


Challenges to Using Large Language Models in Code Generation and Repair

Liliana Pasquale  | University College Dublin and Lero—The SFI Research Centre for Software

Antonino Sabetta  | SAP

Marcelo d'Amorim  | North Carolina State University

Péter Hegedűs  | University of Szeged and FrontEndART Ltd.

Mehdi Tarrit Mirakhorli | University of Hawaii at Manoa

Hamed Okhravi  | Massachusetts Institute of Technology Lincoln Laboratory

Mathias Payer | EPFL

Awais Rashid | University of Bristol

Joanna C. S. Santos | University of Notre Dame

Jonathan M. Spring | Cybersecurity and Infrastructure Security Agency

Lin Tan | Purdue University

Katja Tuma | Vrije Universiteit Amsterdam

Large language models (LLMs) hold great promise in solving many challenges arising from software complexity, including the possibility of automating code generation and repair. Although we cannot deny the groundbreaking nature of LLM-based code repair, we must be realistic in positioning current results.

This “Building Security In” column explores the challenges in using LLMs for automated code generation and program repair.

Introduction

Joanna C. S. Santos and Mathias Payer: With the latest advances in LLMs, artificial intelligence (AI)-based code development assistants are increasingly part of day-to-day software development. A recent study (<https://tinyurl.com/3kub3awn>) of 500 U.S.-based developers showed that 92% use AI coding assistants for work and personal use. The increased productivity perceived by developers

partly explains this fast, widespread adoption; AI helps them automate repetitive tasks so that they can focus on higher-level challenging tasks.¹

Péter Hegedűs, and Lin Tan: Automated program repair (APR) aims to generate source code to fix software defects and vulnerabilities automatically. Research on APR has advanced significantly with generative AI models. Long short-term memory models achieved notable success in generating complex, syntactically correct code after training on extensive source code datasets. LLMs further improved APR. Since they are pretrained on an enormous amount of natural language text and source code, they also offer an out-of-the-box solution for code

repair. Recent studies^{2,3,10,11} show that LLMs can fix issues in the code, such as defects, vulnerabilities, and code smells. In some cases, code

DISCLAIMER

The views expressed in this document do not necessarily represent the views of the U.S. government or the Cybersecurity and Infrastructure Security Agency. Reference to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply an endorsement, recommendation, or favoring by the U.S. government or the Cybersecurity and Infrastructure Security Agency.

repair is treated as a code generation task with a prompt explicitly instructing the model to fix a problem in a given location.^{2,3}

Hegedűs: While LLMs, such as GPT-4, excel at fixing functional bugs in laboratory environments (i.e., on synthetic or isolated issues), their real-world application, especially when the task is fixing complex, security-related issues, remains limited.⁴

LLMs Generate Vulnerable and Incorrect Code

Awais Rashid: Software professionals are concerned about AI-generated code quality, correctness, and security and the need to scrutinize and validate such code.⁵ This is particularly critical for program repair. The CrowdStrike case has highlighted how errors in a single patch can have a global impact, halting critical services.

Marcelo d'Amorim: There is evidence that LLMs can produce code containing security weaknesses even when the user of the LLM is not malicious.⁶ Prevention and detection are two directions to mitigate this problem. For *prevention*, responsibly disclosing the weaknesses of an LLM to the public encourages the LLM maintainers to curate training datasets actively. Users must know the threats and limitations associated with the versions of the LLMs they are using. LLM maintainers are expected to care about public announcements about weaknesses in their LLM and will address them in subsequent releases. The LVE Repository (<https://lve-project.org/>) is a commendable global initiative in that direction. For *detection*, LLMs can be used to explain the weaknesses identified by third-party tools. Ideally, those explanations should describe the consequences of not taking some action to mitigate the weakness, i.e., counterfactual explanations are likely more

helpful to users. Such explanations should help the distracted trained developer and help to train inexperienced developers.

Hamed Okhravi: Source code often must comply with many other requirements besides functionality. These may include soft and hard real-time constraints, power usage requirements (e.g., for embedded code), and side-channel resilience (e.g., for crypto code), as well as more generic nonfunctional requirements such as readability, maintainability, performance, portability, testability, and modularity. LLM-generated code rarely accounts for these requirements.

LLMs must also understand the underlying platforms to generate the correct code to fix specific bugs.⁹ Platform-specific parameters may include Windows versus Linux file handling, 32-bit versus 64-bit code, Windows versus POSIX threading application programming interface (API), network socket differences, or memory alignment. To successfully repair code, the LLM should be trained on all those platforms, and detailed platform information must be provided when prompting it to repair source code.

Santos and Antonino Sabetta: LLMs have a prompt input and output size threshold (e.g., GPT-4 can take up to 128,000 tokens and generate up to 16,384 tokens). Considering real software systems' sheer complexity and size, these thresholds are insufficient. As such, LLMs may miss the broader context of a project and can generate a limited repair size. Understanding the complete environment in which the code operates (e.g., configuration files, external dependencies, database structures, etc.) is crucial for code generation and repair.

Despite improvements in token counts (e.g., Gemini 1.5 allows up to 1 million tokens), capturing sufficient relevant context may require more than just a large token

capacity. Effective code repair depends on some form of reasoning about the code context, which involves understanding the structure and purpose of the application to generate repairs that align with the codebase's security needs and a variety of technical constraints.

Payer: AI-based assistants must be sufficiently scoped to create correct code, especially in highly optimized environments. Although research has explored integrating LLMs into automated testing, the results only marginally improve on existing methods when incorporating the cost of LLM queries. A more promising application of LLMs is in generating test drivers to target specific functionalities as they can generate and refine drivers to improve code path coverage. While manually written drivers often fall short, LLMs could fill these gaps and enhance API coverage. However, LLM-generated drivers may be flawed or incomplete, potentially leading to false positives and wasted resources.

Hegedűs: The reproducibility of the fixing process is a major challenge as LLM results are nondeterministic. Since prompts can have a major impact on the results, instead of model training, one would need to invest effort into prompt engineering.

Katja Tuma: From experience assessing the effectiveness of LLMs in fixing security misconfigurations in Kubernetes-based applications,⁷ existing tools (Checkov, Datree, and KICS, to name a few) adopt different rules and security policies to identify security misconfigurations. These tools may produce both false positives and negatives. Some configurations (such as allowing network access to a container) might be flagged as insecure, while they are required for the running application to perform its key functionalities (e.g., network monitoring). This can substantially affect the performance of LLMs in fixing security misconfigurations

for K8-based applications. Keeping the human in the loop is essential: For infrastructure-as-code repair with LLMs, first, we need to distinguish among misconfiguration fixes that can (and should) be verified by humans and those that could potentially be automated with limited security risks. Second, we need to establish a common taxonomy of misconfigurations and robustness measures for more effective tool benchmarking and experimental validation. This could help associate a certain level of confidence in the LLM-generated fixes for certain types of misconfigurations and instead leave the (orders of magnitude smaller) remaining set of issues for humans to handle.

Insufficient Training Data and How to Add Software Domain Knowledge

Okhravi: Supervised approaches may be necessary for APR to succeed. To achieve this aim, LLMs must capture a solid notion of vulnerable and secure code to repair code successfully. However, realistic data for vulnerable and secure code samples are insufficient to apply supervised learning. The entire National Vulnerability Database (NVD) contains around 260,000 vulnerabilities at the time of writing. Consider further that not every reported vulnerability has an associated code sample available, and some vulnerabilities in the NVD are too old to be relevant to modern code. As a result, there are often fewer than tens of thousands of vulnerable code samples on which to train an LLM. This is insufficient to ensure the LLM is properly trained to generate only secure code. Recent work in this domain suggests that enriching existing data with additional properties (context, syntax, and semantics) allows one to achieve better accuracy, precision, and recall in distinguishing between vulnerable and secure code.⁸

Santos: Prior work¹¹ examined whether LLMs could repair their generated insecure code. Stark differences exist among the issues LLMs could repair for each programming language. For example, for Python, LLMs can solve issues related to XML validation vulnerabilities but are less capable of solving issues related to the use of a broken or risky cryptographic algorithm (CWE-327), path traversal (CWE-22), and incorrect permission assignment for a critical resource (CWE-732). We also observed that overall, LLMs are more capable of repairing Python code than Java code. These results indicate open challenges in effectively using LLMs to repair insecure code. LLMs are trained with popular languages, especially Python. Consequently, LLMs will struggle to repair insecure code for languages with fewer samples in their training data. Even in cases where the language is well covered, a model generates repairs to insecure code based on historical data. Still, vulnerabilities and secure coding practices continually change as technology evolves. Thus, the precision observed today likely will not be the same tomorrow.

Tan: Another important question is whether adding more data to train deep learning (DL) models, including LLMs, is a promising direction to improve APR techniques. Using increasingly large amounts of data has succeeded in tasks such as speaking a natural language, which may fundamentally differ from coding tasks. Babies learn to speak their mother tongue by mimicking and learning implicitly from what they hear. However, software developers do not simply learn programming and program repair by reading code and patches; they also use logic and reasoning by taking programming, algorithms, and data structure courses. Thus, while adding more data may improve LLMs

for text and other modalities, it may not be the most effective approach for APR tasks. Adding explicit domain knowledge (including but not limited to type rules) to models may be a more effective approach.⁹ On the other hand, models may not need to learn the same way humans do, and the most effective learning approaches for humans may not be the most effective ones for models, suggesting that more data could be more effective.

Recent DL-based program repair techniques provide conflicting results in this respect. For example, KNOD employs a domain-rule distillation technique to explicitly inject domain knowledge including types into the neural network decoding procedure.⁹ Specifically, the domain-rule distillation technique 1) represents syntax and semantics as rules in first-order logic and 2) uses these logic rules to refine the teacher-student probability distributions to guide the model to learn to follow these syntactic and semantic rules. This approach shows that adding domain knowledge explicitly improves the effectiveness of neural networks for program repair. Yet, other studies (e.g., Jiang et al.¹⁰) suggest the opposite. They show that LLMs for code, *without or with fine-tuning*, outperform existing DL-based program repair techniques specially designed for APR to fix software defects. These generic LLMs for code are pretrained with a vast amount of data but are not designed for APR. Since these LLMs are typically trained with more data than existing DL-based APR approaches, the finding suggests that more data could be more effective for improving LLM-based program repair. The next relevant open questions are 1) whether we have more data for DL models to improve APR and code generation and (2) how to add domain knowledge to LLMs effectively.

Limited LLM Accountability and Overreliance

Hegedűs: Another major challenge with LLM-based code repair is the validation of the fixes they produce. It is not always easy to determine if an LLM-generated patch is genuinely good, meaning humans still play an essential role in verifying the correctness of the generated patches.

d'Amorim: A challenge is avoiding hallucinations, which can be especially detrimental to inexperienced developers, who may not realize incoherences in the discourse.

The complementary problem of vulnerability repair can be even more challenging in practice if we consider the possibility of developers accepting plausible patches recommended by an LLM. The possibility of introducing bugs or other vulnerabilities when repairing code is well known in software engineering, but security weaknesses can be more consequential. Developers need to validate the security patches that automated tools generate. However, for small single-hunk patches, which are prevalent, the human cost of reviewing may well dominate the cost of writing the patch. So, the benefit of automated repair in that context is questionable. It is therefore important 1) to focus automated repair on multiple hunk patches, 2) to develop tools capable of explaining the repairs, and 3) to ensure developers validate these patches.

Jonathan M. Spring: Developers need a robust development environment to place more trust in the outputs of an LLM. That means good specification and documentation of the API of the project or module, adequate unit tests, adequate integration tests, repeatable build processes, appropriate program verification techniques to detect specific common classes of vulnerabilities, appropriate testing to check parsing and error handling, and so

on. An organization should have these tools established and working well before moving to automated code repair.

However, there are some critical tasks an LLM cannot do. An LLM cannot take ownership of maintaining a software product that is out of support or is at the end of its life. An LLM cannot automatically write in interoperable, open standards for communication and data formats. Free and open standards will help others (using an LLM tool or not) repair your code after you move on to another project.

With or without LLM assistance, a software vendor should meet the goals of the Cybersecurity and Infrastructure Security Agency's Secure by Design initiative. When a software vendor offers a product on which the engineers use LLM-based code repair, the vendor should provide software transparency and vulnerability management. A system owner or acquisitions team should still ask for a software bill of materials and ask the vendor about their vulnerability disclosure and reporting practices. Vendors should still pledge the organizational work to make software secure by design.

If we demand that software is secure by design, tools such as LLMs for code repair can help software developers meet that demand.

Rashid: Several open questions surround the quality of LLM outputs. Would we see situations where the computer (LLM) says "no repair is needed" when one is required or where it hallucinates one? Similar questions arise about the repair itself. Who will scrutinize and validate the repair, and how, so it does not introduce undesirable side effects, such as impacting other software functionalities or introducing security weaknesses or vulnerabilities?

There is an expectation that the developer's role will change, from the *driver* who writes the code to a *navigator* who will check and

validate the driver's work, that is, the LLM's. However, we also know that automation and reliance on tools erode skills. I am reminded of a problem with my car: The hazard lights kept coming on when parked, draining the battery. Neither the small handheld diagnostic computer (with the repair person) nor the extensive diagnostic rig at the garage could replicate the issue or isolate the fault. The problem kept recurring until a different repair person came out to recharge the battery, used the same handheld diagnostic computer to no effect, gave it some thought, and then noted that it was likely a faulty burglar alarm. He isolated it, and the problem was solved. Even if we use LLMs for code repair, we need skilled software engineers to understand, scrutinize, and validate the outcomes.

Liliana Pasquale: LLMs can generate code that no longer satisfies system requirements or introduces vulnerabilities. Despite this, their growing power has led software engineers to increasingly depend on them, sometimes overly. This overconfidence becomes concerning as developers rely on LLMs for coding and program repair, where accuracy is critical. Existing AI coding assistants should identify the criticality of software development tasks and configure the reliance that developers can place on them accordingly. For example, LLMs can still be useful for several applications where errors can be tolerated. Thus, developers can entirely rely on LLMs to automate simple and repetitive programming tasks in noncritical applications. More complex programming tasks of noncritical applications could require the supervision of a senior software engineer to review the code generated automatically. New and large programming tasks, especially for critical applications, may require using LLMs only to oversee software development activities,

such as generating test cases or performing code reviews.

Mehdi Tarrit Mirakhorli: Code repair generated by LLMs, while often functional, provides no guarantees that the repaired code is free of vulnerabilities, meets specific safety criteria, or truly addresses the underlying requirements. This lack of assurance can be problematic, especially in critical systems where correctness, security, and performance are nonnegotiable. One idea is to use LLMs to generate test cases and validate the repaired or synthesized code. However, a stronger idea is to provide proof of correctness. Since proofs equate with programs, one can deliver an LLM-based approach to generate proofs of correctness automatically using similar programs. We discussed the foundation of shifting toward certified code repair, where LLMs are integrated with formal verification techniques.¹² Based on the theory that proofs can equate with programs, we can think of generating proofs as a task similar to generating code. This theoretical foundation suggests that with appropriate training and fine-tuning, LLMs can be guided to produce not only code repairs but also formal proofs that guarantee the correctness of the generated solutions. In such a transformative approach, along with the code fix, the LLM generates a formal proof that certifies the repaired code satisfies a set of predefined safety or correctness properties, security policies, or design rules. A lightweight verification tool can independently check the proof, ensuring the code fix meets the necessary safety criteria before deployment.

Certified code repair (or synthesis) is foundational for enabling AI autonomously and developing secure and trustworthy systems. Pre-LLMs and through my NSF CAREER award, I focused on the challenges of realizing such a foundational approach where software engineers could focus on the key

engineering tasks of 1) creativity and 2) design, then collaborate with a design synthesis tool to generate low-level code that correctly implements their design choices. While we are closer to such an idea today, there are challenges to achieving it for modern large-scale systems. For instance, generating formal proofs for code repairs can be computationally expensive, especially for large and complex systems. Proof generation requires rigorous formalization of the code's properties and behavior, and ensuring that these properties hold under all conditions can be time-consuming. Also, modern software has many third-party dependencies, adding to the complexity of generating proof of correctness. Fine-tuning LLMs on datasets that include examples of formal methods, symbolic reasoning, and proof generation tasks can help bridge this gap. Integrating language models with formal proof engines could also enhance their capabilities in proof generation.

Opportunities for Software Testing

Santos: LLMs cannot simply be used off the shelf as a foolproof tool to solve the insecure code repair problem. LLMs should *enhance* classic APR techniques rather than fully *replacing* them. Such a hybrid approach has been shown by prior work to help in generating tests.¹³ In that context, LLMs generated more diversified inputs to increase test coverage for an underlying search-based software testing approach.

Payer: Two key areas are certainly human-in-the-loop code completion and the generation of unit tests and fuzzers. Automated testing, particularly fuzzing, has experienced a meteoric rise in popularity, mirroring the growth of LLMs in computer science. Despite its conceptual simplicity, fuzzing effectively uncovers bugs by randomly probing various inputs to expose program

vulnerabilities. A promising application of LLMs is generating test drivers to target specific functionalities¹⁴ as they can create and refine drivers to improve code path coverage. While manually written drivers often fall short, LLMs could fill these gaps and enhance API coverage. However, LLM-generated drivers may be flawed or incomplete, potentially leading to false positives and wasted resources.

A promising use case of LLMs is in the bug-fixing process.³ After a fuzzer detects a bug and generates test inputs to reproduce it, an LLM could assist the developer by iteratively suggesting patches to address the underlying vulnerability. The fuzzer could then explore the patched code to uncover any lingering weaknesses of the patch. This iterative approach, alternating between fuzzers and LLMs, may lower developer involvement and reduce the costs of producing a complete patch. A hybrid approach combining fuzzers, LLMs, and developers could be a promising future direction for integrating LLMs into the bug discovery and remediation cycle. As it neither increases costs nor produces false positives, this approach is likely the most interesting angle for LLMs, but it will require careful customization and optimization.

However, while LLMs offer significant potential for enhancing fuzzing, the baseline approach without LLMs is already highly optimized, and the cost of querying LLMs must be carefully balanced against the potential benefits. LLMs trained on source code and specifications may improve mutation operators and driver generation, but some challenges, such as false positives, remain.

Rashid: “Many people expect advances in artificial intelligence to provide the revolutionary breakthrough that will give

order-of-magnitude gains in software productivity and quality. I do not,” wrote Fred Brooks Jr. in “No Silver Bullet,” his seminal 1986 essay tackling essential and accidental complexity in software engineering.¹⁵

Will LLMs for code repair tasks alleviate essential complexity or exacerbate accidental complexity? Unless we systematically address issues such as correctness, verifiability, and explainability, LLMs will likely add accidental complexity, potentially an order of magnitude, to the task of program repair, thus eroding any gains they may provide.

There are several open questions about the quality of LLM outputs. Time will tell. Let us know what your experience and opinions are. ■

References

1. A. Ziegler et al., “Productivity assessment of neural code completion,” in *Proc. 6th ACM SIGPLAN Int. Symp. Mach. Program. (MAPS)*, New York, NY, USA: Association for Computing Machinery, 2022, pp. 21–29, doi: [10.1145/3520312.3534864](https://doi.org/10.1145/3520312.3534864).
2. H. Joshi, J. C. Sanchez, S. Gulwani, V. Le, G. Verbruggen, and I. Radiček, “Repair is nearly generation: Multilingual program repair with LLMs,” in *Proc. 37th AAAI Conf. Artif. Intell. 35th Conf. Innov. Appl. Artif. Intell. 13th Symp. Educ. Adv. Artif. Intell. (AAAI/IAAI/EAAI)*, 2023, pp. 5131–4140, doi: [10.1609/aaai/v37i4.25642](https://doi.org/10.1609/aaai/v37i4.25642).
3. H. Pearce, B. Tan, B. Ahmad, R. Karri, and B. Dolan-Gavitt, “Examining zero-shot vulnerability repair with large language models,” in *Proc. IEEE Symp. Secur. Privacy (SP)*, Los Alamitos, CA, USA: IEEE Comput. Soc. Press, 2023, pp. 2339–2356, doi: [10.1109/SP46215.2023.10179324](https://doi.org/10.1109/SP46215.2023.10179324).
4. Z. Ságodi, G. Antal, B. Bogenfürst, M. Isztin, P. Hegedűs, and R. Ferenc, “Reality check: Assessing GPT-4 in fixing real-world software vulnerabilities,” in *Proc. 28th Int. Conf. Eval. Assess. Softw. Eng. (EASE)*, New York, NY, USA: Association for Computing Machinery, 2024, pp. 252–261, doi: [10.1145/3661167.3661207](https://doi.org/10.1145/3661167.3661207).
5. J. H. Klemmer et al., “Using AI assistants in software development: A qualitative study on security practices and concerns,” in *Proc. 31st ACM Conf. Comput. Commun. Secur. (CCS)*, New York, NY, USA: Association for Computing Machinery, 2024, pp. 2726–2740, doi: [10.1145/3658644.3690283](https://doi.org/10.1145/3658644.3690283).
6. S. Hamer, M. d’Amorim, and L. Williams, “Just another copy and paste? Comparing the security vulnerabilities of ChatGPT generated code and StackOverflow answers,” in *Proc. IEEE Secur. Privacy Workshops (SPW)*, San Francisco, CA, USA, 2024, pp. 87–94, doi: [10.1109/SPW63631.2024.00014](https://doi.org/10.1109/SPW63631.2024.00014).
7. F. Minna, F. Massacci, and K. Tuma, “Analyzing and mitigating (with LLMs) the security misconfigurations of helm charts from artifact hub,” 2024, *arXiv:2403.09537*.
8. S. Shimmi, A. Rahman, M. Gadde, H. Okhravi, and M. Rahimi, “VulSim: Leveraging similarity of multi-dimensional neighbor embeddings for vulnerability detection,” in *Proc. 33rd USENIX Secur. Symp. (USENIX Secur.)*, Philadelphia, PA, USA, 2024, pp. 1777–1794. [Online]. Available: <https://www.usenix.org/system/files/usenixsecurity24-shimmi.pdf>
9. N. Jiang, T. Lutellier, Y. Lou, L. Tan, D. Goldwasser, and X. Zhang, “KNOD: Domain knowledge distilled tree decoder for automated program repair,” in *Proc. 45th Int. Conf. Softw. Eng. (ICSE)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 1251–1263, doi: [10.1109/ICSE48619.2023.00111](https://doi.org/10.1109/ICSE48619.2023.00111).
10. N. Jiang, K. Liu, T. Lutellier, and L. Tan, “Impact of code language models on automated program repair,” in *Proc. 45th Int. Conf. Softw. Eng. (ICSE)*, Melbourne, Australia, 2023, pp. 1430–1442, doi: [10.1109/ICSE48619.2023.00125](https://doi.org/10.1109/ICSE48619.2023.00125).
11. M. Siddiq, L. Casey, and J. C. S. Santos, “FRANC: A lightweight framework for high quality code generation,” in *Proc. 24th Int. Conf. Source Code Anal. Manipulation (SCAM)*, Piscataway, NJ, USA: IEEE Press, 2024, pp. 106–117.
12. M. Fazelnia, M. Mirakhorli, and H. Bagheri, “Translation titans, reasoning challenges: Satisfiability-aided language models for detecting conflicting requirements,” in *Proc. 39th IEEE/ACM Int. Conf. Automat. Softw. Eng. (ASE)*, New York, NY, USA: Association for Computing Machinery, 2024, pp. 2294–2298, doi: [10.1145/3691620.3695302](https://doi.org/10.1145/3691620.3695302).
13. C. Lemieux, J. P. Inala, S. K. Lahiri, and S. Sen, “CodaMosa: Escaping coverage plateaus in test generation with pre-trained large language models,” in *Proc. 45th Int. Conf. Softw. Eng. (ICSE)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 919–931, doi: [10.1109/ICSE48619.2023.00085](https://doi.org/10.1109/ICSE48619.2023.00085).
14. Y. Lyu, Y. Xie, P. Chen, and H. Chen, “Prompt fuzzing for fuzz driver generation,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, New York, NY, USA: Association for Computing Machinery, 2024, pp. 3793–3807, doi: [10.1145/3658644.3670396](https://doi.org/10.1145/3658644.3670396).
15. F. P. Brooks Jr., “No silver bullet—essence and accident,” in *The Mythical Man-Month, Essays on Software Engineering*, Anniversary Edition, Reading, MA, USA: Addison-Wesley, 1995, pp. 177–203.

Liliana Pasquale is an associate professor at University College Dublin, Dublin D04 V1W8, Ireland, and a funded investigator at Lero—The SFI Research Centre for Software, Dublin D04 V1W8, Ireland. Her research interests include requirements engineering and adaptive systems, focusing on security, privacy, and digital forensics. Pasquale received a Ph.D. in information and software technologies from Politecnico di Milano. She is an associate editor of *IEEE Transactions on Software Engineering*, a department editor of *IEEE Security & Privacy Magazine*, and a member of the ACM

Transactions on Software Engineering and Methodology review board. Contact her at liliana.pasquale@ucd.ie.

Antonino Sabetta is a principal research scientist at SAP, 06254 Mougins, France. His research interests include applications of AI to software security, as well as the security of AI-based software systems. Sabetta received a Ph.D. in computer science and automation engineering from the University of Rome Tor Vergata. He serves as the editor of the “Building Security In” department of *IEEE Security & Privacy Magazine*. Contact him at antonino.sabetta@sap.com.

Marcelo d’Amorim is an associate professor in computer science at North Carolina State University, Raleigh, NC 27695-8206 USA. His research interests include helping developers build correct software and preventing, finding, diagnosing, and repairing software bugs and vulnerabilities. d’Amorim received a Ph.D. from the University of Illinois at Urbana-Champaign and an M.S. and a B.S. from UFPE. Contact him at mdamori@ncsu.edu.

Péter Hegedűs is an assistant professor at the Department of Software Engineering, University of Szeged, H-6720 Szeged, Hungary, and a researcher at FrontEndART Ltd., 744W+XJ Szeged, Hungary. His research interests include software maintainability models, DL applications, source code analysis, and vulnerability detection and prediction. Hegedűs received a Ph.D. in computer science from the University of Szeged. He was a program committee member for the CSMR, MSR, QUATIC, ESEM, and SCAM conferences. He has received various awards and scholarships during his career, including the prestigious Bolyai János research scholarship. Contact him at hpeter@inf.u-szeged.hu.

Mehdi Tarrit Mirakhorli is a faculty member at University of Hawaii at Manoa, Honolulu, HI 96822 USA. His research interests include the broad area of software engineering, including trustworthy software, software assurance, cybersecurity, AI, scientific software development, and software enabled sustainable disposal. Mirakhorli received a Ph.D. in computer science from DePaul University. He has received multiple ACM SIGSOFT Distinguished Paper Awards and was a recipient of the NSF CAREER Award in 2020. Contact him at mehdi23@hawaii.edu.

Hamed Okhravi is a senior researcher at the Massachusetts Institute of Technology (MIT) Lincoln Laboratory, Lexington, MA 02421 USA. His research interests include systems security, security evaluation, and operating systems. Okhravi received a Ph.D. in electrical and computer engineering from the University of Illinois at Urbana-Champaign. He is the recipient of two Best Paper Awards, three R&D 100 Awards, the Federal Laboratory Consortium for Technology Transfer Excellence in Technology Transfer Award, MIT Lincoln Laboratory’s Best Invention and Early Career Technical Achievement Awards, and the National Security Agency’s Best Scientific Cybersecurity Paper Award. He is a Senior Member of IEEE. Contact him at hamed.okhravi@ll.mit.edu.

Mathias Payer is an associate professor at EPFL, CH-1015 Lausanne, Switzerland, where he leads the HexHive group. His research interests include software and system security, particularly protecting applications from vulnerabilities like memory corruption or type violations. Payer received a doctorate in computer science from ETH Zurich. Contact him at mathias.payer@epfl.ch.

Awais Rashid is a professor of cybersecurity; the director of the EPSRC Centre for Doctoral Training in Cyber Security; and the director of the National Research Centre on Privacy, Harm Reduction and Adversarial Influence Online, University of Bristol, Bristol BS8 1UB, U.K. His research interests include security and privacy in large connected infrastructures with a particular focus on cyberphysical systems, software security, and human factors. Rashid received a Ph.D. in computer science from Lancaster University. Contact him at awais.rashid@bristol.ac.uk.

Joanna C. S. Santos is an assistant professor in the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556 USA. Her research interests include the domain of software engineering and security. Santos received a Ph.D. in computing and information sciences from Rochester Institute of Technology. She was the recipient of the Best Paper Award at the 2017 IEEE International Conference on Software Architecture for the paper “Understanding Software Vulnerabilities Related to Architectural Security Tactics: An Empirical Investigation of Chromium, PHP and Thunderbird.” Contact her at joannacss@nd.edu.

Jonathan M. Spring is a senior technical advisor for security at scale at the Cybersecurity and Infrastructure Security Agency, VA 22203 USA. His research interests include providing reliable evidence to support effective cybersecurity policies at various levels of vulnerability management, machine learning, and threat intelligence. Spring received a Ph.D. in computer science and philosophy of science from University College London. Contact him at spring@cisa.dhs.gov.

Lin Tan is a Mary J. Elmore New Frontiers Professor in the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA. Her research interests include software dependability, software-AI synergy, and software text analytics. Tan received a Ph.D. from the University of Illinois, Urbana-Champaign. She was the recipient of an Early Career Academic Achievement Alumni Award from the University of Illinois, Urbana-Champaign; was a Canada Research Chair; and received multiple industry awards including J. P. Morgan AI Faculty Research Awards, Meta/Facebook Research Awards, Google Faculty

Research Awards, and an IBM CAS Research Project of the Year Award. Her papers have received four ACM Distinguished Paper Awards, and one was an *IEEE Micro* Top Pick. She served as a program cochair of the ACM International Conference on the Foundations of Software Engineering in 2024. She was an associate editor of *IEEE Transactions on Software Engineering* (2017–2022) and Springer's *Empirical Software Engineering* journal (2015–2021). She was the ACM SIGSOFT treasurer and an elected member-at-large (2021–2024). Contact her at lintan@purdue.edu.

Katja Tuma is an assistant professor at the Department of Computer Science, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, The Netherlands. Her research interests include the intersection of software engineering, security and AI, and risk analysis. Tuma received a Ph.D. in computer science and engineering from the University of Gothenburg. She is the founder and coordinator of Hack4Her, the national women-focused hackathon; the coordinator of the national working group on AI for security and security for AI; and a coorganizer of the international workshop DeMeSSA. Contact her at k.tuma@vu.nl.

Call for Articles

IEEE Pervasive Computing

seeks accessible, useful papers on the latest peer-reviewed developments in pervasive, mobile, and ubiquitous computing. Topics include hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, and systems considerations, including deployment, scalability, security, and privacy.

Author guidelines:

www.computer.org/mc/pervasive/author.htm

Further details:

pervasive@computer.org
www.computer.org/pervasive



Digital Object Identifier 10.1109/MSEC.2025.3548711