

ECE444: Software Engineering

Software Engineering Research


Shurui Zhou



The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

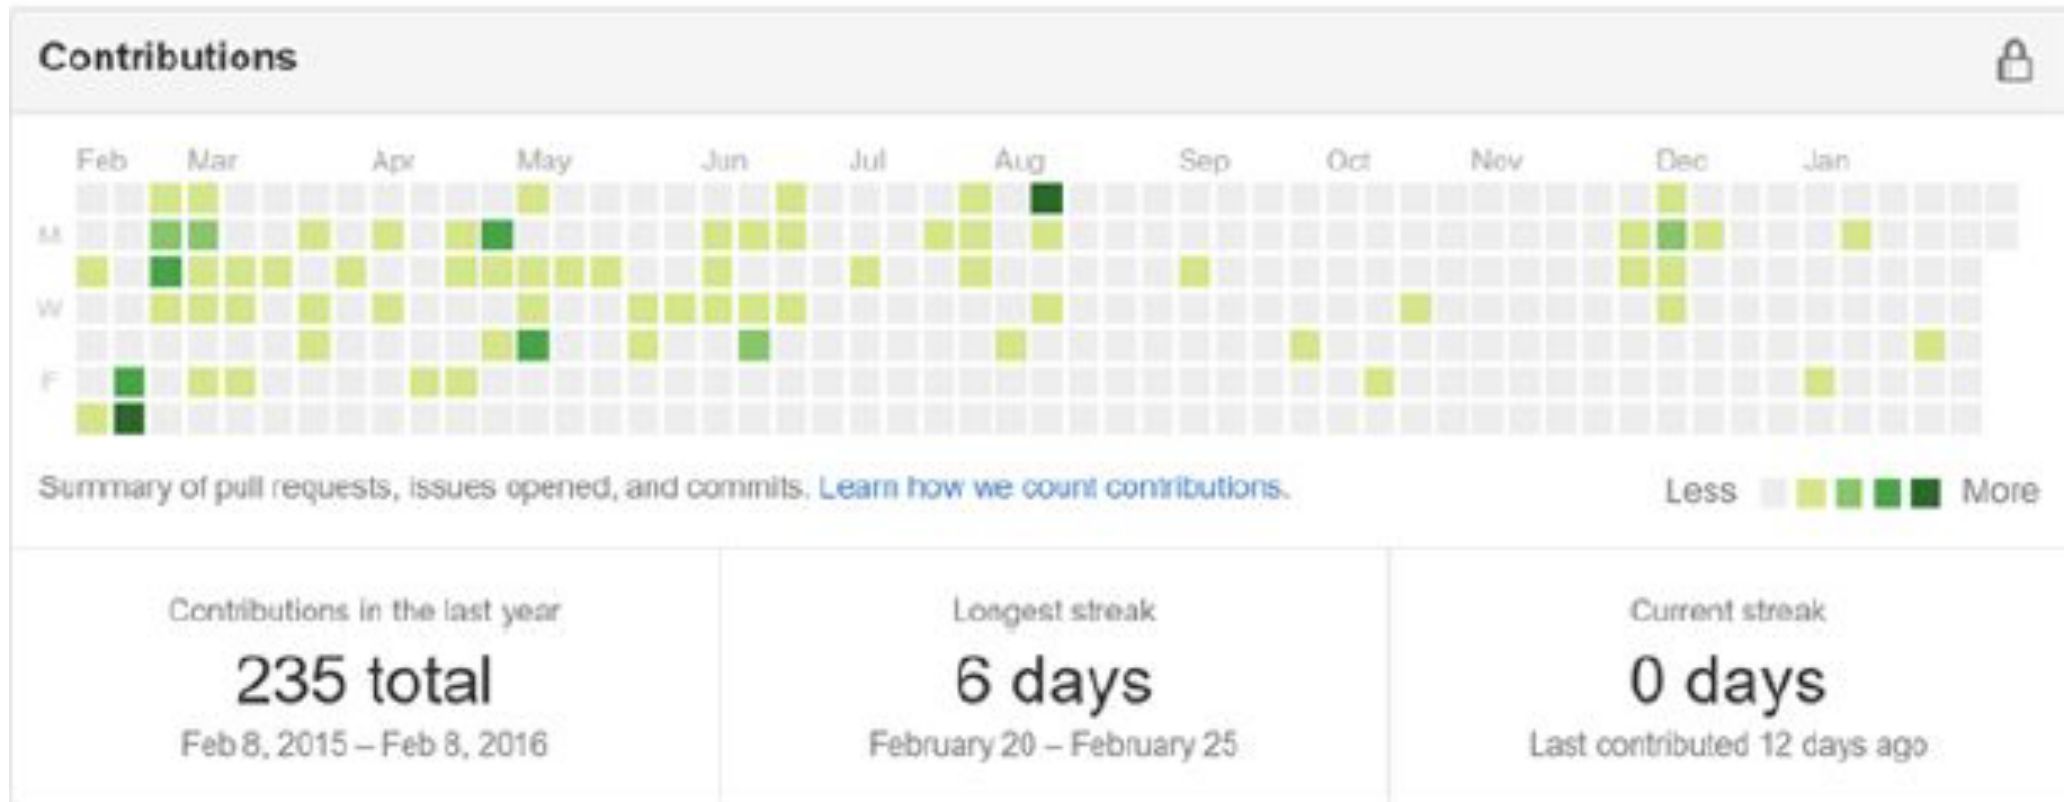
(Competing) concerns in SE...

- **Code:** faster, cheaper, more features, more reliable/secure
- **Developers:** more productive, more skilled, happier, better connected
- **Organizations/communities:** attract/retain contributors, encourage a participatory culture, increase value






“Measuring programming progress by lines of code is like measuring aircraft building progress by weight.”



Contributing graphs considered harmful (Hanselman)


<https://www.hanselman.com/>

Software Engineering Design Space



Research success?

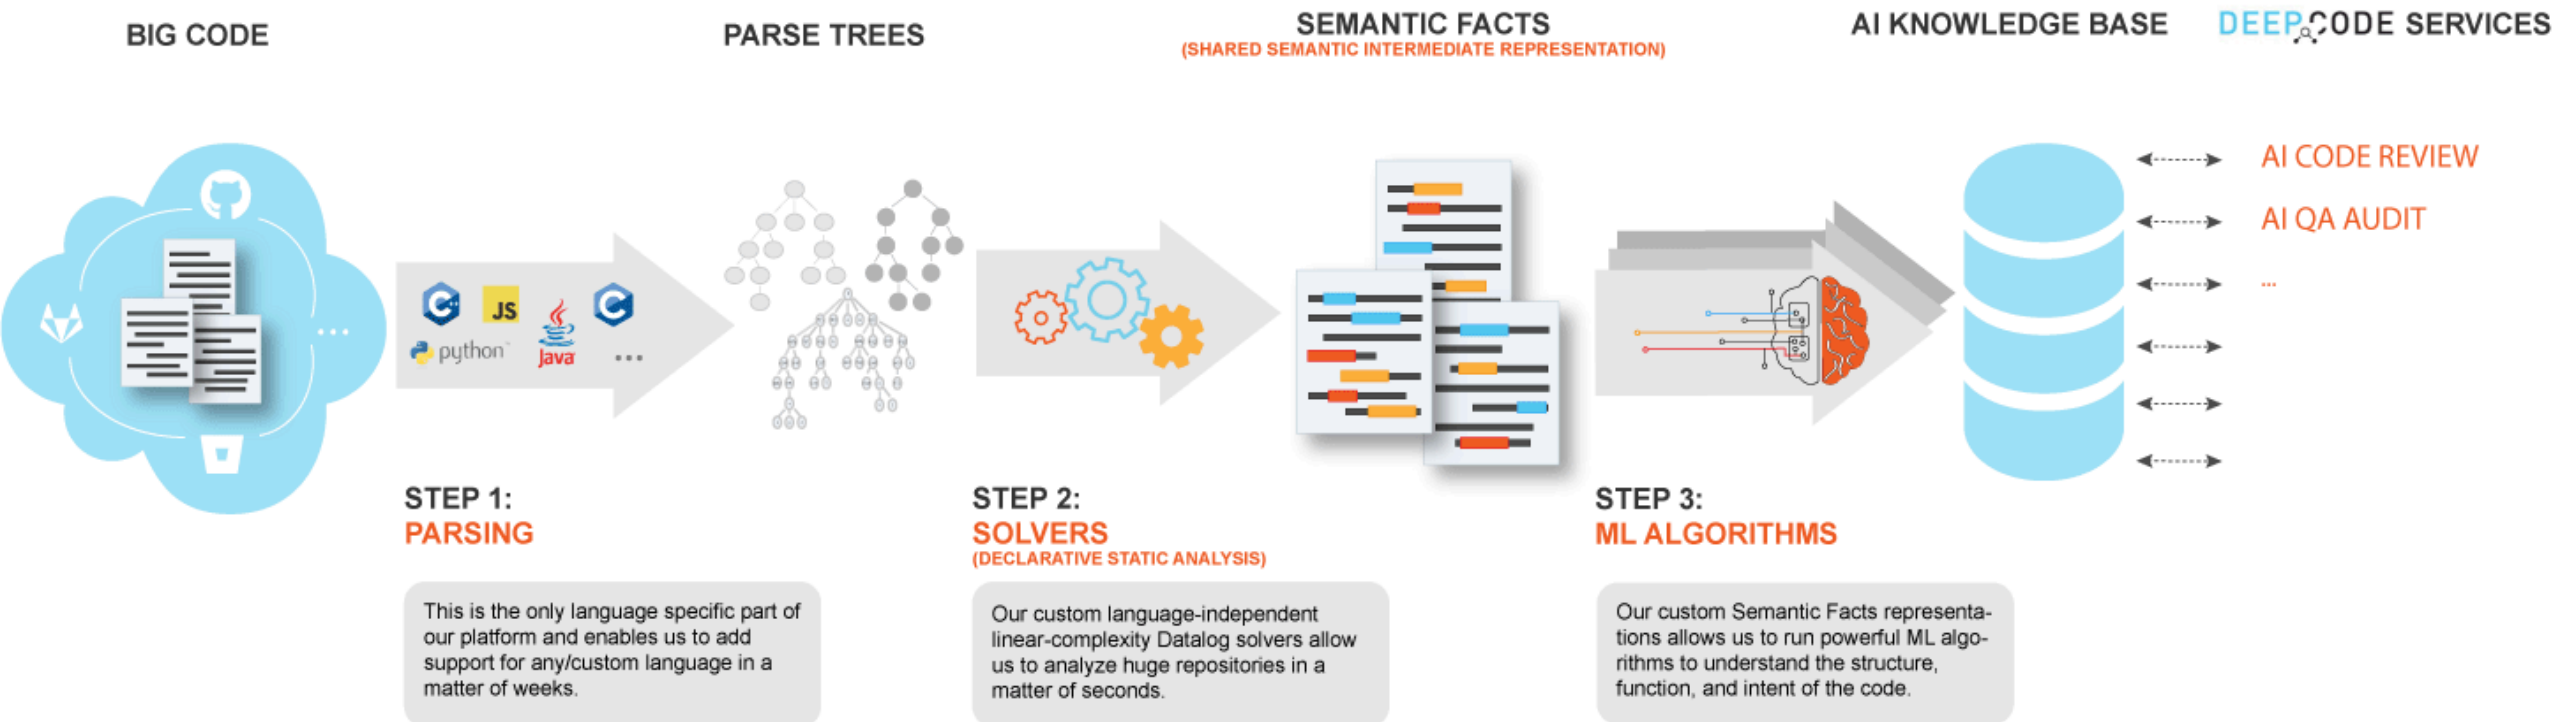





Success practice transfer stories from research

- Automated testing (Facebook)
- Code review tools (Microsoft)
- Software Analytics (Hassan et al.)

DeepCode from ETH




kite from MIT&Stanford



Code Faster.
Stay in Flow.

Kite adds AI powered code completions to your code editor, giving developers superpowers.

 Download for Free

```
1 import os
2 import sys
3
4 def count_py_files_in_repos(dirname):
5     i|
```

kite.com

1968 NATO Conference on Software Engineering

- international experts on computer software who agreed on defining best practices for software grounded in the application of engineering.



“Academic software engineering research has been a backwater primarily staffed by those interested in theory, with a tenuous connection to practical software development.”

- Lack of **industrial relevance** (doesn't scale or solve industry problems) [Briand]
- **Poor replication** of software engineering studies [Menzies et al.]
- **Poor actionability** (practitioners know which modules are buggy...)
- **Perils of mining** software repositories [Kaliyamvakou, German et al.]
- Lack of focus on **human/social aspects** [Storey et al.]

**Evidence-based
Software Engineering**
based on the publicly available data

Derek M. Jones

ISBN: 978-1-8382913-0-3
Publisher: Knowledge Software, Ltd
Released: November 8, 2020

The content of this book is licensed under the [Creative Commons Attribution-ShareAlike](#)

4.0 International License (CC BY-SA 4.0)





ESEC/FSE 2020

Sun 8 - Fri 13 November 2020 Sacramento, California, United States

Attending ▾ Program ▾ Tracks ▾ Organization ▾ Search Series ▾

Sign i




Sacramento Skyline at Night

ASE 2020

Mon 21 - Fri 25 September 2020 Melbourne, Australia

Attending ▾ Tracks ▾ Organization ▾ Search Series ▾



Melbourne City

○○●

Photo by satanoid



29th IEEE International Requirements Engineering Conference

Notre Dame, South Bend, USA
September 20-24, 2021

ISSTA 2020

Sat 18 - Wed 22 July 2020

Attending ▾

Sponsorship ▾

Program ▾

MSR 2020

Mon 29 - Tue 30 June 2020

Attending ▾

Travel Support ▾

Program ▾

Tracks ▾

Organization ▾

🔍 Search

Series ▾



What metrics are the **best predictors of failures**?

If I increase **test coverage**, will that actually increase software quality?

What is the **data quality** level used in empirical studies and how much does it actually matter?

Are there any **metrics that are indicators of failures** in both Open Source and Commercial domains?

I just submitted a **bug report**.
Will it be fixed?

How can I tell if a piece of software will have **vulnerabilities**?

Should I be writing **unit tests** in my software project?

Do **cross-cutting concerns** cause defects?

Is strong **code ownership** good or bad for software quality?

Does **Test Driven Development** (TDD) produce better code in shorter time?

Does **Distributed/Global software development** affect quality?

© Microsoft Corporation

Mining Software Repository

Interpreting Cloud Computer Vision Pain-Points: A Mining Study of Stack Overflow

Alex Cummaudo
ca@deakin.edu.au

Applied Artificial Intelligence Inst.
Deakin University
Geelong, Victoria, Australia

Rajesh Vasa

Scott Barnett

SAFE: A Simple Approach for Feature Extraction from App Descriptions and App Reviews

John
john.grund

The GHTorrent project

Welcome to the G

Follow [@ghtorren](#)

Software Documentation Issues Unveiled

Emad Aghajani*, Csaba Nagy*, Olga Lucero Vega-Márquez[†]

Mario Linares-Vásquez[†], Laura Moreno[‡], Gabriele Bavota*, Michele Lanza*

**Software Institute, Università della Svizzera italiana (USI), Switzerland*

[†]Systems and Computing Engineering Department, Universidad de los Andes, Colombia

[‡]Department of Computer Science, Colorado State University, USA

Beyond the Code: Mining Self-Admitted Technical Debt in Issue Tracker Systems

Laerte Xavier

ASERG Group - Department of Computer Science
Federal University of Minas Gerais (UFMG)
Belo Horizonte, Brazil
laertexavier@dcc.ufmg.br

Fabio Ferreira

Center of Informatics - Federal Institute
of the Southeast of Minas Gerais
Barbacena, Brazil
fabio.ferreira@ifsudestemg.edu.br

Rodrigo Brito

ASERG Group - Department of Computer Science

Marco Tulio Valente

ASERG Group - Department of Computer Science

Refactor `:comment_personal_snippet` to `:create_note`

Follow-up to https://dev.gitlab.org/gitlab/gitlabhq/merge_requests/2794

We should probably keep this confidential until that MR is merged to master and issue is made public.

The `:create_note` permission is being checked on the Noteable when replying to email notifications. The previous MR adds the `:create_note` permission to `ProjectSnippetPolicy`.

This is a duplicate of an existing `:comment_personal_snippet` permission. We should refactor uses of `:comment_personal_snippet` to use the common `:create_note` permission instead.

Related issues ⓘ 0

Related merge requests ⓘ 1

Remove the `comment_personal_snippet` permission !27999

11.11

Milestone

11.11

Time tracking

No estimate or time spent

Due date

None

Accepting merge requests

Plan [DEPRECATED]

backend

types:plan permissions

sn

technical debt

Requirement

Can a Conversation Paint a Picture?

Mining Requirements in Software Forums

James Tizard, Hechen Wang, Lydia Yohannes, Kelly Blincoe

University of Auckland, New Zealand

{jtiz003,hwan531}@aucklanduni.ac.nz, l.yohannes@web.de, k.blincoe@auckland.ac.nz

ARdoc Classes	Mapped Forum Classes
Problem Discovery	Apparent bug
Feature Request	Feature request
Information Seeking	Question on application Help seeking Requesting more information Question on background
Information Giving	Application guidance User setup Praise for application Dispraise for application Application usage Attempted solution Acknowledgement of resolution

Detecting Bad Smells in Use Case Descriptions

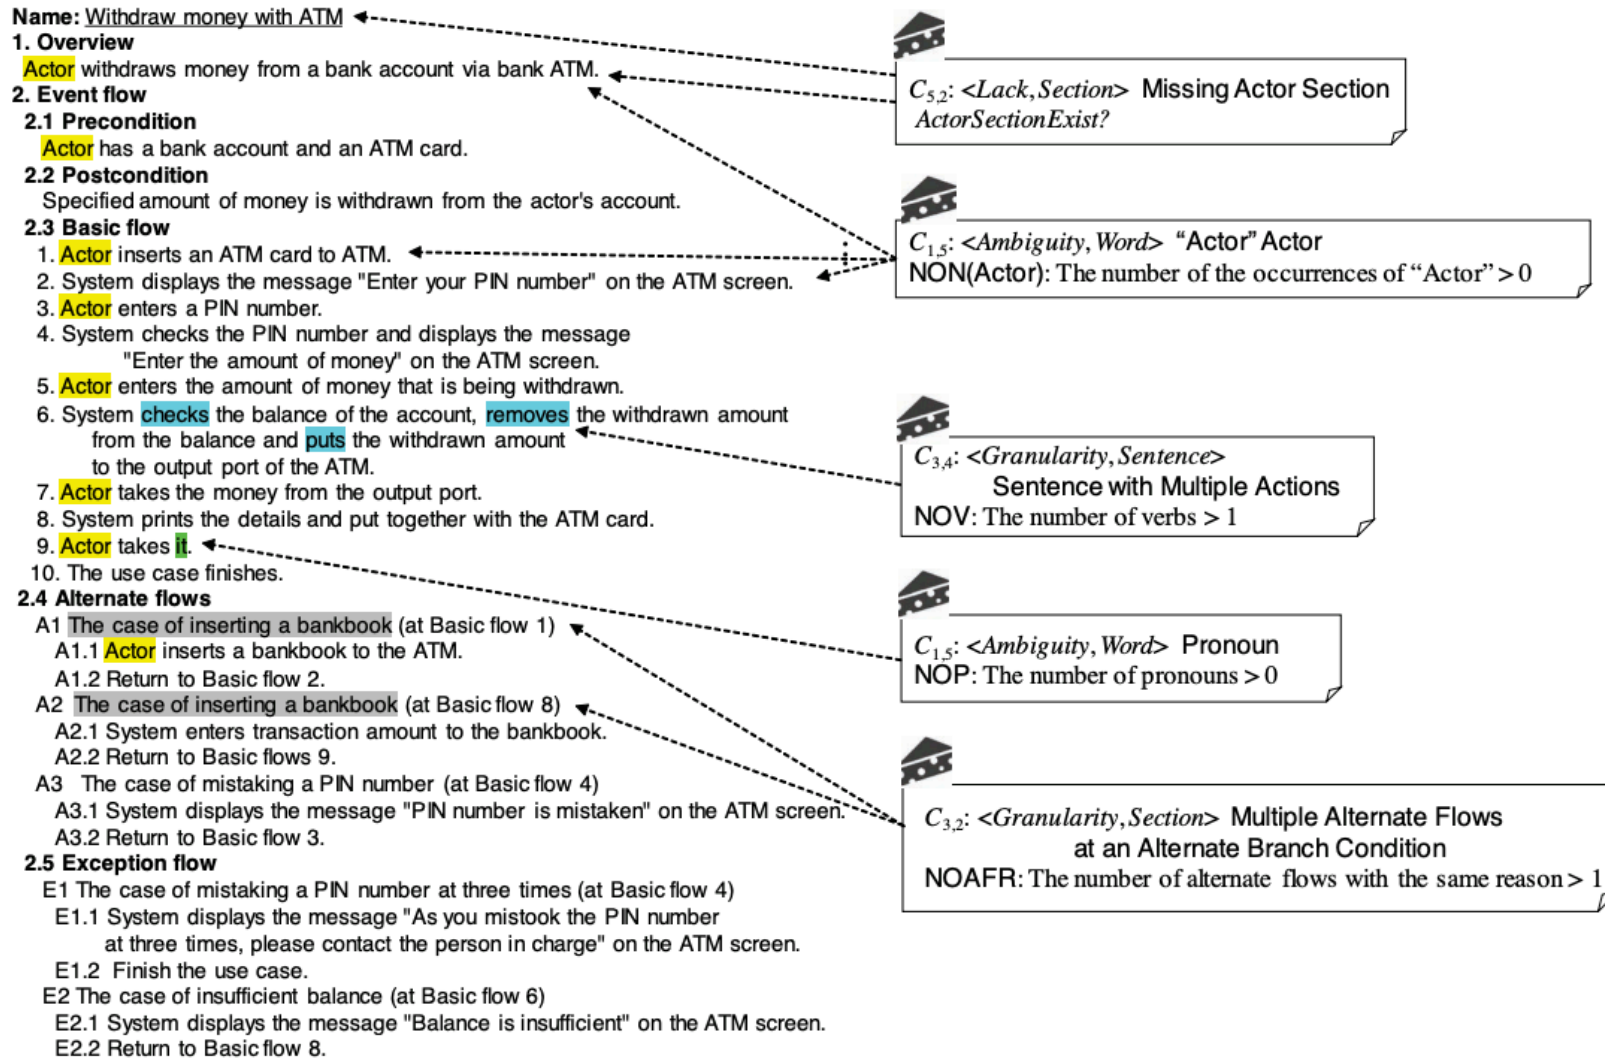


Fig. 6. Example result of the smell detection.

NERO: A Text-based Tool for Content Annotation and Detection of Smells in Feature Requests

Fangwen Mu^{1,2}, Lin Shi^{1,2*}, Wei Zhou¹, Yuanzhong Zhang¹, Huixia Zhao¹

TABLE II
EXPLANATIONS AND DETECTION METHODS OF SMELLS

Smell Category	Smell Name	Explanation	Detection method
Ambiguous	Vagueness	Vagueness occurs whenever a statement admits borderline cases, such as appropriate, clear, significant, etc.	Keyword glossary, Lemmatization
	Weakness	Weakness occurs when the feature requests use words with weak semantic content and little emotional color, such as could, may, might, etc.	Keyword glossary, Lemmatization
	Generality	Generality occurs when the sentence contains words that identify a certain type of object, and no modifiers limit its scope, such as flow, access, data, interface, etc.	Keyword glossary, Lemmatization, Dependency parsing
	Coordination ambiguity	Coordination ambiguity occurs when the use of coordinating conjunctions leads to multiple potential interpretations of a sentence.	POS tagging, Regular expression
	Referential ambiguity	Referential ambiguity occurs when an anaphor (e.g. it, that, which, etc.) can take its reference from more than one element, each playing the role of the antecedent.	POS tagging, Regular expression
	Passive voice	Passive voice occurs when the passive voice is used in the feature requests.	Dependency parsing, Regular expression
Incomplete	Missing condition	Missing condition occurs when the sentence contains an if clause expressing the condition, but there is no corresponding else/otherwise clause.	Keyword glossary, Lemmatization
	Missing description	Missing description occurs when the sentence contains omitted-meaning words, such as as defined, to be completed, to be determined, etc.	Regular expression
Unintelligible	Unreadability	Unreadability occurs when the sentences in one feature request are too long or not smooth.	GPT2 LM, Coleman-Liau formula
	Partial Content	Partial Content occurs when the feature requests lack any of the five semantic annotations (except Trivia) mentioned in the content annotation. We assume that feature requests with more different content annotations will deliver more diverse information.	Weighted analysis

NERO: A Text-based Tool for Content Annotation and Detection of Smells in Feature Requests

Fangwen Mu^{1,2}, Lin Shi^{1,2*}, Wei Zhou¹, Yuanzhong Zhang¹, Huixia Zhao¹

¹Laboratory for Internet Software Technologies, Institute of Software Chinese Academy of Sciences, Beijing, China.

²University of Chinese Academy of Sciences, Beijing, China.

{fangwen, zhouwei, yuanzhong, huixia}@itechs.iscas.ac.cn, {shilin}@iscas.ac.cn

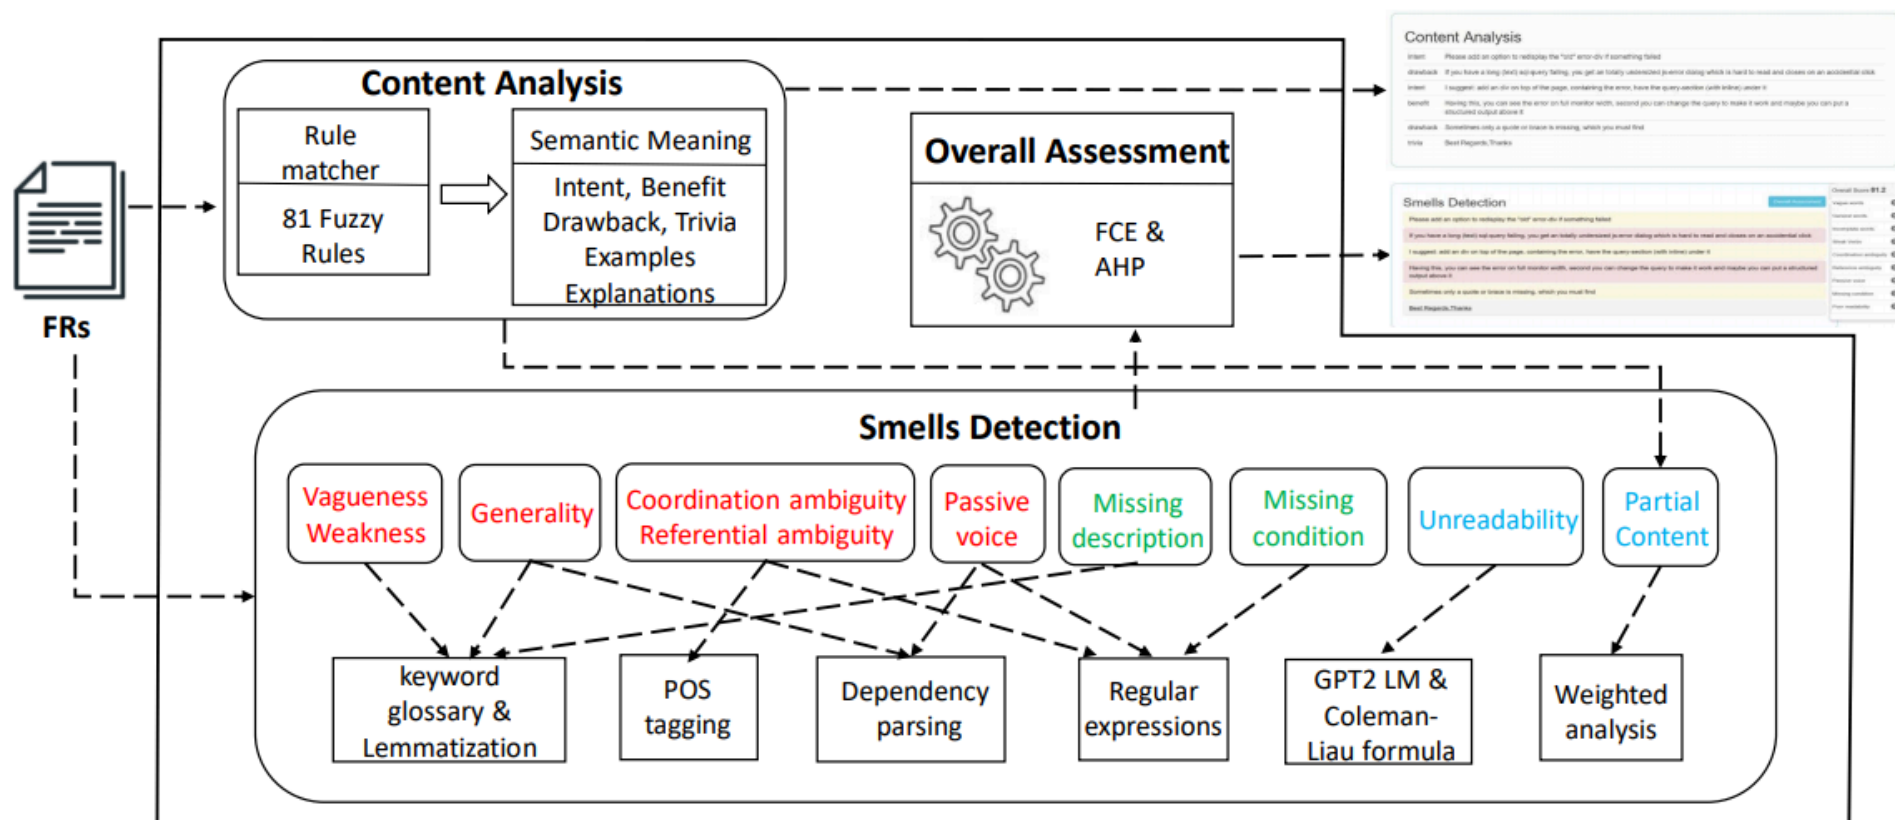


Fig. 1. The Overview of NERO

Documentation

On Automatically Generating Commit Messages via Summarization of Source Code Changes

Luis Fernando Cortés-Coy¹, Mario Linares-Vásquez², Jairo Aponte¹, Denys Poshyvanyk²

Automatic Documentation Generation via Source Code Summarization of Method Context

Paul W. McBurney and Collin McMillan

ARENA: An Approach for the Automated Generation of Release Notes

Laura Moreno, *Member, IEEE*, Gabriele Bavota, *Member, IEEE*, Massimiliano Di Penta, *Member, IEEE*, Rocco Oliveto, *Member, IEEE*, Andrian Marcus, *Member, IEEE*, and Gerardo Canfora

Traceability

A Novel Approach to Tracing Safety Requirements and State-Based Design Models

Mounifah Alenazi*
University of Cincinnati
Cincinnati, Ohio
alenazmh@mail.uc.edu

Nan Niu
University of Cincinnati
Cincinnati, Ohio
nan.niu@uc.edu

Juha Savolainen
Danfoss Drives A/S
Gråsten, Denmark
juha.savolainen@danfoss.com

Establishing Multilevel Test-to-Code Traceability Links

Robert White
University College London
London, UK

Jens Krinke
University College London
London, UK

Raymond Tan
University College London
London, UK

Improving the Effectiveness of Traceability Link Recovery using Hierarchical Bayesian Networks

Kevin Moran
William & Mary
Williamsburg, VA, USA
kpmoran@cs.wm.edu

David N. Palacio
William & Mary
Williamsburg, VA, USA
dnaderp@cs.wm.edu

Carlos Bernal-Cárdenas
William & Mary
Williamsburg, VA, USA
cebernal@cs.wm.edu

Daniel McCrystal
William & Mary
Williamsburg, VA, USA
dmc@cs.wm.edu

Code Review

Mitigating Turnover with Code Review Recommendation: Balancing Expertise, Workload, and Knowledge Distribution

Ehsan Mirsaeedi

Department of Computer Science and
Software Engineering

Concordia University, Montréal, Québec, Canada
s_irsae@encs.concordia.ca

Peter C. Rigby

Department of Computer Science and
Software Engineering

Concordia University, Montréal, Québec, Canada
peter.rigby@concordia.ca

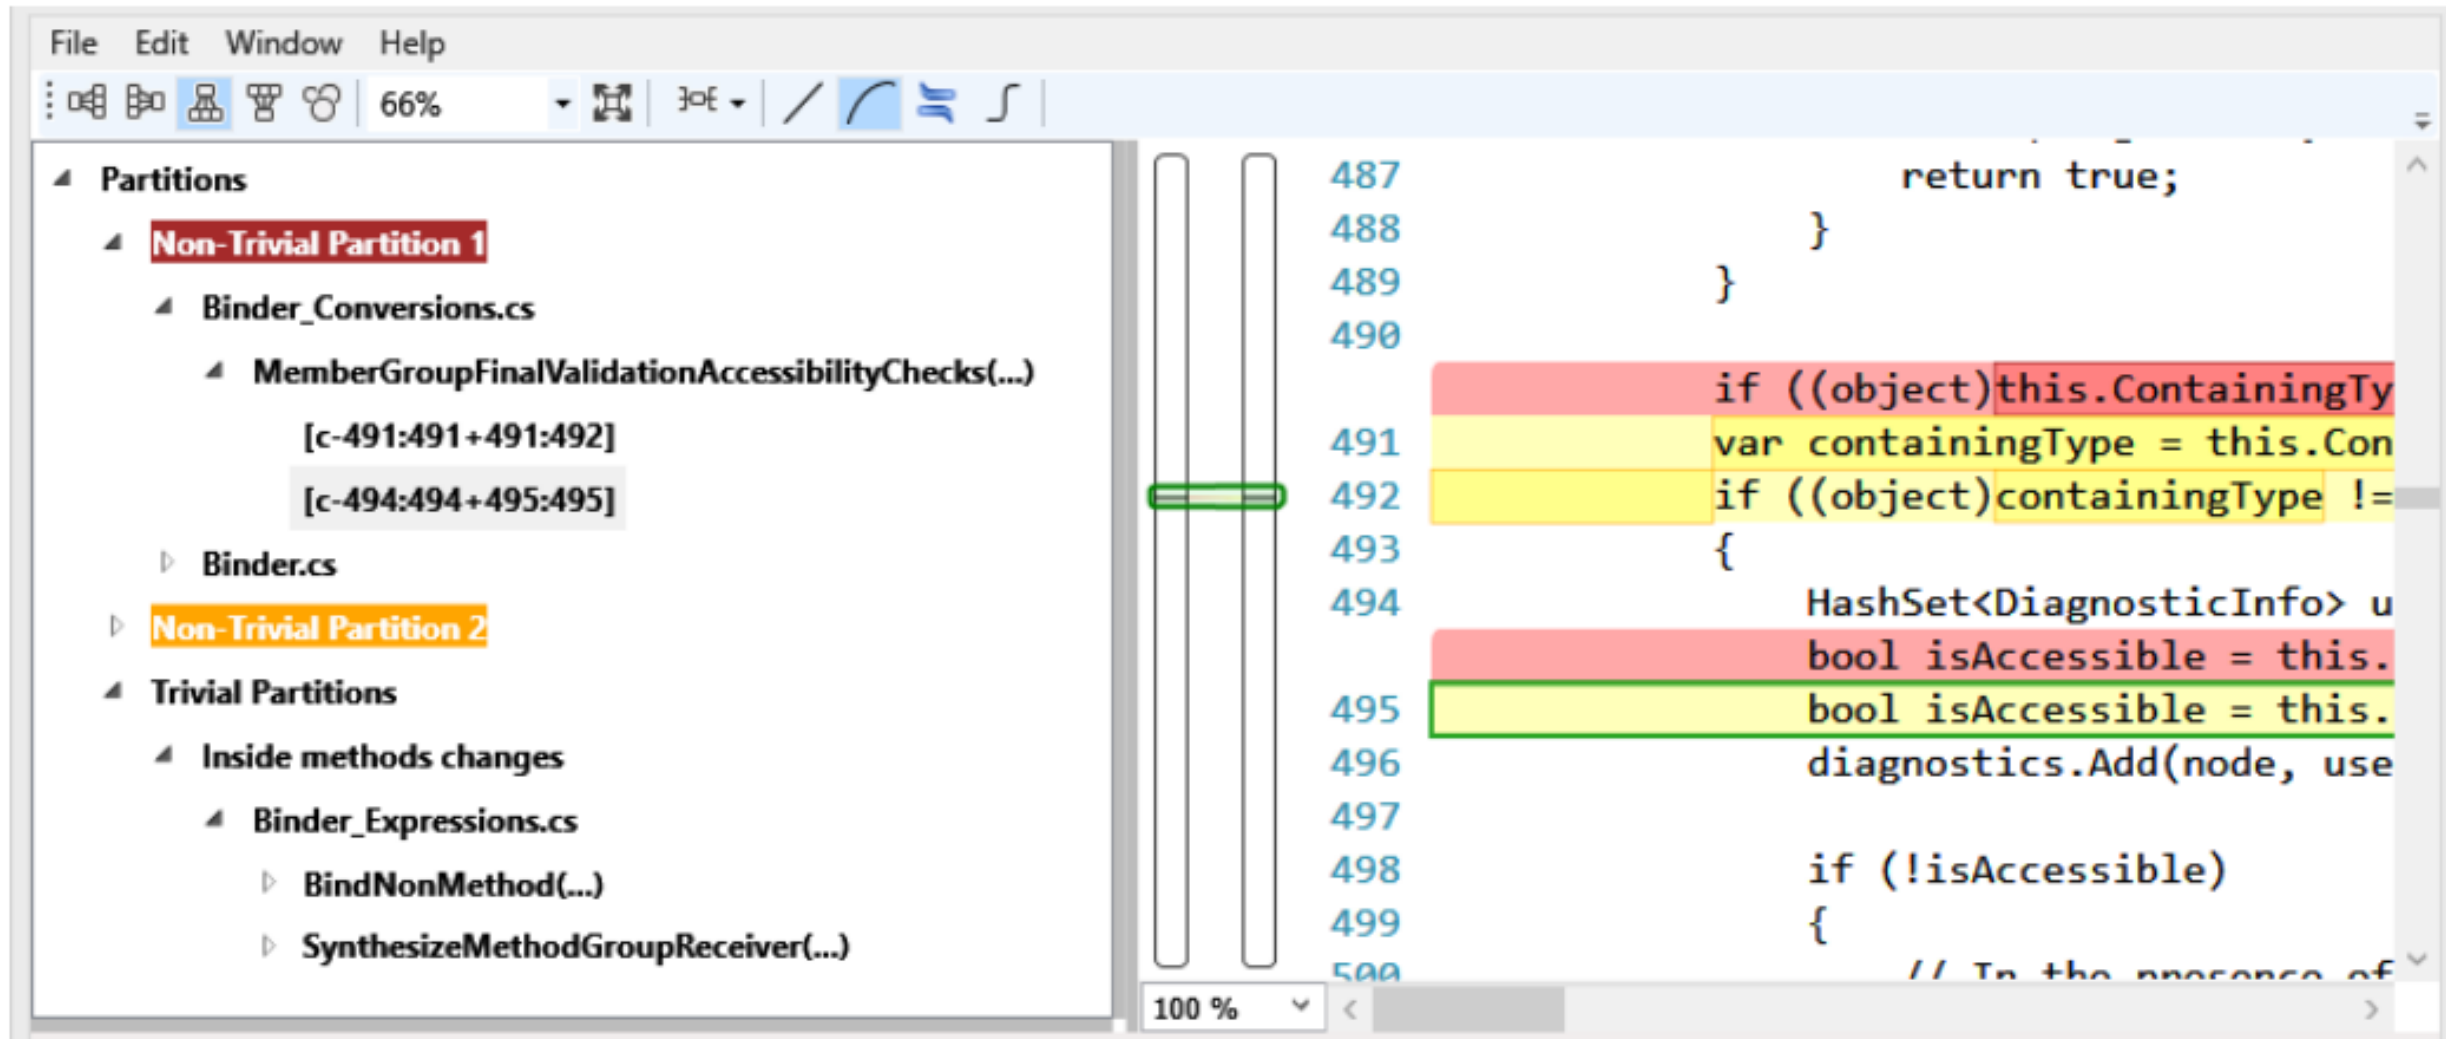
Who Should Review My Code?

A File Location-Based Code-Reviewer Recommendation Approach for Modern Code Review

Patanamon Thongtanunam*, Chakkrit Tantithamthavorn*, Raula Gaikovina Kula[†],
Norihiro Yoshida[‡], Hajimu Iida*, Ken-ichi Matsumoto*

*Nara Institute of Science and Technology, [†]Osaka University, [‡]Nagoya University, Japan
{patanamon-t, chakkrit-t, matumoto}@is.naist.jp, iida@itc.naist.jp, raula-k@ist.osaka-u.ac.jp, yoshida@ertl.jp

Helping Developers Help Themselves: Automatic Decomposition of Code Review Changesets



Deployment

FastLane: Test Minimization for Rapidly Deployed Large-scale Online Services

Adithya Abraham Philip, Ranjita Bhagwan, Rahul Kumar, Chandra Sekhar Maddila and Nachiappan Nagappan
Microsoft Research

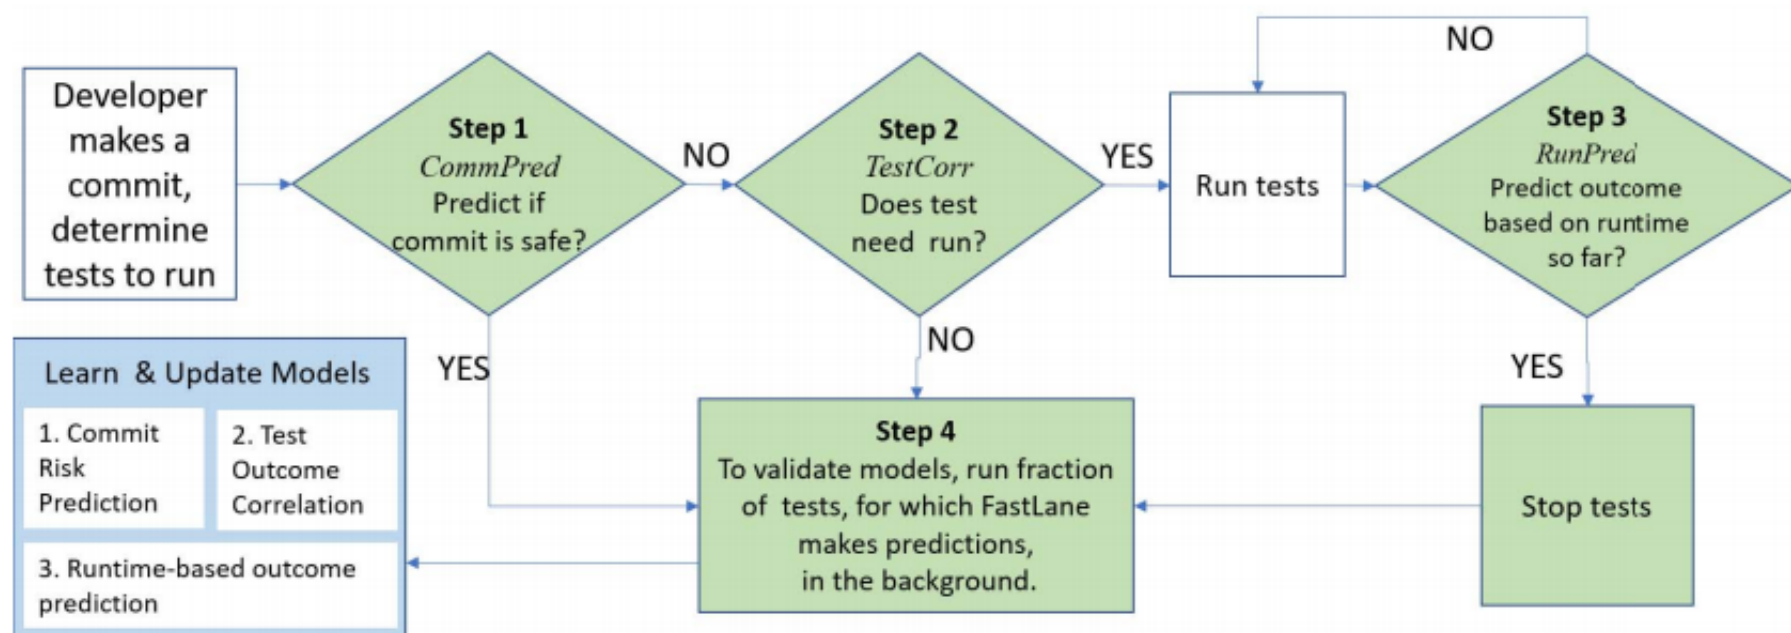


Fig. 1. The integrated algorithm and FastLane test prediction flow. This includes all three types of predictions FastLane performs (green boxes): 1. commit risk prediction, 2. test outcome-based correlation, and 3. runtime-based outcome prediction. The blue box captures the learning functionality used to continuously train FastLane.

Productivity

Characterizing Software Developers by Perceptions of Productivity

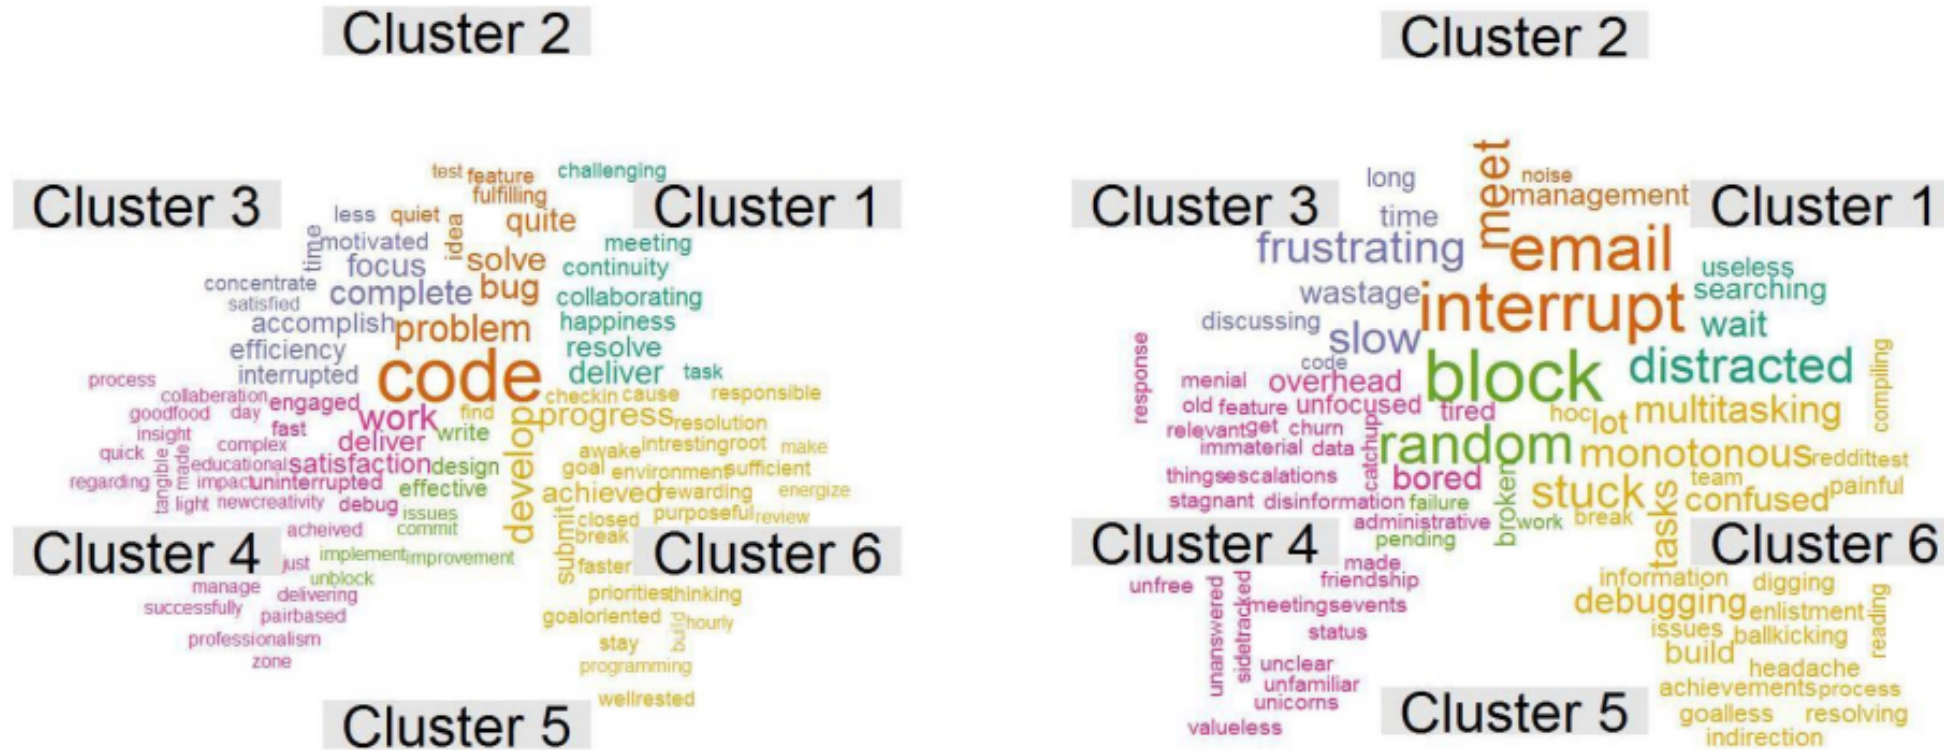


Fig. 1: Comparing the clusters with respect to words that developers associate with *productive* (left, Q1) and *unproductive* work days (right, Q2). Terms in **turquoise** are related to Cluster 1, **orange** to Cluster 2, **purple** to Cluster 3, **pink** to Cluster 4, **green** to Cluster 5, and **gold** to Cluster 6. The size of a term corresponds to the difference between the maximum relative frequency and the average relative frequency of the word across the six clusters.

Do Developers Discover New Tools On The Toilet?

Emerson Murphy-Hill Edward K. Smith* Caitlin Sadowski Ciera Jaspan Collin Winter*
Google, LLC *Bloomberg* *Google, LLC* *Google, LLC* *Waymo*
 emersonm@google.com esmith404@bloomberg.net supertri@google.com ciera@google.com collinwinter@waymo.com

Matthew Jorde Andrea Knight Andrew Trenk Steve Gross
Google, LLC *Google, LLC* *Google, LLC* *Google, LLC*
 majorde@google.com aknight@google.com atrenk@google.com stevegross@google.com

- To increase awareness and adoption of software tools and practices, Google uses a technique called “Testing on the Toilet”, or TotT for short
- Evaluation of the effectiveness of TotT
- **Hypothesis:** Testing on the Toilet increases usage of advertised developer tools.
- **Case Study:** CausalImpact, a Bayesian statistical technique that was developed to evaluate the impact of advertising on website traffic

Testing on the Toilet Presents... Healthy Code on the Commode



Automatic formatting for C++

by Daniel Jasper in Munich



Are you tired of hitting space and backspace more often than anything else while coding? Are you **annoyed by fighting over parameter and comment alignment in code reviews?**

Consistent formatting allows readers to quickly scan and interpret code, dedicating their attention to what the code does and how it works. Without this consistency, effort is wasted parsing the wide variety of personal styles code might follow. However, **keeping your code formatting nice and shiny is not a good task for humans**. Luckily, we now have clang-format, which can do this tedious task for you.

Clang-format produces both readable and Google style-compliant code:

```
$ cat file.cc
int a; // clang-format can ..
int bbb;           // .. align trailing comments.
#define UNDERSTAND_MULTILINE_MACROS int cc; int d;
LOG(INFO) << ".. align operators\n" << ".. and many more things";
$ clang-format file.cc -style Google
int a;    // clang-format can ..
int bbb;  // .. align trailing comments.
#define UNDERSTAND_MULTILINE_MACROS
    int cc;
    int d;
LOG(INFO) << ".. align operators\n"
          << ".. and many more things";
```

Conveniently **integrating with your editor**, you can format the current statement or a selected **region** (available for vim, emacs and eclipse - go/clang-format). You can also reformat unified diffs, e.g. in a CitC client, by:

```
$ g4 diff -du0 | /usr/lib/clang-format/clang-format-diff.py
```

In addition to making the editor-based code development faster and more fun, **consistently using clang-format provides other advantages:**

- Code reviewers don't even need to consider whether all your spaces are correct
- Source files become fully **machine editable**, e.g. for API maintenance

So, give it a try and see how much fun it is to just type everything into a single line and let clang-format do the rest. If you encounter clang-format messing up the formatting, e.g. producing style guide violations, please file a bug on go/clang-format-bug.

[clang-format](http://go/clang-format)
 Learn how to use clang-format in your workflow.
<http://go/clang-format>

Find out more: go/CodeHealth

[Scythe](http://go/scythe)
 Want to see your dead code and automatically get rid of it?
<http://go/scythe>

Read all TotTs online: <http://tott>

FLOSS Participants' Perceptions about Gender and Inclusiveness: A Survey

Investigating the Effects of Gender Bias on GitHub

Nasif Imtiaz¹, Justin Middleton¹, Joymallya Chakraborty¹, Neill Robson¹, Gina Bai¹, and Emerson Murphy-Hill^{*2}

¹Department of Computer Science, North Carolina State University

²Google, LLC

{n.imtiaz, jmiddle1, jchakra, nrobson, gbai21}@ncsu.edu, emersonm@google.com

Engineering Gender-Inclusivity into Software: Tales from the Trenches

Claudia Hilderbrand, Christopher Perdriau, Lara Letaw, Jillian Emard, Zoe Steine-Hanson, Margaret Burnett, Anita Sarma[†]

Oregon State University
Corvallis, Oregon, USA

{minic, perdriac, letawl, emardj, steinehz, burnett, sarmaa}@oregonstate.edu

Clone Detection

Code Clone Categorization

- Type-1 clones – Identical code fragments but may have some variations in whitespace, layout, and comments
- Type-2 clones – Syntactically equivalent fragments with some variations in identifiers, literals, types, whitespace, layout and comments
- Type-3 clones – Syntactically similar code with inserted, deleted, or updated statements
- Type-4 clones – Semantically equivalent, but syntactically different code

Key points of Code Clone

- Pros
 - Increase performance
 - Code inlining vs. function call
 - Increase program readability
- Cons
 - Increase maintenance cost
 - If one code fragment contains a bug and gets fixed, all its clone peers should be always fixed in similar ways.
 - Increase code size

Detecting Strategies

- Text matching
- Token sequence matching
- Graph matching

Collaboration

Welc

The ACM C
internation
place where
latest in int
field of HCI
the first tim

Our theme
discoveries
CHI as a pla
This gives r

Please see
online virtu

An aerial photograph of a coastline, possibly a fjord or a bay, with steep, rocky cliffs and a body of water. The entire image is overlaid with a strong pink and red color filter, giving it a surreal, artistic appearance. The text of the conference title is centered over the lower half of the image.

The 24th ACM Conference on Computer-Supported Cooperative Work and Social Computing

I Interest

Game

Predicting Developers' Negative Feelings about Code Review

Carolyn D. Egelman¹, Emerson Murphy-Hill¹, Elizabeth Kammer¹, Margaret Morrow Hodges²,
Collin Green¹, Ciera Jaspan¹, James Lin¹

How Software Practitioners Use Informal Local Meetups to Share Software Engineering Knowledge

How to Hackathon: Socio-technical Tradeoffs in Brief, Intensive Collocation

Erik H. Trainer, Arun Kalyanasundaram, Chalalai Chaihirunkarn, James D. Herbsleb
Institute for Software Research
Carnegie Mellon University

Scaling Open Source Communities: An Empirical Study of the Linux Kernel

Xin Tan
Department of Computer Science and
Technology, Peking University

Minghui Zhou*
Department of Computer Science and
Technology, Peking University

Brian Fitzgerald
Lero—the Irish Software Research
Centre, University of Limerick