
CONFERENCE PROGRAM

ASE 2026

12–16 October 2026

Hochstraße 3, Munich, Germany

Contents

Monday 12 October 2026

09:00–10:00	Forum 6	TRUST	4
	Forum 7	Harness4GenUI	4
	Forum 9	ASE New-Faculty Symposium	4
	Forum 10	ASE Doctoral Symposium	4
	Forum 11	JAWs 2	4
	Forum 13	APR	4
	Forum 15	RASE	4
	Boardroom	RSSR	5
10:00–10:30	Forum 12	Break	5
10:30–12:30	Forum 6	TRUST	5
	Forum 7	Harness4GenUI	5
	Forum 9	ASE New-Faculty Symposium	5
	Forum 10	ASE Doctoral Symposium	6
	Forum 11	JAWs 2	6
	Forum 13	APR	6
	Forum 15	RASE	6
	Boardroom	RSSR	7
14:00–16:00	Forum 7	Harness4GenUI	7
	Forum 9	ASE New-Faculty Symposium	7
	Forum 10	ASE Doctoral Symposium	7
	Forum 11	JAWs 2	7
	Forum 13	APR	8
	Forum 15	RASE	8
	Boardroom	CONFLANG	8
16:30–18:00	Forum 7	RSSR	9
	Forum 9	ASE New-Faculty Symposium	9
	Forum 10	ASE Doctoral Symposium	9
	Forum 11	JAWs 2	9
	Forum 15	RASE	9
	Boardroom	CONFLANG	9

Tuesday 13 October 2026

08:30–09:00	Plenary Room (Forum 1-3)	ASE Opening	10
09:00–10:00	Plenary Room (Forum 1-3)	Trust the Delta: Rethinking Software Validation in the AI Era	10
10:30–12:30	Forum 1	AI and Software Engineering: AI4SE 1	10
	Forum 2	AI and Software Engineering: Autonomous and Self-Adapting Systems	10
	Forum 2	Security and Other Non-Functional Properties: Security and Privacy 7	11
	Forum 6	Human and Social Aspects: Program Comprehension and Visualization	11
	Forum 7	Testing and Analysis: Regression and Model-Based Testing 1	11
	Forum 8	Maintenance and Evolution: Refactoring and Reengineering 1	11
	Forum 15	Security and Other Non-Functional Properties: Reliability and Availability 1	12
	Forum 16	Security and Other Non-Functional Properties: Security and Privacy 1	12
14:00–16:00	Forum 1	AI and Software Engineering: AI4SE 2	13
	Forum 2	Testing and Analysis: Regression and Model-Based Testing 2	13
	Forum 2	Testing and Analysis: Regression and Model-Based Testing 4	13
	Forum 2	Testing and Analysis: Regression and Model-Based Testing 2	13
	Forum 2	Human and Social Aspects: Human-Computer Interface	13
	Forum 2	Testing and Analysis: Regression and Model-Based Testing 2	14
	Forum 6	Security and Other Non-Functional Properties: Security and Privacy 2	14
	Forum 7	Security and Other Non-Functional Properties: Performance 1	14
	Forum 8	Testing and Analysis: Automated Program Repair and Synthesis 1	15
	Forum 15	Formal Aspects of Software Engineering: Specification Languages 1	15
	Forum 16	Maintenance and Evolution: API Design and Management	15
16:30–17:30	Forum 1	AI and Software Engineering: SE4AI 1	16
	Forum 2	Testing and Analysis: Fuzzing and Mutation Testing 1	16
	Forum 6	Software Analytics: Software Metrics	16
	Forum 7	Security and Other Non-Functional Properties: Reliability and Availability 2	16
	Forum 8	Human and Social Aspects: Software Engineering Processes 1	17
	Forum 15	Requirements and Design: Software Architecture and Design	17
	Forum 16	Security and Other Non-Functional Properties: Security and Privacy 3	17
18:00–21:30	Plenary Room (Forum 1-3)	Conferene Reception	18

Wednesday 14 October 2026

09:00–10:00	 Plenary Room (Forum 1-3)	Beyond the Hype: What AI Really Did to Software Engineering	19
10:30–12:30	 Forum 1	AI and Software Engineering: AI4SE 3	19
	 Forum 2	Testing and Analysis: Automated Program Repair and Synthesis 2	19
	 Forum 6	Security and Other Non-Functional Properties: Security and Privacy 4	20
	 Forum 7	Maintenance and Evolution: Refactoring and Reengineering 2	20
	 Forum 8	Human and Social Aspects: Software Engineering Processes 2	20
	 Forum 15	Requirements and Design: Requirements and Traceability	21
	 Forum 16	Testing and Analysis: Regression and Model-Based Testing 3	21
14:00–16:00	 Forum 1	AI and Software Engineering: SE4AI 2	22
	 Forum 2	Testing and Analysis: Static and Dynamic Analysis 1	22
	 Forum 6	Software Analytics: Mining Software Repositories 1	22
	 Forum 7	Testing and Analysis: Regression and Model-Based Testing 4	23
	 Forum 8	AI and Software Engineering: AI4SE 4	23
	 Forum 15	Security and Other Non-Functional Properties: Reliability and Availability 3	23
	 Forum 16	Security and Other Non-Functional Properties: Security and Privacy 5	24
16:30–17:30	 Plenary Room (Forum 1-3)	ASE Awards	24
18:30–22:30	 Festsaal	Conference Dinner	24

Thursday 15 October 2026

09:00–10:00	 Plenary Room (Forum 1-3)	Lost in Automated Specifications: A Time and Space Odyssey in System Testing	25
10:30–12:30	 Forum 1	AI and Software Engineering: AI4SE 5	25
	 Forum 2	Security and Other Non-Functional Properties: Security and Privacy 6	25
	 Forum 6	Testing and Analysis: Regression and Model-Based Testing 5	26
	 Forum 7	Security and Other Non-Functional Properties: Dependability and Safety 1	26
	 Forum 8	Testing and Analysis: Fuzzing and Mutation Testing 2	26
	 Forum 15	Formal Aspects of Software Engineering: Software Validation and Verification	27
	 Forum 16	Maintenance and Evolution: Debugging and Fault Localization	27
14:00–16:00	 Forum 1	AI and Software Engineering: SE4AI 3	27
	 Forum 2	Testing and Analysis: Static and Dynamic Analysis 2	28
	 Forum 6	Security and Other Non-Functional Properties: Security and Privacy 7	28
	 Forum 7	Security and Other Non-Functional Properties: Performance 2	29
	 Forum 8	AI and Software Engineering: AI4SE 6	29
	 Forum 15	Formal Aspects of Software Engineering: Formal Methods and Model Checking 1	29
	 Forum 16	Testing and Analysis: Fuzzing and Mutation Testing 3	30
16:30–17:30	 Forum 1	Human and Social Aspects: Human-Computer Interface	30
	 Forum 2	Software Analytics: Mining Software Repositories 2	30
	 Forum 6	Security and Other Non-Functional Properties: Security and Privacy 8	31
	 Forum 7	Formal Aspects of Software Engineering: Specification Languages 2	31
	 Forum 8	Testing and Analysis: Regression and Model-Based Testing 6	31
	 Forum 15	Formal Aspects of Software Engineering: Formal Methods and Model Checking 2	31
	 Forum 16	Security and Other Non-Functional Properties: Dependability and Safety 2	32

Friday 16 October 2026

08:45–10:00	 Forum 7	SECUTE	33
	 Forum 8	KLEE	33
	 Forum 10	LiveWrite	33
	 Forum 11	CPAchecker	33
	 Forum 13	MAS-GAIN	33
	 Forum 14	SWGeno	33
	 Forum 15	ASYDE	33
	 Forum 16	POVC	34
10:30–12:30	 Forum 7	SECUTE	34
	 Forum 8	KLEE	34
	 Forum 10	LiveWrite	34
	 Forum 11	CPAchecker	34
	 Forum 13	MAS-GAIN	35
	 Forum 14	SWGeno	35
	 Forum 15	ASYDE	35
	 Forum 16	POVC	35
14:00–16:00	 Forum 8	KLEE	36
	 Forum 9	ATIQSER	36
	 Forum 10	LiveWrite	36
	 Forum 11	CPAchecker	36
	 Forum 16	POVC	36
16:30–17:35	 Forum 8	KLEE	37
	 Forum 10	LiveWrite	37
	 Forum 11	CPAchecker	37
	 Forum 16	POVC	37

All times are local time (Europe/Berlin).

Program

Monday

12 October 2026

101 talks

Forum 6

09:00–09:58

TRUST · 5 talks

- 09:00 TRUST 26 Opening
- 09:10 Software Project Management with LLM-Based Automation: Coordination, Validation, and Governance in Practice
Ronnie de Souza Santos, Italo Santos, Cleyton Magalhaes
- 09:22 RigorBench: Benchmarking Engineering Process Discipline in Autonomous AI Coding Agents
Meher Bhaskar Madiraju, Meher Sai Preetam Madiraju
- 09:34 Closed-World Trustworthy Copilots for Legal and Regulatory Text
Monalisha Ojha, Pragya Shrey Mishra
- 09:46 From Regulation to Innovation: Mining High-Impact AI Patents Grounded in the EU AI Act and Korea's AI Basic Act
Jimin Kim, Ho Jun Kang, Jaehyun Lee

Forum 7

09:00–10:00

Harness4GenUI · 2 talks

- 09:00 Introduction & Icebreaker
- 09:15 GenUI and User Centered Design and their Relationship with Automation Design and Engineering of Interactive Systems and Technologies
Philippe Palanque

Forum 9

09:00–10:00

ASE New-Faculty Symposium · 2 talks

- 09:00 Opening and Introductions
Marcel Böhme
- 09:30 Keeping the Flame Alive: Reflections on Building a Meaningful Research Career
Lionel Briand

Forum 10

09:00–10:00

ASE Doctoral Symposium · 2 talks

- 09:00 Opening
Tegan Brennan, Malte Mues
- 09:15 The Devil's Guide to Doing your PhD
Andreas Zeller

Forum 11

09:00–10:00

JAWs 2 · 1 talk

- 09:00 Editors in Conversation: What Makes a Strong TSE and TOSEM Article?

Forum 13

09:00–10:00

APR · 1 talk

- 09:00 What Are We Really Repairing? Lessons from Cloud Failures
Minghua Ma

Forum 15

09:00–10:00

RASE · 2 talks

- 09:00 Workshop Opening

- 09:15 Combining Large-Scale Datasets, Program Analysis, and AI for Reliable and Trustworthy Software Analysis: Lessons from Android Apps
Jacques Klein

Boardroom

09:00–10:00

RSSR · 2 talks

- 09:00 Opening
Marc Carwehl, Sinem Getir Yaman
- 09:10 Can Our Systems Survive Their Own Evolution?
Morgan Ericsson

10:00–10:30 Break

ASE Catering · Forum 12

Forum 6

10:30–12:30

TRUST · 10 talks

- 10:30 AgentChaosBench: Benchmarking External Fault Detection and Localization in Agentic Systems
Chenkai Zhang, Yiran Li, Yifang Tian, Michail Bachras, Hans-Arno Jacobsen
- 10:42 Can Open-Weight LLMs Predict Test Outcomes from Traces? A Cross-Subject Boundary Study on Tests4Py
Giuliano Ricardo Rasper, Tural Mammadov, Marius Smytcek, Andreas Zeller
- 10:54 Dynamic Deception: When Pedestrians Team Up to Fool Autonomous Cars
Masoud Jamshidiyan Tehrani, Marco Gabriel, Jinhan Kim, Paolo Tonella
- 11:06 Explaining LLM Adaptation Decisions in Self-Adaptive Systems via Faithful Input Attribution
Balint Mate, Vincenzo Scotti, Diego Perez-Palacin, Raffaella Mirandola
- 11:18 Metamorphic Stress Testing of Vision-Language Components for Time-Critical Autonomous Systems
Arman Chhetri, Yinxi Liu
- 11:30 Leveraging Large Language Models for Automated Discriminatory Test Generation
Sadia Afrin Mim, Fatema Tuz Zohra, Brittany Johnson
- 11:42 Fairness Guardrails for Autonomous Agents: A Trustworthiness-by-Design Vision from Social Cybersecurity
Maria Teresa Baldassarre, Vita Santa Barletta, Vito Bavaro, Ronnie de Souza Santos
- 11:50 On Resilience in Multi-agent Systems for Software Engineering
Sandra Mitrovic, Cezar Sas, Matteo Salani
- 11:58 Automated Counterfactual Scenario Generation for Fairness Assessment of Large Language Models: The ScenGen Framework
Alessandra Parziale, Antonio D'Auria, Daniele Pio Scaparra, Valeria Pontillo
- 12:10 Open Discussion and Closing

Forum 7

10:30–12:30

Harness4GenUI · 5 talks

- 10:30 What Endures When Everything Evolves? AFS-UI and the Search for Stable Foundations in Generative UI
Robert Mao
- 11:15 Toward Frontier-Quality Declarative UI Generation at Small-Model Cost
Yingxiang Yang, Weihang Xiao, Ben Bullough, Tushar Deshpande, Niresh Agarwal
- 11:35 The Driver Was Never the Problem: Cross-Platform Test Portability as a Semantic Contract
Simarjot Khanna
- 11:55 Prompt-to-Feature Attribution for Explainable Software Generation
Radoslaw Klimek
- 12:15 Interactive Activity

Forum 9

10:30–12:30

ASE New-Faculty Symposium · 4 talks

- 10:30 Research, People, and Me
Sukyoung Ryu
- 11:00 What Counts and What Matters: Careers, Criteria, and What is Actually Worth Doing
Alexander Pretschner
- 11:30 Panel 1: Reflections and Mentorship
Sukyoung Ryu, Alexander Pretschner, Lionel Briand

12:00 2-on-1: Mentorship Round 1*Sukyoung Ryu, Manuel Rigger, Arie van Deursen, Lionel Briand, Alexander Pretschner, Andreas Zeller***Forum 10**

10:30–12:30

ASE Doctoral Symposium · 3 talks**10:30 Automated Explanation-Based Testing and Improvement of Trustworthiness in AI-Based Systems***Lam Nguyen Tung***11:10 Studying, Identifying, and Fixing Hidden Technical Debt in AI-Intensive Cyber-Physical Systems***Beena Beena***11:50 Efficient Compiler Auto-tuning via Optimization Space Reduction***Mingxuan Zhu***Forum 11**

10:35–12:27

JAWs 2 · 12 talks**10:35 Architecting Fairness-Aware LLM Interactions***Alessandra Parziale, Valeria Pontillo, Fabio Palomba, Patrizio Pelliccione***10:40 Structured Generative Autonomy: Combining LLM Flexibility with Assured Robotic Mission Engineering***Martina De Sanctis, Gianluca Filippone, Patrizio Pelliccione, Sofiia Prokofieva***10:45 Verification as an Architectural Layer for LLM Agents: A V-Model Design, and a Pilot Study of Its Deterministic Core***Ali Afoud, Jie JW Wu***10:50 Model-Based Agentic Software Engineering***James C. Davis, Kelechi G. Kalu, Huiyun Peng, Parth Vinod Patil***10:55 Toward Bounding Coding-Agent Autonomy in Software Maintenance: A Pilot Study of Boundary Contracts***Heeyong Eun, Seungyeon Chung, Hyejee Park, Youngeun Lee, Sukhyun Hong, Hyewon Jung***11:00 Journal Clinic I: Consultation Tables (Wise Crowds)****11:32 Pre- and Post-Processing Pipeline for Local Small LMs in Requirements Engineering****11:37 When the Judge Is the Bottleneck: Measuring Where an Adversarial NL-to-LTL Loop Actually Fails***Xuan-Bach Le, Tuan Phat Tran-Truong***11:42 Trustworthy Software Project Generation: A Case Study with an Interactive Theorem Prover***Jian Fang, Yingfei Xiong***11:47 Contract-in-the-Loop: Context-Aware API Misuse Detection***Ragib Shahariar Ayon, Shabbir Ahmed***11:52 Skill or Threat? Risks and Scanner Gaps in the LLM Skill Ecosystem***Xinghao Gao, Xinhe Zhang, Yanxiang Mao, Jiongchi Yu, Xiaoyu Zhang, Ming Hu, Geguang Pu***11:57 Journal Clinic II: Consultation Tables (Wise Crowds)****Forum 13**

10:30–12:20

APR · 5 talks**10:30 Remediation Experience as a First-Class Citizen in Automated Program Repair***Alex Alibrahim, Alexander Rage, Andrii Vozniuk, Ashan Tharindu, David Godoy, Harry Anderson, Isuru Perera, Jacek Poreda, Joseph Wang, Juncheng Loy, Manuel Vianello, Paulo Canelas, Rasmus Vahelaan, Rodolfo Biazzi, Scott Bell, Silpa Vipparthi, Tom Van Braband, Vasileios Naskos, Vojtech Suchy***11:00 Test Cases Considered Harmful***Steven P. Reiss***11:20 FitGen: Validation and Guided Repair of LLM Generated Executable Fitness Functions for Search Based Software Testing***Md Al Muzahid Nayim, Faezeh Rajabi Kouchi, Amit Kumar Sikder, Myra Cohen***11:40 DafnyFix: Single-Transformation Automated Repair in Dafny***Isabel Amaral, Álvaro F. Silva, João F. Ferreira, Alexandra Mendes***12:00 Whose Oracle Is It? The Authority Problem in Deployed Agentic Program Repair***Bardia Mohammadi, Laurent Bindschaedler***Forum 15**

10:30–12:15

RASE · 6 talks**10:30 Toward Defensible System Behavior for AI-Generated Software Artifacts***Radoslaw Klimek*

- 10:45 **TaxVerif: Verification-Oriented COBOL Tax-Code Migration via Invariant Extraction and Semantic Validation**
Vadim Tynchenko, Igor Masich, Dmitry Martysyuk, Yury Kostromin, Aleksei Borodulin, Ivan Malashin, Vladimir Nelyub, Andrei Gantimurov
- 11:05 **NORMA: Automated AI Act Compliance Assessment for AI Training Datasets**
Elisabetta Arba, Allegra Berardi, Bedilia Estrada Torres, Francesco Leotta, Flavia Monti, Massimo Mecella
- 11:20 **Can We Trust Strong Teachers? Auditing Knowledge Distillation for Binary Code Similarity Detection**
Minh Khang Van, Yunru Wang, Johannes Kinder
- 11:40 **Whose Specification Did You Test? The Hidden Interpretation Bias in LLM-Generated Test Suites**
Md Mainul Islam
- 11:55 **How Grammar-Based Synthetic Data Shapes Behavior Learning: An Empirical Study of Input Generation Strategies**
Tural Mammadov, Alireza Kheradmand, Marius Smytcek, Andreas Zeller

Boardroom

10:30–12:20

RSSR · 4 talks

- 10:30 **Model-centric approaches to assuring AI-based autonomy**
Simon Burton
- 11:20 **Break-out groups: preparation**
Marc Carwehl, Sinem Getir Yaman
- 11:40 **A Resilience-Oriented Orchestration Framework for Mission Continuity in Multi-Agent Autonomous Systems**
Anjali Santhosh, Massimo Tivoli, Marco Autili
- 12:00 **R²: Repairability-Preserving Recursive Repair for Multimodal Agentic Systems**
Dominic Dabish

12:30–14:00 Lunch Break

Forum 7

14:00–15:50

Harness4GenUI · 5 talks

- 14:00 **Demo Session**
- 14:05 **Generative UI for TaskGenAgent**
Razieh Akbarilalaei
- 14:10 **GenAI for App Feedback**
Ali Ebrahimi Pourasad
- 14:30 **Panel Discussion**
Ali Ebrahimi Pourasad, Jieshan Chen, Silvia Abrahão, Robert Mao
- 15:30 **Closing**

Forum 9

14:00–16:00

ASE New-Faculty Symposium · 4 talks

- 14:00 **Career development and leadership growth**
Arie van Deursen
- 14:30 **On the Tenure Track: Lessons Learned So Far**
Manuel Rigger
- 15:00 **On Impact in SE Research**
Andreas Zeller
- 15:30 **Panel 2: Growth and Career**
Andreas Zeller, Arie van Deursen, Manuel Rigger

Forum 10

14:00–16:00

ASE Doctoral Symposium · 3 talks

- 14:00 **Evaluating and improving Quality of Experience in virtual environments**
Lara Ihlbrock
- 14:40 **Understanding Human-AI Teaming for Software Engineering Practices**
Jiessie Tie
- 15:20 **Human-centered LLM-based Agents to Support Agile Task Management in Software Engineering**
Razieh Akbarilalaei

Forum 11

JAWs 2 · 11 talks

14:04–15:50

- 14:04 Context Engineering for Uplift-Risk Decision Support: An Evidence-Packaging Study of Agentic Triage in Firefox
Emre Dinc, Eray Tüzün, Mohammad Hamdaqa, Moataz Chouchen, Marco Castelluccio, Suhaib Mujahid, Connor Sheehan
- 14:09 Do Agent Skills Improve Agentic Code Review?
Arda Sarıtaş, Hüseyin Özgür Kamalı, Eray Tüzün
- 14:14 Modernitic: Experience-Driven Agentic Modernization for Scalability-Ready Microservices
Majid Dashtbani, Ladan Tahvildari
- 14:19 From Behavioral Differences to Structured Feedback: Human Review in AI-Assisted Legacy Modernization
Tobias Böhm, Dick Carstens, Tom Schoonbaert, Cornel Klein, Bart Carpels, Andreas Biesdorf
- 14:24 Journal Clinic III: Consultation Tables (Wise Crowds)
- 14:55 Detecting Secret Agents
Romain Robbes, Théo Matricon, Thomas Degueule, Andre Cavalcante Hora, Stefano Zacchirolı, François Charoy
- 15:00 Loop Engineering: Building Blocks, Adoption, and Impact
Jai Lal Lulla, Vahram Nersesyan, Seyedmoein Mohsenimofidi, Christoph Treude, Sebastian Baltes
- 15:05 Beyond Authorship: A Research Vision for Developer Criticality in AI-Assisted Software Engineering
Fereshteh Vedadi, İlke Latifoğlu, Eray Tüzün
- 15:10 BurnRiSc: Toward Non-Invasive Burnout Screening in Open Source from Public Repository Signals
Timofey Sanko, Yuan Tian, Mariam Guizani
- 15:15 Technostress in the Age of AI: A Preliminary Study with Software Professionals
Ronnie de Souza Santos, Italo Santos, Cleyton Magalhaes, Zixuan Feng
- 15:20 Journal Clinic IV: Consultation Tables (Wise Crowds)

Forum 13

APR · 4 talks

14:00–15:20

- 14:00 Execution Features and Their Relations to Faults: A Replication on Defects4J
Marius Smytzek
- 14:20 Failure Explanations: The Missing Abstraction between Debugging and Repair
Marius Smytzek
- 14:40 AnchorRepair: A Reference-Guided Dual-Agent Framework for Code Translation Repair
Jiajun Sun, Yaocai Zhao, Fengjie Li, Jiajun Jiang
- 15:00 RAWL: Retrieval-Augmented Smart Contract Repair using Weisfeiler-Lehman-Based Retrieval
Chihiro Kado, Tatsuhiro Tsuchiya

Forum 15

RASE · 6 talks

14:00–15:45

- 14:00 Software Engineering Governance: Foundations for Trustworthy Automated Software Engineering
Carlos Mario Braga Ortuño, Manuel Serrano, Eduardo Fernandez-Medina
- 14:20 HermesAngler: Enabling Dynamic Instrumentation of iOS React Native Apps
Martin Renze, Fabian Scherf
- 14:35 Formal Analysis of Software Bills of Materials with Alloy
Sabato Nocera, Giuseppe Scanniello
- 14:50 Minimizing Vulnerability Prevalence in Maven Dependencies via Answer Set Programming
Christian Macho, Laura Hofer, Martin Pinzger
- 15:10 When Log Parsing Removes What Matters: Detecting Failures in Distributed GPU Workloads
Boyuan Cui, Yuyang Zeng, Yifang Tian, Hans-Arno Jacobsen
- 15:25 When Do Inference-Cost Stress Tests Generalize? A Cross-Model and Cross-Configuration Measurement Study
Arman Chhetri, Yinxi Liu

Boardroom

CONFLANG · 3 talks

14:00–16:00

- 14:00 A brief history of security configuration
Patrick Stöckle
- 15:20 Poliac: Production Observability for LLM-based Infrastructure as Code
Haitam El Hayani, Jolan Philippe, Stéphanie Challita

15:40 How to represent CUE terms as graphs
Eloi Perdereau, Jacques Noyé

16:00–16:30 *Coffee Break*

Forum 7

16:30–18:00

RSSR · 4 talks

- 16:30 AIEngOrchestrator-HITL: Human-in-the-Loop Process Intelligence for Resilient AI Systems
Razan Abualsaud
- 17:00 MAVLink Proxy Defense with Intrusion Detection, Safety Quarantine, and Dual-View Deception for UAV Command-and-Control
Utkarsh Balu Lubal, Muhammad Abdul Basit Ur Rahim, Muhammad Rizwan Abid
- 17:30 Break-out groups: summary
Sinem Getir Yaman, Marc Carwehl
- 17:50 Closing
Sinem Getir Yaman, Marc Carwehl

Forum 9

16:30–17:30

ASE New-Faculty Symposium · 2 talks

- 16:30 2-on-1: Mentorship Round 2
Sukyoung Ryu, Manuel Rigger, Arie van Deursen, Alexander Pretschner, Lionel Briand, Andreas Zeller
- 17:00 Summary and Closing
Marcel Böhme

Forum 10

16:30–17:30

ASE Doctoral Symposium · 3 talks

- 16:30 Safeguarding the Development and Deployment of Agentic AI Applications in the Enterprise Context
Rui Yang
- 17:10 Ask us Anything
Tegan Brennan, Malte Mues, Stefan Brunthaler, Vadim Zaytsev
- 17:25 Closing
Tegan Brennan, Malte Mues

Forum 11

16:30–17:30

JAWs 2 · 1 talk

- 16:30 From Workshop Paper to Journal Article: Minimum Specifications and Next Steps

Forum 15

16:30–17:30

RASE · 2 talks

- 16:30 Discussion Panel
- 17:20 Awards & Workshop Closing - chaired by RASE Chairs and AI4I-IAS (Valeria Morra)

Boardroom

16:30–17:30

CONFLANG · 3 talks

- 16:30 Scopeannon: A Static Analyzer of Ansible's Scope Ambiguities
Olivia Proust, Hélène Coullon, Jolan Philippe, Frederic Louergue
- 16:50 Verifying the Stability of Terraform Executions
Eloi Perdereau, Daniel Sokolowski, Jolan Philippe, Hélène Coullon, Guido Salvaneschi
- 17:10 Déjà vu? Redundant Dependencies in Terraform
Carlos Felgueiras, Luís W. Barbosa, Nuno Saavedra, Alexandra Mendes, João F. Ferreira

Tuesday

13 October 2026

152 talks

08:30–09:00 ASE Opening

Stefan Winter, Shiva Nejati, Justyna Petke

ASE Organization · Plenary Room (Forum 1-3)

09:00–10:00 Trust the Delta: Rethinking Software Validation in the AI Era

Michael Pradel

ASE Keynotes · Plenary Room (Forum 1-3)

10:00–10:30 Coffee Break

Forum 1

10:30–12:30

AI and Software Engineering: AI4SE 1

ASE Research Papers · 8 talks

- 10:30 How Reasoning Shapes Social Bias in LLM-Generated Code?
Weifeng Sun, Jieke Shi, Zhou Yang, Yuchen Chen, Hongyan Li, Meng Yan, David Lo
- 10:45 Think Before You Code: Dual Reasoning for the NLSafety–Utility Trade-Off in LLM Code Generation
Honghao Tan, Haibo Wang, Shin Hwei Tan
- 11:00 Exact-IaC: Execution-Driven Exact Knowledge Injection for Reliable IaC Synthesis
Hua Ge, Minxue Pan
- 11:15 SwiftGram: Dynamic Lexical Forking for Precise and Efficient Grammar-Constrained Decoding
Shuyi Ling, Liangyi Kang, Jie Liu, Shuai Wang, Zhirou Ma, Dan Ye
- 11:30 Formal Verification-Integrated Reinforcement Learning for Automated ORM Schema Synthesis
Md Rashedul Hasan, Hamid Bagheri
- 11:45 AlphaCu: A Transformation-Driven Synthesis Framework for LLM-Based GPU Kernel Generation
Yangjie Zhou, Xing Ma, Zihan Liu, Wang Qing, Weihao Cui, Sun Wu, Hanjing Wang, Binhang Qi, Jingwen Leng, Yun Lin, Minyi Guo, Jin Song Dong
- 12:00 MultiKernelBench: A Multi-platform Benchmark for Kernel Generation
Zhongzhen Wen, zhangyinghui, Zhong Li, Zhongxin Liu, Linna Xie, Tian Zhang
- 12:15 RealisticTritonBench: A Benchmark for Triton-Kernel Generation in Real-World AI Frameworks
Jinjun Huang, Zhongzhen Wen, Tongtong Xu, Meng Yan, Xin Xia, Zhongxin Liu

Forum 2

10:30–12:15

AI and Software Engineering: Autonomous and Self-Adapting Systems

ASE Research Papers · 8 talks

- 10:30 ProbGuard: Proactive Runtime Monitoring for LLM Agent Safety via Probabilistic Prediction
Haoyu Wang, Chris Poskitt, Jiali Wei, Jun Sun
- 10:45 What Breaks When LLMs Code? Characterizing Operational Safety Failures of Agentic Code Assistants
Alif Al Hasan, Sumon Biswas
- 11:00 Maintaining Control States for Initialization-Aware GRPO in Industrial LLM Workflows
Chen Liu, Zhenzhao Yuan, Fengyuan Fan, Limei Che, qiao qian, Qingyuan Liang, Zeyu Sun, Lu Zhang
- 11:15 From Plan to Action: How Well Do Agents Follow the Plan?
Shuyang Liu, Saman Dehghan, Jatin Ganhotra, Martin Hirzel, Reyhaneh Jabbarvand
- 11:30 Safeguarding LLM Agents from Misalignment through Provenance Analysis
Yining She, Yiliang Liang, Eunsuk Kang
- 11:45 Introducing Safe LLM-Agent Execution via Typed Tool Mediation and World-Effect Evaluation
Jun Zhou, Tao Jia

- 11:55 **FeatX: Editing Software by Editing Features for Repository-Level Code Evolution**
Xutian Li, Yifeng Zhu, Xianlin Zhao, Yanzhen Zou, Lu Zhang, Bing Xie
- 12:05 **Graphectory Viewer: A Tool for Process-Centric Analysis of Agentic Software Trajectories**
Charlie Jyu, Shuyang Liu, Reyhaneh Jabbarvand

Forum 2

12:15–12:30

Security and Other Non-Functional Properties: Security and Privacy 7

ASE Research Papers · 1 talk

- 12:15 **Assessing the adoption of security policies by developers in terraform across different cloud providers**
Alexandre Verdet, Mohammad Hamdaqa, Leuson Da Silva, Foutse Khomh

Forum 6

10:30–12:30

Human and Social Aspects: Program Comprehension and Visualization

ASE Research Papers · 8 talks

- 10:30 **Issue-Oriented Agent-Based Framework for Automated Review Comment Generation**
Shuochuan Li, Dong Wang, Patanamon Thongtanunam, Zan Wang, Jiuqiao Yu, Junjie Chen
- 10:45 **ExplainBench: Evaluating Code Explanations from Agents**
Zhiyuan Pan, Sungmin Kang, Imam Nur Bani Yusuf, Abhik Roychoudhury
- 11:00 **Automating Low-Risk Code Review at Meta: RADAR, Risk Calibration, and Review Efficiency**
Audris Mockus, Peter C Rigby, Nachiappan Nagappan
- 11:15 **RepoProbe: Benchmarking Architecture-Aware Repository Comprehension with Checklists**
Yuxi Yang, Alyssa Wu, Ji Luo, Richeng Xuan, Zhichao Hu, Yuhong Liu, Zhen Qin
- 11:30 **We Must Have Missed This Comment: Detecting and Repairing Stale Function References in Linux Kernel Comments**
Kexin Sun, Yunbo Lyu, Xutong Ma, Hongyu Kuang, Ratnadira Widyasari, He Zhang, Xiaoxing Ma, Julia Lawall, David Lo
- 11:45 **How Developer Coreness Influences the Patch-Review Process: A Mixed-Method Study**
Christian Hecht, Thomas Bock, Ralf Ramsauer, Wolfgang Maurer, Sven Apel
- 12:00 **On the Reliability of Code Comprehension Proxies**
Erfan Arvan, Nadeeshan De Silva, Oscar Chaparro, Martin Kellogg
- 12:15 **Same Scrutiny, More Time: Eye Tracking Insights into Reviewing LLM-Labelled Code**
Ranim Khojah, Francisco Gomes de Oliveira Neto, Mazen Mohamad, Julian Frattini, Philipp Leitner

Forum 7

10:30–12:30

Testing and Analysis: Regression and Model-Based Testing 1

ASE Research Papers · 9 talks

- 10:30 **Documentation vs. Code Patterns: What Drives LLM-Based Exception Oracle Generation?**
Soneya Binta Hossain, Matthew B Dwyer, Tasfia Tasnim
- 10:45 **DynCA: An Effective Algorithm for 3-Wise Combinatorial Interaction Testing in Highly Configurable Systems**
Ruizhi Shi, Shuangyu Lyu, Renyu Yang, Wei Wu, Chanjuan Liu, Chunming Hu, Chuan Luo
- 11:00 **Escaping the Self-Repair Trap: Improving Test Oracle Generation via Dual-Context Awareness**
Kefan Li, Hongyue Yu, Yuan Yuan
- 11:15 **JSTestCraft: Addressing Context Deficits in JavaScript Unit Test Generation via Agentic Multi-Level Contextual Analysis**
Yiyang Liu, Yanjie Zhao, Haoyu Wang
- 11:30 **RustGTC: Constraint-Guided Test Generation for Rust APIs via Comment Analysis and Code Context**
Yichi Zhang, Jinhao Huang, Yang Feng
- 11:45 **EvoLLM: High-Coverage, Readable, and Token-Efficient Unit Test Generation for Industrial Java Codebases**
Seokhyeon Moon, Myungho Lee, Jihun Kim, Yoon-Chan Jhi, Jihoon Cho, Hakjoo Oh
- 12:00 **SplitCA: An Effective Tool for Generating 4-wise Covering Arrays for Large-Scale Highly Configurable Systems**
Shuangyu Lyu, Ruizhi Shi, Yuechen Duan, Chunming Hu, Chuan Luo
- 12:10 **PITMuS: A Tool for Automated Bug Dataset Generation via Source-Level Mutant Reconstruction**
Tasfia Tasnim, Soneya Binta Hossain
- 12:20 **UltraInstinctVR: A Model-Based Tool for Automated System Testing of VR Applications**
Gerry Longfils, Maxime Cauz, Arnaud Blouin, Xavier Devroey

Forum 8

10:30–12:30

Maintenance and Evolution: Refactoring and Reengineering 1

ASE Research Papers · 8 talks

- 10:30 If You Build It, Will They Come? Lessons from 20 Years of Industrial Use of a Code Transformation DSL
Johan Fabry, Vadim Zaytsev
- 10:45 PCodeTrans: Translate Decompiled Pseudocode to Trace-Level Equivalent Source Code
Yuxin Cui, Zeyu Gao, Siliang Qin, Yuanda Wang, Jiaming Wang, xizeLuo, Peng Zhou, Chao Zhang
- 11:00 Beyond Text Matching: Towards Reference-Free Evaluation for Human-Oriented Binary Reverse Engineering
Xiuwei Shang, Li Hu, Xiao Jiang, Jieke Shi, Junda He, Zhou Yang, Shaoyin Cheng, Guoqiang Chen, Weiming Zhang, David Lo
- 11:15 The CodeInverter Suite: Structure- and Data-Aware Binary Decompilation with Efficient LLMs
Peipei Liu, Jian Sun, Rongkang Sun, Li Chen, Zhaoteng Yan, Xiaoling Zhang, Dawei Wang, Dapeng Sun, Peizheng Zhang, Dan Li
- 11:30 HyBinMAS: LLM-Powered Hybrid Analysis for Function Name Recovery of Stripped Binaries
Chenxi Wang, Lei Xu, Xincheng He, Gaozhan Yu, Ye Wang
- 11:45 BDiff: Block-Aware and Accurate Text-Based Code Differencing
Yao Lu, Wanwei Liu, Tanghaoran Zhang, Kang Yang, Yang Zhang, Wenyu Xu, Longfei Sun, Xinjun Mao, Shuzheng Gao, Michael Lyu
- 12:00 Needle in the Repo: Diagnosing Maintainability Failures in AI-Generated Repository Edits
Haichao Zhu, Qian Zhang, Jiyuan Wang, Zhaorui Yang, Yuxin Qiu
- 12:15 Code Transformation Rule Synthesis using LLMs: Potential and Limits
Axel Allain, Aymeric Blot, Djamel Eddine Khelladi, Mathieu Acher

Forum 15

10:30–12:30

Security and Other Non-Functional Properties: Reliability and Availability 1

ASE Research Papers · 8 talks

- 10:30 TELLER: Non-intrusive Cross-Layer Root-Cause Analysis for LLM Inference
Ruilin Xu, Junyi Li, Pengfei Chen, Zongxuan Xie
- 10:45 When Does AI Actually Help in Incident Response? Identifying Good First Messages in Cloud Service Incidents
Minghua Ma, Rujia Wang, Chetan Bansal, Saravan Rajmohan, Yingnong Dang, Hongyu Zhang
- 11:00 Expanding Guard Reach: Online Learning-Driven Incremental Kernel Panic Diagnosis
Yicheng Sui, Xiaotian Wang, Hailin Zhang, Shenglin Zhang, Binhong Jiang, Heng Zheng, Yongqian Sun, Yuzhi Zhang, Yang Zhang, Xin Wu, Liangyan Peng, Linlin Han, Feng Wang, Xiaozhou Liu, Yu Zhang
- 11:15 GALA: Graph-Augmented LLM Agents for Root Cause Analysis and Incident Response in Microservices
Yifang Tian, Yaming Liu, Zichun Chong, Zihang Huang, Yiran Li, Hans-Arno Jacobsen
- 11:30 OpsAgent: An Evolving Multi-agent System for Incident Management in Microservices
Yu Luo, Jiamin Jiang, Jingfei Feng, Lei Tao, Qingliang Zhang, Xidao Wen, Yongqian Sun, Shenglin Zhang, Tong Liu, Wenjie Zhang, Dan Pei
- 11:45 ActionNex: A Production-Grade System for Next-Best-Action Recommendation in Cloud Outage Management
Zhenfeng Lin, Haoji Hu, Ming Hao, Xuchao Zhang, Yudi He, Oleg Kulygin, Hatay Tuna, Chetan Bansal, Ze Li, Murali Chintalapati, Sheila Jiang, Salman Zafar, Angie Anderson, Ryan zhang
- 12:00 ReLog: Execution-Aware Logging with Runtime Feedback for LLM-Oriented Debugging
Xin Wang, Yang Feng, Xiaoqian Jiao, Yang Zhang, Zhenhao Li, Zishuo Ding
- 12:15 Towards Reliable Alarm Flood Reduction in Data Center via Grammar-Guided Verifiable LLM Reasoning
Jiacheng Liu, Guocong Lin, Jichao Yan, Huiwen Zheng, Hongli Ma, Changhao Zhou, Zhiyi Yu, Yuanchun Li, Yunxin Liu, Ding Li, Tao Xie, Yao Guo

Forum 16

10:30–12:30

Security and Other Non-Functional Properties: Security and Privacy 1

ASE Research Papers · 9 talks

- 10:30 An Empirical Study of Data Access Practice in Android AR Apps to Understand User Privacy Risks
Sabbir Hussain Meraj, Long Trac, Xiaoyin Wang, Xusheng Xiao, Wei Wang
- 10:45 Software Engineering Meets the EU AI Act: Artifacts, Patterns, and Future Directions
Sridhar Chimalakonda, Karimi Upendra
- 10:55 When Extensions Lie: Large-Scale Study of Privacy Policy Compliance in Browser Extensions
Zilun Wang, Zeyang Zhuang, Mingxue Zhang, Wei Meng, Michael Lyu
- 11:10 VeriSBOM: Secure and Verifiable SBOM Sharing via Zero-Knowledge Proofs
Gianpietro Castiglione, Shahriar Ebrahimi, Narges Khakpour

- 11:25 Evaluating Inference-Time Defenses against Package Hallucination in LLM-Generated Code
Alberick Euraste Djire, Iyiola E. Olatunji, Melissa Tessa, Earl T. Barr, Jacques Klein, Tegawendé F. Bissyandé
- 11:40 Post-quantum Cryptography in the Wild: Assessing the Readiness of Open-Source Ecosystems
Tongxin Yuan, Zhanpeng Liu, Jiashuo Liang, Zhuosheng Zhang, Gongshen Liu, Yang Yu, Guancheng Li
- 11:55 Execution-as-Configuration: Security Smells in Model Configuration Artifacts
Mohammed Latif Siddiq, Prince Noah Johnson, Joanna C. S. Santos
- 12:10 AIGen: Automating AI Bill of Materials Generation through Hybrid MLOps Integration
Federica Pepe, Daniele Bifulco, Costantino Martignetti, Aureliano D'Amici, Fabiano Izzo, Damian Andrew Tamburri, Massimiliano Di Penta
- 12:20 AIRBORNE: Augmented fine-grained software Bill Of materials generator for stack overflow
Sabato Nocera, Massimiliano Di Penta, Simone Romano, Giuseppe Scanniello

12:30–14:00 Lunch Break

Forum 1

14:00–16:00

AI and Software Engineering: AI4SE 2

ASE Research Papers · 8 talks

- 14:00 ExeCRE: Execution-Consistency Guided Reliability Estimation for Self-Correcting Code Generation
Yiru Dong, Richong Zhang, Fanshuang Kong, Si Chen
- 14:15 Translation Tag Team: Formal Rules and LLMs Translate More Macros Together Than Apart
Brent Pappas, Joseph Zalusky, Zachary Burkett, Paul Gazzillo
- 14:30 Regression Accumulation in Multi-turn LLM Programming Conversations
Andie Huang, Lin Ma, Amjed Tahir, Qian Zhang, Liwen Xiao, Lysa Xiao
- 14:45 When Comments Mislead: An Empirical Study of Outdated Comments in LLM-Based Code Completion
Zhikun Xia, Yuan Huang, Xiangping Chen, Zibin Zheng
- 15:00 Towards Reliable C-to-Rust Translation with Rule-Guided Reasoning and Reinforcement Learning
Feng Luo, Jiachen Liu, Cuiyun Gao, jiafeng, Kui Liu
- 15:15 Instruction Alignment for Binary Code Representation Learning
Huaijin Wang, Shuai Wang
- 15:30 Post-hoc Attention Steering of Large Language Models for Robust Code Understanding under Obfuscation
Xiaokai Rong, Aashish Yadavally, Tien N. Nguyen
- 15:45 Lost in Translation, Found in Summary: Summary-Driven Supervision for Cross-Language Code Retrieval
Zhijie Liu, Yisi Liu, Linxi Jiang, Yutian Tang, Lili Xiao, Zhipeng Xue

Forum 2

14:00–14:15

Testing and Analysis: Regression and Model-Based Testing 2

ASE Research Papers · 1 talk

- 14:00 Piece by Piece: Automating Combination Interaction GUI Testing via Planning and Dual Memory
Zhe Liu, Oujin Wang, Junjie Wang, Chunyang Chen, Mengzhuo Chen, Boyu Wu, Yuekai Huang, Qing Wang

Forum 2

14:15–14:30

Testing and Analysis: Regression and Model-Based Testing 4

ASE Research Papers · 1 talk

- 14:15 AgentExecutor: Partial Code Execution via Agentic Context Generation
Junkai Chen, Chengran Yang, Xing Hu, Zhenhao Li, Xin Xia, David Lo

Forum 2

14:30–14:45

Testing and Analysis: Regression and Model-Based Testing 2

ASE Research Papers · 1 talk

- 14:30 Evaluating the Maintainability of ViewModel Tests using Multi-modal DSLs: An Industrial Mixed-Methods Study
Mario Fuksa, Sandro Speth, Steffen Becker

Forum 2

14:45–14:55

Human and Social Aspects: Human-Computer Interface

ASE Research Papers · 1 talk

14:45 AutoFLC: An LLM-Based Flowchart Generation Tool for Legacy Code Repository

Yixing Luo, Longxing Yang, Xiaofeng Li, Bin Gu, Zhi Jin

Forum 2

15:00–16:00

Testing and Analysis: Regression and Model-Based Testing 2

ASE Research Papers · 4 talks

15:00 Using Large Language Models to Test Voice User Interfaces

Emanuela Guglielmi, Angelica Spina, Gabriele Bavota, Rocco Oliveto, Simone Scalabrino

15:15 VUI Testing of VPA Apps via Behavior Model-Enhanced LLM Agents

Suwan Li, Lei Bu, Shangqing Liu, Guangdong Bai, Fuman Xie, Kai Chen, Chang Yue

15:30 Doc2Test: An Industrial Multi-agent Framework for Document-Driven End-to-End Selenium Test Generation

Andrea Lops, Antonio Ferrara, Fedelucio Narducci, Azzurra Ragone, Michelantonio Trizio

15:45 Model-Guided Property-Based Testing of WeChat Pay at Billion-User Scale

Xiangchen Shen, Yiting Wang, Ting Su, Jingjing Liang, Jingling Sun, Xixian Liang, Haiying Sun, Xinjie Xu, Haochuan Lu, Yuetang Deng, Pengcheng Wang, Geguang Pu, Zhendong Su, John Hughes

Forum 6

14:00–16:00

Security and Other Non-Functional Properties: Security and Privacy 2

ASE Research Papers · 8 talks

14:00 Mind the Gap: Do Widely-Used SAST Tools Really Cover the Language-Specific Top CWEs?

zhilin li, Yiran Nie, Xianglong Qi, Jiwei Yan, Jun Yan

14:15 Perception, Reasoning, and Action: A ReAct-Based Vulnerability Detection Framework with Repository-Level Reasoning

Xin-Cheng Wen, Hanyue Luo, Cuiyun Gao, Hao Zhu, Yang Xiao, Ge Li

14:30 One Is Not Enough: The Untold Story of Multiple Security Patches for One Vulnerability

Fangyuan Zhang, Lyuye Zhang, Lingling Fan, Chengwei Liu, Yinan Li, Liang Huang, Yang Liu, Zheli Liu, Sen Chen

14:45 Are Reasoning Traces Faithful in LLM-Based Code Vulnerability Detection?

Wang Jialai, Juncheng Lu, Songtao Yang, Bingyi He, Zeyu Gao, Chao Zhang, Shuai Wang

15:00 Harnessing Uncertainty in Code Language Models: Lessons from Vulnerability Detection

Haodong Li, Xiao Cheng, Xudong Wang, Zhihao Guo, Haoyu Wang

15:15 UntrustVul: Automated Untrustworthy Alert Identification in Vulnerability Detection Models

Lam Nguyen Tung, Xiaoning Du, Neelofar Neelofar, Aldeida Aleti

15:30 No Isolated Evolution, No Blind Search: Memory-Guided Prompt Co-evolution for Industrial Multi-agent Vulnerability Detection

Jie Xu, Zihan Wu, Yun Peng, Chun Yong Chong, Linghan Meng, Kui Liu, Xiaohua Jia

15:45 AgenticSCR: An Autonomous Agentic Secure Code Review for Immature Vulnerabilities Detection

Wachiraphan (Ping) Charoenwet, Kla Tantithamthavorn, Patanamon Thongtanunam, Hong Yi Lin, Minwoo Jeong, Ming Wu

Forum 7

14:00–16:00

Security and Other Non-Functional Properties: Performance 1

ASE Industry Showcase · 8 talks

14:00 Stop When It Matters: Detectability-Guided Microbenchmarking for Performance Regression Testing

Zongxiong Chen, Derui Zhu, Jinfu Chen, Kundi Yao, Alexander Pretschner, Weiyi Shang, Manfred Hauswirth, Sonja Schimmler

14:15 Multi-level Code Optimization via Mixture of Prompts

Yun Peng, Jun Wan, Jiakun Liu, Shuzheng Gao, David Lo, Xiaoxue Ren

14:30 Less Is More: Tuning Configurable Systems with Imperfect Fidelity

Yulong Ye, Miqing Li, Tao Chen

14:45 Characterizing and Mitigating Productivity Loss in Large-Scale Model Training: An Empirical Study

Dinghao Xue, Yanxiang Chen, Junhong Liu, Yepeng Zhang, Guangsen Ni, Menghao Zhang, Tianyu Wo, Zheng Zheng, Chunming Hu, Jin Ouyang, Renyu Yang

15:00 Understanding the Performance-Effectiveness Trade-Offs of AddressSanitizer in Real-World Programs: An Empirical Study

Jiaxin Xi, Cheng Wen, Zhiwei Lin, Yuandao Cai, Jie Su, Bin Yu, Zhuohua Li, Shengchao Qin

15:15 Li-WaC: A Lightweight WebAssembly Container for Independently Upgradable IoT Applications in Xiaomi Vela

Bin Liao, Canfeng Zheng, Shiwen Shan, Qi Huang, Junjie Dong, Xiang Xiao, Yaoyao Gu, Yuxin Su

15:30 DBRepro: Automated Database Synthesis via a Hybrid Constraint-Solving Approach for Reproducing Slow Queries

zhaoyang zhang, Shuang Liu, Dengfeng Xu, Wei Lu, Jianquan Leng, Sheng Du, Xiaoyong Du

15:45 WiseSword: An Automated Computation Graph Partitioning Tool for Industrial DLRM Deployment on Ascend NPU

Ke Cheng, Yuhan Zhang, Bo Pan, Jiawei Zhu, Qidong Zhang, Yuting Yan, Zhengyong Zhang, Yibo Jin, Sheng Zhang

Forum 8

14:00–16:00

Testing and Analysis: Automated Program Repair and Synthesis 1

ASE Research Papers · 9 talks

- 14:00 RegoMender: A Multi-agent Framework for Automated Rego Policy Repair
Ruijie Zhang, Huaxiao Yin, Li Hu, Fanqi Kong, Haozhe Liang, Shaoyin Cheng, Kun Zhang, Bibo Tu
- 14:15 DesignFix: Context-Augmented Repair of Design Rule Violations
Mainul Hossain, Thomas D. LaToza
- 14:30 ArkEval: Benchmarking and Evaluating Automated Code Repair for ArkTS
bang xie, Senjian Zhang, Zhiyuan Peng, Wei Chen, Xin Yin, Chenhao Ying, Yuan Luo
- 14:45 LaMAR: Latent Multi-agent Collaboration via KV-Cache Communication for Automated Program Repair
Rui Jiao, Yue Zhang, Zhexiong Feng, Jinku Li
- 15:00 From Security Findings to Verified Fixes: Automated Repair of Infrastructure-as-Code Misconfigurations
Ying Li, Tamjid Rahat, Yingjun Lyu, Omer Tripp, George Argyros
- 15:15 Catching Developers in the Flow: Low-Latency Agentic Program Repair at Google Scale
Celal Ziftci, Spencer Greene, Ray Liu, Livio Dalloro, Lorenzo Dini
- 15:30 TraceGate: A Framework for Policy-Controlled Failure Evidence in LLM-Assisted Program Repair
Nicolas Schuler, Vincenzo Scotti, Raffaella Mirandola
- 15:40 Is Three the Magic Number? An Empirical Evaluation of LLM-Based Repair Loops
Tobias Kiecker, Eik Reichmann, Hosung Kang, Gabin An, Lars Grunske
- 15:50 ReproBreak: A Dataset of Reproducible Web Locator Breaks
Thiago Santos de Moura, Leon Adamietz, Samra Mehboob, Yannic Noller

Forum 15

14:00–16:00

Formal Aspects of Software Engineering: Specification Languages 1

ASE Research Papers · 9 talks

- 14:00 RFC2TLA+: Extracting and Verifying Formal Models from RFC Documents using Continuous LLM Feedback
Guozhen Ding, Kexin Li, Ilya Grishchenko, David Lie
- 14:15 Exploring Development Methods for Reactive Synthesis Specifications
Dor Ma'ayan, Shahar Maoz, Jan Oliver Ringert
- 14:30 TLCrafter: Peeling the Layers of Natural Language to Reveal Temporal Logic Specifications
Cheng Wen, Wenjun Ke, Xiao Liang, Zhi Ma, Xiaoyong Xue, Shengchao Qin, Cong Tian
- 14:40 ACSLBench: A Verified C/ACSL Corpus and Benchmark with Compositional Call Chains for Formal Specification Synthesis
Hu Junjie, Cheng Wen, Bin Yu, Jialun Cao, Dugang Liu, Weidi Sun, Haokun Li, Shengchao Qin, Cong Tian
- 14:50 Can Formal Specifications Be Synthesized from Tests Alone?
Tianhai Liu, Maximilian Müller, Tobias Hey, Vitus Lüntzel, Muhammad Minhas, Anne Koziolk, Bernhard Beckert
- 15:00 Formula Completion Suggestions for Alloy Models with Selection Guidance
Mohammad Nurullah Patwary, Allison Sullivan
- 15:15 Beyond Benchmarks: A Case Study of LLM-Generated Verus Specification Failures on Asterinas Vostd
Dugang Liu, Ding Wang, Chao Jiang, Cheng Wen, Zhiwu Xu, Zhong Ming, Yuwei Liu, Xinyi Wan, Lin Huang
- 15:30 All or Nothing: How Library Type Annotations Affect Client-Side Errors in Python
Eric Asare, Luca Di Grazia, Sarah Nadi
- 15:45 MiniPoly: Automatic Extraction of Efficient JavaScript Polyfill from Language Specification
Jungwoong Kim, Youngmin Cho, Jihyeok Park

Forum 16

14:00–16:00

Maintenance and Evolution: API Design and Management

ASE Research Papers · 8 talks

- 14:00 Understanding the Code and Architectural Smells Generated by LLMs and Autonomous Agents
Yuecai Zhu, Nikolaos Tsantalis, Peter Rigby
- 14:15 ABACI: Automated Downstream Bug Resolution via Semantic-Aware Bisection and Convergent Backporting
Runzhe Wang, Yuanhang Zhou, Qingming Su, Yong Yang, Xiaohai Shi, Tao Ma
- 14:30 HyTri: Hybrid Triage of CI Failures via LLM Guided Semantic Reasoning and Change Attribution
Lior Broide, Argaman Mordoch, Roni Stern

- 14:45 **Beyond Syntactic Differencing: Semantic-Aware Scheduling for Cross-Configuration Incremental Builds**
Xiaolong Dai, Shanshan Li, Jun Lyu, Xintao Niu
- 15:00 **Implicit, Yet Impactful: Understanding Hidden Dependencies in Java Projects**
Lyuye Zhang, Chengwei Liu, Fangyuan Zhang, Yiran Zhang, Yuan Zhou, Yang Liu
- 15:15 **DepUpgrade: Automating Dependency Upgrade through State-Path Exploration**
Yifan An, Xiangxi Ma, Wentong Tian, Xuanqi Wang, Qingao Dong, Xiang Gao, Hailong Sun
- 15:30 **DAGGER: Dynamic Agentic Graph-Guided Efficient Reasoner for Container Image Dependency Optimization**
Shikhar Vashistha, Neetesh Kumar
- 15:45 **Byam: Fixing Breaking Dependency Updates with Large Language Models**
Frank Reyes Garcia, May Mahmoud, Federico Bono, Sarah Nadi, Benoit Baudry, Martin Monperrus

16:00–16:30 *Coffee Break*

Forum 1

16:30–17:30

AI and Software Engineering: SE4AI 1

ASE Research Papers · 4 talks

- 16:30 **Tensor-Based Batch Fuzzing with Adaptive Perturbation Scaling for Deep Neural Networks**
Guanqin Zhang, Yulei Sui
- 16:45 **AgentChaos: Chaos Engineering for Agent Systems via Programmatic Fault Injection**
Gou Tan, Zhensu Sun, Jieke Shi, Ting Zhang, Zilong He, Qingfu Wu, Shuai Liang, Weifeng Sun, Junda He, Pengfei Chen, Chuanfu Zhang, Lwin Khin Shar, David Lo
- 17:00 **How Composite Metamorphic Relations Enhance Test Effectiveness of DNN Testing: An Empirical Study**
Huayao Wu, Peng Wang, Shengyou Hu, Xintao Niu, Changhai Nie, Tsong Yueh Chen
- 17:15 **When Ambiguity Meets Atypicality: Dual-Perspective Test Input Prioritization for DNNs**
Haoran Li, Shihai Wang, Bin Liu, Jialuo Chen, Wenjing Zhu, Yu Liu, Tengfei Shi, Shudi Guo

Forum 2

16:30–17:30

Testing and Analysis: Fuzzing and Mutation Testing 1

ASE Research Papers · 4 talks

- 16:30 **TDiFf: Detecting Bugs in DataFrame Systems via Transferred DBMS Test Cases**
Jiaxin Hu, Shaowei Chen, Rongxin Wu
- 16:45 **Fuzz4DB: A Practice of LLM-Agent-Guided Fuzzing for Database Feature-Level Delta Testing**
Chunling Qin, Yong Hu, Xiao Zhang, Jinchuan Chen, Fangtao Gu, Baoxun Wang, Yuxing Chen, Anqun Pan, Lixiong Zheng
- 17:00 **KQFuzz: Knowledge-Guided Fuzzing for Quantum Libraries via Large Language Models**
Fuyuan Xia, Qixin Zhang, Chenhao Ying, Haojin Zhu, Shuai Wang, Yuan Luo, Pingchuan Ma, Yuxuan Du
- 17:15 **Finding Typing-Related Bugs in JVM Implementations**
Yanzhou Lu, Junjie Chen, Yingquan Zhao, Tianchang Gao, Zan Wang

Forum 6

16:30–17:30

Software Analytics: Software Metrics

ASE NIER · 5 talks

- 16:30 **Less Noise, More Signal: the DRR Effect for Better Optimizations of a Range of SE Tasks**
Andre Lustosa, Tim Menzies
- 16:45 **How Well Do LLMs Generate Taxonomies in the SE Domain? A Multi-perspective Evaluation Framework**
Sota Nakashima, Yuta Ishimoto, Masanari Kondo, Tao Xiao, Yasutaka Kamei
- 17:00 **The Ground Is Shifting: A Reflection on the Foundations of Software Measurement**
Thomas Bock, Audris Mockus, Bogdan Vasilescu
- 17:10 **ECLAIR: A Causally-Grounded AI Framework for Scientific Discovery in Empirical Software Engineering**
Alejandro Velasco, Daniel Rodriguez-Cardenas, Dipin Khatri, David Nader Palacio, Denys Poshyvanyk
- 17:20 **When Elo Lies: Hidden Biases in Codeforces-Based Evaluation of Large Language Models in Practice**
Shenyu Zheng, Ximing Dong, Xiaoshuang Liu, Chun Yong Chong, Gustavo A. Oliva, Dayi Lin, Boyuan Chen, Shaowei Wang, Ahmed E. Hassan

Forum 7

16:30–17:30

Security and Other Non-Functional Properties: Reliability and Availability 2

ASE Research Papers · 4 talks

- 16:30 **ARMOR: A Robust Self-Supervised Framework for Root Cause Analysis in Microservices under Missing Modality**
Wenzhuo Qian, Hailiang Zhao, Ziqi Wang, Zhipeng Gao, Jiayi Chen, Zhiwei Ling, Shuiguang Deng
- 16:45 **Source-Free Detection and Impact Analysis of Compiler Optimization Problems in Mobile Applications**
Han Hu, Xiaoheng Xie, Bo Sun, Jian Gu, Gang Fan, Li Li
- 17:00 **CAPMix: Robust KPI Anomaly Detection for AIOps in Noisy and Dynamic Environments**
Xudong Mou, Rui Wang, Tiejun Wang, Zexin Wu, Fangda Guo, Jie Sun, Shiru Chen, Penghao Zhang, Tiezi Zhang, Tianyu Wo, Hao Peng, Chunming Hu, Xudong Liu, Renyu Yang
- 17:15 **OpsLens: Runbook-Free Root Cause Analysis via Multi-agent Collaboration and Tool Management**
Haiyu Huang, Jiewei Lyu, Gou Tan, Pengfei Chen, Zhihan Jiang, Guangba Yu, Penghui Mi, Fuqiang Wang, Jian Han, Chongkang Tan, Michael Lyu

Forum 8

16:30–17:30

Human and Social Aspects: Software Engineering Processes 1

ASE Tools and Datasets · 6 talks

- 16:30 **FlightRecorder: Tracing Fine-Grained Development with AI Assistants**
Tobias Dick, Benedict Bliem, Kallistos Weis, Alisa Carla Welter, Sven Apel
- 16:40 **Do Not Copy/Paste: Soft Barriers for Copying in AI-Assisted Programming**
Iyiola E. Olatunji, Alberick Euraste Djire, Jacques Klein, Tegawendé F. Bissyandé
- 16:50 **Prompt Coach: An Empirical Evaluation of an Agentic Tutor for Learning Prompt Engineering in Software Development**
Rohit Mehra, Kapil Singi, Vikrant Kaulgud, Vibhu Saujanya Sharma, Swapnajeet Choudhury, Swati Sharma, Adam P. Burden, Majd Sakr
- 17:00 **Agents That Teach: Towards Designing Incidental Learning Back into AI-Assisted Software Development**
Rohit Mehra, Samdyuti Suri, Prithviraj K. Tagadinamani, Kapil Singi, Vikrant Kaulgud, Adam P. Burden
- 17:10 **ScratchLog+: Live Analytics of Learners' Exercise Progress**
Benedikt Fein, Gordon Fraser
- 17:20 **VidTutorAssistant: Automating Responses to Programming Tutorial Questions**
Ahmad Tayeb, Sonia Haiduc, Mohammad D. Alahmadi

Forum 15

16:30–17:30

Requirements and Design: Software Architecture and Design

ASE Industry Showcase · 5 talks

- 16:30 **MicroAgent: Context-Augmented Multi-agent Framework for Automatic Microservice Decomposition**
Zishan Su, Junjie Huang, Shiwen Shan, chenxingyan, Hui Zeng, Yuxin Su, Yanlin Wang, Michael Lyu
- 16:45 **MicroWeaver: A Constraint-Guided Structure–Semantic Aware Approach to Automated Microservice Decomposition**
Huitian Xue, Kairui Zhang, Chenxi Zhang, Zheng Li, Chuan Luo, Yilong Yang, Tianyu Wo, Chunming Hu, Yanliang Li, Zhou Yang, Penghao Zhang, Tiezi Zhang, Renyu Yang
- 17:00 **The ARDoCo Tool Landscape: REST API, TraceView, and TraceViz for Architecture Traceability**
Jan Keim, Dominik Fuchß, Sophie Corallo, Tobias Hey, Julian Winter, Kevin Feichtinger
- 17:10 **DAL: Decoupled Aspect-Based Layering of I/O Processing for Automotive Software**
Muhammet USLU, Atif Imran Penkar, Robert Rasche, Olaf Spinczyk, Marc Fessler, Michael Unterreiner
- 17:20 **CAS2UML: A Handwritten Sketch-to-PlantUML Dataset for Class and Activity Diagrams**
Simon Scholz, Mersedeh Sadeghi

Forum 16

16:30–17:30

Security and Other Non-Functional Properties: Security and Privacy 3

ASE Research Papers · 5 talks

- 16:30 **Assessing the Cross-Version Applicability of Java Library Vulnerability Exploits**
Zirui Chen, Qi Zhan, Jiayuan Zhou, Xing Hu, Xin Xia, Xiaohu Yang
- 16:45 **RAVEN: Measuring Detection Instability in Anti-virus Engines under Semantics-Preserving Perturbations**
Younghoon Ban, Beomjin Jin, Doowon Kim, Hyoungshick Kim, Haehyun Cho
- 17:00 **JavaVulBench: A Java Vulnerability Benchmark with Realistic Splits, a Unified Multi-backend Harness, and a Leakage-Aware Evaluation Mode**
Norbert Szolnoki Sándor, Gabor Antal
- 17:10 **AgentSysFilter: A Multi-agent Analysis Approach for Scenario-Aware Syscall Limitation**
Dongyang Zhan, Xingchen Yan, Zhaofeng Yu, Yangfan Guo, tan kai, Langtong Zhang

17:20 **Vulnerability Tracking using Normalized Scope+Offset**
Julian Thome, Hua Yan, Lucas Charles, Craig Smith, Jason Leasure

17:30–18:00 *Coffee Break*

18:00–21:30 Conferene Reception
ASE Catering · Plenary Room (Forum 1-3)

Wednesday

14 October 2026

119 talks

09:00–10:00 Beyond the Hype: What AI Really Did to Software Engineering

Joost Noppen

ASE Keynotes · Plenary Room (Forum 1-3)

10:00–10:30 Coffee Break

Forum 1

10:30–12:30

AI and Software Engineering: AI4SE 3

ASE Research Papers · 9 talks

- 10:30 PAIChecker: Uncovering and Checking PR-Issue Misalignment in SWE-Bench-Like Benchmarks
Manyi Wang, Junjielong Xu, Pinjia He
- 10:45 REAP: Automatic Curation of Coding Agent Benchmarks from Interactive Production Usage
Smriti Jha, Matteo Paltenghi, Chandra Sekhar Maddila, Vijayaraghavan Murali, Shubham Ugare, Satish Chandra
- 11:00 SkillMOO: Multi-objective Optimization of Agent Skills for Software Engineering
Jingzhi Gong, Ruizhen Gu, Zhiwei Fei, Yazhuo Cao, Lukas Twist, Alina Geiger, Shuo Han, Dominik Sobania, Federica Sarro, Jie M. Zhang
- 11:10 RELY-SE: A Dashboard for Measuring Operational Reliability of AI Coding Agents
Akilesh P, Shivadharshan Sankar, Sridhar Chimalakonda, Vibhu Saujanya Sharma, Vikrant Kaulgud
- 11:20 Automated Table Reproduction via Code Generation
Doehyun Baek, Michael Pradel
- 11:35 Free Is More: Empowering Software Agents via Issue-Free Trajectory Learning and Entropy-Aware RLVR Training
Xin-Cheng Wen, Binbin Chen, Haoxuan Lan, Hang Yu, Peng Di, Cuiyun Gao
- 11:50 LLM Agents Can See Code Repositories
Dongjian Ma, Silin Chen, yufei yang, Yuling Shi, Yanfu Yan, Xiaodong Gu
- 12:05 ReCodeAgent: A Multi-agent Workflow for Language-Agnostic Translation and Validation of Large-Scale Repositories
Ali Reza Ibrahimzada, Brandon Paulsen, Daniel Kroening, Reyhaneh Jabbarvand
- 12:20 K-COBAGent: A Knowledge-Driven Multi-agent Framework for COBOL Modernization
Kavyasri Gajula, Phanindra Kodela, Sridhar Chimalakonda

Forum 2

10:30–12:30

Testing and Analysis: Automated Program Repair and Synthesis 2

ASE Research Papers · 8 talks

- 10:30 Kozuchi Agent: A Language-Agnostic Open-Weight Agent for Software Repair
Mehdi Bahrami, Kosaku Kimura, Satoshi Munakata, Satoshi Nakashima, Yu Ishikawa, Kosuke Maeda, Nao Soma, Kenichi Kobayashi, Keisuke Miyazaki, Keizo Kato, Shigeki Fukuta, Tatsuo Kumano, Nobutaka Imamura, Kevin Musgrave, Shahbaz Abdul Khader, Kwun Ho Ngan, Joe Townsend, Fayas Asharindavida, Matthieu Parizy, Akira Sakai, Yuma Ichikawa, Yang Zhao, Michiaki Takizawa, Taku Fukui, Hiroki Ohtsuji, Wei-Peng Chen, Hiro Kobashi
- 10:45 InfraAPR: An Agent-Based Repair Framework for Cloud Infrastructure-as-Code Programs
Quan Chen, Zhao Liu, Chunfang Li, YiFan Zhang, Ning Hu, Qing Liao
- 11:00 MultiFixer: A Coordinator-Proposer Based Multi-agent Framework for Fixing Multi-hunk Bugs
hai chuan hu, Chunrong Fang, ye shang, Jiawei Liu, Weifeng Sun, Guoqing Xie, Chenxing Zhong, Qianjun Zhang
- 11:15 From Guessing to Seeing: Enhancing LLM-Based Program Repair via Trace-Guided Multi-strategy Debate
Jiaqing Wu, Tong Wu, Manqing Zhang, Yunwei Dong, Bo Shen
- 11:30 Probe to Generate: Program Variant-Guided Test Augmentation for Repository-Level Repair Benchmarks
Chenglin Li, Yisen Xu, Zehao Wang, Shin Hwei Tan, Tse-Hsun (Peter) Chen
- 11:45 A11YRepair: Bridging Web Accessibility Barriers via Knowledge-Enhanced Divide-and-Conquer Repair
Kai Huang, Ling Zhu, Jian Zhang, Xiaofei Xie, Chunyang Chen

- 12:00 PrevaRank: Ranking Plausible Patches by Historic Feature Frequencies
Shifat Sahariar Bhuiyan, Abhishek Tiwari, Yu Pei, Carlo A. Furia
- 12:15 PSearch: Search-Based Patch Generation in the Era of LLM-Based Automated Program Repair
hai chuan hu, ye shang, Weifeng Sun, Qunjun Zhang

Forum 6

10:30–12:30

Security and Other Non-Functional Properties: Security and Privacy 4

ASE Research Papers · 8 talks

- 10:30 GraphQLer: Enhancing GraphQL Security with Context-Aware API Testing
Omar Tsai, Jianing Li, Tsz Tung Cheung, Lejing Huang, Hao Zhu, Jianrui Xiao, Iman Sharafaldin, Mohammad A. Tayebi
- 10:45 When Compression Becomes an Attack Surface: Black-Box Attacks on Prompt-Compressed LLM Agents
Zesen Liu, Zhixiang Zhang, Yuchong Xie, Dongdong She
- 11:00 Self-Discovering Security Oracles: Meta-learning Vulnerability Detection Strategies through Adversarial Self-Play
Md Rashedul Hasan, Hamid Bagheri
- 11:15 To Think or Not to Think: Evaluating LLM Reasoning and Agents in Vulnerability Detection
Hua Beng Tan, Ratnadira Widyasari, Merve Astekin, Arda Goknil, Hasan Sozer, Yan Naing Tun, Erik Johannes Husom, Lwin Khin Shar, Lingxiao Jiang
- 11:30 Decompiling the Web: Automated Semantic Recovery of Post-compilation Abstraction Leaks
Shriyans Sudhi, Yinxi Liu
- 11:45 Context-Enhanced Vulnerability Detection Based on Large Language Models
Yixin Yang, Bowen Xu, Xiang Gao, Hailong Sun
- 12:00 CLEAR: Causal Context-Based Agentic Reasoning for Vulnerability Detection
Sungju Yun, Sijune hwang, Yeonjoon Lee, Kyungtae Kang, Sungbin Park
- 12:15 Learning to Triage Vulnerability Reports from Program Analysis: An Empirical Study in Node.js
Ronghao Ni, Aidan Z.H. Yang, Min-Chien Hsu, Nuno Sabino, Limin Jia, Ruben Martins, Darion Cassel, Kevin Cheang

Forum 7

10:30–12:30

Maintenance and Evolution: Refactoring and Reengineering 2

ASE Research Papers · 9 talks

- 10:30 Structure-Guided Semantic Deobfuscation of Android Identifiers
Mingyang Chen, Zhentao Xie, Yaqi Gao, Shuang Li, Wenrui Diao
- 10:45 Latent Reuse in Agent Skills: Multi-modal Clone Detection at Ecosystem Scale
Jiaying Zhu, Lyuye Zhang, Wenbo Guo, Yang Liu
- 11:00 On the Effects of Customized Configurations of Static Code Analysis Tools: A Prospective Cohort Study of SonarQube Cloud
Sabato Nocera, Sira Vegas, Giuseppe Scanniello
- 11:15 A Unified Model for Cross-Domain Clone Detection via Model Merging
Palash Ranjan Roy, Banani Roy, Kevin Schneider, Chanchal K. Roy
- 11:30 LLM-Based Agents for Identifying Bug-Introducing Commits
Niklas Risse, Marcel Böhme
- 11:45 A Smart Assistant for Debugging
Steven P. Reiss, Mohamed Salem
- 11:55 LintCFG: DSL-Driven Linter Configuration Generation for Coding Standards
Yi Tang, zejun zhang, Zhenchang Xing, Minxue Pan, Tian Zhang, Xuandong Li
- 12:05 SmellCC: A Tool for Automated Code Smells Remediation
Xiaoting Zhang, Yujie Zhang, Zhipeng Gao, Xing Hu, Xin Xia
- 12:15 Smelly-shot is All You Need: An Empirical Study to Compare In-Context Learning versus Fine-Tuning for Code Smell Detection
Nawaf Alomari, Moussa Redah, Ahmad Ashraf, Mohammad Alshayeb

Forum 8

10:30–12:30

Human and Social Aspects: Software Engineering Processes 2

ASE Research Papers · 7 talks

- 10:30 A Longitudinal Analysis of Software Development Evolution in a Financial Organization
Wentao Chen, Zehong Zhang, Huiqun Yu, Guisheng Fan, Huan Wang, Kaiwen Zhi
- 10:45 Fragmented Markers, Mixed Results: A Systematic Review of AI Coding Assistants and Developer Productivity
Annemarie Wittig, Alina Mailach, Norbert Siegmund

- 11:00 **Command Coding: A Methodological Framework for AI-Accelerated Software Development**
Thomas S. Villani, Martin Kellogg
- 11:15 **To Ban or Not to Ban? How Open Source Projects Govern GenAI Contributions**
Wenhao Yang, Runzhi He, Minghui Zhou
- 11:45 **"Should I Give Up Now?" Investigating LLM Pitfalls in Software Engineering**
Jiessie Tie, Bingsheng Yao, Tianshi Li, Hongbo Fang, Syed Ishtiaque Ahmed, Dakuo Wang, Shurui Zhou
- 12:00 **A Low-Cost Human-in-the-Loop Investigation of Toxicity on GitHub at Scale**
Rahat Rizvi Rahman, Mia Mohammad Imran, Kostadin Damevski
- 12:15 **Programming by Chat: A Large-Scale Behavioral Analysis of 11,579 Real-World AI-Assisted IDE Sessions**
Ningzhi Tang, Chaoran Chen, Zihan Fang, Gelei Xu, Maria Dhakal, Yiyu Shi, Collin McMillan, Yu Huang, Toby Jia-Jun Li

Forum 15

10:30–12:30

Requirements and Design: Requirements and Traceability

ASE Research Papers · 9 talks

- 10:30 **Can Large Language Models Decompose User Stories into Tasks? Exploring the Role of Prompting Strategies and Models**
Delina Ly, Fatma Başak Aydemir, Fabiano Dalpiaz
- 10:45 **Bridging Stakeholder and Product Requirements: An Empirical Study of Requirement Engineering in the Automotive Industry**
Zixu Wang, Shengcheng Yu, Zhenchang Xing, Tobias Wenzel, Chunyang Chen
- 11:00 **From Backlog Items to Security Guidance: Towards Continuous Security Compliance**
Ignacio García Núñez, Florian Angermeir, Fabiola Moyón
- 11:15 **A Scalable Rule-Based Deep Reinforcement Learning Framework for the Next Release Problem**
Hemant Gudaparthi, Nan Niu, Prem Swaroopaanda Ramalingam, Laxman Chowdary Kakara
- 11:30 **Neuro-symbolic Requirements Elicitation: Utilizing Formal Verification Counterexamples as Contextual Prompts**
Okba Tibermacine, Chouki Tibermacine
- 11:45 **ReqEvolve: User-Oriented Software Self-Evolution through Automatic Requirement Interpretation**
Md Asif Iqbal Fahim, Alessio Ferrari
- 12:00 **Specifying the Delegated-Autonomy Boundary: Requirements Engineering for Agentic AI**
Chetan Arora, Andreas Vogelsang, Abbi Sharma
- 12:10 **End-to-End Regulatory Compliance Evolution Pipeline: From Legal Requirements to Developer Tests**
Carolin Brandt, Sallam Abualhaija
- 12:20 **Agent-HAnS: Maintaining Feature and Code Traceability in Agentic Development**
Johan Martinson, Jacob Bengtsson, Moltas Hultin, Kevin Hermann, Thorsten Berger

Forum 16

10:30–12:30

Testing and Analysis: Regression and Model-Based Testing 3

ASE Research Papers · 9 talks

- 10:30 **Scenario-Guided LLM-based Mobile App GUI Testing**
Shengcheng Yu, Yuchen Ling, Chunrong Fang, Quan Zhou, Yi Zhao, Chunyang Chen, Shaomin Zhu, Zhenyu Chen
- 10:45 **Repository-Aware Metamorphic Relation Generation for Augmented Reality Applications using Large Language Models**
Dibyendu Brinto Bose, Jiawei Qin, Chris Brown
- 11:00 **PropGen: Automated Property Generation for Property-Based Testing of Mobile Apps**
Yiheng Xiong, Shiwen Song, Bo Ma, Ting Su, Xiaofei Xie
- 11:15 **DragonCrawl: A Generative, Intent-Based Framework for Scalable Mobile End-to-End Testing**
Sowjanya Puligadda, Mengdie Zhang, Ali Zamani, Dhruva Dixith Kurra, Eric Chen, Juan Marcano
- 11:30 **Testing IFrame Rendering in Web Browsers**
Dahyeon Park, Yuseok Jeon, Seongil Wi, Mijung Kim
- 11:45 **LDMDroid: Leveraging LLMs for Detecting Data Manipulation Errors in Android Apps**
Xiangyang Xiao, Huaxun Huang, Rongxin Wu
- 12:00 **UIXPOSE: Screen-Conditioned Runtime Discrepancy Analysis for Security-Aware Android GUI Exploration**
Amirmohammad Pasdar, Thuan Pham, Toby Murray
- 12:10 **E-CoDrive: A Co-Simulation Framework for Testing Energy-Critical Driving Scenarios**
Manfredi Napolitano, Alessandra Somma, Alessio Gambi, Andrea Stocco, Nicola Mazzocca
- 12:20 **Cam2Sim: Neural Scenario Reconstruction for Closed-Loop Autonomous Driving Simulation**
Davide Jannussi, Stefano Carlo Lambertenghi, Constantin Carste, Andrea Stocco

12:30–14:00 *Lunch Break***Forum 1**

14:00–16:00

AI and Software Engineering: SE4AI 2

ASE Research Papers · 8 talks

- 14:00 **Benchmarking Contextual Understanding for In-Car Conversational Systems**
Philipp Habicht, Lev Sorokin, Abdullah Saydemir, Ken Friedl, Andrea Stocco
- 14:15 **Understanding Bugs in Modern Agentic Frameworks: A Study of Symptoms, Root Causes, and Triggering Conditions**
Xiaowen Zhang, Hannuo Zhang, Shin Hwei Tan
- 14:30 **DiffGAN: A Test Generation Approach for Differential Testing of Deep Neural Networks for Image Analysis**
Zohreh Aghababaeian, Manel Abdellatif, Lionel Briand, Ramesh S
- 14:45 **Search-Based Testing of Vision Language Models for In-Car Scene Understanding**
Lev Sorokin, Chen Yang, Ken Friedl, Andrea Stocco
- 15:00 **When Knowledge Changes: Metamorphic Testing of RAG Systems with Mutations**
Jinhan Kim, Samuele Pasini, Paolo Tonella
- 15:15 **Automated Assertion Generation and Regression Testing for Machine Learning Notebooks**
Yingao (Elaine) Yao, Vedant Nimje, Varun Viswanath, Saikat Dutta
- 15:30 **False-Positive Bug Reports in Deep Learning Compilers: Stages, Root Causes, and Mitigation**
Lili Huang, Qingchao Shen, Dong Wang, Yunping Wu, Meng Wang, Junjie Chen
- 15:45 **CIPHunter: Detecting Configuration-Induced Prediction Instability in Deep Learning Frameworks**
Yanzhou Mu, Shuo Meng, Mijung Kim, Xiang Chen, Chunrong Fang, Zhenyu Chen, Juan Zhai

Forum 2

14:00–16:00

Testing and Analysis: Static and Dynamic Analysis 1

ASE Research Papers · 8 talks

- 14:00 **Every Scenario Matters: Scenario-Guided White-Box Testing for Static Code Checkers Powered by LLMs**
Jun Liu, Liqiang Ji, Jiwei Yan, Linjie Pan, Jinsong Ren, Jun Yan, Jian Zhang
- 14:15 **Intent-Aware Defect Pattern Extraction from Singular Examples**
Jiachen Han, Fengjie Li, Jiajun Jiang, Ruihang Fan, Yingfei Xiong, Linjie Pan, Biliang Wang, Junjie Chen
- 14:30 **Conflict Extraction in Probabilistic Datalog Analyses**
Siyu Chen, Chungha Sung, Xuyang Li, Jingbo Wang
- 14:45 **Large Language Model Enabled Symbolic Execution for Automated Functional Analysis**
Meixi Liu, Zhenbang Chen, Xudong Wang, Ziran He, Jiachen Gu, Minghui Chen, Lei Wang, Wei Dong
- 15:00 **Orbis: Guiding Symbolic Execution Techniques to Maximize Option-Related Branch Coverage**
Minjong Kim, Sungjae Hwang, Sooyoung Cha
- 15:15 **SIMD-Accelerated Sparse Bit-Vectors for Pointer Analysis**
Zhaoyang Tan, Peisen Yao, Kui Ren
- 15:30 **IncSFS: Incremental Full-Sparse Flow-Sensitive Pointer Analysis for C/C++**
Kunlin Liu, Zhenbang Chen, Piya Zu, Yide Du, Ji Wang
- 15:45 **PELSE: Asymmetric-Evidence-Guided MCU Peripheral Model Synthesis**
Guohao Wu, Chunlin Wang, Rui Yao, Feng Hong, Hongliang Liang, Qiuping Yi

Forum 6

14:00–16:00

Software Analytics: Mining Software Repositories 1

ASE Research Papers · 8 talks

- 14:00 **Characterizing the Landscape of Open-Source Satellite Software**
Jinfeng Wen, Qi Liang, Yuehan Sun, Federica Sarro, Ao Zhou, Xuanzhe Liu, Shangguang Wang
- 14:15 **Recovering Revisions of Pull Requests with Altered History**
Gengyi Sun, Georges Aaron RANDRIANAINA, Tao Xiao, Yasutaka Kamei, Shane McIntosh
- 14:30 **An Exploratory Study on LLM-Generated Code and Comments in Code Repositories**
Yongyi Ji, Jiaji Wang, Yi Zhou, Fuxiang Chen, Hongji Yang
- 14:45 **RealAIGC: Towards More Realistic Evaluation of AI-Generated Code Detection in Real-World Code Repository**
Wentao Chen, Yufei Chen, Huiqun Yu, Guisheng Fan, Kaiwen Zhi, Peite Guo, Jianyuan Yu
- 15:00 **Looks Good, but Does It Run? A Large-Scale Empirical Study of the Executability of Example Code in Model Cards**
Simeng Yao, Jialin Zhao, Yang Zhang, Tun Li, Tao Wang, Changrong Xie, Zezhou Tang, Yiwen Wu

- 15:15 **Observing Fine-Grained Changes in Jupyter Notebooks During Development Time**
Sergey Titov, Konstantin Grotov, Cristina Sarasua, Yaroslav Golubev, Dhivyabharathi Ramasamy, Alberto Bacchelli, Abraham Bernstein, Timofey Bryksin
- 15:30 **ATLAS: Agentic Taxonomy of LArge-Scale Software Ecosystems**
Junyi Lu, Mengyao Lyu, Wu Jiahui, Lei Yu, Chengwei Liu, Fengjun Zhang, Li Yang, Chun Zuo, Yang Liu
- 15:45 **SustainSR: Interpretable OSS Sustainability Prediction via Symbolic Regression**
Yang Deng, Mustafa Misir, Jinfu Chen, Jifeng Xuan

Forum 7

14:15–16:00

Testing and Analysis: Regression and Model-Based Testing 4

ASE Research Papers · 8 talks

- 14:15 **TDRepro: A Neurosymbolic Approach to Reproducing Timing-Dependent Flaky Test Failures**
Shanto Rahman, Talank Baral, August Shi, Wing Lam
- 14:30 **Synthesizing File-Level Data for Unit Test Generation with Chain-of-Thoughts via Self-Debugging**
Ziyue Hua, Tianyu Chen, Yeyun Gong, Shuai Lu, Peng Cheng, Qinglin Zhu, Yibo He, Yingjie Fu, Wenpin Jiao, Wei Yang, Tao Xie
- 14:45 **Feature-Focused Test Generation**
José Antonio Zamudio Amaya, Gaetano Sapia, Alexi Turcotte, David Benavides, Marcel Böhme, Andreas Zeller
- 15:00 **Hallucination to Consensus: Multi-Agent LLMs for End-to-End JUnit Test Generation**
Xu Qinghua, Guancheng Wang, Lionel Briand, Kui Liu
- 15:15 **FlakeFighters: A PyTest Plugin to Detect Flaky Test Failures**
Michael Foster, Owain Parry, Michael Hilton, Gregory Kapfhammer, Phil McMinn
- 15:25 **multiplex: A Modular LLM-Based Mutation Framework**
Megan Maton, Gregory Kapfhammer, Phil McMinn
- 15:35 **Don't Retest Everything: Auditable Regression Scoping for Safety-Critical Systems**
Francesco Basciani, Daniele Masti, Patrizio Pelliccione, Alberto Petrucci, Franco Raimondi, Emerson Sales, Paolo Serri
- 15:50 **TORI: A Tool for Oracle Quality Assessment**
Facundo Molina, Nazareno Aguirre, Alessandra Gorla

Forum 8

14:00–16:00

AI and Software Engineering: AI4SE 4

ASE Industry Showcase · 10 talks

- 14:00 **Why Is This on My Bill? Documentation-Grounded Rule Learning for Itemized Bill Explanation**
Martin Schäff, Muhammad Numair Mansur, John Konstantinides, Willem Visser, Jason Lefever
- 14:15 **Empirical Computation: Prompting versus Programming**
Eric Tang, Jing Liu, Marcel Böhme
- 14:25 **BESSER-NN: Visual Neural Network Design and Automated Code Generation with the BESSER Low-Code Platform**
Nadia Daoudi, Armen Sulejmani, Jordi Cabot
- 14:35 **Distributional Program Analysis: Treating Large Language Models as Program Samplers**
Tuan Phat Tran-Truong, Xuan-Bach Le
- 14:45 **UIBenchKit: A Unified Toolkit for Design-to-Code Model Evaluation**
Chinh T. Le, Trevor Yee Siang Ong, Jingyu Xiao, Yuxuan Wan, Yintong Huo
- 14:55 **IP-Preserving Enterprise System Modernization via LLM-Driven Specification Extraction and Multi-target Code Generation**
Andreas Biesdorf, Eckhard Seibert, Nicole Wengatz, Bart Carpels, Tom Schoonbaert
- 15:05 **SecChain: Operationalizing Security Guidelines for Multi-agent Smart Contract Generation**
Dianxiang Sun, Liming Nie, Yang Liu
- 15:20 **Smart Contract Synthesis via Multi-modal Specifications**
Tanglin Chen, Haoxian Chen, Huilin Xiang, Yuepeng Wang
- 15:35 **MergeSE: Post-Hoc Model Merging for Software Engineering Tasks without Retraining**
Palash Ranjan Roy, Banani Roy, Kevin Schneider, Chanchal K. Roy
- 15:45 **Requirements-Augmented Generation for Trustworthy Acceptance Testing of LLM-Based Software**
Fanyu Wang, Chetan Arora, Zhenping Xie, Yonghui Liu, Kla Tantithamthavorn, Aldeida Aleti, Siwei Jiang

Forum 15

14:00–16:00

Security and Other Non-Functional Properties: Reliability and Availability 3

ASE Research Papers · 8 talks

- 14:00 **HERO: Hypothesis-Centered Root-Cause Analysis for Microservice Incidents**
Jiewei Lyu, Junquan Yi, Shu Liang, Pengfei Chen, Long Pan
- 14:15 **Observability of web applications: an industrial case study**
Anas Shatnawi, Bachar Rima, Zakarea AL SHARA, Gabriel Darbord, Abdelhak-Djamel Seriai, Christophe Bortolaso
- 14:30 **Enhancing Trace-Based Root Cause Analysis for Microservice Systems via Code Change Understanding**
Min Zhang, Chenxi Zhang, Senyu Xie, Shihong Chen, Lei Wu, Xin Peng
- 14:45 **Bifrost: Empowering Pretrained Language Model with Fallibility Representation for Log-Based Fault Diagnosis**
Minghua He, Tong Jia, Lingzhe Zhang, Chiming Duan, Xinlong Zhao, Leyi Pan, Cheng Wang, Kangjin Wang, Yinghao Yu, Liping Zhang, Yifan Wu, Ying Li
- 15:00 **LLM-Assisted Joint Ticket and Log Analysis for Incident Triage in Intelligent and Connected Vehicles**
Ruwei Fu, Shenglin Zhang, Wenwei Gu, Weiguo Li, Yongqian Sun, Dan Pei
- 15:15 **AlarmClaw: Context-Enriched Alarm Management with Category-/Severity-Aware Incident Graphs**
Siyu Yu, Meizhen Li, Jiacheng Yang, Yifan Wu, Ning Wang, Zhaoxing Chen, Ming Liu, Xinchu Ren, Xincheng Ren, Chan Li, Tong Jia, Xiang Zhang, Ying Li
- 15:30 **Smart Brain: Semantic Anomaly Detection for Operational Time Series in Large Scale Service Systems**
Hang Cui, Zexin Wang, Jingjing Li, Juncheng Hu, Haotian Si, Cenjie Hu, Quan Zhou, Yongchang Hu, Lei Han, Dan Pei, Changhua Pei, Gaogang Xie
- 15:45 **KRCA: An Efficient Root Cause Analysis System in Hyper-scale Microservice Systems via Agentic AI**
Jiamin Jiang, Jingfei Feng, Yu Luo, Qingliang Zhang, Yongqian Sun, Wenwei Gu, Shenglin Zhang, Tianyu Cui, Yao Wu, Jielong Huang, Nan Qi, Dan Pei

Forum 16

14:00–16:00

Security and Other Non-Functional Properties: Security and Privacy 5

ASE Research Papers · 8 talks

- 14:00 **GraphLedger: Repository-Level Vulnerability Detection via Graph-Based Compression and Assumption Validation**
Thi-Hong-Cuc Le, Hoang Quoc Bao Hua, Xuan-Bach Le
- 14:15 **Breaking Customized LLMs for Coding: Automated Red Teaming for Instruction Backdoor Attacks**
Yuchen Chen, Wei Cheng, Yuan Xiao, Weisong Sun, Chunrong Fang, Yang Liu, Zhenyu Chen, Baowen Xu
- 14:30 **IAM Policy Autopilot: Static Analysis for Policy Generation from Application Code**
Antonio Filieri, Luke Kennedy, Kevin Luo, Matt Luttrell, Adrian Palacios, Akash Panda, Neha Rungta, Matthias Schlaipfer, Karan Jit Singh, Anvesh Tanuku, Nick Winans, Diana Yin, Weiben Zhang
- 14:45 **A Comprehensive Evaluation of Code Language Models for Security Patch Detection**
Nils Loose, Joseph Bienhüls, Kristoffer Hempel, Felix Mächtle, Thomas Eisenbarth
- 15:00 **LLMs as Hackers: Autonomous Linux Privilege Escalation Attacks**
Andreas Happe, Aaron Kaplan, Jürgen Cito
- 15:15 **"Impossible to Hide Secret ...": Uncovering Security and Privacy Issues in LLM-Native IDEs**
Mostafijur Rahman Akhond, Md Afif Al Mamun, Gias Uddin, Song Wang
- 15:30 **ReSEDA: Automatic and Precise Program-Level ReDoS Triage via Interprocedural Taint Analysis and Agentic Exploit Generation**
Rongchen Li, Weihao Su, Chengyao Peng, hong huang, Haiming Chen, Guiyi He
- 15:45 **Benchmarking Automated Security Patch Backporting: How Far Are We?**
jincheng yang, Yulong Fu, Chengwei Liu, Lyuye Zhang, Fangyuan Zhang, Bingyang Ren, Yang Liu, Hui Li

16:00–16:30 *Coffee Break***Plenary Room (Forum 1-3)**

16:30–17:30

ASE Awards · 2 talks

- 16:30 **ASE Award**
- 17:00 **Usage, costs, and benefits of continuous integration in open-source projects**
Michael Hilton, Timothy Tunnell, kai huang, Darko Marinov, Danny Dig

17:30–18:30 *Break***18:30–22:30 Conference Dinner**

ASE Catering · Festsaal

Thursday

15 October 2026

146 talks

09:00–10:00 Lost in Automated Specifications: A Time and Space Odyssey in System Testing

Myra Cohen

ASE Keynotes · Plenary Room (Forum 1-3)

10:00–10:30 Coffee Break

Forum 1

10:30–12:30

AI and Software Engineering: AI4SE 5

ASE Research Papers · 8 talks

- 10:30 DepWareTrans: Dependency-Aware Incremental Repository Migration across Co-executable Languages
Sivajeet Chand, Alexander Pretschner, Steve Haupt, Derui Zhu, Sushant Kumar Pandey
- 10:45 Aligning User Requirement with Large Language Model in Code Generation
Jia Li, Ruiqi Bai, Yangkang Luo, Yiran Zhang, Wentao Yang, Zeyu Sun, Tiankuo Zhao, Dongming Jin, Lei Li, Zhi Jin
- 11:00 Pattern over Pixels: Measuring Pattern Completion Bias in Multimodal Code Generation
Khai Nguyen, Oscar Chaparro, Antonio Mastropaolo
- 11:15 Customizing an LLM for Enterprise Software Engineering
Aditya Kini, Satish Chandra, Milad Hashemi, Saksham Thakur, Aditya Pandey, Vincent Nguyen, Marc Brockschmidt, Franjo Ivančić, Danny Tarlow, Parthasarathy Ranganathan, Petros Maniatis, Ahmed Omran, Zaheer Abbas, Anita Gergely, Martin Sevenich, Gufeng Zhang, Amy Hua, Alexander Frömmgen
- 11:30 From Vibe to Spec: An Empirical Study to Demystify Specifications in Code Generation
Ruirui He, Hanqing Lu, Lezhi Ma, Kai Ye, Qiong Wu, Shangqing Liu, Lei Bu, Xuandong Li
- 11:45 AEG-Edit: Bridging API Evolution and Code Semantics for Accurate API Usage in Code Generation
Yitong Lin, Haoye Tian, Dehai Zhao, Chao Ni, Xiaohu Yang, Xiaoxue Ren
- 12:00 Effective and Efficient Context Retrieval via Partial Dependency Graph for Repository-Level Code Generation
Zhongxin Liu, Zhonghao Jiang, Zhifan Ye, Haoye Wang, Jiakun Liu, Xiaoxue Ren
- 12:15 Towards Fully Automated Medical Imaging Code Generation via Validation-Based Context Engineering
Zixiao Zhao, Jing Sun, Zhe Hou, Cheng-Hao Cai, Qian Liu, Mengze Li, Zijian Zhang, Jin Song Dong

Forum 2

10:30–12:30

Security and Other Non-Functional Properties: Security and Privacy 6

ASE Research Papers · 9 talks

- 10:30 An empirical analysis of vulnerability detection tools for solidity smart contracts
Francesco Salzano, Cosmo Kevin Antenucci, Simone Scalabrino, Giovanni Rosa, Rocco Oliveto, Remo Pareschi
- 10:45 SmartPoC: Validating Smart Contract Vulnerability Findings through PoC Generation and Execution
Longfei Chen, Ruibin Yan, Taiyu Wong, Yiyang Chen, Wang Jialai, Chao Zhang
- 11:00 Zero-Knowledge VMs: Security and Usability in Practice
Jinan Jiang, Jinzhao Chu, Pengyu Xue, Xinghao Peng, Haoran Qin, Xiapu Luo
- 11:15 GoLiSA: An Advanced Static Analysis Tool for Detecting Potential Risks and Vulnerabilities in Hyperledger Fabric Smart Contracts
Luca Olivieri, Luca Negrini, Vincenzo Arceri, Pietro Ferrara, Fausto Spoto, Agostino Cortesi
- 11:25 MalSkills: Detecting Malicious Skills in the Agentic Supply Chain via Neuro-symbolic Reasoning
Shenao Wang, HeJunjie, Yanjie Zhao, Yayi Wang, Kan Yu, Haoyu Wang
- 11:40 LARA: Static ReDoS Detection for Regular Expressions with Lookahead via Matching Automata
Mengxi Wang, Chengyao Peng, Weihao Su, Guiyi He, Rongchen Li, hong huang, Haiming Chen
- 11:55 How Your Credentials Are Leaked by LLM Agent Skills: An Empirical Study
Zhihao Chen, Ying Zhang, Yi Liu, Gelei Deng, Yuekang Li, Yanjun Zhang, Jianting Ning, Leo Zhang, Lei Ma, Zhiqiang Li
- 12:10 AIALIB: A Threat Library for AI-Generated Code
Ali Aljaberi, Jian Kang, Youcheng Sun

12:20 Cochise: A Reference Harness for Autonomous Penetration Testing
Andreas Happe, Jürgen Cito

Forum 6

10:30–12:30

Testing and Analysis: Regression and Model-Based Testing 5

ASE Research Papers · 9 talks

- 10:30 FaultForge: A Business-Semantics-Driven Cross-Layer Fault Injection Framework for Microservice
Yongxin Yang, Yang Yang, Anni Jiang, Mingyuan Yang, Peng Yu
- 10:40 DT-Drive: A Tool for Deterministic Replay-Based Testing and Debugging of Autonomous Driving Systems
Sanjeetha Pennada, Matthew Leach, Olek Osikowicz, Donghwan Shin
- 10:50 Continuous On-Commit Certification Testing of Railway Control Software: An Experience Report
Jacopo Maltagliati, Salvatore De Simone, Giovanni Denaro
- 11:00 Real-World Perturbation Testing of Autonomous Driving Systems
Stefano Carlo Lambertenghi, Matthias Weil, Andrea Stocco
- 11:15 Property-Based Testing of Wafer Handling Robots via Model-Based Planning and Fuzzing
Ruiyang Xu, Jingjing Liang, Guoyue Zheng, Geguang Pu, Ting Su
- 11:30 Test Suite Generation for Simulink Models: An Experience with Reactis
Andrea Bombarda, Nunzio Marco Bisceglia, Marco Saini, Federico Formica, Angelo Gargantini, Claudio Menghi
- 11:45 Where Does Balance Break? Boundary Discovery for Game Balance Testing under a Finite Simulation Budget
Hiroki Mukai, Yusaku Kato, Norihiro Yoshida, Erina Makihara, Katsuro Inoue
- 12:00 QDoctor: Probe-Based Testing for Quantum Programs
Qichen Wang, Shangzhou Xia, Xiaoyu Guo, Jianjun Zhao
- 12:15 Quantum Squeeziness: An Information Theoretical Metric for Quantum Software Testability
Avner Bensoussan, Hector Menendez, Mohammad Reza Mousavi

Forum 7

10:30–12:30

Security and Other Non-Functional Properties: Dependability and Safety 1

ASE Research Papers · 9 talks

- 10:30 Not In My Git Yard: Catching Backdoors at Commit and Release Time
Dimitri Kokkonis, Michaël Marcozzi, Stefano Zacchiroli
- 10:45 IntoAgent: A Practical Framework for Automated Detection and Triggering of Integer Overflow Vulnerabilities
Yuelin Wang, Jiongchi Yu, Xiaofei Xie, Yaohui Sun, Tianyu Shi, Yanbang Sun, Qiang Hu, Junjie Wang
- 11:00 Precise and Efficient Static Data Race Detection
Pei Wang, Zhihang Sun, Fei He
- 11:15 From Thread to Pool: Modeling and Static Detection of Thread Pool Misuses in Java Programs
Baoquan Cui, Yucheng Qiu, Rui Han, Hang Gao, Jian Zhang
- 11:30 SafeDep: Detecting State Mutation Unsoundness in Rust via Dependency Slicing and LLM Reasoning
Zihao Rao, Yilin Chen, Chengjun Chen, Hui Xu
- 11:45 SAF: An AI-Agent-Ready and Browser-Accessible Static Analysis Framework for LLVM IR
Yuekang Li, Wei Li, Wei Song, Yi Liu, Gelei Deng
- 11:55 Trace-Driven Automated PoC Synthesis for On-Chain Attacks
Xing Su, Hao Wu, Hanzhong Liang, Yunlin Jiang, Yuxi Cheng, Yating Liu, Fengyuan Xu, Sheng Zhong
- 12:10 MASSIV: A Hybrid RAG for Automated Defeater Extraction in Autonomous Driving Safety Cases
Tihomir Rohlinger, Daniel Ratiu, Stefan Wagner
- 12:20 xDECAF: An Extensible Data Flow Diagram Analysis Framework for Information Security
Benjamin Arp, Felix Schwickerath, Alexander Vogt, Tom Hüller, Nils Niehues, Nicolas Boltz

Forum 8

10:30–12:30

Testing and Analysis: Fuzzing and Mutation Testing 2

ASE Research Papers · 8 talks

- 10:30 Learning the Lexical Structure of Black-Box Systems by Parsing Systematic String Edit Mutations
Moeketsi Raselimo, Lars Grunske, Bernd Fischer
- 10:45 Planner-Fuzz: Overcoming Diagnostic Saturation in UAV Planners via Boundary-Oriented Fuzzing
Rongbo Chen, Chengcheng Zhao, Peng Cheng, Jiming Chen, Taegyu Kim
- 11:00 SpecTrum: Specification-Guided Differential Fuzzing for Ethereum Consensus Clients
Seokhun Jeong, Gyeongmin Dan, Sukyoung Ryu, Sungjae Hwang
- 11:15 Breaking the Isolation: Coordinated Multi-agent Fuzzing for Smart Contracts with Multi-dimensional Objective Learning
Ruichao Liang, Ruochen Cao, Jing Chen, Kun He, Yebo Feng, Yue Xue, Cong Wu

- 11:30 **RamFuzz: LLM-Guided Greybox Fuzzing for Spatial Memory Corruption via Valid Range Violation**
Shangzhi Xu, Wei Song, Yuekang Li, Nan Sun, Muhammad Ejaz Ahmed, Willy Susilo, Benjamin Turnbull, Xiao Cheng, Siqi Ma
- 11:45 **Maximal Format-Free Data Repair**
Zijian Luo, Xi Wu, Hong Jin Kang, Alan Fekete, Rahul Gopinath
- 12:00 **CtxFuzz: Context-Aware Directed Fuzzing via Runtime Fuzzer-Agent Cooperation**
Jiho Kim, Dongkwan Kim, HyungSeok Han, Youngjoon Kim, Soyeon Park, Sangwoo Ji, Joshua Wang, Dae R. Jeong, Taesoo Kim
- 12:15 **From Function to Repository: Towards Understanding and an Agentic Approach for Repo-Level Directed Input Generation**
Haoyu Wang, Ming Wen, Haoran Yan, Jiayi Huang, Xuanhua Shi, Hai Jin

Forum 15

10:30–12:30

Formal Aspects of Software Engineering: Software Validation and Verification

ASE Research Papers · 9 talks

- 10:30 **Generating Loop Invariants for Deductive Verification using Iterator Profiles**
Keisuke Nishimura, Jean-Pierre Lozi, Julia Lawall
- 10:45 **Clarity: From Formal System Designs to Verified RL Controllers**
Austin O'Quinn, Conor Snedeker, Max Taylor, Lance Joneckis, Christopher Stewart
- 11:00 **ControlsDSL: A Language for Verifiable Cloud Configuration Controls**
Miyree Alair, Shrutarshi Basu, Sean Cai, Loris D'Antoni, Julio Delgado Jr., Antonio Filieri, Andrew Gacek, Anand Goyal, Sebastiaan Joosten, Varad Panch, Sorawee Porncharoenwase, Vishwanath Raman, Niloofar Razavi, Neha Rungta, Mike Schliep, Preethi Sekaran, Chase Johnson
- 11:15 **Certified Program Synthesis with a Multi-modal Verifier**
Yueyang Feng, Dipesh Kafle, Vladimir Gladshtein, Vitaly Kurin, George Pirlea, Qiyuan Zhao, Peter Müller, Ilya Sergey
- 11:30 **KVerus: Scalable and Resilient Formal Verification Proof Generation for Rust Code**
Yuwei Liu, Xinyi Wan, Yanhao Wang, Minghua Wang, Lin Huang, Tao Wei
- 11:45 **Automated Lemma Discovery in Agentic Program Verification**
Huan Zhao, Haoxin Tu, Zhengyao Liu, Martin C. Rinard, Abhik Roychoudhury
- 12:00 **ProofPulse: Interactive Proof Coverage Analysis for Dafny**
Álvaro F. Silva, Ruben Martins, Alexandra Mendes
- 12:10 **DSpec2Test: Specification-Driven Test Generation in Dafny**
Sofia Vieira Pinto, Álvaro F. Silva, João Pascoal Faria, Alexandra Mendes
- 12:20 **QUTest: A Native Testing Framework for Quantum Programs**
José Campos

Forum 16

10:30–12:15

Maintenance and Evolution: Debugging and Fault Localization

ASE Journal First · 7 talks

- 10:30 **Efficient Black-Box Fault Localization for System-Level Test Code Using Large Language Models**
Ahmadreza Saboor Yaraghi, Golnaz Gharachorlu, Sakina Fatima, Lionel Briand, Ruiyuan Wan, Ruifeng Gao
- 10:45 **COBug: A Bug Localization Pipeline for COBOL: Does It Work? An Exploratory Study**
Kavyasri Gajula, Phanindra Kodela, Sridhar Chimalakonda
- 11:00 **Improving IR-based Bug Localization with Semantics-Driven Query Reduction**
Asif Mohammed Samir, Masud Rahman
- 11:15 **BugMentor: Generating Answers to Follow-up Questions from Software Bug Reports using Structured Information Retrieval and Neural Text Generation**
Usmi Mukherjee, Masud Rahman
- 11:30 **MERA: A Multi-expert LLM Framework for Severity Assessment of User Tickets in Mobile OS Beta Testing**
Yongqian Sun, Yimin Zuo, Qingliang Zhang, Bowen Hao, Wenwei Gu, Jiaqi Luan, Sixu Zhou, Shenglin Zhang, Dan Pei
- 11:45 **Disassembly-Augmented Root Cause Analysis for C/C++ Segmentation Faults in SAP HANA**
Jihoon Jung, Chansong You, Jingun Hong, Thomas Bach, Gabin An
- 12:00 **One Sentence Can Kill the Bug: Auto-replay Mobile App Crashes from One-sentence Overviews**
Yuchao Huang, Junjie Wang, Zhe Liu, Mingyang Li, Song Wang, Chunyang Chen, Yuanzhe Hu, Qing Wang

12:30–14:00 Lunch Break

Forum 1

14:00–16:00

AI and Software Engineering: SE4AI 3

ASE Research Papers · 8 talks

- 14:00 RAGVA: Engineering retrieval augmented generation-based virtual assistants in practice
Rui Yang, Michael Fu, Kla Tantithamthavorn, Chetan Arora, Lisa Vandenhurk, Joey Chua
- 14:15 Localizing Data Annotation Faults in Semantic Segmentation Systems
Pengyi Yang, Teng Long, Zhiwu Xu, Linhai Jing
- 14:30 Targeted Deep Learning System Boundary Testing
Oliver Weiß, Amr Abdellatif, Xingcheng Chen, Giorgi Merabishvili, Vincenzo Riccio, Severin Kacianka, Andrea Stocco
- 14:45 Tangent: An Empirical Study of Testing Practices for LLM-Based Agent Applications
Rangeet Pan, Tyler Stennett, Divya Sankar, Bridget McGinn, Alessandro Orso, Raju Pavuluri, Saurabh Sinha, Maja Vukovic
- 15:00 XMutant: XAI-based fuzzing for deep learning systems
Xingcheng Chen, Matteo Biagiola, Vincenzo Riccio, Marcelo d'Amorim, Andrea Stocco
- 15:15 Optimization-Free Repair of Code LLMs: A Closed-Form Fisher-Guided Bayesian Approach
Zhiyu Duan, Chong Bian, iainzhang, Shunkun Yang
- 15:30 Evolution-Aware Regression Test Prioritization of ML-Enabled Systems using Gradient-Based Behavior Vectors
Eunho Cho, Donghwan Shin, In-Young Ko
- 15:45 Not as Sweet by Another Name: An Empirical Study of Format Robustness in LLM Document Workflows
Xiaoyu Zhang, Xianyun Cheng, Tianlin Li, Yuwei Zheng, Yue Yang, Yang Liu

Forum 2

14:00–16:00

Testing and Analysis: Static and Dynamic Analysis 2

ASE Research Papers · 8 talks

- 14:00 Testing Static Analyzers via Semantic-Preserving Mutators Learned from Real-World Refactoring Practice
Meilin Li, Kaixuan Li, Zifan Xie, Shiyu Qiu, Ming Wen, Maolin Sun, Hongyu Zhang
- 14:15 PL4SA: Optimized Partial Library Selection for Efficient Static Analysis
Guohao Feng, Yifei Lu, Minxue Pan
- 14:30 Verifier-in-the-Loop LLM Solving of Heap Constraints for Concolic Execution
Shaoran Xia, Leyi Cheng, Caihua Dong, Dongdong She, Bo Wang, Xin Peng, Zhen Dong
- 14:45 Top-Down Multi-agent LLM System with Runtime Data for MISRA C Violation Detection
Ahcheong Lee, Yeongbin Kang, Sechang Jang, Heechan Yang, Moonzoo Kim
- 15:00 Tracing the Invisible: Semantic Message Flow Discovery via Data Contracts in Real-World Distributed Systems
youlong chen, Jie Lu, Mingtao Huang, Chenghang Shi, Yongheng Huang, Haofeng Li, Mengna Ma, Yong Liu, Qinfen Hao, Lian Li
- 15:15 GPU-Accelerated Belief Propagation for Program Analysis
Haoyu Feng, Xin Zhang
- 15:30 TasmScan: Continuation-Aware Taint Analysis for TVM Bytecode with Savelist Abstraction
Yixuan Liu, Yin Wu, Yi Li
- 15:45 Safe Multi-site Binary Instrumentation of Running x86-64 Programs
Jihun Baek, Sanghoon Han, Hyungon Moon

Forum 6

14:00–15:45

Security and Other Non-Functional Properties: Security and Privacy 7

ASE Research Papers · 7 talks

- 14:00 S2CAgent: A Self-Learning Agent Knowledge Graph Framework for Software Supply Chain Analysis
Youfei Huang, Zhong Li, Minxue Pan, Xuandong Li
- 14:15 Closing the Gap: Automated Discovery of Secure Dockerfile Reference Standards via Semantic Clustering in Enterprise Inner Source
Jessica Hösl, Benedikt Hofmann, Patrick Stöckle
- 14:30 How Effective Are NPM Malicious Package Detectors? A Large-Scale Empirical Study
Wenbo Guo, Zhongwen Chen, Zhengzi Xu, Ming Kang, Shiwen Song, Chengwei Liu, Yijia Xu, Weisong Sun, Yang Liu
- 14:45 Defensive Capability Analysis for JavaScript Libraries
Wenyuan Xu, Anders Møller
- 15:00 When Dependencies Depend: A Large-Scale Study on Conditional Dependencies and Their Security Implications in C/C++
Yifeng Di, Juewei Zhang, Tianyi Zhang

- 15:15 A Longitudinal Study of Android Apps Signing Key Protection
Mark Meng, Qing Zhang, Weirao Lu, Chunyang Chen
- 15:30 GANADI: Uncovering C/C++ OSS Reuse Genealogies via Pivotal Function-Based Clustering to Enhance Supply Chain Security
Dongyeon Kim, Seunghoon Woo, Heejo Lee

Forum 7

14:00–16:00

Security and Other Non-Functional Properties: Performance 2

ASE Research Papers · 9 talks

- 14:00 EffiHolmes: Differential Profiling-Guided Repository Level Time Inefficiency Fix Localization
Haowen Yang, Yun Peng, Zishuo Ding
- 14:15 Assessing Uncertainty in Performance Modeling: Conformal Prediction vs. Bayesian Regression
Stefan Jahns, Johannes Dorn, Max Weber, Sven Apel, Norbert Siegmund
- 14:30 Beyond Single-Level PGO: Understanding and Harnessing Cross-Level Optimization Effects in Industrial Systems
Yong Liu, Yifan Cheng, Yunfeng Cui, Ting Chen, Jiachi Chen, Zhihong Xue, Chong Hu, Shengyu Cheng
- 14:45 Baseline: A Plugin-Oriented, Hardware-Agnostic GPU Benchmarking Library
Côme Eyraud, Maxime Lamothe
- 14:55 CppPerf: An Automated Pipeline and Dataset for Performance-Improving C++ Commits
Tommy Ho, Khashayar Etemadi, Zhendong Su
- 15:05 EnergyTrackr: A Modular Energy Regressions Detection Tool
François Bechet, Jérôme Maquoi, Luís Cruz, Benoît Vanderoose, Xavier Devroey
- 15:15 SimServing: Native-Execution Simulation for Evolution-Resilient LLM Serving Configuration Tuning
Yujie Huang, Jiazhen Gu, Zhihan Jiang, Michael Lyu
- 15:30 Discovering Performance Archetypes: Critical-Path-Aware Pattern Analysis and Regression Detection
Kaveh Shahedi, Heng Li, Maxime Lamothe, Foutse Khomh
- 15:45 Bridging the Gap between Intent and Impact: An Empirical Study of GPU Optimizations in Deep Learning Frameworks
Siqi Wang, Fanjing Luo, Xing Hu, Xinyu Wang, Xin Xia

Forum 8

14:00–16:00

AI and Software Engineering: AI4SE 6

ASE Research Papers · 8 talks

- 14:00 Doc2Feat-Bench: Evaluating Documentation-Driven Feature Addition
Zhonghao Jiang, Le Deng, Jialun Cao, Michael Pradel, Zhongxin Liu
- 14:15 From Runnable Code to Shippable Applications: Test-Driven Development for Full-Stack Web Application Generation
Yuxuan Wan, Tingshuo Liang, Jiakai Xu, Jingyu Xiao, Yintong Huo, Michael Lyu
- 14:30 DCAS: Decoupling CLI Agent Scaffolding to Internalize Planning across Scaffolds
Kishanthan Thangarajah, Boyuan Chen, Ahmed E. Hassan
- 14:45 Generation of Web Apps with Agentic IDEs: An Empirical Assessment
Manuel Marceca, Maria Teresa Rossi, Leonardo Mariani
- 15:00 Revisiting LLMs on New Feature Implementation in Real-World Software Development Practices
Jia Li, Tingxuan Huang, Tiankuo Zhao, Yiming Wang, Yiran Zhang, Anshun Hu, Kechi Zhang, Lecheng Wang, Zinan Sheng, Kui Liu, Ge Li, Zhi Jin
- 15:15 Compressing Code Context for LLM-Based Issue Resolution
Haoxiang Jia, Earl T. Barr, Sergey Mechtaev
- 15:30 RACE-Bench: A Reasoning-Augmented Benchmark for Repository-Level Code Agents on Feature Addition
Shuhan Liu, Zhiyi Zhao, Xing Hu, Kui Liu, Xiaohu Yang, Xin Xia
- 15:45 Debugging the Debuggers: Failure-Anchored Structured Recovery for Software Engineering Agents
Chenyu Zhao, Shenglin Zhang, Yihang Lin, Wenwei Gu, Zhimin Chen, Yongqian Sun, Dan Pei, Chetan Bansal, Saravan Rajmohan, Minghua Ma

Forum 15

14:00–16:00

Formal Aspects of Software Engineering: Formal Methods and Model Checking 1

ASE Research Papers · 9 talks

- 14:00 Specification-Guided Synthesis of Deadlock-Free Communication Protocol Refinements with Large Language Models
Yang Li, Ping Hou, Nobuko Yoshida

- 14:15 **Kani: A Model Checker for Rust**
Remi Delmas, Ziad Hassan, Qinheping Hu, Rahul Kumar, Felipe R. Monteiro, Thanh Nguyen, Adrian Palacios, Celina Val, Michael Tautschnig, Justus Adam, Daniel Schwartz-Narbonne, Carolyn Zech
- 14:30 **Leveraging System-Level Observations to Inform Bayesian Learning of Model Parameters for Quantitative Verification**
Simos Gerasimou, Xingyu Zhao
- 14:45 **Visual Consistency in Formal Modeling: An Empirical Evaluation**
Yiliang Liang, Avinash Palliyil, Eunsuk Kang, Joshua Sunshine
- 15:00 **Mining Tactics for Automated Theorem Proving**
Jian Fang, Yixun Yao, Yingfei Xiong
- 15:15 **Quick Bug Detection through Black-Box Checking: A Systematic Evaluation**
Bram Pellen, María Belén Rodríguez, Frits Vaandrager, Petra van den Bos
- 15:30 **Towards Diagnostic Explainability in Workflow Verification via Shapley Attribution**
Radosław Klimek, Jakub Błażowski
- 15:40 **VerusSeek: Retrieval-Augmented LLM-Based Proof Synthesis for Rust Programs**
Yuchen Zhang, Cheng Wen, Jialun Cao, Dugang Liu, Zhiwu Xu, Yuwei Liu, Shengchao Qin
- 15:50 **MIMOS-Tools: An Integrated Toolchain for Model-Based Design of Embedded Systems**
Wang Yi, Chengzi Huang, Behnam Khodabandeloo, Duc Anh Nguyen

Forum 16

14:00–16:00

Testing and Analysis: Fuzzing and Mutation Testing 3

ASE Research Papers · 9 talks

- 14:00 **Mutation-Guided Unit Test Generation with a Large Language Model**
Guancheng Wang, Xu Qinghua, Lionel Briand, Kui Liu
- 14:15 **Marv: An Interactive Tool for Visualising and Reviewing the Results of Mutation Analysis**
Daniel Wells, Zolán Lévai, Phil McMinn
- 14:25 **ConfFuzz: Parameter-Aware Greybox Fuzzing for Configurable Cloud Systems**
Shuai Wang, Hao Wang, Darko Marinov, Tianyin Xu, Yongle Zhang
- 14:40 **ExcavatorFuzz: LLM-Guided Path-Difficulty-Driven Directed Greybox Fuzzing**
Kaiyu Xie, Siyuan Li, Jiaqian Peng, Yun Yang, Yaowen Zheng, Kai Cheng, Jiayu Zhang, Shizhao Tian, Iiyubo, Hongsong Zhu
- 14:55 **Hybrid Fuzzing in Input Generation Logic**
Siwei Wei, Shihao Zhu, Ruijie Meng, Yan Cai
- 15:10 **Accurate and Low-Cost Residual Risk Assessment via Sampled Profiling and Structure-Aware Coverage Amplification**
Seongmin Lee, Işıl Özgü, Marcel Böhme, Miryung Kim
- 15:25 **Gray-Box Fuzzing in Local Space**
Martin Jonáš, Jan Strejcek, Marek Trtík
- 15:40 **InsightQL: A Unified Code Database for Fuzz Blocker Analysis**
Wentao Gao, Renata Borovica-Gajic, Sang Kil Cha, Michael Fu, Tian Qiu, Thuan Pham
- 15:50 **Mutation Testing of Simulink Cyber-Physical System Models: Challenges and Solutions in Practice**
Murat Kavak, Onur Kilincceker, Serge Demeyer, Kevin Vandenbroucke, Halim Abdurrahman Ceylan

16:00–16:30 *Coffee Break***Forum 1**

16:30–17:30

Human and Social Aspects: Human-Computer Interface

ASE Research Papers · 4 talks

- 16:30 **Toward Inclusive Programming Support for Novice Programmers with Dyslexia: Insights from a Comparative Study**
Zihan Fang, Janice Chung, Ruijia Chen, Marcia Barnes, Yuhang Zhao, Yu Huang
- 16:45 **A Selective Undo/Redo Algorithm for Dependency-Constrained Operation Histories Based on Change Propagation**
Xiaoxiong Fan, Hou Tam, Hanxi Zhu, Song-Hai Zhang
- 17:00 **Towards Multilingual AI-Assisted Software Engineering**
Shivadharshan Sankar, Anirudh Arrepu, Sridhar Chimalakonda
- 17:20 **Agentic Code Review in the Terminal: A Trajectory-Level Analysis of Behavior, Cost, and Human-Alignment**
Wachiraphan (Ping) Charoenwet, Kla Tantithamthavorn, Patanamon Thongtanunam, Hong Yi Lin, Minwoo Jeong, Ming Wu

Forum 2

16:30–17:30

Software Analytics: Mining Software Repositories 2

ASE Research Papers · 4 talks

- 16:30 From Exploration to Revelation: App-Level Context-Aware Deceptive Pattern Detection for Mobile Applications
Jieshan Chen, Zhen Wang, Jiamou Sun, Zhenchang Xing, Qinghua Lu, Qing Huang, Xiwei (Sherry) Xu, Liming Zhu
- 16:45 A Few Pages of Markdown: Committed AI Configuration and Lower Quality Cost after Coding-Agent Adoption
Yegor Denisov-Blanch, Shyam Agarwal, Pavel Azaletskiy, Hao He, Rylan Schaeffer, Brando Miranda, Bogdan Vasilescu, Sanmi Koyejo
- 17:00 How Quantum Bugs Live and Die: A Lifecycle-Based Empirical Study of Bugs in Quantum Software
Yasai Shi, Xiangxin Meng, Xiangjie Huang, Jian Zhang, Tianyu Wo, Xu Wang
- 17:15 CodeFault: Predicting Fault Risks from Code Changes via Multi-modal Learning
Yifan Xiao, Shijie Li, Yu Huang

Forum 6

16:30–17:30

Security and Other Non-Functional Properties: Security and Privacy 8

ASE Research Papers · 4 talks

- 16:30 Evidence-Grounded Logic Vulnerability Detection for Smart Contracts via Runtime Context Reasoning
Guang Gong, Yucong Jiang, Rong Jian, Xiaotao Feng, Zhihan Wei, Qi Li, Ke Xu
- 16:45 Airdrop Scams in Practice: An Industry-Oriented Taxonomy, Benchmark, and Tool Audit
Ningyu He, Xiaohui Hu, Haoyu Wang
- 17:00 PoCo: Agentic Proof-of-Concept Exploit Generation for Smart Contracts
Vivi Andersson, Sofia Bobadilla, Harald Hobbelhaugen, Martin Monperrus
- 17:15 Demystifying Solana Bots: From GitHub Blueprints to On-Chain Fingerprints
Xiaoye Zheng, Yujing Chen, Minghao Wu, David Lo, Difan Xie, Daoyuan Wu, Xiaohu Yang, Zhiyuan Wan

Forum 7

16:30–17:30

Formal Aspects of Software Engineering: Specification Languages 2

ASE Research Papers · 4 talks

- 16:30 Barista: Synthesizing Typestate Specifications with LLM Agents
Catarina Gamboa, Paulo Canelas, Ricardo Costa, Márcio Caetano, Jonathan Aldrich, Alcides Fonseca
- 16:45 SpecFSM: Extracting and Repairing Finite State Machines from Protocol Specification Documents
Xiangdong Li, DaWei Huang, Jingjing Guan, Hui Li
- 17:00 Diagnosing Violations of State-based Specifications in iCFTL
Cristina Stratan, Claudio Mandrioli, Domenico Bianculli
- 17:15 Refine2Diff: Detecting Protocol Specification–Implementation Inconsistencies via Specification-Driven Code Refinement
Yuekun Wang, Lili Quan, Xiaoferi Xie

Forum 8

16:30–17:30

Testing and Analysis: Regression and Model-Based Testing 6

ASE Research Papers · 4 talks

- 16:30 PWCC: Probability-Weighted Combinatorial Coverage for Risk-Informed Scenario Testing of Autonomous Systems
Borhaneddine Hamadou, Chokri Mraidha, Nicholas Matragkas
- 16:45 Constrained Co-evolutionary Differential Metamorphic Testing for Autonomous Systems with an Interpretability Approach
Hossein Yousefzadeh, Shenghui Gu, Lionel Briand, Ali Nasr
- 17:00 Simulator Ensembles for Trustworthy Autonomous Driving Systems Testing
Lev Sorokin, Matteo Biagiola, Andrea Stocco
- 17:15 TrackTest: Trajectory Generation and Prioritization for Testing Multiple Object Tracking System
Xinyu Gao, Haoxin Chen, Shuoxiao Zhang, Enyi Tang, Hengrui Xu, Haoxiang Tian, Minghui Wei, An Guo

Forum 15

16:30–17:30

Formal Aspects of Software Engineering: Formal Methods and Model Checking 2

ASE Research Papers · 4 talks

- 16:30 VARIES: Verification Harness Synthesis and Efficient Scheduling for Unsoundness Detection in Rust Libraries
Huan Li, Xing Hu, Xin Xia, Xinyu Wang

- 16:45 **Sound and Efficient Statistical Model Checking for Probabilities and Bounded Rewards**
Hao Bu, Lin Huang, Tao Wei, Jingyi Wang
- 17:00 **FlowChecker: One Command to Model Check Control-Flow Rules in Industrial Java**
Qiuye Wang, Liqing Cao, Wei Hu, Dengcheng He
- 17:15 **Inferring the Shape of Data Frames in R Programs using Abstract Interpretation**
Oliver Gerstl, Florian Sihler, Matthias Tichy

Forum 16

16:30–17:30

Security and Other Non-Functional Properties: Dependability and Safety 2

 **ASE Industry Showcase** · 4 talks

- 16:30 **Understanding Real-World Ordering-Related Concurrency Bugs under Weak-Memory and Cross-Architecture Settings**
Shaohao Wang, Cheng Wen, Yuandao Cai, Shengchao Qin, Jie Su, Mengda He, Xiaoxue Ma, Cong Tian
- 17:00 **Deployment Risk Assessment using Diff-Aware Features: A Case Study at Prime Video**
Mayur Premkumar Kurup, Hyunjae Suh, Swathi Vaidyanathan, Pranesh Vyas, Srinidhi Madabhushi, Yegor Silyutin
- 17:10 **ThreatCraft: Automated Attack Scenario Generation via Hybrid Rule-Based and LLM-Driven Reasoning**
Dohee Kang, Jiwon Kwak, Geunwoo Baek, Seungjoo Kim
- 17:20 **The EVERest Dataset for Secure Software Engineering**
Sophie Corallo, Debora Grupp, Dominik Fuchß, Jan Keim, Frederik Reiche, Tobias Hey, Anne Koziol

Friday

16 October 2026

75 talks

Forum 7

09:00–10:00

SECUTE · 2 talks

- 09:00 Welcome
Emanuele Iannone
- 09:10 Language Models as Penetration Testing Agents: Surprising Capabilities and Ethical Dimensions
Jürgen Cito

Forum 8

08:45–10:00

KLEE · 2 talks

- 08:45 Opening Remarks
Cristian Cadar
- 09:00 Deriving Test Oracles for Verification Infrastructure
Maria Christakis

Forum 10

09:00–10:00

LiveWrite · 2 talks

- 09:00 Welcome
Carmen Armenti, Tim Menzies
- 09:40 Livewrite in Action!

Forum 11

09:00–10:00

CPAchecker · 2 talks

- 09:00 Opening Remarks
Dirk Beyer
- 09:05 Bi-Abductive Analysis of Code with Pointers and Advanced Low-Level Operations
Tomas Vojnar

Forum 13

09:00–10:00

MAS-GAIN · 4 talks

- 09:00 MAS-GAIN Opening Session
Giacomo Valente
- 09:05 Decompose to Conquer: A Multi-Agent LLM Pipeline for Requirements Engineering and Software Design
Mesfin Tibebe, Mohamed Nabeel
- 09:30 SAGE: Source-Aligned Governance for Human-Agent Requirements Clarification in Multi-Agent Software Engineering
Tingting Feng, Meifang Xue, Jian Wang, Zhongyan Du
- 09:45 A Workflow-Driven Multi-Agent Architecture for Requirements Engineering
Radoslaw Klimek

Forum 14

09:00–10:00

SWGeno · 4 talks

- 09:00 The Gene Is Not the Gate: Conceptualizing the Semantic Genome of Quantum Software
- 09:15 From Code to Knowledge: Enhancing Coding Agents via Reverse-Engineered Software Knowledge
- 09:30 Reachability Is Only the Beginning: Toward Effect-Aware Vulnerability Analysis for AI Agents
- 09:45 ForkOps: Operating Fleets of Generated Software Variants

Forum 15

09:00–10:00

ASYDE · 4 talks

- 09:00 Opening

- 09:15 **Refinement for Correct Iterative Model-Based Systems Engineering of Cyber-Physical Systems**
Luca Brodo, Giuseppe Scalora, Stefan Henkler
- 09:29 **Systematic evaluation of security flaws in low-level code generated by LLMs**
Fynn Janson, Peter Barth, Eric Heim
- 09:43 **Towards a Vertically Integrated Compiler Infrastructure to Improve Solver Time during Symbolic Execution**
Sören Tempel, Christian Dietrich

Forum 16

09:00–10:00

POVC · 1 talk

- 09:00 **LLM-Driven Requirements Engineering: Automating Dependencies, Regulatory Tracing, and Impact Analysis**
Lionel Briand

10:00–10:30 *Coffee Break***Forum 7**

10:30–12:30

SECUTE · 7 talks

- 10:30 **Using LLM Agents for Security Testing of Compiled Mobile Applications: A Case Study on a Real-World Compiled iOS App**
Fabian Scherf, Florian Magin
- 10:45 **The Role of Prompt Patterns in LLM-Based Smart Contract Vulnerability Detection**
Gerardo Iuliano, Dario Di Nucci
- 11:00 **AutoDriver: Purely Source-based Driver Generation for Standard Library Fuzzing Pipelines**
Vincent Ahlrichs, Florian Kasten, Jonas Bogenberger, Dieter Schuster, Julian Horsch
- 11:15 **Artifacts to Actors: Assessing the Maturity of Software Supply Chain Security Measurement Practices**
Alexis Butler, Dan O'Keeffe, Santanu Dash
- 11:30 **Sentry: Towards a Sound Dynamic Analysis Framework for Hybrid Android Applications**
Jyoti Prakash, Mathias Ransgaard Yde, Miguel Campusano, Abhishek Tiwari
- 11:40 **The Present and Future of Software Security Testing (Live Collaborative Activity)**
- 12:25 **Closing**
Emanuele Iannone

Forum 8

10:30–12:30

KLEE · 9 talks

- 10:30 **Automatic Unit-Test Generation Based on Symbolic Execution**
Tigran Avagyan, Levon Shahbazyan, Hovhannes Movsisyan, Hayk Aslanyan
- 10:35 **Co-Designing Compiler Infrastructure and Symbolic Execution to Improve SMT-LIB Encoding**
Sören Tempel, Christian Dietrich
- 10:40 **EAR: Debugging Libraries through API Call Traces**
Tongjun Zhang, Cristian Cadar
- 10:45 **Interprocedural Symbolic Static Analysis for Go**
Varvara Dvortsova, Alexey Borodin
- 10:50 **FeatMaker: Automated Feature Engineering for Search Strategy of Symbolic Execution**
Jaehan Yoon, Sooyoung Cha
- 11:10 **Empc: Effective Path Prioritization for Symbolic Execution with Path Cover**
Shuangjie Yao, Dongdong She
- 11:30 **Enhancing Symbolic Execution of Programs for Control Flow Path Feasibility Analysis**
Hovhannes Movsisyan, Hripsime Hovhannisyan, Hayk Aslanyan, Tigran Avagyan
- 11:50 **Enhancing Symbolic Execution with Self-Configuring Parameters**
Minjong Kim, Sooyoung Cha
- 12:10 **Guiding Symbolic Execution with Static Analysis and LLMs for Vulnerability Discovery**
Md Shafiuzzaman, Achintya Desai, Wenbo Guo, Tefvik Bultan

Forum 10

10:30–12:30

LiveWrite · 1 talk

- 10:30 **Livewrite in Action!**

Forum 11

10:30–12:30

CPAchecker · 4 talks

- 10:30 Dataflow Analyses in CPAchecker
Manuel Pietsch
- 11:00 Framework Optimizations for Distributed Summary Synthesis
Matthias Kettl
- 11:30 Interpolation-Based Termination Analysis in CPAchecker
Marek Jankola
- 12:00 SV-LIB support in CPAchecker
Marian Lingsch-Rosenfeld

Forum 13

10:30–12:15

MAS-GAIN · 5 talks

- 10:30 Towards Declarative Composition of Heterogeneous AI Agents
Antonio Marcos Alves Morais, Paulo Henrique, Felip Manyà
- 10:55 Sufficiently Advanced Agents Are Indistinguishable from Erlang: An Operational Mapping from OTP to Multi-Agent LLM Systems
Mick Darling
- 11:20 Does Architecture Matter? Comparing Multi-Agent Workflows for Filtering and Prioritizing Static Analysis Findings
Linus Bennin, Oshando Johnson, Eric Bodden
- 11:45 Towards Reliable Agentic Infrastructure Provisioning with LLMs
Yhon Castañeda, Camilo Escobar-Velásquez
- 12:10 MAS-GAIN Closing Session and Final Wrap-up
Giacomo Valente

Forum 14

10:30–12:15

SWGeno · 5 talks

- 10:30 RepoTune: Co-Evolving Vulnerability Detection with a Repository's Genome
- 10:45 The Bad AI Bubble. Implications for Security and Software Engineering Research
- 11:30 Same Task, Different Fate: Population Structure in Trajectory-Based Analyses of Coding Agents
- 11:45 Seeing the Forest: A Taxonomy-Guided Analysis and Improvement of Malicious PyPI Package Detectors
- 12:00 Vulnerability Proof-of-Concepts in Software Supply Chain: Real-World Usage Scenarios from the Perspective of Migration

Forum 15

10:30–12:30

ASYDE · 7 talks

- 10:30 Constrained Automated Requirements Formalization for Rigorous Analysis with Alloy
Viktoria Koscinski, Mehdi Mirakhorli
- 10:47 Dead-End Ratio: A Behavioral Model Quality Metric for Constraint-Aware Runtime Verification
Natalia Dybczak, Jagoda Flejmer, Radoslaw Klimek
- 11:04 STAFF: Towards Translating Formal Verification Benchmarks Between Formal Specification Languages
Merlijn Sevenhuijsen
- 11:21 Controller Mimicking for On-The-Fly DES Reachability Control
Florencia Zanollo, Fuyuki Ishikawa, Sebastian Uchitel, Victor Braberman
- 11:38 From Statistical Code Generation to Correct-by-Construction AI Programming: A Basis-Function Composition Framework
Yu Wu, Yuchen Yang
- 11:55 WanderVision: Scenario-Based GUI Testing for Smart TVs via Vision-Language Models
Sihyun Ahn, Juyeon Yoon, Taehyoung Kim, Eunki Song, Shin Yoo
- 12:12 Discussion and closing

Forum 16

10:30–12:30

POVC · 4 talks

- 10:30 Prompts Matter: On the Impact of Code Smell Remediation with Large Language Models
Nadim Hagatulah, Markus Borg, Emma Söderberg
- 11:00 Geometry-Aware Prompt Context for Model Consistency Management
Monalisha Ojha, Shilpi Gupta, Rahul Sharma

- 11:30 **Agent4RE: A Self-refining Multi-agent Framework for End-to-End Software Requirements Engineering and Benchmarking**
Yongjian Tang, Linhan Li, Thomas Runkler
- 12:00 **Prompt Pipelines as Engineered Artifacts: Regression-Testing a Multi-Agent Threat-Modeling Pipeline**
Eldar Mametov, Andrey Sadovykh, Eugene Zouev, Alberto Sillitti

12:30–14:00 *Lunch Break*

Forum 8

14:00–16:00

KLEE · 6 talks

- 14:00 **Defusing Logic Bombs in Symbolic Execution with LLM-Generated Ghost Code**
Dimitrios Stamatios Bouras, Sergey Mechtaev
- 14:20 **Enhancing Symbolic Execution with Machine-Checked Safety Proofs**
David Trabish, Shachar Itzhaky
- 14:40 **Agentic Concolic Execution**
Zhengxiong Luo, Huan Zhao, Dylan Wolff, Cristian Cadar, Abhik Roychoudhury
- 15:00 **LLM-Driven Generation of Summaries for Symbolic Execution**
Frederico Ramos, João Matos, Jorge Hristovsky, Pedro Adão, Ruben Martins, José Fragoso Santos
- 15:20 **A Case Study in Firmware Verification: Applying Formal Methods to Intel TDX Module**
Dirk Beyer, Po-Chun Chien, Bo-Yuan Huang, Nian-Ze Lee, Thomas Lemberger
- 15:40 **Triggering the Hidden: Oracle Design Meets KLEE for Backdoor Detection**
Michaël Marcozzi

Forum 9

14:00–16:00

ATIQSER · 5 talks

- 14:00 **Welcome to ATIQSER 2026**
Jürgen Cito, Ben Hermann
- 14:10 **What we learned in 15 years of artifact evaluation**
Ben Hermann
- 15:00 **doi-checker: A Tool for Validating References of Research Papers**
Dirk Beyer, Marian Lingsch-Rosenfeld
- 15:20 **ATIBA: Grounded Integrity and Quality Checking for Research Papers**
Veli Karakaya, Semih Çağlar, Yusuf Yiğit Korkmaz, Eray Tüzün
- 15:40 **RefVerifier: Semi-Automated Reference Claim Verification for Scientific Manuscripts**
Stefania Mocan, Florian Angermeir, Mark Kreitz

Forum 10

14:00–16:00

LiveWrite · 1 talk

- 14:00 **Livewrite in Action!**

Forum 11

14:00–16:00

CPAchecker · 4 talks

- 14:00 **Distributed Summary Synthesis Multi-processing**
Akshay Warrier
- 14:30 **Trace Equations for Program Verification**
Gidon Ernst
- 15:00 **Leveraging Function Contracts for Blame Analysis**
Nicola Thoben
- 15:30 **ACF: Debugging Long-Running Systems with Tracing and Unit Test Execution**
Pavel Parizek

Forum 16

14:00–16:00

POVC · 4 talks

- 14:00 **VibeCheck - Assessing the Quality of LLM-Generated Unit Tests: A Multi-Agent Empirical Study Across Heterogeneous Repositories**
Anika Tabassum, Intesum Mushahid, Md. Fahim Arefin, Tarannum Shaila Zaman

- 14:30 **TurnCost: Turns, Not Observations—Cache-Aware Loop Engineering for Coding Agents**
Pramod Kumar, Raj Kumar Pandey
- 15:00 **Towards Trace-Guided Automated Program Repair**
David Georg Reichelt, Shinhyung Yang, Wilhelm Hasselbring
- 15:30 **Ranking Is Not Authority: Executable Context-Routing Contracts for Coding Agents**
Nikhil Mahesh

16:00–16:30 *Coffee Break*

Forum 8 16:30–17:35

KLEE · 2 talks

- 16:30 **The Role of SAT Solving in Automated Reasoning**
Armin Biere
- 17:30 **Closing Remarks**
Cristian Cadar

Forum 10 16:30–17:30

LiveWrite · 1 talk

- 16:30 **Debriefing**
Carmen Armenti, Dario Di Nucci, Tim Menzies, Klaus Schmid

Forum 11 16:30–17:30

CPAchecker · 3 talks

- 16:30 **New Developments in JavaSMT 6 and Beyond**
Daniel Baier
- 17:00 **Continuous Modernization of CPAchecker**
Philipp Wendler
- 17:25 **Discussion and Closing**
Dirk Beyer

Forum 16 16:30–17:00

POVC · 1 talk

- 16:30 **Your Prompt Is Only Half the Prompt: Silent Chat-Template Drift in Open-Weight Model Repositories**
Simarjot Khanna



At JetBrains Research we are dedicated to advancing the state of the art in computer science and software engineering. We are a team of researchers and engineers exploring innovative solutions at the intersection of research and software development. In collaboration with academic and industry experts, we contribute cutting-edge technology and scientific research to developers worldwide.



Come see our talk

Observing Fine-Grained Changes
in Jupyter Notebooks During
Development Time

Wednesday 14th Oct 15:15 - 15:30,
Forum 6



Check out our academic licences



Come see us at our stand near the registration area!

Sponsors and Supporters of ASE 2026

Sponsors



Association for
Computing Machinery



IEEE

Supporters

PLATINUM



amazon
automated
reasoning



JETBRAINS

GOLD



HUAWEI

IN-KIND

