={decoration={markings,mark=at position .5 with {
    \arrow[>=latex]{>}}},postaction=decorate}
]
  \pgfmathsetlengthmacro{\r}{3cm}
  \pgfmathsetmacro{\f}{.7}
  \pgfmathsetlengthmacro{\arcradius}{.8cm}
  \pgfmathsetlengthmacro{\dotradius}{.6cm}
  \pgfmathsetlengthmacro{\arclabelradius}{1cm}
  \pgfmathsetmacro{\incidentangle}{asin(\f)}
  \coordinate (O) at (0, 0);
  \coordinate (A) at (\incidentangle:\r);
  \draw (O) circle (\r);
  \draw[ray] (A -| \r*3, 0) -- (A);
  \draw[gray] (O) -- ($ (O)!1.5!(A) $) node[pos=1.05] {$n$};
  \drawarcdelta{(A)}{0}{\incidentangle}{\arcradius-1pt}
  \drawlabeledarcdelta{(A)}{0}{\incidentangle}{\arcradius+1pt}
    {$i$}{\arclabelradius}
  \foreach \index/\color in {1.32/red, 1.34/blue} {
    \pgfmathsetmacro{\refractedangle}{asin(sin(\incidentangle) / \index)}
    \pgfmathsetmacro{\angleindrop}{180 - 2*\refractedangle}
    \coordinate (A') at (\incidentangle+\angleindrop:\r);
    \coordinate (A'') at (\incidentangle+2*\angleindrop:\r);
    \begin{scope}[opacity=.5, color=\color]
      \draw[ray] (A) -- (A');
      \draw[ray] (A') -- (A'');
    \end{scope}
  }

```

Slicing Non-existing(?) Semantic

- Slicing Picture Description Languages:
 - “There is this source code; it compiles into this diagram. Which lines affect this particular sub-image?”
- The semantics we want to preserve while slicing is defined ***visually!***
- No existing slicing technique can do this

Visual-ORBS

Yoo, Binkley, and Eastman (<https://doi.org/10.1016/j.jss.2016.04.009>)



Source

```
\begin{tikzpicture}[xscale=-1,
  ray/.style={decoration={markings,mark=at position .5 with {
    \arrow[>=latex]{>}}},postaction=decorate}
]
\pgfmathsetlengthmacro{\r}{3cm}
\pgfmathsetmacro{\f}{.7}
\pgfmathsetlengthmacro{\arcradius}{.8cm}
\pgfmathsetlengthmacro{\dotradius}{.6cm}
\pgfmathsetlengthmacro{\arclabelradius}{1cm}
\pgfmathsetmacro{\incidentangle}{asin(\f)}
\coordinate (O) at (0, 0);
\coordinate (A) at (\incidentangle:\r);
\draw (O) circle (\r);
\draw[ray] (A -| \r*3, 0) -- (A);
...
```

Slice

```
\begin{tikzpicture}[xscale=-1,
  ray/.style={decoration={markings,mark=at position .5 with {
    \arrow[>=latex]{>}}},postaction=decorate}
]
\pgfmathsetlengthmacro{\r}{3cm}
\pgfmathsetmacro{\f}{.7}
\pgfmathsetlengthmacro{\arcradius}{.8cm}
\pgfmathsetlengthmacro{\arclabelradius}{1cm}
\pgfmathsetmacro{\incidentangle}{asin(\f)}
\coordinate (O) at (0, 0);
\coordinate (A) at (\incidentangle:\r);
\draw[ray] (A -| \r*3, 0) -- (A);
\draw[gray] (O) -- ($ (O)!1.5!(A) $) node[pos=1.05] {$n$};
\drawarcdelta{(A)}{0}{\incidentangle}{\arcradius-1pt}
\drawlabeledarcdelta{(A)}{0}{\incidentangle}{\arcradius+1pt}
{$i$}{\arclabelradius}
\end{tikzpicture}
```

Criterion (sub-image)



(Failure) Attribution in Agents

- Agent Trajectories: a sequence of reasoning and action steps taken by an agent to achieve a given task.
- Failure Attribution: identifying the root cause step that is responsible for the failure of the task.
- Trajectory Minimisation: reducing the trajectory to the bare minimum required to achieve the task.
- On the surface, both can be considered in the context of DD/slicing: attribution is leaning towards DD, minimisation is leaning towards slicing.
- Do you see any problem? :)

Lack of Concrete Dependence Relations in Trajectories

- Suppose you have the following agent trajectory, which eventually produces a correct patch for the given issue.
- Any trajectory that includes **Write Patch** will succeed, so DD will remove all other action steps, including necessary predecessors!
- Ironically, we need formally defined notion of dependencies for natural language reasoning, whereas deleting and observing worked for programming languages 😎



Issue
Report

Read File A


Write Patch

Run Test

Read File B

Write Patch

Run Test




Task
Success

Fault Localisation

- We will examine three major approaches:
 - Information-Retrieval based FL
 - Spectrum-based FL
 - Mutation-based FL

Let's have a look at a real bug report



https://bugzilla.mozilla.org/show_bug.cgi?id=1001903

Firefox Bug 1001903

- After creating a folder in the bookmark toolbar and a bookmark inside the folder, mouse left-click successfully opens the folder but touch-click fails to open it: it opens but immediately closes itself.
- While we do not know the internals of Firefox, it is easy to guess which parts of the overall system is most likely to contain this fault.

Firefox Bug 1001903

- After **creating** a **folder** in the **bookmark toolbar** and a **bookmark** inside the **folder**, mouse left-click successfully opens the folder but **touch-click** fails to **open** it: it **opens** but immediately **closes** itself.
- While we do not know the internals of Firefox, it is easy to guess which parts of the overall system is most likely to contain this fault.

IR based Fault Localisation

- Information Retrieval techniques help a user to quickly obtain resources relevant to an information need, from a large collection of information resources.
- Fault localisation can be thought of as finding a resource (a program element) that is relevant to an information need (the reported symptoms of the failure), from a large collection of information resources (the entire system).
 - The bug report becomes our query.
 - The entire source code becomes our collection of documents.
 - Fault localisation is to find the source code that matches the bug report the best.

How do we do the matching?

- There are many ways to represent documents in IR: here, we are going to look at one of the basic model, called Vector Space Model (VSM).
- VSM allows us to represent queries and documents as vectors.
- This, in turn, allows us to calculate distances between our query and documents easily.

Vector Space Model

- First, define the vocabulary, a set of words that are meaningful to us.
- Given a vocabulary with N terms, we represent both the query and the documents as vectors:

$$d_j = (w_{1,j}, w_{2,j}, \dots, w_{N,j})$$

$$q = (w_{1,q}, w_{2,q}, \dots, w_{N,q})$$

- The dimensionality of these vectors is equal to the number of terms in the vocabulary, N .
- If the term w_2 appears in d_j , $w_{2,j}$ is a non-zero number; otherwise, it is a zero. There are many ways to set the non-zero number: we are going to study *tf-idf*.

Tf-Idf

- Term frequency of a term t in document d , $tf(t, d)$, is simply the number of times t appears in d .
 - Inverse document frequency of a term t in a set of documents D , $idf(t, D)$, is the logarithmically scaled inverse fraction of the documents that contain the term t .
 - A higher $idf(t, D)$ means that only a few documents in D contains the term t : in other words, t is not common.
 - A lower $idf(t, D)$ means that most of the documents in D contain the term t : in other words, t is very common.
- $$idf(t, D) = \log \frac{N}{|\{d \in D : t \in d\}|}$$

Tf-Idf

- Tf-Idf is simply the product of $tf(t, d)$ and $idf(t, D)$.
 - A high $tfidf(t, d, D)$ means that t happens specifically in d , but not commonly in D . In other words, t is unique to d , and not other documents in D .
 - A low $tfidf(t, d, D)$ means that either t does not appear frequently in d , or if it does then it appears frequently in most of the documents in D . In other words, t is not unique to d .

$$tfidf(t, d, D) = tf(t, d) \cdot idf(t, D)$$

Preparing bug reports for IR based Fault Localisation

- We follow the standard text cleansing process used in many IR and machine learning applications.
 - Tokenisation: we break down the bug report to a list of word tokens.
 - Remove punctuation: for example, “file’s” becomes “file”
 - Case normalisation: for example, “File” and “FILE” all become “file”
 - Stop word filtering: some words are extremely common in English (e.g., “a”, “an”, “and”...) and best left out. These words are called stop words.
 - Stemming: reduce each word to its root (e.g., “opens” and “opening” both become “open”).
- There are many widely used libraries that will perform all of these steps.



Stemming Applied

Preparing code for IR based Fault Localisation

- The standard text cleansing process also allies to source code.
- With source code, stop words should also include the reserved keywords of the used programming language.
- In addition, we usually apply additional normalisation to identifiers:
 - Camel case breakdown: for example, divide “openDialog” to “open” and “dialog”.
 - Snake case breakdown: for example, divide “create_bookmarks” to “create” and “bookmarks”

VSM and Distance

- After cleansing, using a predefined vocabulary, we can calculate $tfidf(t, d, D)$ for each t (word) and d (the query or a source code file). This converts the query and each source code file to a vector.
- Given two vector, one representing the query and another representing a source code file, the distance between two vectors can be calculated using the cosine distance:

$$\cos(d_j, q) = \frac{d_j \cdot q}{\|d_j\| \|q\|} = \frac{\sum_{i=1}^N w_{i,j} w_{i,q}}{\sqrt{\sum_{i=1}^N w_{i,j}^2} \sqrt{\sum_{i=1}^N w_{i,q}^2}}$$

- Choose the file that has the shortest distance to the query, and it is more likely to contain the fault that is responsible for the symptom described in the query (i.e., the bug report).


IR Based Fault Localisation

- Strengths
 - Simple and available: many existing work on IR already improved the accuracy of textual queries - we get free ride!
 - Intuitive, easy to understand.
- Weaknesses
 - Requires a (very good and detailed) bug report.
 - There is inherent limit to accuracy, because the document to be matched (i.e., the unit of localisation) needs to be of certain size (otherwise lexical similarity becomes more random).

Spectrum Based Fault Localisation

- Intuition: If there is a faulty line, executing it is likely to increase the probability of a failure. Similarly, not executing the faulty line is likely to decrease the probability of a failure.
- This decision process
 - requires only test results and structural coverage, but
 - is essentially a statistical inference, and does not give any guarantee
- Question: how do we formalise the above intuition in a computable form?

Spectrum Based Fault Localisation



Spectrum Based Fault Localisation

- Program Spectrum: for each structural unit (i.e. statements or branches), summarise the test result AND coverage into a tuple of the following four numbers
 - ep: # of test cases that execute this unit and pass
 - ef: # of test cases that execute this unit and fail
 - np: # of test cases that do not execute this unit and pass
 - nf: # of test cases that do not execute this unit and fail

Spectrum Based Fault Localisation

Structural Elements	Test	Test	Test	Spectrum				Tarantula	Rank
	t_1	t_2	t_3	e_p	e_f	n_p	n_f		
s_1	●			1	0	0	2	0.00	9
s_2	●			1	0	0	2	0.00	9
s_3	●			1	0	$\frac{e_f}{e_f + n_f}$	2	0.00	9
s_4	●			1	0	$\frac{e_f + n_f}{e_f + n_f}$	2	0.00	9
s_5	●			1	0	$\frac{e_p}{e_p + n_p}$	2	0.00	9
s_6	●		●	1	0	0	1	0.33	4
s_7 (faulty)	●		●	0	2	1	0	1.00	1
s_8	●	●		1	1	0	1	0.33	4
s_9	●	●	●	1	2	0	0	0.50	2
Result	P	F	F						

State of the Art c.a. 2010

$$\frac{e_f}{e_f + n_f + e_p}$$

$$\frac{2e_f}{e_f + n_f + e_p}$$

$$\frac{2(e_f + n_p)}{2(e_f + n_p) + e_p + n_f}$$

$$\frac{e_f}{e_f + n_p + 2(e_p + n_f)}$$

$$\frac{e_f}{e_f + 2(n_f + e_p)}$$

Over 30 formulæ in the literature: none guaranteed to perform best for all types of faults

$$\frac{e_f + n_p}{n_f + e_p}$$

$$\frac{e_f}{e_f + n_f + e_p + n_p}$$

$$\frac{2e_f}{2e_f + n_f + e_p}$$

$$\frac{1}{2}\left(\frac{e_f}{e_f + n_f} + \frac{e_f}{e_f + e_p}\right)$$

$$\frac{\frac{e_f}{e_f + n_f}}{\frac{e_p}{e_p + n_p} + \frac{e_f}{e_f + n_f}}$$

$$\frac{e_f + n_p - n_f - e_p}{e_f + n_f + e_p + n_p}$$


Name	Formula	Name	Formula
ER1 _a	$\begin{cases} -1 & \text{if } n_f > 0 \\ n_p & \text{otherwise} \end{cases}$	ER1 _b	$ef - \frac{ep}{ep + np + 1}$
ER5 _a	$ef - \frac{ef}{ep + np + 1}$	ER5 _b	$\frac{ef}{ef + nf + ep + np}$
ER5 _c	$\begin{cases} 0 & \text{if } ef < F \\ 1 & \text{otherwise} \end{cases}$	GP ₂	$2(ef + \sqrt{ep + np}) + \sqrt{ep}$
Ochiai	$\frac{ef}{(ef + nf)(ef + ep)}$	GP ₃	$\sqrt{ ef^2 - \sqrt{ep} }$
Jaccard	$\frac{ef}{ef + nf + ep}$	GP ₁₃	$ef(1 + \frac{1}{2ep + ef})$
AMPLE	$ \frac{ef}{F} - \frac{ep}{P} $	GP ₁₉	$ef\sqrt{ ep - ef + F - P }$
Hamann	$\frac{ef + np - ep - nf}{P + F}$	Tarantula	$\frac{\frac{ef}{ef + nf}}{\frac{ef}{ef + nf} + \frac{ep}{ep + np}}$
Dice	$\frac{2ef}{ef + ep + nf}$	RusselRao	$\frac{ef}{ep + ef + np + nf}$
M1	$\frac{ef + np}{nf + ep}$	SørensenDice	$\frac{2ef}{2ef + ep + nf}$
M2	$\frac{ef}{ef + np + 2nf + 2ep}$	Kulczynski1	$\frac{ef}{nf + ep}$
Hamming	$ef + np$	Kulczynski2	$\frac{1}{2}(\frac{ef}{ef + nf} + \frac{ef}{ef + ep})$
Goodman	$\frac{2ef - nf - ep}{2ef + nf + ep}$	SimpleMatching	$\frac{ef + np}{ep + ef + np + nf}$
Euclid	$\sqrt{ef + np}$	RogersTanimoto	$\frac{ef + np}{ef + np + 2nf + 2ep}$
Wong1	ef	Sokal	$\frac{2ef + 2np}{2ef + 2np + nf + ep}$
Wong2	$ef - ep$	Anderberg	$\frac{ef}{ef + 2ep + 2nf}$
Wong3	$ef - h, h = \begin{cases} ep & \text{if } ep \leq 2 \\ 2 + 0.1(ep - 2) & \text{if } 2 < ep \leq 10 \\ 2.8 + 0.01(ep - 10) & \text{if } ep > 10 \end{cases}$		
Ochiai2	$\frac{ef np}{\sqrt{(ef + ep)(nf + np)(ef + np)(ep + nf)}}$		
Zoltar	$\frac{ef}{ef + ep + nf + \frac{10000nfep}{ef}}$		

Parnin & Orso, ISSTA 2011

- Title: “Are Automated Debugging Techniques Actually Helping Programmers?”
- Based on human studies, authors say “no, not really”!
 - Some of the widely accepted assumptions, such as perfect bug understanding, may not hold in reality.
 - **BIGGEST IMPACT:** “percentage will not cut it” - later studies move away from the Expense metric.

These are formulas - can we evolve them?

Yoo 2012 (https://doi.org/10.1007/978-3-642-33119-0_18)



The Competition

- We choose 9 formulæ from Naish et al. 2011:
- Op2 is the known best.
- Jaccard, Tarantula, Ochiai are widely studied in SE.
- Wong & AMPLE are recent additions.

$$Op1 = \begin{cases} -1 & \text{if } n_f > 0 \\ n_p & \text{otherwise} \end{cases} \quad Op2 = e_f - \frac{e_p}{e_p + n_p + 1}$$

$$Jaccard = \frac{e_f}{e_f + n_f + e_p}$$

$$Tarantula = \frac{\frac{e_f}{e_f + n_f}}{\frac{e_p}{e_p + n_p} + \frac{e_f}{e_f + n_f}}$$

$$AMPLE = \left| \frac{e_f}{e_f + n_f} - \frac{e_p}{e_p + n_p} \right|$$


$$Ochiai = \frac{e_f}{\sqrt{(e_f + n_f) \cdot (e_f + e_p)}}$$

$$Wong1 = e_f \quad Wong2 = e_f - e_p$$

$$Wong3 = e_f - h, h = \begin{cases} e_p & \text{if } e_p \leq 2 \\ 2 + 0.1(e_p - 2) & \text{if } 2 < e_p \leq 10 \\ 2.8 + 0.001(e_p - 10) & \text{if } e_p > 10 \end{cases}$$

Results


- **Green**: GP outperforms the other.
- **Orange**: GP exactly matches the other.
- **Red**: The other outperforms GP.



4 Unix tools w/ 92 faults: 20 for training, 72 for evaluation.

Results

- GP completely outperforms Ochiai, Tarantula, Wong 1 & 2, and mostly outperforms AMPLE.
- Op1, Jaccard, and Wong 3 are tough to beat.
- Op2 is very good but it is not impossible to do better.



4 Unix tools w/ 92 faults: 20 for training, 72 for evaluation.

Evolved Formulæ

ID	Refined Formula	ID	Refined Formula
GP01	$e_f(n_p + e_f(1 + \sqrt{e_f}))$	GP16	$\sqrt{e_f^{\frac{3}{2}} + n_p}$
GP02	$2(e_f + \sqrt{n_p}) + \sqrt{e_p}$	GP17	$\frac{2e_f + n_f}{e_f - n_p} + \frac{n_p}{\sqrt{e_f}} - e_f - e_f^2$
GP03	$\sqrt{ e_f^2 - \sqrt{e_p} }$	GP18	$e_f^3 + 2n_p$
GP04	$\sqrt{ \frac{n_p}{e_p - n_p} - e_f }$	GP19	$e_f \sqrt{ e_p - e_f + n_f - n_p }$
GP05	$\frac{(e_f + n_p)\sqrt{e_f}}{(e_f + e_p)(n_p n_f + \sqrt{e_p})(e_p + n_p)\sqrt{ e_p - n_p }}$	GP20	$2(e_f + \frac{n_p}{e_p + n_p})$
GP06	$e_f n_p$	GP21	$\sqrt{e_f + \sqrt{e_f + n_p}}$
GP07	$2e_f(1 + e_f + \frac{1}{2n_p}) + (1 + \sqrt{2})\sqrt{n_p}$	GP22	$e_f^2 + e_f + \sqrt{n_p}$
GP08	$e_f^2(2e_p + 2e_f + 3n_p)$	GP23	$\sqrt{e_f}(e_f^2 + \frac{n_p}{e_f} + \sqrt{n_p} + n_f + n_p)$
GP09	$\frac{e_f \sqrt{n_p}}{n_p + n_p} + n_p + e_f + e_f^3$	GP24	$e_f + \sqrt{n_p}$
GP10	$\sqrt{ e_f - \frac{1}{n_p} }$	GP25	$e_f^2 + \sqrt{n_p} + \frac{\sqrt{e_f}}{\sqrt{ e_p - n_p }} + \frac{n_p}{(e_f - n_p)}$
GP11	$e_f^2(e_f^2 + \sqrt{n_p})$	GP26	$2e_f^2 + \sqrt{n_p}$
GP12	$\sqrt{e_p + e_f + n_p - \sqrt{e_p}}$	GP27	$\frac{n_p \sqrt{(n_p n_f - e_f)}}{e_f + n_p n_f}$
GP13	$e_f(1 + \frac{1}{2e_p + e_f})$	GP28	$e_f(e_f + \sqrt{n_p} + 1)$
GP14	$e_f + \sqrt{n_p}$	GP29	$e_f(2e_f^2 + e_f + n_p) + \frac{(e_f - n_p)\sqrt{n_p e_f}}{e_p - n_p}$
GP15	$e_f + \sqrt{n_f + \sqrt{n_p}}$	GP30	$\sqrt{ e_f - \frac{n_f - n_p}{e_f + n_f} }$

Theoretical Analysis of SBFL

- For all the gory details:
 - X. Xie, T. Y. Chen, F.-C. Kuo, and B. Xu. A theoretical analysis of the risk evaluation formulas for Spectrum-Based Fault Localization. ACM Transactions on Software Engineering Methodology, 22(4):31:1– 31:40, October 2013.
 - X. Xie, F.-C. Kuo, T. Y. Chen, S. Yoo, and M. Harman. Provably optimal and human-competitive results in SBSE for Spectrum Based Fault Localisation. In G. Ruhe and Y. Zhang, editors, Search Based Software Engineering, volume 8084 of Lecture Notes in Computer Science, pages 224–238. Springer Berlin Heidelberg, 2013.

Crash Course into Our Proof System

- Given a test suite and a formula R , a program $PG = \{s_1, s_2, \dots, s_n\}$ can be divided into the following three subsets:
 - S^R_B : set of elements that are ranked higher than the faulty element, s_f
 - S^R_F : set of elements that are tied to the faulty element, s_f
 - S^R_A : set of elements that are ranked lower than the faulty element, s_f

$$S^R_B = \{s_i \in S \mid R(s_i) > R(s_f), 1 \leq i \leq n\}$$

$$S^R_F = \{s_i \in S \mid R(s_i) = R(s_f), 1 \leq i \leq n\}$$


$$S^R_A = \{s_i \in S \mid R(s_i) < R(s_f), 1 \leq i \leq n\}$$

Crash Course into Our Proof System

- Formula R_1 *dominates* formula R_2 ($R_1 \rightarrow R_2$) if, for any combination of program, faulty statement, test suite, and a consistent tie-breaking scheme, $E_1 \leq E_2$
 - That is, $S_B^{R_1} \subseteq S_B^{R_2} \wedge S_A^{R_2} \subseteq S_A^{R_1}$
- Formula R_1 is *equivalent* to formula R_2 ($R_1 \leftrightarrow R_2$) if, for any combination of program, faulty statement, test suite, and a consistent tie-breaking scheme, $E_1 = E_2$
 - That is, $S_B^{R_1} = S_B^{R_1} \wedge S_F^{R_1} = S_F^{R_1} \wedge S_A^{R_1} = S_A^{R_1}$
- Finally, $R_1 \leftrightarrow R_2$ iff $R_1 \rightarrow R_2$ and $R_2 \rightarrow R_1$

Crash Course into Our Proof System

Statement Ranking



Formula X

Formula Y

To show that Y dominates X,
we show that:

$$S_B^Y \subseteq S_B^X \wedge S_A^X \subseteq S_A^Y$$

(assuming that we break ties
in F sets consistently)

Equivalence is defined as:

$$X \leftrightarrow Y \iff X \rightarrow Y \wedge Y \rightarrow X$$

The Current Maximal Formulas

	Name	Formula expression
ER1'	Naish1	$\begin{cases} -1 & \text{if } e_f < F \\ P - e_p & \text{if } e_f = F \end{cases}$
	Naish2	$e_f - \frac{e_p}{e_p + n_p + 1}$
	GP13	$e_f \left(1 + \frac{1}{2e_p + e_f} \right)$
ER5	Wong1	e_f
	Russel & Rao	$\frac{e_f}{e_f + n_f + e_p + n_p}$
	Binary	$\begin{cases} 0 & \text{if } e_f < F \\ 1 & \text{if } e_f = F \end{cases}$
GP02		$2(e_f + \sqrt{n_p}) + \sqrt{e_p}$
GP03		$\sqrt{ e_f^2 - \sqrt{e_p} }$
GP19		$e_f \sqrt{ e_p - e_f + n_f - n_p }$

SBFL

- Strengths
 - Only requires what is already there: coverage and test results
 - Relatively intuitive
- Weaknesses
 - Single formulas are usually limited. In fact,
 - There exists a theoretical proof that no single formula works the best against all faults
 - Does not work against omission faults
 - Does not work well against multiple faults

Mutation-based FL

Mutation Based Fault Localisation

- How do we use mutants to localise a fault? Mutants are injected faults themselves!!
- Consider this: what would happen if you mutate a program you know to be faulty?

Case 1: Mutating Correct Statements

Equivalent

P F

New Fault

P- F+

Mask

P+ F-

CASE 2: mutating Faulty Statement

(Partial) Fix

P + F.

(New) Fault

P? F?

Equivalent

P F

Mask

P + F.

Hypotheses

- An arbitrary mutation operator applied to a correct statement is likely to introduce a new fault
- An arbitrary mutation operator applied to a faulty statement is either likely to keep the program still faulty or, even better, (partially) fix the program
- The majority of statements in a faulty program is correct; we detect the faulty one by observing the anomaly from our hypotheses

MUSE (Moon et al., 2014)

$$\mu(s) = \frac{1}{|mut(s)|} \sum_{m \in mut(s)} \left(\frac{|f_P(s) \cap p_m|}{|f_P|} - \alpha \cdot \frac{|p_P(s) \cap f_m|}{|p_P|} \right)$$

Proportion of test cases that mutant m turns from fail to pass

Average over all mutation applied to statement s


Proportion of test cases that mutant m turns from pass to fail

$$\alpha = \frac{f2p}{|mut(P)| \cdot |f_P|} \cdot \frac{|mut(P)| \cdot |p_P|}{p2f}$$

Ahead of Time MBFL

SIMFL (Kim et al., ISSRE 2021)

- Perform mutation analysis (i.e., inject faults into different locations, and record which tests fail)
- Using this information, learn the reverse relationship. The ML model is designed to answer this question: “if these test cases failed, where is the mutant?”
- When real faults occur, give the test outcome to the trained model, and ask where the mutant (=real fault) is.
- Promising results (113 out of 203 faults ranked at the top, method level FL on Java)



Hybrid/Learn-to-Rank FL

Hybrid SBFL Approaches

- Primarily based on the understanding that a single SBFL formula is very much limited. How can we add more information?
- Hybridisation
 - Use multiple formulas at the same time, or
 - Use SBFL formulas with some additional input features
- As we accept more diverse input, fault localisation became a learning-to-rank problem, instead of a design-a-technique problem.

Learn-to-Rank

ML for ranking models

- Pointwise L2R: consider individual datapoints as if standard classification or regression - rank them based on their score (either softmax, or direct regression). Any linear/logistic regression or other decision trees.
- Pairwise L2R: formulate the problem as classification of pairs, i.e., given (A, B), learn to predict whether A ranks higher than B or vice versa - can be expensive ($O(n^2)$). Ranking SVM, RankNet, etc.
- Listwise L2R: learn to produce the entire permutation. LambdaRank, LambdaMART, etc.
- FL is a unique application in that we are only interested in the rank of the faulty program element, not any of others. Consequently, many use point wise.

Xuan & Monperrus (ICSME 2014)

- Simultaneously use 25 SBFL formulas, using the weighted-sum method.
- Use a learning-to-rank technique to train the linear point wise ranking model (i.e. to optimise the weights that yield best ranks for training set faults)

Learning to Combine Multiple Ranking Metrics for Fault Localization

Jifeng Xuan
INRIA Lille - Nord Europe
Lille, France
jifeng.xuan@inria.fr

Martin Monperrus
University of Lille & INRIA
Lille, France
martin.monperrus@univ-lille1.fr

Abstract—Fault localization is an inevitable step in software debugging. Spectrum-based fault localization applies a ranking metric to identify faulty source code. Existing empirical studies on fault localization show that there is no optimal ranking metric for all the faults in practice. In this paper, we propose *MULTRIC*, a learning-based approach to combining multiple ranking metrics for effective fault localization. In *MULTRIC*, a suspiciousness score of a program entity is a combination of existing ranking metrics. *MULTRIC* consists two major phases: learning and ranking. Based on training faults, *MULTRIC* builds a ranking model by learning from pairs of faulty and non-faulty source code. When a new fault appear, *MULTRIC* computes the final ranking with the learned model. Experiments are conducted on 5386 seeded faults in ten open-source Java programs. We empirically compare *MULTRIC* against four widely-studied metrics and three recently-proposed metrics. Our experimental results show that *MULTRIC* localizes faults more effectively than state-of-art metrics, such as Tarantula, Ochiai, and Ample.

I. INTRODUCTION

Debugging is an expensive task in software development. *Fault localization* eases debugging by automatically analyzing bugs to find out the root cause – a specific location in source code – of the problem. *Spectrum-based fault localization* (also called *coverage-based fault localization*) is a class of fault localization approaches, which leverages the execution traces of test cases to predict the likelihood – called *suspiciousness scores* – of source code locations to be faulty. In spectrum-based fault localization, a faulty program is instrumented for collecting the running traces; then a *ranking metric* is applied to compute suspiciousness scores for *program entities* (such as methods [25], [24], basic blocks [31], and statements [12], [2]). The most suspicious entities are given to a developer for further analysis [21]. The developer manually analyzes the source code, from the most suspicious location and downward, to confirm the inferred root cause of the bug.

The ideal fault localization ranking metric would always rank the faulty source code entity at the top, by giving highest suspiciousness score. However, there is no such metric, and

[12], Ochiai [2], Jaccard [2], and Ample [4]). Most of these metrics are manually and analytically designed based on assumptions on programs, test cases, and their relationship with faults [16]. To our knowledge, only the work by Wang et al. [27] considers the combination of multiple ranking metrics. They propose search-based algorithms to form such a combination in fault localization. In this paper, we propose the combination approach of multiple ranking metrics based on machine learning. Our idea is to leverage the diversity of ranking metrics by automatically combining multiple metrics.

We propose *MULTRIC*, a learning-based approach that combines multiple fault localization ranking metrics. *MULTRIC* consists in two major phases: learning and ranking. In the learning phase, a ranking model is learned based on the ranks of program entities in training faulty programs. When a new fault happens, this is the ranking phase, in which the model computes weights for combining metrics and generating the final ranking. In other words, *MULTRIC* determines the final ranking of program entities based on a combination of various suspiciousness scores.

Fault localization cannot be transferred into a general classification problem because the ranks of program entities are hard to label as classes in machine learning. To solve this problem, we use a learning-to-rank method. Instead of classes in classification, a learning-to-rank method models orders between faulty entities and non-faulty entities and optimizes the model to satisfy those orderings. In other words, the ranking of program entities is converted to indicate whether an actual faulty one ranks above a non-faulty one. The learning-to-rank method is used to to combine 25 existing ranking metrics.

Experiments are conducted on 5386 seeded faults in ten open-source Java programs. We empirically compare *MULTRIC* against four widely-studied ranking metrics (Tarantula [12], Ochiai [2], Jaccard [2], and Ample [4]) and three recently-proposed ranking metrics (Naish1 [18], Naish2 [18], and GP13 [33]). Experimental results show that *MULTRIC* can localize faults more effectively than the ranking metrics.

Le et al. (ISSTA 2016)

- In addition to 35 SBFL formulas, *Savant* use invariant difference:
- Use Daikon to infer method invariants twice: with passing tests, and with failing test. For the faulty method, two sets of invariant will tend to be more different.

A Learning-to-Rank Based Fault Localization Approach using Likely Invariants

Tien-Duy B. Le¹, David Lo¹, Claire Le Goues², and Lars Grunske³

¹School of Information Systems, Singapore Management University, Singapore

²School of Computer Science, Carnegie Mellon University, USA

³Humboldt University of Berlin, Germany

{btdle.2012,davidlo}@smu.edu.sg, clegoues@cs.cmu.edu, grunske@informatik.hu-berlin.de

ABSTRACT

Debugging is a costly process that consumes much developer time and energy. To help reduce debugging effort, many studies have proposed various fault localization approaches. These approaches take as input a set of test cases (some failing, some passing) and produce a ranked list of program elements that are likely to be the root cause of the failures (i.e., failing test cases). In this work, we propose *Savant*, a new fault localization approach that employs a learning-to-rank strategy, using likely invariant *diffs* and suspiciousness scores as features, to rank methods based on their likelihood of being a root cause of a failure. *Savant* has four steps: method clustering and test case selection, invariant mining, feature extraction, and method ranking. At the end of these four steps, *Savant* produces a short ranked list of potentially buggy methods. We have evaluated *Savant* on 357 real-life bugs from 5 programs from the Defects4J benchmark. We find that, on average, *Savant* can identify the correct buggy method for 63.03, 101.72, and 122 bugs at the top 1, 3, and 5 positions in the produced ranked lists. We have compared *Savant* against several state-of-the-art spectrum-based fault localization baselines. We show that *Savant* can successfully locate 57.73%, 56.69%, and 43.13% more bugs at top 1, top 3, and top 5 positions than the best performing baseline, respectively.

CCS Concepts: Software and its engineering → Software testing and debugging

Keywords: Learning to Rank, Program Invariants, Automated Debugging

1. INTRODUCTION

Software systems are often plagued with bugs that compromise system reliability, usability, and security. One of the main tasks involved in fixing such bugs is identifying the associated buggy program elements. Developers can then study the implicated program elements and their context,

and make necessary modifications to resolve the bug. This is a time consuming and expensive process. Many real-world projects receive a large number of bug reports daily [6], and addressing them requires considerable time and effort. Debugging can contribute up to 80% of the total software cost for some projects [50]. Thus, there is a pressing need for automated techniques that help developers debug. This problem has motivated considerable work proposing automated debugging solutions for a variety of scenarios, e.g., [4, 7, 8, 14, 21, 30, 32, 36, 44, 48, 53, 54, 55, 62, 63, 64].

In addition to potentially providing direct developer support, automated debugging approaches are also used by recent work in automated program repair (including [28, 40, 35], and many others). Such tools use automated debugging approaches as a first step to identify likely faulty program elements. These lists guide program repair tools to generate program patches that lead previously-failing tests to pass. The accuracy of automated debugging approaches therefore plays an important role in the effectiveness of program repair tools. Thus, there is a need to improve the effectiveness of automated debugging tools further to support both developers and current program repair techniques.

In this work, we are particularly interested in a family of automated debugging solutions that takes as input a set of failing and passing test cases and then highlights the suspicious program elements that are likely responsible for the failures (failing test cases), e.g., [4, 7, 8, 14, 21, 30, 32, 36, 53, 54, 55, 62, 63]. While these techniques have been shown effective in many contexts, their effectiveness needs to be further improved to localize more bugs more accurately.

We propose a novel technique, *Savant*, for effective automated fault localization. *Savant* uses a learning-to-rank machine learning approach to identify buggy methods from failures by analyzing both classic suspiciousness scores and inferred likely invariants observed on passing and failing test cases. *Savant* is built on three high-level intuitions. First, program elements which follow different invariants when run in failing versus correct executions are suspicious. Second, such program elements are even more suspicious if they are assigned higher suspiciousness scores computed by existing

FL is just DP happening late 🧐

Sohn & Yoo (<https://doi.org/10.1145/3092703.3092717>)

- A simple and naive insight: we already have a technique that claims it can tell us where the faults are - even before we test. Defect Prediction!
- DP predicts where faults will be before they actually happen; FL finds where fails are AFTER they actually happen.
- We simply add three major DP features to learn-to-rank model for FL.
 - Age: how long has the given element existed in the code base?
 - Churn: how frequently has the given element been changed?
 - Complexity: how complex is the given element
- Interestingly, we use genetic programming to evolve a point wise ranking model (i.e., a formula produces ranking scores).

Summary

- Delta Debugging minimises a failure-inducing input, thereby helping the developer to localise the corresponding part in the program.
- IR based FL queries the documents (source code files) using symptoms of the failure (bug report) as the query text.
- SBFL capsulate the intuition about test results and coverage into risk evaluation formula.
- MBFL is based on the idea that mutating an already faulty program can reveal insights about the fault (further damage or partial fix).
- There is a growing body of agentic FL work, but we will revisit them later.