

2025 40th IEEE/ACM International Conference on Automated Software Engineering (ASE 2025)

**Seoul, Republic of Korea
16-20 November 2025**

Pages 1–699

IEEE Catalog Number: CFP25075-POD
ISBN: 979-8-3503-5734-9



Copyright © 2025, IEEE

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:	CFP25075-POD
ISBN (Print-On-Demand):	979-8-3503-5734-9
ISBN (Online):	979-8-3503-5733-2
ISSN:	1938-4300

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

TABLE OF CONTENTS

AlertGuardian: Intelligent Alert Life-Cycle Management for Large-scale Cloud Systems	1
<i>Guangba Yu, Genting Mai, Rui Wang, Ruipeng Li, Pengfei Chen, Long Pan, Ruijie Xu</i>	
Argus: Resilience-Oriented Safety Assurance Framework for End-to-End ADSs	13
<i>Dingji Wang, You Lu, Bihuan Chen, Shuo Hao, Haowen Jiang, Yifan Tian, Xin Peng</i>	
Vul-R2: A Reasoning LLM for Automated Vulnerability Repair	26
<i>Xin-Cheng Wen, Zirui Lin, Yijun Yang, Cuiyun Gao, Deheng Ye</i>	
MIMIC: Integrating Diverse Personality Traits for Better Game Testing Using Large Language Model	39
<i>Yifei Chen, Sarra Habchi, Lili Wei</i>	
Learning from the Past: Real-World Exploit Migration for Smart Contract PoC Generation	52
<i>Kairan Sun, Zhengzi Xu, Kaixuan Li, Lyuye Zhang, Yebo Feng, Daoyuan Wu, Yang Liu</i>	
Propagation-Based Vulnerability Impact Assessment for Software Supply Chains	65
<i>Bonan Ruan, Zhiwei Lin, Jiahao Liu, Chuqi Zhang, Kaihang Ji, Zhenkai Liang</i>	
Debut: Detecting Bundled JavaScript Libraries on Web using Property-Order Graphs	78
<i>Seojin Kim, Sungmin Park, Jihyeok Park</i>	
When AllClose Fails: Round-Off Error Estimation for Deep Learning Programs	91
<i>Qi Zhan, Xing Hu, Yuanyi Lin, Tongtong Xu, Xin Xia, Shanping Li</i>	
Faultseeker: LLM-Empowered Framework for Blockchain Transaction Fault Localization	104
<i>Kairan Sun, Zhengzi Xu, Kaixuan Li, Lyuye Zhang, Yuqiang Sun, Liwei Tan, Yang Liu</i>	
SateLight: A Satellite Application Update Framework for Satellite Computing	117
<i>Jinfeng Wen, Jianshu Zhao, Zixi Zhu, Xiaomin Zhang, Qi Liang, Ao Zhou, Shangguang Wang</i>	
Diagnosing Performance Differences in Model Checkers via Runtime-Guided Problem Generation	129
<i>Yibo Dong, Yicong Xu, Wenjing Deng, Yu Chen, Xiaoyu Zhang, Jianwen Li, Chengyu Zhang, Geguang Pu</i>	
LongCodeZip: Compress Long Context for Code Language Models	141
<i>Yuling Shi, Yichun Qian, Hongyu Zhang, Beijun Shen, Xiaodong Gu</i>	
Code-DiTing: Automatic Evaluation of Code Generation without References or Test Cases	154
<i>Guang Yang, Yu Zhou, Xiang Chen, Wei Zheng, Xing Hu, Xin Zhou, David Lo, Taolue Chen</i>	
An Empirical Study of Knowledge Transfer in AI Pair Programming	166
<i>Alisa Welter, Niklas Schneider, Tobias Dick, Kallistos Weis, Christof Tinnes, Marvin Wyrich, Sven Apel</i>	
Wired for Reuse: Automating Context-Aware Code Adaptation in IDEs via LLM-Based Agent	178
<i>Taiming Wang, Yanjie Jiang, Chunhao Dong, Yuxia Zhang, Hui Liu</i>	
"My productivity is boosted, but ..." Demystifying Users' Perception on AI Coding Assistants	191
<i>Yunbo Lyu, Zhou Yang, Jieke Shi, Jianming Chang, Yue Liu, David Lo</i>	
Finding Bugs in WebAssembly Interface Type Binding Generators	204
<i>Ethan Stanley, Eric Eide</i>	
R ³ -Bench: Reproducible Real-world Reverse Engineering Dataset for Symbol Recovery	217
<i>Muzhi Yu, Zhengran Zeng, Wei Ye, Jinan Sun, Xiaolong Bai, Shikun Zhang</i>	

Relia: Accelerating the Analysis of Cloud Access Control Policies	228
<i>Dan Wang, Peng Zhang, Zhenrong Gu, Weibo Lin, Shibiao Jiang, Zhu He, Xu Du, Longfei Chen, Jun Li Xiaohong Guan</i>	
Interaction2Code: Benchmarking MLLM-based Interactive Webpage Code Generation from Interactive Prototyping	241
<i>Jingyu Xiao, Yuxuan Wan, Yintong Huo, Zixin Wang, Xinyi Xu, Wenxuan Wang, Zhiyao Xu, Yuhang Wang Michael R. Lyu</i>	
Defects4C: Benchmarking Large Language Model Repair Capability with C/C++ Bugs	254
<i>Jian Wang, Xiaofei Xie, Qiang Hu, Shangqing Liu, Jiongchi Yu, Jiaolong Kong, Yi Li</i>	
Evolution-Aware Heuristics for GR(1) Realizability Checking	266
<i>Dor Ma'ayan, Shahar Maoz, Jan Oliver Ringert</i>	
Do LLMs Generate Useful Test Oracles? An Empirical Study with an Unbiased Dataset	278
<i>Davide Molinelli, Luca Di Grazia, Alberto Martin-Lopez, Michael D. Ernst, Mauro Pezzè</i>	
Loupe: End-to-End Learning of Loop Unrolling Heuristics for Abstract Interpretation	291
<i>Maykel Mattar, Michele Alberti, Valentin Perrelle, Salah Sadou</i>	
Advancing Binary Code Similarity Detection via Context-Content Fusion and LLM Verification	304
<i>Chaopeng Dong, Jingdong Guo, Shouguo Yang, Yi Li, Dongliang Fang, Yang Xiao, Yongle Chen, Limin Sun</i>	
PrefGen: A Preference-Driven Methodology for Secure Yet Gas-Efficient Smart Contract Generation	317
<i>Zhiyuan Peng, Xin Yin, Zijie Zhou, Chenhao Ying, Chao Ni, Yuan Luo</i>	
LogMoE: Lightweight Expert Mixture for Cross-System Log Anomaly Detection	330
<i>Jiaxing Qi, Zhongzhi Luan, Shaohan Huang, Carol Fung, Yuchen Wang, Aibin Wang, Hongyu Zhang Hailong Yang, Depei Qian</i>	
Risk Estimation in Differential Fuzzing via Extreme Value Theory	342
<i>Rafael Baez, Alejandro Olivas, Nathan K Diamond, Marcelo Frias, Yannic Noller, Saeid Tizpaz-Niari</i>	
Enhancing LLMs with Staged Grouping and Dehallucination for Header File Decomposition	355
<i>Yue Wang, Jiaxuan Sun, Yanzhen Zou, Bing Xie</i>	
Automated Repair of Ambiguous Problem Descriptions for LLM-Based Code Generation	367
<i>Haoxiang Jia, Robbie Morris, He Ye, Federica Sarro, Sergey Mechtaev</i>	
Towards More Accurate Static Analysis for Taint-Style Bug Detection in Linux Kernel	380
<i>Haonan Li, Hang Zhang, Kexin Pei, Zhiyun Qian</i>	
Democratizing the Cryptocurrency Ecosystem by Just-In-Time Transformation of Mining Programs	393
<i>Wei Liu, Zhenhua Li, Feng Qian, Feiyu Jin, Hao Lin, Yannan Zheng, Bo Xiao, Xiaokang Qin, Tianyin Xu</i>	
From Sparse to Structured: A Diffusion-Enhanced and Feature-Aligned Framework for Coincidental Correctness Detection	406
<i>Huan Xie, Chunyan Liu, Yan Lei, Zhenyu Wu, Jinping Wang</i>	
ProfMal: Detecting Malicious NPM Packages by the Synergy between Static and Dynamic Analysis	419
<i>Yiheng Huang, Wen Zheng, Susheng Wu, Bihuan Chen, You Lu, Zhuotong Zhou, Yiheng Cao, Xiaoyu Li Xin Peng</i>	

Why AI Agents Still Need You: Findings from Developer-Agent Collaborations in the Wild	432
<i>Aayush Kumar, Yasharth Bajpai, Sumit Gulwani, Gustavo Soares, Emerson Murphy-Hill</i>	
TensorGuard: Gradient-Based Model Fingerprinting for LLM Similarity Detection and Family Classification	445
<i>Zehao Wu, Yanjie Zhao, Haoyu Wang</i>	
Uncovering Prompt Elements: Cloning System Prompts from Behavioral Traces	457
<i>Yi Qian, Fei Peng, Hao Wu, Ligeng Chen, Bing Mao</i>	
Comprehend, Imitate, and then Update: Unleashing the Power of LLMs in Test Suite Evolution	469
<i>Tangzhi Xu, Jianhan Liu, Yuan Yao, Cong Li, Feng Xu, Xiaoxing Ma</i>	
VRExplorer: A Model-based Approach for Semi-Automated Testing of Virtual Reality Scenes	482
<i>Zhengyang Zhu, Hong-Ning Dai, Hanyang Guo, Zeqin Liao, Zibin Zheng</i>	
RSFuzz: A Robustness-Guided Swarm Fuzzing Framework Based on Behavioral Constraints	495
<i>Ruoyu Zhou, Zhiwei Zhang, Haocheng Han, Xiaodong Zhang, Zehan Chen, Jun Sun, Yulong Shen, Dehai Xu</i>	
Provable Fairness Repair for Deep Neural Networks	508
<i>Jianan Ma, Jingyi Wang, Qi Xuan, Zhen Wang</i>	
GlassWing: A Tailored Static Analysis Approach for Flutter Android Apps	521
<i>Xiangyu Zhang, Yucheng Su, Lingling Fan, Miaoying Cai, Sen Chen</i>	
RustAssure: Differential Symbolic Testing for LLM-Transpiled C-to-Rust Code	534
<i>Yubo Bai, Tapti Palit</i>	
DNAFuzz: Descriptor-Aware Fuzzing for USB Drivers	547
<i>Zhengshu Wang, Peng He, Fuchen Ma, Yuanliang Chen, Shuoshuo Duan, Yiyuan Bai, Yu Jiang</i>	
Security Debt in LLM Agent Applications: A Measurement Study of Vulnerabilities and Mitigation Trade-offs	559
<i>Zhuoxiang Shen, Jiarun Dai, Yuan Zhang, Min Yang</i>	
SSR: Safeguarding Staking Rewards by Defining and Detecting Logical Defects in DeFi Staking	571
<i>Zewei Lin, Jiachi Chen, Jingwen Zhang, Zexu Wang, Yuming Feng, Weizhe Zhang, Zibin Zheng</i>	
EditFusion: Resolving Code Merge Conflicts via Edit Selection	584
<i>Changxin Wang, Lei Xu, Rundong Wang, Yiming Ma, Weifeng Zhang</i>	
Automatic Fixing of Missing Dependency Errors	597
<i>Jun Lyu, He Zhang, Lanxin Yang, Yue Li, Chenxing Zhong, Manuel Rigger</i>	
RustRepoTrans: Repository-level Context Code Translation Benchmark Targeting Rust	610
<i>Guangsheng Ou, Mingwei Liu, Yuxuan Chen, Yanlin Wang, Xin Peng, Zibin Zheng</i>	
Enhancing LLM's Ability to Generate More Repository-Aware Unit Tests Through Precise Context Injection	623
<i>Xin Yin, Chao Ni, Xinrui Li, Liushan Chen, Guojun Ma, Xiaohu Yang</i>	
Finding Bugs in MLIR Compiler Infrastructure via Lowering Space Exploration	636
<i>Jingjing Liang, Shan Huang, Ting Su</i>	

BinStruct: Binary Structure Recovery Combining Static Analysis and Semantics	648
<i>Yiran Zhang, Zhengzi Xu, Zhe Lang, Chengyue Liu, Yuqiang Sun, Wenbo Guo, Chengwei Liu, Weisong Sun Yang Liu</i>	
United We Stand: Towards End-to-End Log-based Fault Diagnosis via Interactive Multi-Task Learning	661
<i>Minghua He, Chiming Duan, Pei Xiao, Tong Jia, Siyu Yu, Lingzhe Zhang, Weijie Hong, Jin Han, Yifan Wu Ying Li, Gang Huang</i>	
Triangle: Empowering Incident Triage with Multi-Agent	674
<i>Zhaoyang Yu, Aoyang Fang, Minghua Ma, Jaskaran Singh Walia, Chaoyun Zhang, Shu Chi, Ze Li Murali Chintalapati, Xuchao Zhang, Rujia Wang, Chetan Bansal, Saravan Rajmohan, Qingwei Lin Shenglin Zhang, Dan Pei, Pinjia He</i>	
Navigating the Labyrinth: Path-Sensitive Unit Test Generation with Large Language Models	687
<i>Dianshu Liao, Xin Yin, Shidong Pan, Chao Ni, Zhenchang Xing, Xiaoyu Sun</i>	
LogAction: Consistent Cross-system Anomaly Detection through Logs via Active Domain Adaptation	700
<i>Chiming Duan, Minghua He, Pei Xiao, Tong Jia, Xin Zhang, Zhewei Zhong, Xiang Luo, Yan Niu Lingzhe Zhang, Siyu Yu, Yifan Wu, Weijie Hong, Ying Li, Gang Huang</i>	
When Control Flows Deviate: Directed Grey-box Fuzzing with Probabilistic Reachability Analysis	713
<i>Peihong Lin, Pengfei Wang, Xu Zhou, Wei Xie, Xin Ren, Kai Lu</i>	
Improving LLM-based Log Parsing by Learning from Errors in Reasoning Traces	726
<i>Jialai Wang, Juncheng Lu, Jie Yang, Junjie Wang, Zeyu Gao, Chao Zhang, Zhenkai Liang, Ee-Chien Chang</i>	
Watson: A Cognitive Observability Framework for the Reasoning of LLM-Powered Agents	739
<i>Benjamin Rombaut, Sogol Masoumzadeh, Kirill Vasilevski, Dayi Lin, Ahmed E. Hassan</i>	
Speculative Automated Refactoring of Imperative Deep Learning Programs to Graph Execution	752
<i>Raffi Khatchadourian, Tatiana Castro Vélez, Mehdi Bagherzadeh, Nan Jia, Anita Raja</i>	
DualFuzz: Detecting Vulnerability in Wi-Fi NICs through Dual-Directional Fuzzing	765
<i>Yuanliang Chen, Fuchen Ma, Yanyang Zhao, Yuanyi Li, Yu Jiang</i>	
DrainCode: Stealthy Energy Consumption Attacks on Retrieval-Augmented Code Generation via Context Poisoning	778
<i>Yanlin Wang, Jiadong Wu, Tianyue Jiang, Mingwei Liu, Jiachi Chen, Chong Wang, Ensheng Shi, Xilin Liu Yuchi Ma, Zibin Zheng</i>	
Hit The Bullseye On The First Shot: Improving LLMs Using Multi-Sample Self-Reward Feedback for Vulnerability Repair	791
<i>Rui Jiao, Yue Zhang, Jinku Li, Jianfeng Ma</i>	
Let the Code Speak: Incorporating Program Dynamic State for Better Method-Level Fault Localization	804
<i>Yihao Qin, Shangwen Wang, Bo Lin, Xin Peng, Sheng Ouyang, Liqian Chen, Xiaoguang Mao</i>	
Execution-Aware Program Reduction for WebAssembly via Record and Replay	816
<i>Doehyun Baek, Daniel Lehmann, Ben L. Titzer, Sukyoung Ryu, Michael Pradel</i>	
Repairing Leaks in Resource Wrappers	828
<i>Sanjay Malakar, Michael D. Ernst, Martin Kellogg, Manu Sridharan</i>	
PseudoFix: Refactoring Distorted Structures in Decompiled C Pseudocode	841
<i>Gangyang Li, Xiuwei Shang, Shaoyin Cheng, Junqi Zhang, Li Hu, Xu Zhu, Weiming Zhang, Nenghai Yu</i>	

DLBench: A Comprehensive Benchmark for SQL Translation with Large Language Models	854
<i>Li Lin, Hongqiao Chen, Qinglin Zhu, Liehang Chen, Linlong Tang, Rongxin Wu</i>	
An Empirical Study of Python Library Migration Using Large Language Models	867
<i>Mohayeminul Islam, Ajay Kumar Jha, May Mahmoud, Ildar Akhmetov, Sarah Nadi</i>	
AppBDS: LLM-Powered Description Synthesis for Sensitive Behaviors in Mobile Apps	880
<i>Zichen Liu, Xusheng Xiao</i>	
LineBreaker: Finding Token-Inconsistency Bugs with Large Language Models	893
<i>Hongbo Chen, Yifan Zhang, Xing Han, Tianhao Mao, Huanyao Rong, Yuheng Zhang, XiaoFeng Wang</i> <i>Luyi Xing, Xun Chen, Hang Zhang</i>	
Profile Coverage: Using Android Compilation Profiles to Evaluate Dynamic Testing	906
<i>Jakob Bleier, Felix Kehr, Jürgen Cito, Martina Lindorfer</i>	
Lares: LLM-driven Code Slice Semantic Search for Patch Presence Testing	919
<i>Siyuan Li, Yaowen Zheng, Hong Li, Jingdong Guo, Chaopeng Dong, Chunpeng Yan, Weijie Wang</i> <i>Yimo Ren, Limin Sun, Hongsong Zhu</i>	
Exploring Static Taint Analysis in LLMs: A Dynamic Benchmarking Framework for Measurement and Enhancement	932
<i>Haoran Zhao, Lei Zhang, Keke Lian, Fute Sun, Bofei Chen, Yongheng Liu, Zhiyu Wu, Yuan Zhang</i> <i>Min Yang</i>	
Demystifying OpenZeppelin’s Own Vulnerabilities and Analyzing Their Propagation in Smart Contracts	945
<i>Han Liu, Daoyuan Wu, Yuqiang Sun, Shuai Wang, Yang Liu, Yixiang Chen</i>	
iKnow: an Intent-Guided Chatbot for Cloud Operations with Retrieval-Augmented Generation	958
<i>Junjie Huang, Yuedong Zhong, Guangba Yu, Zhihan Jiang, Minzhi Yan, Wenfei Luan, Tianyu Yang, Rui Ren</i> <i>Michael R. Lyu</i>	
AlignCoder: Aligning Retrieval with Target Intent for Repository-Level Code Completion	971
<i>Tianyue Jiang, Yanli Wang, Yanlin Wang, Daya Guo, Ensheng Shi, Yuchi Ma, Jiachi Chen, Zibin Zheng</i>	
Coverage-Based Harmfulness Testing for LLM Code Transformation	983
<i>Honghao Tan, Haibo Wang, Diany Pressato, Yisen Xu, Shin Hwei Tan</i>	
Not Every Patch is an Island: LLM-Enhanced Identification of Multiple Vulnerability Patches	996
<i>Yi Song, Dongchen Xie, Lin Xu, He Zhang, Chunying Zhou, Xiaoyuan Xie</i>	
Automated Inline Comment Smell Detection and Repair with Large Language Models	1008
<i>Hatice Kübra Çağlar, Semih Çağlar, Eray Tüzün</i>	
Programmers’ Visual Attention on Function Call Graphs During Code Summarization	1020
<i>Samantha McLoughlin, Zachary Karas, Robert Wallace, Aakash Bansal, Collin McMillan, Yu Huang</i>	
An LLM-based multi-agent framework for agile effort estimation	1032
<i>Thanh-Long Bui, Hoa Khanh Dam, Rashina Hoda</i>	
PromFuzz: Leveraging LLM-Driven and Bug-Oriented Composite Analysis for Detecting Functional Bugs in Smart Contracts	1044
<i>Xingshuang Lin, Qinge Xie, Binbin Zhao, Yuan Tian, Saman Zonouz, Na Ruan, Jiliang Li, Raheem Beyah</i> <i>Shouling Ji</i>	

LOSVER: Line-Level Modifiability Signal-Guided Vulnerability Detection and Classification	1057
<i>Doha Nam, Jongmoon Baik</i>	
A Large Scale Study of AI-based Binary Function Similarity Detection Techniques for Security Researchers and Practitioners	1070
<i>Jingyi Shi, Yufeng Chen, Yang Xiao, Yuekang Li, Zhengzi Xu, Sihao Qiu, Chi Zhang, Keyu Qi, Yeting Li</i> <i>Xingchu Chen, Yanyan Zou, Yang Liu, Wei Huo</i>	
DebCovDiff: Differential Testing of Coverage Measurement Tools on Real-World Projects	1083
<i>Wentao Zhang, Jinghao Jia, Er kai Yu, Darko Marinov, Tianyin Xu</i>	
ZendDiff: Differential Testing of PHP Interpreter	1095
<i>Yuancheng Jiang, Jianing Wang, Qiang Liu, Yeqi Fu, Jian Mao, Roland H. C. Yap, Zhenkai Liang</i>	
BCFuzz: Bytecode-Driven Fuzzing for JavaScript Engines	1107
<i>Jiming Wang, Chenggang Wu, Jikai Ren, Yuhao Hu, Yan Kang, Xiaojie Wei, Yuanming Lai, Mengyao Xie</i> <i>Zhe Wang</i>	
CoorLog: Efficient-Generalizable Log Anomaly Detection via Adaptive Coordinator in Software Evolution	1119
<i>Pei Xiao, Chiming Duan, Minghua He, Tong Jia, Yifan Wu, Jing Xu, Gege Gao, Lingzhe Zhang</i> <i>Weijie Hong, Ying Li, Gang Huang</i>	
GUIFuzz++: Unleashing Grey-box Fuzzing on Desktop Graphical User Interfacing Applications	1132
<i>Dillon Otto, Tanner Rowlett, Stefan Nagy</i>	
Hypergraph Neural Network-based Multi-Granular Root Cause Localization for Microservice Systems	1143
<i>Yaxiao Li, Lu Wang, Chenxi Zhang, Qingshan Li, Siming Rong, Baiyang Wen, Xuyang Li, Kun Ma</i> <i>Quanwei Du, KeYang Li, Lingfeng Pan, Xinyue Li, Mingxuan Hui</i>	
Seeing is Fixing: Cross-Modal Reasoning with Multimodal LLMs for Visual Software Issue Repair	1156
<i>Kai Huang, Jian Zhang, Xiaofei Xie, Chunyang Chen</i>	
When Autonomous Vehicle Meets V2X Cooperative Perception: How Far Are We?	1169
<i>An Guo, Shuoxiao Zhang, Enyi Tang, Xinyu Gao, Haomin Pang, Haoxiang Tian, Yanzhou Mu, Wu Wen</i> <i>Chunrong Fang, Zhenyu Chen</i>	
Backdoors in Code Summarizers: How Bad Is It?	1182
<i>Chenyu Wang, Zhou Yang, Yaniv Harel, David Lo</i>	
AdaptEval: A Benchmark for Evaluating Large Language Models on Code Snippet Adaptation	1195
<i>Tanghaoran Zhang, Xinjun Mao, Shangwen Wang, Yuxin Zhao, Yao Lu, Jin Zhang, Zhang Zhang</i> <i>Kang Yang, Yue Yu</i>	
Bridging Natural Language and Formal Specification—Automated Translation of Software Requirements to LTL via Hierarchical Semantics Decomposition Using LLMs	1208
<i>Zhi Ma, Cheng Wen, Zhexin Su, Xiao Liang, Cong Tian, Shengchao Qin, Mengfei Yang</i>	
RFCAudit: AI Agent for Auditing Protocol Implementations Against RFC Specifications	1221
<i>Mingwei Zheng, Chengpeng Wang, Xuwei Liu, Jinyao Guo, Shiwei Feng, Xiangyu Zhang</i>	
RFCScope: Detecting Logical Ambiguities in Internet Protocol Specifications	1234
<i>Mrigank Pawagi, Lize Shao, Hyeonmin Lee, Yixin Sun, Wenxi Wang</i>	

Protecting Source Code Privacy When Hunting Memory Bugs	1247
<i>Jielun Wu, Bing Shui, Hongcheng Fan, Shengxin Wu, Rongxin Wu, Yang Feng, Baowen Xu, Qingkai Shi</i>	
What's DAT Smell? Untangling and Weaving the Disjoint Assertion Tangle Test Smell	1260
<i>Monil Narang, Hang Du, James A. Jones</i>	
SMTgazer: Learning to Schedule SMT Algorithms via Bayesian Optimization	1273
<i>Chuan Luo, Shaoke Cui, Jianping Song, Xindi Zhang, Wei Wu, Chanjuan Liu, Shaowei Cai, Chunming Hu</i>	
Agentic Specification Generator for Move Programs	1286
<i>Yu-Fu Fu, Meng Xu, Taesoo Kim</i>	
AMPLE: Fine-grained File Access Policies for Server Applications	1299
<i>Seyedhamed Ghavamnia, Julien Vanegue</i>	
Detecting Semantic Clones of Unseen Functionality	1312
<i>Konstantinos Kitsios, Francesco Sovrano, Earl T. Barr, Alberto Bacchelli</i>	
Efficient and Verifiable Proof Logging for MaxSAT Solving	1325
<i>Raoul Van Doren, Timos Antonopoulos, Ruzica Piskac</i>	
FGit: Fault-Guided Fine-Tuning for Code Generation	1338
<i>Lishui Fan, Zhongxin Liu, Haoye Wang, Lingfeng Bao, Xin Xia, Shanping Li</i>	
Destabilizing Neurons to Generate Challenging Neural Network Verification Benchmarks	1351
<i>Linhan Li, ThanhVu Nguyen</i>	
SATORI: Static Test Oracle Generation for REST APIs	1364
<i>Juan C. Alonso, Alberto Martin-Lopez, Sergio Segura, Gabriele Bavota, Antonio Ruiz-Cortés</i>	
Learning Project-wise Subsequent Code Edits via Interleaving Neural-based Induction and Tool-based Deduction	1377
<i>Chenyan Liu, Yun Lin, Yuhuan Huang, Jiaxin Chang, Binhang Qi, Bo Jiang, Zhiyong Huang, Jin Song Dong</i>	
NATE: A Network-Aware Testing Enhancer for Network-Related Fault Detection in Android Apps	1390
<i>Yuanhong Lan, Shaoheng Cao, Yifei Lu, Minxue Pan, Xuandong Li</i>	
WEST: Specification-Based Test Generation for WebAssembly	1403
<i>Dongjun Youn, Wonho Shin, Sukyoung Ryu</i>	
VeriExploit: Automatic Bug Reproduction in Smart Contracts via LLMs and Formal Methods	1415
<i>Chenfeng Wei, Shiyu Cai, Yiannis Charalambous, Tong Wu, Sangharatna Godbole, Lucas Cordeiro</i>	
Rechecking Recheck Requests in Continuous Integration: An Empirical Study of OpenStack	1427
<i>Yelizaveta Brus, Rungroj Maipradit, Earl T. Barr, Shane McIntosh</i>	
Hierarchical Knowledge Injection for Improving LLM-based Program Repair	1440
<i>Ramtin Ehsani, Esteban Parra, Sonia Haiduc, Preetha Chatterjee</i>	
VERT: Polyglot Verified Equivalent Rust Transpilation with Large Language Models	1453
<i>Aidan Z.H. Yang, Yoshiki Takashima, Brandon Paulsen, Josiah Dodds, Daniel Kroening</i>	
ScaleCirc: Scaling the Analysis over Circom Circuits	1464
<i>Jinan Jiang, Haoran Qin, Xiapu Luo</i>	

Aligning LLMs to Fully Utilize the Cross-file Context in Repository-level Code Completion	1477
<i>Jia Li, Hao Zhu, Huanyu Liu, Xianjie Shi, He Zong, Yihong Dong, Kechi Zhang, Siyuan Jiang, Zhi Jin Ge Li</i>	
CoTune: Co-evolutionary Configuration Tuning	1490
<i>Gangda Xiong, Tao Chen</i>	
Sifting Truth from Coincidences: A Two-Stage Positive and Unlabeled Learning Model for Coincidental Correctness Detection	1503
<i>Chunyan Liu, Huan Xie, Yan Lei, Zhenyu Wu, Jinping Wang</i>	
Improving NLSAT for Nonlinear Real Arithmetic	1516
<i>Zhonghan Wang</i>	
Finding Insecure State Dependency in DApps via Multi-Source Tracing and Semantic Enrichment	1529
<i>Jingwen Zhang, Yuhong Nan, Wei Li, Kaiwen Ning, Zewei Lin, Zitong Yao, Yuming Feng, Weizhe Zhang Zibin Zheng</i>	
Forcrat: Automatic I/O API Translation from C to Rust via Origin and Capability Analysis	1541
<i>Jaemin Hong, Sukyoung Ryu</i>	
It's Not Easy Being Green: On the Energy Efficiency of Programming Languages	1553
<i>Nicolas van Kempen, Hyuk-Je Kwon, Dung Tuan Nguyen, Emery D. Berger</i>	
Belief Propagation with Local Structure and Its Applications in Program Analysis	1566
<i>Yiqian Wu, Yifan Chen, Yingfei Xiong, Xin Zhang</i>	
Cryptbara: Dependency-Guided Detection of Python Cryptographic API Misuses	1578
<i>Seogyong Cho, Seungeun Yu, Seunghoon Woo</i>	
IMUFuzzer: Resilience-based Discovery of Signal Injection Attacks on Robotic Aerial Vehicles	1591
<i>Sudharssan Mohan, Kyeongseok Yang, Zelun Kong, Yonghwi Kwon, Junghwan Rhee, Tyler Summers Hongjun Choi, Heejo Lee, Chung Hwan Kim</i>	
Beyond Static GUI Agent: Evolving LLM-based GUI Testing via Dynamic Memory	1603
<i>Mengzhuo Chen, Zhe Liu, Chunyang Chen, Junjie Wang, Yangguang Xue, Boyu Wu, Yuekai Huang Libin Wu, Qing Wang</i>	
Terminator: Enabling Efficient Fuzzing of Closed-Source GUI Programs by Automatic Coverage-Guided Termination	1616
<i>Jonas Zabel, Philip Kolvenbach, Steven Arzt</i>	
Characterizing Multi-Hunk Patches: Divergence, Proximity, and LLM Repair Challenges	1629
<i>Noor Nashid, Daniel Ding, Keheliya Gallaba, Ahmed E. Hassan, Ali Mesbah</i>	
Fact-Aligned and Template-Constrained Static Analyzer Rule Enhancement with LLMs	1642
<i>Zongze Jiang, Ming Wen, Ge Wen, Hai Jin</i>	
When Faster Isn't Greener: The Hidden Costs of LLM-Based Code Optimization	1655
<i>Tristan Coignon, Clément Quinton, Romain Rouvoy</i>	
Leveraging Mixture-of-Experts Framework for Smart Contract Vulnerability Repair with Large Language Model	1667
<i>Hang Yuan, Xizhi Hou, Lei Yu, Li Yang, Jiayue Tang, Jiadong Xu, Yifei Liu, Fengjun Zhang, Chun Zuo</i>	

Uncovering Discrimination Clusters: Quantifying and Explaining Systematic Fairness Violations	1680
<i>Ranit D. Akash, Ashish Kumar, Verrya Monjezi, Ashutosh Trivedi, Gang Tan, Saeid Tizpaz-Niari</i>	
Demystifying Cookie Sharing Risks in WebView-based Mobile App-in-app Ecosystems	1693
<i>Miao Zhang, Shenao Wang, Guilin Zheng, Yanjie Zhao, Haoyu Wang</i>	
ADPerf: Investigating and Testing Performance in Autonomous Driving Systems	1705
<i>Tri Minh-Triet Pham, Diego Elias Costa, Weiyi Shang, Jinqiu Yang</i>	
Fixing Broken Graphs: LLM-Powered Automatic Code Optimization for DNN Programs	1718
<i>Haotian Wang, Yicheng Sui, Yudong Xie, Yicong Liu, Yufei Sun, Changqing Shi, Yuzhi Zhang</i>	
Function Clustering-Based Fuzzing Termination: Toward Smarter Early Stopping	1731
<i>Liang Ding, Wenzhang Yang, Yinxing Xue</i>	
Cross2OH: Enabling Seamless Porting of C/C++ Software Libraries to OpenHarmony	1744
<i>Qian Zhang, Tsz-On Li, Ying Wang, Li Li, Shing-Chi Cheung</i>	
HybridSIMD: A Super C++ SIMD Library with Integrated Auto-tuning Capabilities	1756
<i>Haolin Pan, Xulin Zhou, Mingjie Xing, Yanjun Wu</i>	
Demystifying Cross-Language C/C++ Binaries: A Robust Software Component Analysis Approach	1768
<i>Meiqiu Xu, Ying Wang, Wei Tang, Xian Zhan, Shing-Chi Cheung, Hai Yu, Zhiliang Zhu</i>	
Detecting Various DeFi Price Manipulations with LLM Reasoning	1781
<i>Juantao Zhong, Daoyuan Wu, Ye Liu, Maoyi Xie, Yang Liu, Yi Li, Ning Liu</i>	
Multi-dimensional Assessment of Crowdsourced Testing Reports via LLMs	1794
<i>Yue Wang, Yuan Zhao, Shengcheng Yu, Zhenyu Chen</i>	
ARG: Testing Query Rewriters via Abstract Rule Guided Fuzzing	1807
<i>Dawei Li, Yuxiao Guo, Qifan Liu, Jie Liang, Zhiyong Wu, Jingzhou Fu, Chi Zhang, Yu Jiang</i>	
On the (In)Security of Non-resettable Device Identifiers in Custom Android Systems	1819
<i>Zikan Dong, Liu Wang, Guoai Xu, Haoyu Wang</i>	
PEACE: Towards Efficient Project-Level Efficiency Optimization via Hybrid Code Editing	1831
<i>Xiaoxue Ren, Jun Wan, Yun Peng, Zhongxin Liu, Ming Liang, Dajun Chen, Wei Jiang, Yong Li</i>	
MCTS-Refined CoT: High-Quality Fine-Tuning Data for LLM-Based Repository Issue Resolution	1844
<i>Yibo Wang, Zhihao Peng, Ying Wang, Zhao Wei, Hai Yu, Zhiliang Zhu</i>	
Can Mamba Be Better? An Experimental Evaluation of Mamba in Code Intelligence	1856
<i>Shuo Liu, Jacky Keung, Zhen Yang, Zhenyu Mao, Yicheng Sun</i>	
ORFuzz: Fuzzing the "Other Side" of LLM Safety - Testing Over-Refusal	1869
<i>Haonan Zhang, Dongxia Wang, Yi Liu, Kexin Chen, Jiashui Wang, Xinlei Ying, Long Liu, Wenhai Wang</i>	
Coding-Fuse: Efficient Fusion of Code Pre-Trained Models for Classification Tasks	1881
<i>Yu Zhao, Lina Gong, Zhiqiu Huang, Yuchen Jin, Mingqiang Wei</i>	
Diplomatist: What Do Cross-language Dependencies Reflect Software Ecosystem Health?	1894
<i>Fanyi Meng, Ying Wang, Chun Yong Chong, Hai Yu, Zhiliang Zhu</i>	
Automated Detection of Web Application Navigation Barriers for Screen Reader Users	1906
<i>Shubhi Jain, Syed Fatiul Huq, Ziyao He, Sam Malek</i>	

SemGuard: Real-Time Semantic Evaluator for Correcting LLM-Generated Code	1919
<i>Qinglin Wang, Zhihong Sun, Ruyun Wang, Tao Huang, Zhi Jin, Ge Li, Chen Lyu</i>	
Defects4Log: Benchmarking LLMs for Logging Code Defect Detection and Reasoning	1931
<i>Xin Wang, Zhenhao Li, Zishuo Ding</i>	
Efficient Understanding of Machine Learning Model Mispredictions	1943
<i>Martin Eberlein, Jürgen Cito, Lars Grunske</i>	
VRTestSniffer: Test Smell Detector for Virtual Reality (VR) Software Projects	1956
<i>Faraz Gurramkonda, Avishak Chakroborty, Bruce Maxim, Mohamed Wiem Mkaouer, Foyzul Hassan</i>	
Non-termination Witnesses and Their Validation	1969
<i>Zsófia Ádám, Paulína Ayaziová, Levente Bajczi, Dirk Beyer, Marek Jankola, Marian Lingsch-Rosenfeld Jan Strejček</i>	
Automated Generation of Issue-Reproducing Tests by Combining LLMs and Search-Based Testing	1982
<i>Konstantinos Kitsios, Marco Castelluccio, Alberto Bacchelli</i>	
Have We Solved Access Control Vulnerability Detection in Smart Contracts? A Benchmark Study	1995
<i>Han Liu, Daoyuan Wu, Yuqiang Sun, Shuai Wang, Yang Liu</i>	
Explainable Fault Localization for Programming Assignments via LLM-Guided Annotation	2008
<i>Fang Liu, Tianze Wang, Li Zhang, Zheyu Yang, Jing Jiang, Zian Sun</i>	
Interpretable Vulnerability Detection Reports	2021
<i>Claudia Mamede, José Campos, Claire Le Goues, Rui Abreu</i>	
Characterizing and Repairing Color-Related Accessibility Issues in Android Apps	2033
<i>Jiahao Gu, Huaxun Huang</i>	
Demystifying the Evolution of Neural Networks with BOM Analysis: Insights from a Large-Scale Study of 55,997 GitHub Repositories	2046
<i>Xiaoning Ren, Yuhang Ye, Xiongfei Wu, Yueming Wu, Yinxing Xue</i>	
FirmProj: Detecting Firmware Leakage in IoT Update Processes via Companion App Analysis	2058
<i>Wenzhi Li, Jialong Guo, Jiongyi Chen, Fan Li, Yujie Xing, Yanbo Xu, Shishuai Yang, Wenrui Diao</i>	
Why Is My Transaction Risky? Understanding Smart Contract Semantics and Interactions in the NFT Ecosystem	2071
<i>Yujing Chen, Xuanming Liu, Zhiyuan Wan, Zuobin Wang, David Lo, Difan Xie, Xiaohu Yang</i>	
Breaking the Traffic Barrier: Unveiling Multi-Format of Protocols via Autonomous Program Exploration	2083
<i>Dingzhao Xue, Yibo Qu, Bowen Jiang, Xin Chen, Shuaizong Si, Shichao Lv, Zhiqiang Shi, Limin Sun</i>	
Spec2Code: Mapping Protocol Specification to Function-Level Code Implementation	2096
<i>Yuekun Wang, Lili Quan, Xiaofei Xie, Junjie Wang, Jianjun Chen</i>	
Using Fourier Analysis and Mutant Clustering to Accelerate DNN Mutation Testing	2109
<i>Ali Ghanbari, Sasan Tavakkol</i>	
PAT-Agent: Autoformalization for Model Checking	2122
<i>Xinyue Zuo, Yifan Zhang, Hongshu Wang, Yufan Cai, Zhe Hou, Jing Sun, Jin Song Dong</i>	

Towards Generalizable Instruction Vulnerability Prediction via LLM-Enhanced Code Representation	2134
<i>Bao Wen, Jingjing Gu, Jingxuan Zhang, Yang Liu, Pengfei Yu, Yanchao Zhao</i>	
Don't Mess with Bro's Cheese! An Empirical Study of Resource Conflict in Android Multi-window	2146
<i>Chenkai Guo, Huimin Zhao, Tianhong Wang, Naipeng Dong, Qingqing Dong, Jiarui Che, Yaqiong Qiao</i> <i>Xiangyang Luo, Zheli Liu</i>	
FlakyGuard: Automatically Fixing Flaky Tests at Industry Scale	2158
<i>Chengpeng Li, Farnaz Behrang, August Shi, Peng Liu</i>	
Spinner: Detecting Locking Violations in the eBPF Runtime	2171
<i>Priya Govindasamy, Joseph Bursey, Hsin-Wei Hung, Ardalan Amiri Sani</i>	
Altered Histories in Version Control System Repositories: Evidence from the Trenches	2184
<i>Solal Rapaport, Laurent Pautet, Samuel Tardieu, Stefano Zacchiroli</i>	
Generating Failure-Based Oracles to Support Testing of Reported Bugs in Android Apps	2196
<i>Jack Johnson, Junayed Mahmud, Oscar Chaparro, Kevin Moran, Mattia Fazzini</i>	
LspFuzz: Hunting Bugs in Language Servers	2209
<i>Hengcheng Zhu, Songqiang Chen, Valerio Terragni, Lili Wei, Yepang Liu, Jiarong Wu, Shing-Chi Cheung</i>	
LLM-Powered Multi-Agent Collaboration for Intelligent Industrial On-Call Automation	2222
<i>Ruowei Fu, Yang Zhang, Zeyu Che, Xin Wu, Zhenyu Zhong, Zhiqiang Ren, Shenglin Zhang, Feng Wang</i> <i>Yongqian Sun, Xiaozhou Liu, Kexin Liu, Yu Zhang</i>	
Enhancing LLM to Decompile Optimized PTX to Readable CUDA for Tensor Programs	2235
<i>Xinyu Sun, Fugen Tang, Yu Zhang, Han Shen, Chengru Song, Di Zhang</i>	
Mixture-of-Experts Low-Rank Adaptation for Multilingual Code Summarization	2248
<i>Tianchen Yu, Li Yuan, Hailing Huang, Jiexin Wang, Yi Cai</i>	
Automated Repair of OpenID Connect Programs	2261
<i>Tamjid Al Rahat, Yanju Chen, Yu Feng, Yuan Tian</i>	
Latra: A Template-Based Language-Agnostic Transformation Framework for Effective Program Reduction	2274
<i>Zhenyang Xu, Yiran Wang, Yongqiang Tian, Mengxiao Zhang, Chengnian Sun</i>	
Faster Runtime Verification during Testing via Feedback-Guided Selective Monitoring	2286
<i>Shinhae Kim, Saikat Dutta, Owolabi Legunsen</i>	
FastCoder: Accelerating Repository-level Code Generation via Efficient Retrieval and Verification	2299
<i>Qianhui Zhao, Li Zhang, Fang Liu, Xiaoli Lian, Qiaoyuanhe Meng, Ziqian Jiao, Zetong Zhou, Jia Li</i> <i>Lin Shi</i>	
Effective Code Membership Inference for Code Completion Models via Adversarial Prompts	2312
<i>Yuan Jiang, Zehao Li, Shan Huang, Christoph Treude, Xiaohong Su, Tiantian Wang</i>	
SPICE: An Automated SWE-Bench Labeling Pipeline for Issue Clarity, Test Coverage, and Effort Estimation	2325
<i>Gustavo A. Oliva, Gopi Krishnan Rajbahadur, Aaditya Bhatia, Haoxiang Zhang, Yihao Chen, Zhilong Chen</i> <i>Arthur Leung, Dayi Lin, Boyuan Chen, Ahmed E. Hassan</i>	

On the Correctness of Software Merge	2338
<i>Akira Mori, Masatomo Hashimoto</i>	
Incremental Program Analysis in the Wild: An Empirical Study on Real-World Program Changes	2350
<i>Xizao Wang, Xiangrong Bin, Lanxin Huang, Shangqing Liu, Jianhua Zhao, Lei Bu</i>	
Polyglot: An Extensible Framework to Benchmark Code Translation with LLMs	2363
<i>Marco Vieira, Priyam Ashish Shah, Bhavain Shah, Rrezarta Krasniqi</i>	
Interleaved Learning and Exploration: A Self-Adaptive Fuzz Testing Framework for MLIR	2376
<i>Zeyu Sun, Jingjing Liang, Weiyi Wang, Chenyao Suo, Junjie Chen, Fanjiang Xu</i>	
FailMapper: Automated Generation of Unit Tests Guided by Failure Scenarios	2388
<i>Ruiqi Dong, Zehang Deng, Xiaogang Zhu, Xiaoning Du, Huai Liu, Shaohua Wang, Sheng Wen, Yang Xiang</i>	
LLMs for Automated Unit Test Generation and Assessment in Java: The AgoneTest Framework	2401
<i>Andrea Lops, Fedelucio Narducci, Azzurra Ragone, Michelantonio Trizio, Claudio Bartolini</i>	
DeepExploitor: LLM-Enhanced Automated Exploitation of DeepLink Attack in Hybrid Apps	2414
<i>Zhangyue Zhang, Lei Zhang, Zhibo Zhang, Yongheng Liu, Zhemin Yang, Yuan Zhang, Min Yang</i>	
An Agent-based Evaluation Framework for Complex Code Generation	2427
<i>Xinchen Wang, Ruida Hu, Pengfei Gao, Chao Peng, Cuiyun Gao</i>	
Token Sugar: Making Source Code Sweeter for LLMs through Token-Efficient Shorthand	2440
<i>Zhensu Sun, Chengran Yang, Xiaoning Du, Zhou Yang, Li Li, David Lo</i>	
Advancing Automated Ethical Profiling in SE: a Zero-Shot Evaluation of LLM Reasoning	2452
<i>Patrizio Miglierini, Mashal Afzal Memon, Marco Autili, Paola Inverardi</i>	
Exact Inference for Quantum Circuits: A Testing Oracle for Quantum Software Stacks	2465
<i>Kanguk Lee, Jaemin Hong, Sukyoung Ryu</i>	
Evaluating and Improving Framework-based Parallel Code Completion with Large Language Models	2478
<i>Ke Liu, Qinglin Wang, Xiang Chen, Guang Yang, Yigui Feng, Gencheng Liu, Jie Liu</i>	
The Fault in our Stats	2491
<i>Alexi Turcotte, Neev Nirav Mehta</i>	
Unit Test Update through LLM-Driven Context Collection and Error-Type-Aware Refinement	2504
<i>Yuanhe Zhang, Zhiquan Yang, Shengyi Pan, Zhongxin Liu</i>	
On Automating Configuration Dependency Validation via Retrieval-Augmented Generation	2516
<i>Sebastian Simon, Alina Mailach, Johannes Dorn, Norbert Siegmund</i>	
Automated Insertion of Flushes and Fences for Persistency	2529
<i>Yutong Guo, Weiyu Luo, Brian Demsky</i>	
Soleker: Uncovering Vulnerabilities in Solana Smart Contracts	2542
<i>Kunsong Zhao, Yunpeng Tian, Zuchao Ma, Xiapu Luo</i>	
ACTaint: Agent-Based Taint Analysis for Access Control Vulnerabilities in Smart Contracts	2555
<i>Huarui Lin, Zhipeng Gao, Jiachi Chen, Xiang Chen, Xiaohu Yang, Lingfeng Bao</i>	
DRIFT: Debug-based Trace Inference for Firmware Testing	2568
<i>Changming Liu, Alejandro Mera, Meng Xu, Engin Kirda</i>	

WingMuzz: Blackbox Testing of IoT Protocols via Two-dimensional Fuzzing Schedule	2580
<i>Xiaogang Zhu, Enze Dai, Xiaotao Feng, Shaohua Wang, Xin Xia, Sheng Wen, Kwok-Yan Lam, Yang Xiang</i>	
On the Robustness Evaluation of 3D Obstacle Detection Against Specifications in Autonomous Driving	2593
<i>Tri Minh-Triet Pham, Bo Yang, Jinqiu Yang</i>	
SE-Jury: An LLM-as-Ensemble-Judge Metric for Narrowing the Gap with Human Evaluation in SE	2606
<i>Xin Zhou, Kisub Kim, Ting Zhang, Martin Weyssow, Luís F. Gomes, Guang Yang, Kui Liu, Xin Xia</i> <i>David Lo</i>	
EfficientEdit: Accelerating Code Editing via Edit-Oriented Speculative Decoding	2619
<i>Peiding Wang, Li Zhang, Fang Liu, Yinghao Zhu, Wang Xu, Lin Shi, Xiaoli Lian, Minxiao Li, Bo Shen</i> <i>An Fu</i>	
AutoFid: Adaptive and Noise-Aware Fidelity Measurement for Quantum Programs via Circuit Graph Analysis	2631
<i>Tingting Li, Ziming Zhao, Jianwei Yin</i>	
Using Active Learning to Train Predictive Mutation Testing with Minimal Data	2644
<i>Miklós Borsi</i>	
QuanBench: Benchmarking Quantum Code Generation with Large Language Models	2657
<i>Xiaoyu Guo, Minggu Wang, Jianjun Zhao</i>	
How Big is the Automaton? Certified Lower Bounds on the Size of Presburger DFAs	2670
<i>Nicolas Amat, Pierre Ganty, Alessio Mansutti</i>	
Understanding Resource Injection Vulnerabilities in Kubernetes Ecosystems	2682
<i>Defang Bo, Jie Lu, Feng Li, Jingting Chen, Jinchen Wang, Chendong Yu, Yeting Li, Wei Huo</i>	
From Characters to Structure: Rethinking Real-Time Collaborative Programming Models	2695
<i>Leon Freudenthaler, Bernhard Taufner, Karl Michael Göschka</i>	
State Field Coverage: A Metric for Oracle Quality	2707
<i>Facundo Molina, Nazareno Aguirre, Alessandra Gorla</i>	
PALM: Synergizing Program Analysis and LLMs to Enhance Rust Unit Test Coverage	2720
<i>Bei Chu, Yang Feng, Kui Liu, Hange Shi, Zifan Nan, Zhaoqiang Guo, Baowen Xu</i>	
Tephra: Principled Discovery of Fuzzer Limitations	2733
<i>Vasil Sarafov, David Markvica, Stefan Brunthaler</i>	
HFuzzer: Testing Large Language Models for Package Hallucinations via Phrase-based Fuzzing	2746
<i>Yukai Zhao, Menghan Wu, Xing Hu, Xin Xia</i>	
LLM-Based Identification of Null Pointer Exception Patches	2759
<i>Tahir Ullah, Waseem Akram, Fiza Khaliq, Hui Liu</i>	
Root Cause Analysis of RISC-V Build Failures via LLM and MCTS Reasoning	2772
<i>Weipeng Shuai, Jie Liu, Zhirou Ma, Liangyi Kang, Zehua Wang, Shuai Wang, Dan Ye, Hui Li, Wei Wang</i> <i>Jiaxin Zhu</i>	
Understanding Feature Request Practice on GitHub via a Large-Scale Empirical Study	2783
<i>Jiajun Li, Wenhua Yang, Minxue Pan, Yu Zhou</i>	

Which Is Better For Reducing Outdated and Vulnerable Dependencies: Pinning or Floatingf	2795
<i>Imranur Rahman, Jill Marley, William Enck, Laurie Williams</i>	
A Multi-Modality Evaluation of the Reality Gap in Autonomous Driving Systems	2808
<i>Stefano Carlo Lambertenghi, Mirena Flores Valdez, Andrea Stocco</i>	
PoliCond: Condition-Aware Ontology-Driven LLMs for Privacy Policy Contradiction Analysis	2821
<i>Yalin Feng, Yifei Lu, Minxue Pan</i>	
Reflective Unit Test Generation for Precise Type Error Detection with Large Language Models	2834
<i>Chen Yang, Ziqi Wang, Yanjie Jiang, Lin Yang, Yuteng Zheng, Jianyi Zhou, Junjie Chen</i>	
Understanding Software Engineering Agents: A Study of Thought-Action-Result Trajectories	2846
<i>Islem Bouzenia, Michael Pradel</i>	
SolContractEval: A Benchmark for Evaluating Contract-Level Solidity Code Generation	2858
<i>Zhifan Ye, Jiachi Chen, Zhenzhe Shao, Lingfeng Bao, Xiaohu Yang, Zhongxin Liu</i>	
LLMPort: Cross-file Patch Porting via Task Decomposition and Self-correction	2870
<i>Bofei Chen, Lei Zhang, Peng Deng, Nan Wang, Haoyu Xu, Mingda Guo, Yuan Zhang, Min Yang</i>	
DiffFix: Incrementally Fixing AST Diffs via Context and Type Information	2882
<i>Guofeng Zeng, Chang-Ai Sun, Kai Gao, Huai Liu</i>	
Better Safe than Sorry: Preventing Policy Violations through Predictive Root-Cause-Analysis for IoT Systems	2894
<i>Michael Norris, Syed Rafiul Hussain, Gang Tan</i>	
Metamorphic Testing for Audio Content Moderation Software	2906
<i>Wenxuan Wang, Yongjiang Wu, Junyuan Zhang, Shuqing Li, Yun Peng, Wenting Chen, Shuai Wang, Michael R. Lyu</i>	
Automated Combinatorial Test Generation for Alloy	2919
<i>Agustín Borda, Germán Regis, Nazareno Aguirre, Marcelo Frias, Pablo Ponzio</i>	
The Cost of Downgrading Build Systems : A Case Study of Kubernetes	2932
<i>Gareema Ranjan, Mahmoud Alfadel, Gengyi Sun, Shane McIntosh</i>	
EPSO: A Caching-Based Efficient Superoptimizer for BPF Bytecode	2945
<i>Qian Zhu, Yuxuan Liu, Ziyuan Zhu, Shangqing Liu, Lei Bu</i>	
When Does Wasm Malware Detection Fail? A Systematic Analysis of Their Robustness to Evasion	2957
<i>Taeyoung Kim, Sanghak Oh, Kiho Lee, Weihang Wang, Yonghwi Kwon, Sanghyun Hong, Hyoungshick Kim</i>	
Vulnerability-Affected Versions Identification: How Far Are We?	2970
<i>Xingchu Chen, Chengwei Liu, Jialun Cao, Yang Xiao, Xinyue Cai, Yeting Li, Jingyi Shi, Tianqi Sun, Haiming Chen, Wei Huo</i>	
LAURA: Enhancing Code Review Generation with Context-Enriched Retrieval-Augmented LLM	2983
<i>Yuxin Zhang, Yuxia Zhang, Zeyu Sun, Yanjie Jiang, Hui Liu</i>	
Your Build Scripts Stink: The State of Code Smells in Build Scripts	2996
<i>Mahzabin Tamanna, Yash Chandrani, Matthew Burrows, Brandon Wroblewski, Laurie Williams, Dominik Wermke</i>	

Mockingbird: Efficient Excessive Data Exposures Detection via Dynamic Code Instrumentation	3009
<i>Chenxiao Xia, Jiazheng Sun, Jun Zheng, Yu-an Tan, Hongyi Su</i>	
RealisticCodeBench: Towards More Realistic Evaluation of Large Language Models for Code Generation	3021
<i>Xiao Yu, Haoxuan Chen, Lei Liu, Xing Hu, Jacky Wai Keung, Xin Xia</i>	
Issue Localization via LLM-Driven Iterative Code Graph Searching	3034
<i>Zhonghao Jiang, Xiaoxue Ren, Meng Yan, Wei Jiang, Yong Li, Zhongxin Liu</i>	
Clarifying Semantics of In-Context Examples for Unit Test Generation	3046
<i>Chen Yang, Lin Yang, Ziqi Wang, Dong Wang, Jianyi Zhou, Junjie Chen</i>	
Who's to Blame? Rethinking the Brittleness of Automated Web GUI Testing from a Pragmatic Perspective	3058
<i>Haonan Zhang, Kundi Yao, Zishuo Ding, Lizhi Liao, Weiye Shang</i>	
Amur: Fixing Multi-Resource Leaks Guided by Resource Flow Analysis	3071
<i>Jinyoung Kim, Eunseok Lee</i>	
Algernon: A Flag-Guided Hybrid Fuzzer for Unlocking Hidden Program Paths	3083
<i>Peng Deng, Lei Zhang, Jingqi Long, Wenzheng Hong, Zhemin Yang, Yuan Zhang, Donglai Zhu, Min Yang</i>	
Agents in the Sandbox: End-to-End Crash Bug Reproduction for Minecraft	3095
<i>Eray Yapağcı, Yavuz Alp Sencer Öztürk, Eray Tüzün</i>	
LLM-Assisted Synthesis of High-Assurance C Programs	3108
<i>Prasita Mukherjee, Minghai Lu, Benjamin Delaware</i>	
How Can Infrastructure as Code Accelerate Data Center Bring-ups? A Case Study at ByteDance	3121
<i>Xianhao Jin, Yifei Feng, Yufei Gao, Yongning Hu, Jie Huang, Kun Xia, Luchuan Guo</i>	
JSIdentify-V2: Leveraging Dynamic Memory Fingerprinting for Mini-Game Plagiarism Detection	3133
<i>Zhihao Li, Chaozheng Wang, Zongjie Li, Xinyong Peng, Qun Xia, Haochuan Lu, Ting Xiong, Shuzheng Gao Cuiyun Gao, Shuai Wang, Yuetang Deng, Huafeng Ma</i>	
Practical Escape of Exploration Tarpits for Mini-Game Testing in an Industrial Setting	3145
<i>Yuan Cao, Dezhi Ran, Haochuan Lu, Chao Guo, Xuran Hao, Zhuoru Chen, Ting Xiong, Yuetang Deng Tao Xie</i>	
Industry Practice of LLM-Assisted Protocol Fuzzing for Commercial Communication Modules	3157
<i>Qiang Fu, Changjian Liu, Yuan Ding, Chao Fan, Yulai Fu, Yuhan Chen, Ying Fu, Ronghua Shi, Fuchen Ma Heyuan Shi</i>	
HarmoBridge: Bridging ArkTS and C/C++ for Cross-Language Static Analysis on HarmonyOS	3168
<i>Jiale Wu, Jiapeng Deng, Yanjie Zhao, Li Li, Haoyu Wang</i>	
APKARMOR: Low-Cost Lightweight Anti-Decompilation Techniques for Android Apps	3180
<i>Jiayang Liu, Yanjie Zhao, Pengcheng Xia, Haoyu Wang</i>	
KAIOPS: A Platform Solution of End-to-End Multi-Modal AIOps for AI Training at Scale	3192
<i>Zeying Wang, Junhong Liu, Penghao Zhang, Xiaoyang Sun, Xu Wang, Tianyu Wo, Chunming Hu Chengru Song, Jin Ouyang, Renyu Yang</i>	

iCodeReviewer: Improving Secure Code Review with Mixture of Prompts	3204
<i>Yun Peng, Kisub Kim, Linghan Meng, Kui Liu</i>	
Should We Evaluate LLM Based Security Analysis Approaches on Open Source Systems?	3216
<i>Kohei Dozono, Jonas Engesser, Benjamin Hummel, Tobias Roehm, Alexander Pretschner</i>	
Tron: Fuzzing Linux Network Stack via Protocol-System Call Payload Synthesis	3227
<i>Qiang Zhang, Yifei Chu, Yuheng Shen, Jianzhong Liu, Heyuan Shi, Yu Jiang, Wanli Chang</i>	
TrioXpert: An Automated Incident Management Framework for Microservice System	3239
<i>Yongqian Sun, Yu Luo, Xidao Wen, Yuan Yuan, Xiaohui Nie, Shenglin Zhang, Tong Liu, Xi Luo</i>	
Metrics Driven Reengineering and Continuous Code Improvement at Meta	3251
<i>Audris Mockus, Peter C Rigby, Rui Abreu, Anatoly Akkerman, Yogesh Bhootada, Payal Bhuptani</i> <i>Gurnit Ghardhora, Lan Hoang Dao, Chris Hawley, Renzhi He, Sagar Krishnamoorthy, Sergei Krauze</i> <i>Jianmin Li, Anton Lunov, Dragos Martac, Francois Morin, Neil Mitchell, Venus Montes, Maher Saba</i> <i>Matt Steiner, Andrea Valori, Shanchao Wang, Nachiappan Nagappan</i>	
Context-Sensitive Pointer Analysis for ArkTS	3262
<i>Yizhuo Yang, Lingyun Xu, Mingyi Zhou, Li Li</i>	
Minuku: Detecting Diverse Display Issues in Mobile Apps with Small-scale Dataset	3274
<i>Yongxiang Hu, Ke Liu, Hailiang Jin, Shiyu Guo, Juxing Yuan, Xin Wang, Yangfan Zhou</i>	
Context-Aware CodeLLM Eviction for AI-assisted Coding	3287
<i>Kishanthan Thangarajah, Boyuan Chen, Shi Chang, Ahmed E. Hassan</i>	
LogPilot: Intent-aware and Scalable Alert Diagnosis for Large-scale Online Service Systems	3299
<i>Zhihan Jiang, Jinyang Liu, Yichen Li, Haiyu Huang, Xiao He, Tieying Zhang, Jianjun Chen, Yi Li, Rui Shi</i> <i>Michael R. Lyu</i>	
Streamlining Acceptance Test Generation for Mobile Applications Through Large Language Models: An Industrial Case Study	3311
<i>Pedro Luís Fonseca, Bruno Lima, João Pascoal Faria</i>	
An Empirical Study on UI Overlap in OpenHarmony Applications	3322
<i>Farong Liu, Mingyi Zhou, Li Li</i>	
Multi-Modal Requirements Data-based Acceptance Criteria Generation using LLMs	3334
<i>Fanyu Wang, Chetan Arora, Yonghui Liu, Kaicheng Huang, Chakkrit Tantithamthavorn, Aldeida Aleti</i> <i>Dishan Sambathkumar, David Lo</i>	
Prompt-with-Me: in-IDE Structured Prompt Management for LLM-Driven Software Engineering	3346
<i>Ziyou Li, Agnia Sergeyuk, Maliheh Izadi</i>	
Bridging Research and Practice in Simulation-based Testing of Industrial Robot Navigation Systems	3357
<i>Sajad Khatiri, Francisco Eli Viña Barrientos, Maximilian Wulf, Paolo Tonella, Sebastiano Panichella</i>	
A Characterization Study of Bugs in LLM Agent Workflow Orchestration Frameworks	3369
<i>Ziluo Xue, Yanjie Zhao, Shenao Wang, Kai Chen, Haoyu Wang</i>	
AdaptiveGuard: Towards Adaptive Runtime Safety for LLM-Powered Software	3381
<i>Rui Yang, Michael Fu, Chakkrit Tantithamthavorn, Chetan Arora, Gunel Gulmammadova, Joey Chua</i>	

Unlocking Reproducibility: Automating re-Build Process for Open-Source Software	3392
<i>Behnaz Hassanshahi, Trong Nhan Mai, Benjamin Selwyn Smith, Nicholas Allen</i>	
Out of Distribution Detection in Self-adaptive Robots with AI-powered Digital Twins	3403
<i>Erbilin Isaku, Hassan Sartaj, Shaikat Ali, Beatriz Sanguino, Tongtong Wang, Guoyuan Li, Houxiang Zhang Thomas Peyrucain</i>	
Data Dependency-Aware Code Generation from Enhanced UML Sequence Diagrams	3415
<i>Wenxin Mao, Zhitao Wang, Long Wang, Sirong Chen, Cuiyun Gao, Luyang Cao, Ziming Liu, Qiming Zhang Jun Zhou, Zhi Jin</i>	
Automated Proactive Logging Quality Improvement for Large-Scale Codebases	3426
<i>Yichen Li, Jinyang Liu, Junsong Pu, Zhihan Jiang, Zhuangbin Chen, Xiao He, Tieying Zhang, Jianjun Chen Yi Li, Rui Shi, Michael R. Lyu</i>	
From Redundancy to Efficiency: Exploiting Shared UI Interactions towards Efficient LLM-Based Testing	3438
<i>Xuan Wang, Yingchuan Wang, Yongxiang Hu, Yu Zhang, Hailiang Jin, Shiyu Guo, Juxing Yuan Yangfan Zhou</i>	
Automated Prompt Generation for Code Intelligence: An Empirical study and Experience in WeChat	3450
<i>Kexing Ji, Shiyun Fu, Cuiyun Gao, Yujia Chen, Zezhou Yang, Chaozheng Wang, Yuetang Deng</i>	
From Technical Excellence to Practical Adoption: Lessons Learned Building an ML-Enhanced Trace Analysis Tool	3462
<i>Kaveh Shahedi, Matthew Khouzam, Heng Li, Maxime Lamothe, Foutse Khomh</i>	
Evaluating Large Language Models for Functional and Maintainable Code in Industrial Settings: A Case Study at ASML	3475
<i>Yash Mundhra, Max Valk, Maliheh Izadi</i>	
BitsAI-Fix: LLM-Driven Approach for Automated Lint Error Resolution in Practice	3486
<i>Yuanpeng Li, Qi Long, Zhiyuan Yao, Jian Xu, Lintao Xie, Xu He, Lu Geng, Xin Han, Yueyan Chen Wenbo Duan</i>	
Adaptive Performance Regression Detection Using A Semi-Supervised Siamese Network	3498
<i>Yongqian Sun, Mengyao Li, Xiao Xiong, Lei Tao, Yimin Zuo, Wenwei Gu, Shenglin Zhang, Junhua Kuang Yu Luo, Huandong Zhuang, Bowen Deng, Dan Pei</i>	
TreeRanker: Fast and Model-Agnostic Ranking System for Code Suggestions in IDEs	3510
<i>Daniele Cipollone, Egor Bogomolov, Arie Van Deursen, Maliheh Izadi</i>	
Evaluating Large Language Models for Time Series Anomaly Detection in Aerospace Software	3522
<i>Yang Liu, Yixing Luo, Xiaofeng Li, Xiaogang Dong, Bin Gu, Zhi Jin</i>	
ErrorPrism: Reconstructing Error Propagation Paths in Cloud Service Systems	3534
<i>Junsong Pu, Yichen Li, Zhuangbin Chen, Jinyang Liu, Zhihan Jiang, Jianjun Chen, Rui Shi, Zibin Zheng Tieying Zhang</i>	
LLM-assisted Industrial-Scale Differential Testing of Package Incompatibilities in Linux Distributions	3546
<i>Yuhao Yang, Chijin Zhou, Runzhe Wang, Weibo Zhang, Yuheng Shen, Xiaohai Shi, Tao Ma, Chang Gao Zhe Wang, Ying Fu, Heyuan Shi</i>	
Shrunk, Yet Complete: Code Shrinking-Resilient Android Third-Party Library Detection	3557
<i>Jingkun Zhang, Jingzheng Wu, Xiang Ling, Tianyue Luo, Bolin Zhou, Mutian Yang</i>	

Tuning LLM-based Code Optimization via Meta-Prompting: An Industrial Perspective	3569
<i>Jingzhi Gong, Rafail Giavrimis, Paul Brookes, Vardan Voskanyan, Fan Wu, Mari Ashiga, Matthew Truscott Mike Basios, Leslie Kanthan, Jie Xu, Zheng Wang</i>	
Element-Aware Fine-Tuning of Vision-Language Models for Cost-Efficient GUI Testing in an Industrial Setting	3581
<i>Mengzhou Wu, Yuzhe Guo, Yuan Cao, Haochuan Lu, Hengyu Zhang, Xia Zeng, Liangchao Yao Yuetang Deng, Dezhi Ran, Wei Yang, Tao Xie</i>	
AutoPLC: Generating Vendor-Aware Structured Text for Programmable Logic Controllers	3593
<i>Donghao Yang, Aolang Wu, Tianyi Zhang, Li Zhang, Xiaoli Lian, Fang Liu, Yuming Ren, Jiaji Tian Xiaoyin Che</i>	
IntelliTopo: An IaC Generation Service for Industrial Network Topology Construction	3605
<i>Mingyu Shao, Zhao Liu, Weihong Han, Cuiyun Gao, Jiachen Liu, Qing Liao</i>	
RPG: Linux Kernel Fuzzing Guided by Distribution-Specific Runtime Parameter Interfaces	3616
<i>Yuhan Chen, Yuheng Shen, Guoyu Yin, Fan Ding, Runzhe Wang, Tao Ma, Xiaohai Shi, Qiang Fu, Ying Fu Heyuan Shi</i>	
DALEQ - Explainable Equivalence for Java Bytecode	3628
<i>Jens Dietrich, Behnaz Hassanshahi</i>	
M2QCode: A Model-Driven Framework for Generating Multi-Platform Quantum Programs	3639
<i>Xiaoyu Guo, Shinobu Saito, Jianjun Zhao</i>	
SCOPE: Evaluating and Enhancing Permission Explanation Transparency in Mobile Apps	3650
<i>Liu Wang, Tianshu Zhou, Haoyu Wang, Xiyuan Liu, Yi Wang</i>	
The Gold Digger in the Dark Forest: Industrial-Scale MEV Analysis in Ethereum	3661
<i>Ningyu He, Tianyang Chi, Xiaohui Hu, Haoyu Wang</i>	
RepoMasterEval: Evaluating Code Completion via Real-World Repositories	3672
<i>Qinyun Wu, Chao Peng, Pengfei Gao, Ruida Hu, Haoyu Gan, Bo Jiang, Jinhe Tang, Zhiwen Deng Zhanming Guan, Cuiyun Gao, Xia Liu, Ping Yang</i>	
Securing Millions of Decentralized Identities in Alipay Super App with End-to-End Formal Verification	3684
<i>Ziyu Mao, Xiaolin Ma, Lin Huang, Huan Yang, Wu Zhang, Weichao Sun, Yongtao Wang, Jingling Xue Jingyi Wang</i>	
Quantum Machine Learning-based Test Oracle for Autonomous Mobile Robots	3695
<i>Xinyi Wang, Qinghua Xu, Paolo Arcaini, Shaikat Ali, Thomas Peyrucain</i>	
Securing Self-Managed Third-Party Libraries	3707
<i>Xin Zhou, Jinwei Xu, He Zhang, Yanjing Yang, Lanxin Yang, Bohan Liu, Hongshan Tang</i>	
MobileUPReg: Identifying User-Perceived Performance Regressions in Mobile OS Versions	3719
<i>Wei Liu, Yi Wen Heng, Feng Lin, Tse-Hsun Peter Chen, Ahmed E. Hassan</i>	
Thinking Longer, Not Larger: Enhancing Software Engineering Agents via Scaling Test-Time Compute	3730
<i>Yingwei Ma, Yongbin Li, Yihong Dong, Xue Jiang, Yanhao Li, Yue Liu, Rongyu Cao, Jue Chen, Fei Huang Binhua Li</i>	

LogSage: An LLM-Based Framework for CI/CD Failure Detection and Remediation with Industrial Validation	3742
<i>Weiyuan Xu, Juntao Luo, Tao Huang, Kaixin Sui, Jie Geng, Qijun Ma, Isami Akasaka, Xiaoxue Shi Jing Tang, Peng Cai</i>	
Kair: A Statistical and Causal Approach to Pinpointing Stragglers in Distributed Model Training	3754
<i>Yitang Yang, Junhong Liu, Jiapeng Chen, Xiaoyang Sun, Tianyu Wo, Chunming Hu, Chengru Song Jin Ouyang, Renyu Yang</i>	
What Types of Code Review Comments Do Developers Most Frequently Resolve?	3760
<i>Saul Goldman, Hong Yi Lin, Jirat Pasuksmit, Patanamon Thongtanunam, Kla Tantithamthavorn, Zhe Wang Ray Zhang, Ali Behnaz, Fan Jiang, Michael Siers, Ryan Jiang, Mike Buller, Minwoo Jeong, Ming Wu</i>	
Towards Reliable LLM-based Exam Generation Lessons Learned and Open Challenges in an Industrial Project	3766
<i>Renzo Degiovanni, Jordi Cabot</i>	
Acceleration of Automotive Software Development by Retrieval Augmented Integration Test Script Generation	3772
<i>Masashi Mizoguchi, Kentaro Yoshimura, Keita Nakazawa, Yasuomi Sato, Takahiro Iida, Fumio Narisawa</i>	
SGCR: A Specification-Grounded Framework for Trustworthy LLM Code Review	3778
<i>Kai Wang, Bingcheng Mao, Shuai Jia, Yujie Ding, Dongming Han, Tianyi Ma, Bin Cao</i>	
Walk the Talk: Is Your Log-based Software Reliability Maintenance System Really Reliable?	3784
<i>Minghua He, Tong Jia, Chiming Duan, Pei Xiao, Lingzhe Zhang, Kangjin Wang, Yifan Wu, Ying Li Gang Huang</i>	
From Modules to Marketplaces: A Vision for Composable Capability Sharing Across Organizations	3789
<i>Wei-Ji Wang</i>	
Measuring Software Resilience Using Socially Aware Truck Factor Estimation	3794
<i>Alexis Butler, Dan O'Keeffe, Santanu Kumar Dash</i>	
Interaction-Aware Patch Assessment for Multi-Fault Automated Program Repair	3799
<i>Omar I. Al-Bataineh</i>	
Debugging the Undebuggable: Why Multi-Fault Programs Break Debugging and Repair Tools	3804
<i>Omar I. Al-Bataineh</i>	
How Does ChatGPT Make Assumptions When Creating Erroneous Programs?	3809
<i>Sadia Jahan, Xiaoyin Wang</i>	
LLM-based Dynamic Differential Testing for Database Connectors with Reinforcement Learning-Guided Prompt Selection	3814
<i>Ce Lyu, Yanhao Wang, Jie Liang, Minghao Zhao</i>	
Uncovering Systematic Failures of LLMs in Verifying Code Against Natural Language Specifications	3819
<i>Haolin Jin, Huaming Chen</i>	
Is Measurement Enough? Rethinking Output Validation in Quantum Program Testing	3824
<i>Jiaming Ye, Xiongfei Wu, Shangzhou Xia, Fuyuan Zhang, Jianjun Zhao</i>	
The Future of Software Transparency: Bridging Understanding, Measurement, and Practice	3829
<i>Gregorio Dalia, Annibale Panichella, Andrea Di Sorbo, Gerardo Canfora, Corrado A. Visaggio</i>	

Fault Injection for Simulink-based CPS Models: Insights and Future Directions	3834
<i>Drishti Yadav, Claudio Mandrioli, Ezio Bartocci, Domenico Bianculli</i>	
Taming Uncertainty via Automation: Observing, Analyzing, and Optimizing Agentic AI Systems	3840
<i>Dany Moshkovich, Sergey Zeltyn</i>	
Detecting and Repairing Incomplete Software Requirements with Multi-LLM Ensembles	3845
<i>Mohamad Kassab, Marwan AbdElhameed</i>	
Linguistic Theories Coincide with Misformalization in Temporal Logic	3850
<i>Colin S. Gordon</i>	
Exploring Autonomous Agents: A Closer Look at Why They Fail When Completing Tasks	3856
<i>Ruofan Lu, Yichen Li, Yintong Huo</i>	
LLM-Powered Fully Automated Chaos Engineering: Towards Enabling Anyone to Build Resilient Software Systems at Low Cost	3861
<i>Daisuke Kikuta, Hiroki Ikeuchi, Kengo Tajiri</i>	
Unseen Data Detection using Routing Entropy in Mixture-of-Experts for Autonomous Vehicles	3866
<i>Sang In Lee, Donghwan Shin, Jihun Park</i>	
CodeACT-R: A Cognitive Simulation Framework for Human Attention in Code Reading	3871
<i>Yueke Zhang, Zihan Fang, Greg Trafton, Daniel Levin, Kevin Leach, Yu Huang</i>	
Envisioning Intelligent Requirements Engineering via Knowledge-Guided Multi-Agent Collaboration	3876
<i>Jiangping Huang, Dongmin Jin, Weisong Sun, Yang Liu, Zhi Jin</i>	
NovaQ: Improving Quantum Program Testing through Diversity-Guided Test Case Generation	3881
<i>Tiancheng Jin, Shangzhou Xia, Jianjun Zhao</i>	
When Abstraction Breaks Physics: Rethinking Modular Design in Quantum Software	3886
<i>Jianjun Zhao</i>	
Towards Automated Governance: A DSL for Human-Agent Collaboration in Software Projects	3891
<i>Adem Ait, Gwendal Jouneaux, Javier Luis Cánovas Izquierdo, Jordi Cabot</i>	
Simulated Interactive Debugging	3896
<i>Yannic Noller, Erick Chandra, Srinidhi Chandrashekar, Kenny Choo, Cyrille Jegourel, Oka Kurniawan Christopher M. Poskitt</i>	
LLM-Guided Genetic Improvement: Envisioning Semantic Aware Automated Software Evolution	3902
<i>Karine Even-Mendoza, Alexander Brownlee, Alina Geiger, Carol Hanna, Justyna Petke, Federica Sarro Dominik Sobania</i>	
ConfuseTaint: Exploiting Vulnerabilities to Bypass Dynamic Taint Analysis	3907
<i>Yufei Wu, Alexandre Bartel</i>	
Human-In-The-Loop Oracle Learning for Simulation-Based Testing	3912
<i>Ben-Hau Chia, Eunsuk Kang, Christopher S. Timperley</i>	
Tether: A Personalized Support Assistant for Software Engineers with ADHD	3917
<i>Aarsh Shah, Cleyton Magalhaes, Kiev Gama, Ronnie de Souza Santos</i>	
Measuring LLM Code Generation Stability via Structural Entropy	3922
<i>Yewei Song, Tiezhu Sun, Xunzhu Tang, Prateek Kumar Rajput, Tegawendé F. Bissyandé, Jacques Klein</i>	

Multiple Schema-Conformant Declarative Code Generation	3927
<i>Mehant Kammakomati, Srikanth G. Tamilselvam</i>	
Requirements Development and Formalization for Reliable Code Generation: A Multi-Agent Vision	3932
<i>Xu Lu, Weisong Sun, Yiran Zhang, Ming Hu, Cong Tian, Zhi Jin, Yang Liu</i>	
IDBFuzz: Web Storage DataBase Fuzzing with Controllable Semantics	3938
<i>Jingyi Chen, Jinfu Chen, Saihua Cai, Shengran Wang</i>	
STaint: Detecting Second-Order Vulnerabilities in PHP Applications with LLM-Assisted Bi-Directional Static Taint Analysis	3943
<i>Yuchen Ji, Hongchen Cao, Jingzhu He</i>	
Vessel: A Taxonomy of Reproducibility Issues for Container Images	3948
<i>Kevin Pitstick, Alex Derr, Lihan Zhan, Sebastián Echeverría</i>	
Autonomous Agents for Accessibility: Simulating Visual Impairments in Web Interfaces	3953
<i>Juan Diego Yepes-Parra, Camilo Escobar-Velásquez</i>	
Are We SOLID Yet? An Empirical Study on Prompting LLMs to Detect Design Principle Violations	3958
<i>Fatih Pehlivan, Arçin Ülkü Ergüzen, Sahand Moslemi Yengejeh, Mayasah Lami, Anil Koyuncu</i>	
RAML: Toward Retrieval-Augmented Localization of Malicious Payloads in Android Apps	3963
<i>Tiezhu Sun, Marco Alecci, Yewei Song, Xunzhu Tang, Kisub Kim, Jordan Samhi, Tegawendé F. Bissyandé Jacques Klein</i>	
A Secure Mocking Approach towards Software Supply Chain Security	3968
<i>Daisuke Yamaguchi, Shinobu Saito, Takuya Iwatsuka, Nariyoshi Chida, Tachio Terauchi</i>	
AndroFL: Evolutionary-Driven Fault Localization for Android Apps	3973
<i>Vishal Singh, Ravi Shankar Das, Prajwal H G, Subhajit Roy</i>	
VUSC: An Extensible Research Platform for Java-Based Static Analysis	3977
<i>Marc Miltenberger, Steven Arzt</i>	
DESIGNATOR: a Toolset for Automated GAN-enhanced Search-based Testing and Retraining of DNNs in Martian Environments	3981
<i>Mohammed Oualid Attaoui, Fabrizio Pastore</i>	
A Large-Scale Evolvable Dataset for Model Context Protocol Ecosystem and Security Analysis	3985
<i>Zhiwei Lin, Bonan Ruan, Jiahao Liu, Weibo Zhao</i>	
LitterBox ⁺ : An Extensible Framework for LLM-enhanced Scratch Static Code Analysis	3989
<i>Benedikt Fein, Florian Obermüller, Gordon Fraser</i>	
CodeGenLink: A Tool to Find the Likely Origin and License of Automatically Generated Code	3993
<i>Daniele Bifulco, Guido Annicchiarico, Pierluigi Barbiero, Massimiliano Di Penta, Fiorella Zampetti</i>	
PyGress: Tool for Analyzing the Progression of Code Proficiency in Python OSS Projects	3997
<i>Rujiphart Charatvaraphan, Bunradar Chatchaiyadech, Thitirat Sukijprasert, Chaiyong Ragkhitwetsagul Morakot Choetkiertikul, Raula Gaikovina Kula, Thanwadee Sunetnanta, Kenichi Matsumoto</i>	
FlowStrider: Low-Friction Continuous Threat Modeling	4001
<i>Bernd Gruner, Noah Erthel, Clemens-Alexander Brust</i>	

Towards Context-aware Mobile Privacy Notice: Implementation of A Deployable Contextual Privacy Policies Generator	4005
<i>Haochen Gong, Zhen Tao, Shidong Pan, Zhenchang Xing, Xiaoyu Sun</i>	
AgentDroid: A Multi-Agent Tool for Detecting Fraudulent Android Applications	4009
<i>Ruwei Pan, Hongyu Zhang, Zhonghao Jiang, Ran Hou</i>	
XRintTest: An Automated Framework for User Interaction Testing in Extended Reality Applications	4013
<i>Ruizhen Gu, José Miguel Rojas, Donghwan Shin</i>	
ReFuzzer: Feedback-Driven Approach to Enhance Validity of LLM-Generated Test Programs	4017
<i>Iti Shree, Karine Even-Mendoza, Tomasz Radzik</i>	
PrioTestCI: Efficient Test Case Prioritization in GitHub Workflows for CI Optimization	4021
<i>Shubham Vasudeo Desai, Shonil Bhide, Souhaila Serbout, Luciano Marchezan, Wesley K. G. Assunção</i>	
DeepTx: Real-Time Transaction Risk Analysis via Multi-Modal Features and LLM Reasoning	4025
<i>Yixuan Liu, Xinlei Li, Yi Li</i>	
FETT: Fault Injection as an Educational and Training Tool in Cybersecurity	4029
<i>Anaé De Baets, Guillaume Nguyen, Xavier Devroey, Fabian Gilson</i>	
WIBE: Watermarks for generated Images – Benchmarking & Evaluation	4033
<i>Aleksey Yakushev, Aleksandr Akimenkov, Khaled Abud, Dmitry Obydenkov, Irina Serzhenko, Kirill Aistov Egor Kovalev, Stanislav Fomin, Anastasia Antsiferova, Kirill Lukianov, Yury Markin</i>	
ORMorpher: An Interactive Framework for ORM Translation and Optimization	4037
<i>Milan Abrahám, Pavel Koupil</i>	
BuilDroid: A Self-Correcting LLM Agent for Automated Android Builds	4041
<i>Jaehyeon Kim, Rui Rua, Karim Ali</i>	
Chrysalis: A Lightweight Logging and Replay Framework for Metamorphic Testing in Python	4045
<i>Jai Parera, Nathan Huey, Ben Limpanukorn, Miryung Kim</i>	
BASHIRI: Learning Failure Oracles from Execution Features	4049
<i>Marius Smytzek, Martin Eberlein, Tural Mammadov, Lars Grunske, Andreas Zeller</i>	
GUI-ReRank: Enhancing GUI Retrieval with Multi-Modal LLM-based Reranking	4053
<i>Kristian Kolthoff, Felix Kretzer, Alexander Maedche, Simone Paolo Ponzetto, Christian Bartelt</i>	
EyeNav: Accessible Webpage Interaction and Testing using Eye-tracking and NLP	4057
<i>Juan Diego Yepes-Parra, Camilo Escobar-Velásquez</i>	
DSBox: A Data Selection Framework for Efficient Deep Code Learning	4062
<i>Xinyang Liu, Lili Quan, Qiang Hu</i>	
BenGQL: An Extensible Benchmarking Framework for Automated GraphQL Testing	4066
<i>Abenezer Angamo, Marcello Maugeri</i>	
PyTrim: A Practical Tool for Reducing Python Dependency Bloat	4070
<i>Konstantinos Karakatsanis, Georgios Alexopoulos, Ioannis Karyotakis, Foivos Timotheos Proestakis Evangelos Talos, Panos Louridas, Dimitris Mitropoulos</i>	
OSSPREY: AI-Driven Forecasting and Intervention for OSS Project Sustainability	4074
<i>Nafiz Imtiaz Khan, Priyal Soni, Arjun Ashok, Vladimir Filkov</i>	

Evaluating Program Coverage for Code-Model Training	4078
<i>Nandakishore Menon, Diptikalyan Saha</i>	
TrustVis: A Multi-Dimensional Trustworthiness Evaluation Framework for Large Language Models	4082
<i>Ruoyu Sun, Da Song, Jiayang Song, Yuheng Huang, Lei Ma</i>	
Metamorphic Testing of Deep Reinforcement Learning Agents with MDPMorph	4086
<i>Jiapeng Li, Zheng Zheng, Yuning Xing, Daixu Ren, Steven Cho, Valerio Terragni</i>	
APIDA-Chat: Structured Synthesis of API Search Dialogues to Bootstrap Conversational Agents	4090
<i>Zachary Eberhart, Collin McMillan</i>	
Quirx: A Mutation-Based Framework for Evaluating Prompt Robustness in LLM-based Software	4094
<i>Souhaila Serbout</i>	
Training-Control-as-Code: Towards a declarative solution to control training	4098
<i>Padmanabha V. Seshadri, Harikrishnan Balagopal, Mehant Kammakomati, Ashok Pon Kumar Dushyant Behl</i>	
LLMorph: Automated Metamorphic Testing of Large Language Models	4102
<i>Steven Cho, Stefano Ruberto, Valerio Terragni</i>	
StackPlagger: A System for Identifying AI-Code Plagiarism on Stack Overflow	4106
<i>Aman Swaraj, Harsh Goyal, Sumit Chadgal, Sandeep Kumar</i>	
Clara: A Developer's Companion for Code Comprehension and Analysis	4110
<i>Ahmed Adnan, Mushfiqur Rahman, Saad Sakib Noor, Kazi Sakib</i>	
PROXiFY: A Bytecode Analysis Tool for Detecting and Classifying Proxy Contracts in Ethereum Smart Contracts	4114
<i>Ilham Qasse, Mohammad Hamdaqa, Björn Þór Jónsson</i>	
evalSmarT: An LLM-Based Framework for Evaluating Smart Contract Generated Comments	4118
<i>Fatou Ndiaye Mbodji, Mame Marieme C. Sougoufara, Wendkûni A. M. Christian Ouedraogo Alioune Diallo, Kui Liu, Jacques Klein, Tegawendé F. Bissyandé</i>	
Verification and Classification of Exploits for Node.js Vulnerabilities	4122
<i>Sungmin Park</i>	
First-Order Quantified Separator in Alloy Analyzer	4125
<i>One An</i>	
Understanding Uncertainty In LLMs	4128
<i>Chandan Kumar Sah</i>	
Dynamic Testing of GUI Exercises in Headless Environments	4131
<i>Benjamin Schmitz</i>	
Detecting Vulnerabilities from Issue Reports for Internet-of-Things	4134
<i>Sogol Masoumzadeh</i>	
Human-Centered Evaluation of REST API Fuzzing Tools: Bridging Academia and Industry	4137
<i>Fanny Febriani Susilo</i>	
Testing Autonomous Driving Systems Through Blind-Spot Guided Fuzzing	4141
<i>Sali Moussa</i>	

Detecting and Mitigating Inconsistencies Between Code, Documentation and Tests	4145
<i>Tobias Kiecker</i>	
Secure Transaction Semantics: Analysis, Vulnerability Detection, and Attack Modeling	4149
<i>Yixuan Liu</i>	
Improving Quality of LLM Code Generation in Low-Resource Programming Languages via Uncertainty Estimation	4153
<i>Georgii Andriushchenko</i>	