2025 IEEE International Conference on Software Maintenance and Evolution (ICSME 2025)

Auckland, New Zealand
7-12 September 2025

Pages 1-479

IEEE Catalog Number: CFP25079-POD
ISBN: 979-8-3315-9588-3

Copyright © 2025 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP25079-POD
ISBN (Print-On-Demand):	979-8-3315-9588-3
ISBN (Online):	979-8-3315-9587-6
ISSN:	1063-6773

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

2025 IEEE International Conference on Software Maintenance and Evolution (ICSME)

ICSME 2025

Table of Contents

Message from the General Chairs ..... **xviii**
Message from the Program Chairs ..... **xix**
Organizing Committee ..... **xx**
Program Committee ..... **xxii**
Sponsors ..... **xxviii**
Keynotes ..... **xxix**

Research Track

Evaluating the Maintainability of Forward-Porting Vulnerabilities in Fuzzer Benchmarks 1
Timothée Riom (Umeå Universitet, Sweden), Sabine Houy (Umeå Universitet, Sweden), Bruno Kreyssig (Umeå Universitet, Sweden), and Alexandre Bartel (Umeå Universitet, Sweden)

Software Fairness Testing in Practice 13
Ronnie de Souza Santos (University of Calgary, Canada), Matheus de Morais Leça (University of Calgary, Canada), Reydne Santos (UFPE, Brazil), and Cleyton Magalhaes (UFPE, Brazil)

Refactoring Deep Learning Code: A Study of Practices and Unsatisfied Tool Needs 24
Siqi Wang (Zhejiang University, China), Xing Hu (Zhejiang University, China), Bei Wang (Zhejiang University, China), Wenxin Yao (Zhejiang University, China), Xin Xia (Zhejiang University, China), and Xinyu Wang (Zhejiang University, China)

APIDocBooster: An Extract-Then-Abstract Framework for Augmenting API Documentation 36
Chengran Yang (Singapore Management University, Singapore), Jiakun Liu (Singapore Management University, Singapore), Bowen Xu (North Carolina State University, USA), Christoph Treude (Singapore Management University, Singapore), Yunbo Lyu (Singapore Management University, Singapore), Junda He (Singapore Management University, Singapore), Ming Li (Nanjing University, China), and David Lo (Singapore Management University, Singapore)

LLM-SZZ: Novel Vulnerability-Inducing Commit Identification Driven by Large Language Model and CVE Description 48
Siqi Fan (Lanzhou University, China), Xin Liu (Lanzhou University, China), Yingli Zhang (Lanzhou University, China), Yuan Tan (Lanzhou University, China), Luxing Yin (Lanzhou University, China), Zhaorun Chen (University of Chicago, United States), Song Li (Zhejiang University, China), Lei Qiao (Lanzhou University, China), and Rui Zhou (Lanzhou University, China)

SAEL: Leveraging Large Language Models with Adaptive Mixture-of-Experts for Smart Contract Vulnerability Detection 61
Lei Yu (Institute of Software, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Shiqi Cheng (Institute of Software, Chinese Academy of Sciences, China), Zhirong Huang (Institute of Software, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Jingyuan Zhang (Institute of Software, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Chenjie Shen (Institute of Software, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Junyi Lu (Institute of Software, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Li Yang (Institute of Software, Chinese Academy of Sciences, China), Fengjun Zhang (Institute of Software, Chinese Academy of Sciences, China), and Jiajia Ma (Institute of Software, Chinese Academy of Sciences, China)

Harnessing LLMs for Document-Guided Fuzzing of OpenCV Library 73
Bin Duan (The University of Queensland, Australia), Tarek Mahmud (Texas State University, USA), Meiru Che (Central Queensland University, Australia), Yan Yan (University of Illinois Chicago, USA), Naipeng Dong (The University of Queensland, Australia), Dan Dongseong Kim (The University of Queensland, Australia), and Guowei Yang (The University of Queensland, Australia)

Boosting Redundancy-Based Automated Program Repair by Fine-Grained Pattern Mining 85
Fengjie Li (Tianjin University, China), Zijie Zhao (University of Pennsylvania, USA), Zhirui Ye (Westlake University, China), Mengjiao Liu (Tianjin University, China), Bo Wang (Beijing Jiaotong University, China), Jiajun Jiang (Tianjin University, China), Hongyu Zhang (Chongqing University, China), and Junjie Chen (Tianjin University, China)

Code Property Graph Meets Typestate: A Scalable Framework to Behavioral Bug Detection 98
Xingjing Deng (Hangzhou Innovation Institute of Beihang University), Zhengyao Liu (Hangzhou Innovation Institute of Beihang University), Xitong Zhong (Hangzhou Innovation Institute of Beihang University), Shuo Hong (Hangzhou Innovation Institute of Beihang University), Yixin Yang (Hangzhou Innovation Institute of Beihang University), Xiang Gao (Hangzhou Innovation Institute of Beihang University), Xuhui Yan (Huawei Company), and Hailong Sun (Hangzhou Innovation Institute of Beihang University)

Enhanced Vulnerability Localization: Harmonizing Task-Specific Tuning and General LLM Prompting 110
Wentong Tian (Hangzhou Innovation Institute of Beihang University), Yuanzhang Lin (Hangzhou Innovation Institute of Beihang University), Xiang Gao (Hangzhou Innovation Institute of Beihang University), and Hailong Sun (Hangzhou Innovation Institute of Beihang University)

Evaluation of the Language Server Protocol for Static Dependency Analysis 123
Falko Galperin (University of Bremen, Germany), Michel Krause (University of Bremen, Germany), and Rainer Koschke (University of Bremen, Germany)

From Release to Adoption: Challenges in Reusing Pre-Trained AI Models for Downstream Developers 135
Peerachai Banyongrakkul (The University of Melbourne, Australia), Mansooreh Zahedi (The University of Melbourne, Australia), Patanamon Thongtanunam (The University of Melbourne, Australia), Christoph Treude (Singapore Management University, Singapore), and Haoyu Gao (The University of Melbourne, Australia)

Toward Realistic Evaluations of Just-In-Time Vulnerability Prediction 148
Duong Nguyen (Hanoi University of Science and Technology, Viet Nam), Thanh Le-Cong (The University of Melbourne, Australia), Triet Huynh Minh Le (University of Adelaide, Australia), M. Ali Babar (University of Adelaide, Australia), and Quyet-Thang Huynh (Hanoi University of Science and Technology, Viet Nam)

Understanding the Faults in Serverless Computing Based Applications: An Empirical Study 161
Changrong Xie (National University of Defense Technology, China), Yang Zhang (National University of Defense Technology, China), Xinjun Mao (National University of Defense Technology, China), Kang Yang (National University of Defense Technology, China), and Tanghaoran Zhang (National University of Defense Technology, China)

Metamorphic Testing of Large Language Models for Natural Language Processing 174
Steven Cho (University of Auckland, New Zealand), Stefano Ruberto (JRC European Commission, Italy), and Valerio Terragni (University of Auckland, New Zealand)

Adoption and Evolution of Code Style and Best Programming Practices in Open-Source Projects 187
Alvari Kupari (University of Auckland, New Zealand), Nasser Giacaman (University of Auckland, New Zealand), and Valerio Terragni (University of Auckland, New Zealand)

Performance Testing in Open-Source Web Projects: Adoption, Maintenance, and a Change Taxonomy 199
Sergio Di Meglio (University of Naples Federico II, Italy), Luigi Libero Lucio Starace (University of Naples Federico II, Italy), Valeria Pontillo (Software Languages (SOFT) Lab, Vrije Universiteit Brussel, Belgium; Gran Sasso Science Institute (GSSI), Italy), Ruben Opdebeeck (Software Languages (SOFT) Lab, Vrije Universiteit Brussel, Belgium), Coen De Roover (Software Languages (SOFT) Lab, Vrije Universiteit Brussel, Belgium), and Sergio Di Martino (University of Naples Federico II, Italy)

A Jump-Table-Agnostic Switch Recovery on ASTs 211
Steffen Enders (Fraunhofer FKIE, Germany), Eva-Maria C. Behner (Fraunhofer FKIE, Germany), and Elmar Padilla (Fraunhofer FKIE, Germany)

Bridging Solidity Evolution Gaps: An LLM-Enhanced Approach for Smart Contract Compilation Error Resolution 223
Likai Ye (Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security, Zhejiang University, China), Mengliang Li (Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security, Zhejiang University, China), Dehai Zhao (Zhejiang University, China), Jiamou Sun (CSIRO's Data61, Australia), and Xiaoxue Ren (Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security, Zhejiang University, China)

TSGuard: Detecting Logic Bugs in Time Series Management Systems via Time Series Algebra 235
Lingwei Kuang (Nanjing University of Aeronautics and Astronautics, China), Liang Liu (Nanjing University of Aeronautics and Astronautics, China), Wenjing Wang (Nanjing University of Aeronautics and Astronautics, China), Ning Cao (Nanjing University of Aeronautics and Astronautics, China), Shijie Li (Nanjing University of Aeronautics and Astronautics, China), Fan Liu (Nanjing University of Aeronautics and Astronautics, China), Haolong Chen (Nanjing University of Aeronautics and Astronautics, China), and WenJian Liao (Nanjing FiberHome StarrySky Communication Development CN, China)

Onweer: Automated Resilience Testing through Fuzzing 248
Gilles Coremans (Vrije Universiteit Brussel, Belgium) and Coen De Roover (Vrije Universiteit Brussel, Belgium)

OptionFuzz: Fuzzing SMT Solvers with Optimized Option Exploration via Large Language Models 260
Yuhao Peng (Institute of Software Chinese Academy of Sciences; University of Chinese Academy of Sciences), Jingzheng Wu (Institute of Software Chinese Academy of Sciences), Xiang Ling (Institute of Software Chinese Academy of Sciences), Zhiyuan Li (Institute of Software Chinese Academy of Sciences; University of Chinese Academy of Sciences), Tianyue Luo (Institute of Software Chinese Academy of Sciences), and Yanjun Wu (Institute of Software Chinese Academy of Sciences)

AST-Enhanced or AST-Overloaded? The Surprising Impact of Hybrid Graph Representations on Code Clone Detection 271
Zixian Zhang (University of Galway, Ireland) and Takfarinas Saber (University of Galway, Ireland; Lero, Research Ireland Centre for Software)

Let it be Chaos in the Plumbing! Usage and Efficacy of Chaos Engineering in DevOps Pipelines 282
Stefano Fossati (JADS - Eindhoven University of Technology, Netherlands), Damian Andrew Tamburri (University of Sannio & JADS/NXP Semiconductors, Italy), Massimiliano Di Penta (University of Sannio, Italy), and Marco Tonnarelli (JADS - Eindhoven University of Technology, Netherlands)

Are Classical Clone Detectors Good Enough for the AI Era? 295
Ajmain I. Alam (University of Saskatchewan, Canada), Palash R. Roy (University of Saskatchewan, Canada), Farouq Al-omari (Thompson Rivers University, Canada), Chanchal K. Roy (University of Saskatchewan, Canada), Banani Roy (University of Saskatchewan, Canada), and Kevin A. Schneider (University of Saskatchewan, Canada)

Are All Code Reviews the Same? Identifying and Assessing the Impact of Merge Request Deviations 308
Samah Kansab (École1 de technologie supérieure (ÉTS), Canada), Francis Bordeleau (École de technologie supérieure (ÉTS), Canada), and Ali Tizghadam (TELUS, Canada)

Characterizing the System Evolution that is Proposed After a Software Incident 321
Matt Pope (Brigham Young University, USA) and Jonathan Sillito (Brigham Young University, USA)

ADPP: Automated Data-Centric Program Partitioning 333
Xuming Jin (Nanjing University of Aeronautics and Astronautics, China) and Hao Han (Nanjing University of Aeronautics and Astronautics, China)

Social Media Reactions to Open Source Promotions: AI-Powered GitHub Projects on Hacker News 344
Prachnachai Meakpaiboonwattana (Mahidol University, Thailand), Warittha Tarntong (Mahidol University, Thailand), Thai Mekratanavorakul (Mahidol University, Thailand), Chaiyong Ragkhitwetsagul (Mahidol University, Thailand), Pattaraporn Sangaroonsilp (Mahidol University, Thailand), Raula Gaikovina Kula (The University of Osaka, Japan), Morakot Choetkiertikul (Mahidol University, Thailand), Kenichi Matsumoto (Nara Institute of Science and Technology, Japan), and Thanwadee Sunetnanta (Mahidol University, Thailand)

Trust and Verify: Formally Verified and Upgradable Trusted Functions 356
Marcus Birgersson (KTH Royal Institute of Technology, Sweden), Cyrille Artho (KTH Royal Institute of Technology, Sweden), and Musard Balliu (KTH Royal Institute of Technology, Sweden)

Retrieve, Refine, or Both? using Task-Specific Guidelines for Secure Python Code Generation 368
Catherine Tony (Hamburg University of Technology, Germany), Emanuele Iannone (Hamburg University of Technology, Germany), and Riccardo Scandariato (Hamburg University of Technology, Germany)

The Impact of Fine-Tuning Large Language Models on Automated Program Repair 380
Roman Macháček (University of Bern, Switzerland), Anastasiia Grishina (Simula Research Laboratory, Norway), Max Hort (Simula Research Laboratory, Norway), and Leon Moonen (Simula Research Laboratory, Norway)

A Taxonomy of Inefficiencies in LLM-Generated Python Code 393
Altaf Allah Abbassi (Polytechnique Montréal, Canada), Leuson Da Silva (Polytechnique Montréal, Canada), Amin Nikanjam (Huawei Distributed Scheduling and Data Engine Lab, Canada), and Foutse Khomh (Polytechnique Montréal, Canada)

Security Vulnerabilities in Docker Images: A Cross-Tag Study of Application Dependencies 405
Hamid Mohayeji (Eindhoven University of Technology, The Netherlands), Eleni Constantinou (University of Cyprus, Cyprus), and Alexander Serebrenik (Eindhoven University of Technology, The Netherlands)

Automatically Augmenting GitHub Issues with Informative User Reviews 418
Arthur Pilone (Institute of Mathematics and Statistics - University of São Paulo, Brazil), Marco Raglianti (REVEAL @ Software Institute --- USI, Lugano, Switzerland), Michele Lanza (REVEAL @ Software Institute --- USI, Lugano, Switzerland), Fabio Kon (Institute of Mathematics and Statistics - University of São Paulo, Brazil), and Paulo Meirelles (Institute of Mathematics and Statistics - University of São Paulo, Brazil)

Together We are Better: LLM, IDE and Semantic Embedding to Assist Move Method Refactoring .. 430
Abhiram Bellur (University of Colorado), Fraol Batole (Tulane University), Mohammed Raihan Ullah (University of Colorado), Malinda Dilhara (Amazon Web Services), Yaroslav Zharov (JetBrains Research), Timofey Bryksin (JetBrains Research), Kai Ishikawa (NEC Corporation), Haifeng Chen (NEC Laboratories America), Masaharu Morimoto (NEC Corporation), Takeo Hosomi (NEC Corporation), Tien N. Nguyen (University of Texas at Dallas), Hridesh Rajan (Tulane University), Nikolaos Tsantalis (Concordia University), and Danny Dig (University of Colorado)

Syntest-ACR: Automated Crash Reproduction for JavaScript 443
Philip Oliver (Victoria University of Wellington, New Zealand), Jens Dietrich (Victoria University of Wellington, New Zealand), Craig Anslow (Victoria University of Wellington, New Zealand), and Michael Homer (Victoria University of Wellington, New Zealand)

Can LLMs Update API Documentation? 455
Seonah Lee (Gyeongsang National University, Republic of Korea), Jueun Heo (Gyeongsang National University, Republic of Korea), and Katherine R. Dearstyne (University of Notre Dame, USA)

Is Quantization a Deal-Breaker? Empirical Insights from Large Code Models 467
Saima Afrin (William & Mary, USA), Bowen Xu (North Carolina State University, USA), and Antonio Mastropaolo (William & Mary, USA)

Boosting Log Observability in Production Systems through Bytecode-Driven Fault Variable Tracking 480
Taizheng Wang (Hainan University, China), Yutong Wang (Hainan University, China), Wei Chang (Hainan University, China), Chunyang Ye (Hainan University, China), and Hui Zhou (Hainan University, China)

Does Editing Improve Answer Quality on Stack Overflow? A Data-Driven Investigation 492
Saikat Mondal (University of Saskatchewan, Canada) and Chanchal K. Roy (University of Saskatchewan, Canada)

RMGenie: An LLM-Based Agent Framework for Open Source Software README Generation 505
Xing Cui (Institute of Software, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Jingzheng Wu (Institute of Software, Chinese Academy of Sciences), Zhiyuan Li (Institute of Software, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Tianyue Luo (Institute of Software, Chinese Academy of Sciences), and Xiang Ling (Institute of Software, Chinese Academy of Sciences; University of Chinese Academy of Sciences)

Roseau: Fast, Accurate, Source-Based API Breaking Change Analysis in Java 517
Corentin Latappy (Univ. Bordeaux, France), Thomas Degueule (Univ. Bordeaux, France), Jean-Rémy Falleri (Univ. Bordeaux, France), Romain Robbes (Univ. Bordeaux, France), and Lina Ochoa (Eindhoven University of Technology)

An Efficient Android App Debloating Approach Based on Multi-Layer Dependence Graph 529
Hengqin Yang (University of Chinese Academy of Sciences, China; Institute of Software, Chinese Academy of Sciences, China), Jiwei Yan (University of Chinese Academy of Sciences, China; Institute of Software, Chinese Academy of Sciences, China), Jun Yan (University of Chinese Academy of Sciences, China; Institute of Software, Chinese Academy of Sciences, China; Institute of Software, Chinese Academy of Sciences, China), Bin Liang (Renmin University of China, China), and Jian Zhang (University of Chinese Academy of Sciences, China; Institute of Software, Chinese Academy of Sciences, China)

Nüwa: Enhancing MLIR Fuzzing with LLM-Driven Generation and Adaptive Mutation 541
Bocan Cao (Northwest University, China), Weiyuan Tong (Northwest University, China), Zhanyong Tang (Northwest University, China), Zixu Wang (Northwest University, China), Hao Huang (Northwest University, China), and Yuheng Yan (Northwest University, China)

Industry Track

HybridRCA Lightweight Critical-Path-Aware Hybrid Tracing for Root-Cause Analysis in Production Microservices 553
Maryam Ekhlasi (Software Engineer at Ciena, Canada), Arnaud Fiorini (Polytechnique Montreal, Canada), Michel Dagenais (Polytechnique Montreal, Canada), Naser Ezzati-Jivan (Brock University, Canada), and Maxime Lamothe (Polytechnique Montreal, Canada)

Improving Merge Pipeline Throughput in Continuous Integration via Pull Request Prioritization 565
Maximilian Jungwirth (BMW Group, University of Passau, Germany), Martin Gruber (BMW Group, Germany), and Gordon Fraser (University of Passau, Germany)

Levels of Binary Equivalence for the Comparison of Binaries from Alternative Builds 576
Jens Dietrich (Victoria University of Wellington, New Zealand), Tim White (Victoria University of Wellington, New Zealand), Behnaz Hassanshahi (Oracle Labs Australia, Australia), and Paddy Krishnan (Oracle Labs Australia, Australia)

Constraint Discovery for Structured Generation via LLM-Guided SMT Inference 588
Hrishikesh Karmarkar (TCS Research, India), Supriya Agrawal (TCS Research, India), Siddhesh Pagar (TCS Research, India), Vaibhavi Joshi (TCS Research, India), Sagar Verma (TCS Research, India), and Naman Paul (TCS Research, India)

Testing is not Boring: Characterizing Challenge in Software Testing Tasks 594
Davi Gama Hardman (CESAR, Brazil), César França (CESAR School, Brazil), Brody Stuart-Verner (University of Calgary, Canada), and Ronnie de Souza Santos (University of Calgary, Canada)

Automated Code Review using Large Language Models at Ericsson: An Experience Report 602
Shweta Ramesh (Ericsson, India), Joy Bose (Ericsson, India), Hamender Singh (Ericsson, India), Ak Raghavan (Ericsson, India), Sujoy Roy Chowdhury (Ericsson, India), Giriprasad Sridhara (Ericsson, India), Nishrith Saini (Ericsson, India), and Ricardo Britto (Ericsson, India)

A Deep Dive into Retrieval-Augmented Generation for Code Completion: Experience on WeChat . 608
Zezhou Yang (Tencent, China), Ting Peng (Tencent, China), Cuiyun Gao (Chinese University of Hong Kong, China), Chaozheng Wang (Chinese University of Hong Kong, China), Hailiang Huang (Tencent, China), and Yuetang Deng (Tencent, China)

Requirements Ambiguity Detection and Explanation with LLMs: An Industrial Study 620
Sarmad Bashir (RISE Research Institutes of Sweden, Sweden), Alessio Ferrari (University College Dublin, Ireland), Abbas Khan (RISE Research Institutes of Sweden, Sweden), Per Erik Strandberg (Westermo Network Technologies AB, Sweden), Zulqarnain Haider (Alstom Rail AB, Sweden), Mehrdad Saadatmand (RISE Research Institutes of Sweden, Sweden), and Markus Bohlin (Mälardalen University, Sweden)

Efficient Detection of Intermittent Job Failures using Few-Shot Learning 632
Henri Aïdasso (École de technologie supérieure, Canada), Francis Bordeleau (École de technologie supérieure, Canada), and Ali Tizghadam (TELUS, Canada)

AskGraph: A Dependency-Aware Code Assistant Powered by Code Graphs and LLM-Generated Cypher Queries 644
Nan Yang (TNO-ESI, The Netherland), Joseph Reynolds (TNO-ESI, The Netherland), Laurens Prast (TNO-ESI, The Netherland), and Rosilde Corvino (TNO-ESI, The Netherland)

Automated Recovery of Software Product Lines from Legacy Configurable Codebases 656
Tewfik Ziadi (University of Doha for Science & Technology, Qatar), Karim Ghallab (RedFabriQ - Mobioos, France), and Zaak Chalal (RedFabriQ - Mobioos, France)

Governance Matters: Lessons from Restructuring the data.table OSS Project 667
Pedro Oliveira (Northern Arizona University, USA), Doris Amoakohene (data.table Project), Toby Hocking (data.table Project), Marco Gerosa (Northern Arizona University, USA), and Igor Steinmacher (Northern Arizona University, USA)

AI Mentor System: Building a Technical Debt Dashboard for Low Code 679
Alexandre Lemos (OutSystems, Portugal) and Joana Coutinho (OutSystems, Portugal)

Integrating Rules and Semantics for LLM-Based C-to-Rust Translation 685
Feng Luo (Harbin Institute of Technology (Shenzhen), China), Kexing Ji (Harbin Institute of Technology (Shenzhen), China), Cuiyun Gao (Harbin Institute of Technology (Shenzhen), China), Shuzheng Gao (The Chinese University of Hong Kong, China), Jia Feng (Harbin Institute of Technology (Shenzhen), China), Kui Liu (Huawei Software Engineering Application Technology Lab, China), Xin Xia (Zhejiang University, China), and Michael R. Lyu (The Chinese University of Hong Kong, China)

Targeted Test Selection Approach in Continuous Integration 697
Pavel Plyusnin (T-Technologies, Russia), Aleksey Antonov (T-Technologies, Russia), Vasilii Ermakov (T-Technologies, Russia), Alexandr Khaybriev (T-Technologies, Russia), Margarita Kikot (T-Technologies, Russia), Ilseyar Alimova (T-Technologies, Russia), and Stanislav Moiseev (T-Technologies, Russia)

Combining Insights from Multiple Tools to Manage Technical Debt in Industrial C# Projects 709
Simeon Tverdal (SINTEF, Norway), Phu Nguyen (SINTEF, Norway), Arda Goknil (SINTEF, Norway), Antonio Martini (University of Oslo, Norway), Merve Astekin (SINTEF, Norway), Mili Orucevic (Visma AS, Norway), Maren Maritsdatter Kruke (Visma AS, Norway), and Håvard Stranden (AKVA Group, Norway)

Monitoring Continuous Integration Practices in Industry: A Case Study 721
Jadson Santos (Federal University of Rio Grande do Norte, Brazil), Daniel Alencar da Costa (University of Otago, New Zealand), and Uirá Kulesza (Federal University of Rio Grande do Norte, Brazil)

Code Readability in the Age of Large Language Models: An Industrial Case Study from Atlassian 732
Wannita Takerngsaksiri (Monash University, Australia), Chakkrit Tantithamthavorn (Monash University, Australia), Micheal Fu (The University of Melbourne, Australia), Jirat Pasuksmit (Atlassian, Australia), Kun Chen (Atlassian, United States), and Ming Wu (Atlassian, United States)

Understanding Practitioners' Perspectives on Monitoring Machine Learning Systems 743
Hira Naveed (Monash University, Australia), John Grundy (Monash University, Australia), Chetan Arora (Monash University, Australia), Hourieh Khalajzadeh (Deakin University, Australia), and Omar Haggag (Monash University, Australia)

New Ideas and Emerging Results (NIER) Track

Towards Interaction-Aware Validation Oracles for Multi-Fault Program Repair 755
Omar I. Al-Bataineh (Gran Sasso Science Institute, Italy)

MalLoc: Toward Fine-Grained Android Malicious Payload Localization via LLMs 761
Tiezhu Sun (University of Luxembourg, Luxembourg), Marco Alecci (University of Luxembourg, Luxembourg), Aleksandr Pilgun (University of Luxembourg, Luxembourg), Yewei Song (University of Luxembourg, Luxembourg), Xunzhu Tang (University of Luxembourg, Luxembourg), Jordan Samhi (University of Luxembourg, Luxembourg), Tegawendé F. Bissyandé (University of Luxembourg, Luxembourg), and Jacques Klein (University of Luxembourg, Luxembourg)

Can LLMs Write CI? A Study on Automatic Generation of GitHub Actions Configurations 767
Taher A. Ghaleb (Trent University, Canada) and Dulina Rathnayake (Trent University, Canada)

From First Use to Final Commit: Studying the Evolution of Multi-CI Service Adoption 773
Nitika Chopra (Trent University, Canada) and Taher A. Ghaleb (Trent University, Canada)

How Does Test Code Differ from Production Code in Terms of Refactoring? An Empirical Study... 779
Kosei Horikawa (Nara Institute of Science and Technology, Japan), Yutaro Kashiwa (Nara Institute of Science and Technology, Japan), Bin Lin (Hangzhou Dianzi University, China), Kenji Fujiwara (Nara Women's University, Japan), and Hajimu Iida (Nara Institute of Science and Technology, Japan)

Inferring Attributed Grammars from Parser Implementations 785
Andreas Pointner (Advanced Information Systems and Technology, University of Applied Sciences, Austria), Josef Pichler (Department for Software Engineering, University of Applied Sciences, Austria), and Herbert Prähofer (Institut for System Software, Johannes Kepler University Linz, Austria)

Client–Library Compatibility Testing with API Interaction Snapshots 791
Gustave Monce (Univ. Bordeaux, France), Thomas Degueule (Univ. Bordeaux, France), Jean-Rémy Falleri (Univ. Bordeaux, France), and Romain Robbes (Univ. Bordeaux, France)

Rethinking Cognitive Complexity for Unit Tests: Toward a Readability-Aware Metric Grounded in Developer Perception 797
Wendkûuni C. Ouédraogo (University of Luxembourg), Yinghua Li (University of Luxembourg), Xueqi Dang (University of Luxembourg), Xin Zhou (Singapore Management University, Singapore), Anil Koyuncu (Bilkent University, Turkey), Jacques Klein (University of Luxembourg), David Lo (Singapore Management University, Singapore), and Tegawendé F. Bissyandé (University of Luxembourg)

Prompting Matters: Assessing the Effect of Prompting Techniques on LLM-Generated Class Code 803
Adam Yuen (University of Calgary, Canada), John Pangas (University of Calgary, Canada), Mainul Hasan Polash (University of Calgary, Canada), and Ahmad Abdellatif (University of Calgary, Canada)

How Much Can a Behavior-Preserving Changeset Be Decomposed into Refactoring Operations? ... 809
Kota Someya (Institute of Science Tokyo, Japan), Lei Chen (Institute of Science Tokyo, Japan), Michael J. Decker (Bowling Green State University, USA), and Shinpei Hayashi (Institute of Science Tokyo, Japan)

Prioritizing Test Smells: An Empirical Evaluation of Quality Metrics and Developer Perceptions 815
Arif Hasan (University of Dhaka, Bangladesh) and Toukir Ahammed (University of Dhaka, Bangladesh)

Explicit Vulnerability Generation with LLMs: An Investigation Beyond Adversarial Attacks 821
Ahmet Emir Bosnak (Bilkent University, Turkey), Sahand Moslemi (Bilkent University, Turkey), Mayasah Lami (Bilkent University, Turkey), and Anil Koyuncu (Bilkent University, Turkey)

LLMShot: Reducing Snapshot Testing Maintenance via LLMs 827
Ergün Batuhan Kaynak (Bilkent University, Turkey), Mayasah Lami (Bilkent University, Turkey), Sahand Moslemi (Bilkent University, Turkey), and Anil Koyuncu (Bilkent University, Turkey)

A Preliminary Study on Large Language Models Self-Negotiation in Software Engineering 833
Chunrun Tao (Kyushu University, Japan), Honglin Shu (Kyushu University, Japan), Masanari Kondo (Kyushu University, Japan), and Yasutaka Kamei (Kyushu University, Japan)

Detecting Adversarial Prompted AI-Generated Code on Stack Overflow: A Benchmark Dataset and an Enhanced Detection Approach 839
Aman Swaraj (Indian Institute of Technology, Roorkee, India), Krishna Agarwal (Indian Institute of Technology Roorkee, India), Atharv Joshi (Indian Institute of Technology Roorkee, India), and Sandeep Kumar (Indian Institute of Technology Roorkee, India)

Repairing Responsive Layout Failures using Retrieval Augmented Generation 845
Tasmia Zerin (University of Dhaka, Bangladesh), Moumita Asad (University of California, Irvine, USA), B M Mainul Hossain (University of Dhaka, Bangladesh), and Kazi Sakib (University of Dhaka, Bangladesh)

ExpertCache: GPU-Efficient MoE Inference through Reinforcement Learning-Guided Expert Selection ..... N/A
Xunzhu Tang (University of Luxembourg), Tiezhu Sun (University of Luxembourg), Yewei Song (University of Luxembourg), Siyuan Ma (Nanyang Technological University, Singapore), Jacques Klein (University of Luxembourg), and Tegawendé F. Bissyandé (University of Luxembourg)

Doctoral Symposium Track

Toward a Resilience-Oriented Understanding of Unit Test Suites and Refactoring in Software Evolution 857
Daniel Knight (Mississippi State University)

Towards Proficiency Assessment through Code 860
Ruksit Rojpaisarnkit (Nara Institute of Science and Technology, Japan)

The Influence of Code Smell Interactions on Software Maintainability 865
Zushuai Zhang (University of Auckland, New Zealand)

Ensuring Code Integrity in the Era of AI-Assisted Software Development 870
Arthur Pilone (University of São Paulo)

Understanding and Simulating OSS Evolution: A Case Study on PyMC 873
Toru Sugiyama (The Open University of Japan, Japan)

The Impact of Generative AI on Developer Practices, Behavior, and Software Quality 878
Julian Oertel (University of Rostock, Germany)

Bugs in AI-Generated Code - Understanding Bug Patterns and Possible Fix Strategies 881
Ruofan Gao (Massey University, New Zealand)

Towards Just-In-Time, Inclusive Clone Refactoring 884
Palash Roy (University of Saskatchewan, Canada)

Enhancing Infrastructure Maintenance and Evolution through Graph-Based Visualization and Analysis 887
Stefano Fossati (JADS - Eindhoven University of Technology, The Netherlands)

DevSecLogs: AI-Powered, Tamper-Evident Log Intelligence for Secure CI/CD Pipelines 890
Sabbir M. Saleh (University of Western Ontario, Canada)

Tool Demonstration Track

LeakageDetector 2.0: Analyzing Data Leakage in Jupyter-Driven Machine Learning Pipelines 895
Owen Truong (Stevens Institute of Technology, USA), Terrence Zhang (Stevens Institute of Technology, USA), Arnav Marchareddy (Stevens Institute of Technology, USA), Ryan Lee (Stevens Institute of Technology, USA), Jeffery Busold (Stevens Institute of Technology, USA), Michael Socas (Stevens Institute of Technology, USA), and Eman Abdullah AlOmar (Stevens Institute of Technology, USA)

DENIM: Exploring Data Access in Microservices 900
Maxime André (University of Namur, Belgium), Marco Raglianti (Università della Svizzera italiana, Switzerland), Anthony Cleve (University of Namur, Belgium), and Michele Lanza (Università della Svizzera italiana, Switzerland)

XTestGen: Natural Language to Maintainable E2E Test Scripts with LLMs 905
Hiroyuki Kirinuki (Computer and Data Science Labs., NTT, Inc., Japan), Masaki Tajima (Computer and Data Science Labs., NTT, Inc., Japan), and Kei Wakabayashi (Computer and Data Science Labs., NTT, Inc., Japan)

MaRCo: Compatible Version Ranges in Maven 910
Cathrine Paulsen (Delft University of Technology, The Netherlands) and Sebastian Proksch (Delft University of Technology, The Netherlands)

AI-Powered Commit Explorer (APCE) 915
Yousab Grees (Belmont University, USA), Polina Iaremchuk (Belmont University, USA), Ramtin Ehsani (Drexel University, USA), Esteban Parra (Belmont University, USA), Preetha Chatterjee (Drexel University, USA), and Sonia Haiduc (Florida State University, USA)

JDala - A Simple Capability System for Java 920
Quinten Smit (Victoria University of Wellington, New Zealand), Jens Dietrich (Victoria University of Wellington, New Zealand), Michael Homer (Victoria University of Wellington, New Zealand), Andrew Fawcet (Victoria University of Wellington, New Zealand), and James Noble (Creative Research and Programming, New Zealand)

MediumDarwin: LittleDarwin Grows with Performance and Research-Oriented Extensions 925
Sajjad Hesamipour Khelejan (Trinity College Dublin & Research Ireland Lero), Thomas Laurent (Trinity College Dublin & Research Ireland Lero), and Anthony Ventresque (Trinity College Dublin & Research Ireland Lero)

LLMLoop: Improving LLM-Generated Code and Tests through Automated Iterative Feedback Loops.. 930
Ravin Ravi (University of Auckland, New Zealand), Dylan Bradshaw (University of Auckland, New Zealand), Stefano Ruberto (JRC European Commission, Italy), Gunel Jahangirova (King's College London, United Kingdom), and Valerio Terragni (University of Auckland, New Zealand)

CodeWatcher: IDE Telemetry Data Extraction Tool for Understanding Coding Interactions with LLMs 935
Manaal Basha (The University of British Columbia), Aimee M. Ribeiro (Federal University of Para), Jeena Javahar (The University of British Columbia), Cleidson R. B. de Souza (Federal University of Para), and Gema Rodriguez Perez (The University of British Columbia)

LadyBug: A GitHub Bot for UI-Enhanced Bug Localization in Mobile Apps 940
Junayed Mahmud (University of Central Florida, USA), James Chen (University of Central Florida, USA), Terry Achille (University of Central Florida, USA), Camilo Alvarez-Velez (University of Central Florida, USA), Darren Dean Bansil (University of Central Florida, USA), Patrick Ijieh (University of Central Florida, USA), Samar Karanch (University of Central Florida, USA), Nadeeshan De Silva (William & Mary, USA), Oscar Chaparro (William & Mary, USA), Andrian Marcus (George Mason University, USA), and Kevin Moran (University of Central Florida, USA)

VulGuard: An Unified Tool for Evaluating Just-In-Time Vulnerability Prediction Models 945
Duong Nguyen (Hanoi University of Science and Technology, Viet Nam), Manh Tran-Duc (Hanoi University of Science and Technology, Viet Nam), Thanh Le-Cong (The University of Melbourne, Australia), Triet Huynh Minh Le (University of Adelaide, Australia), M. Ali Babar (University of Adelaide, Australia), and Quyet-Thang Huynh (Hanoi University of Science and Technology, Viet Nam)

COB2PY - A Non-AI, Rule-Based COBOL to Python Translator 950
Kowshik Reddy Challa (Indian Institute of Technology, India), M V Sonith (Indian Institute of Technology, India), B S Chiranjeevi (Indian Institute of Technology, India), and Sridhar Chimalakonda (Indian Institute of Technology Tirupati, India)

Author Index