resillion

Beyond The Hype: Navigating The Vulnerabilities Of AI-generated Code

Introduction

As businesses increasingly integrate AI into their operations, they must grapple with a delicate balance - harnessing the transformative power of AI solutions while addressing the myriad of risks that come with it. This critical juncture calls for a thoughtful exploration of both the immense potential and the inherent perils of AI, as we all chart our course into an AI-driven future that will undoubtedly reshape the fabric of business and society. One of the recent and growing trends in software development is AI-generated code.

This whitepaper examines the critical risks associated with AI-generated code, including potential security vulnerabilities and quality issues that could compromise the integrity of software systems. It also highlights the transformative opportunities AI-generated code offers, showcasing how this technology can present new opportunities for testing processes and enhancing code quality.

Understanding AI-generated code

AI code generation uses machine learning models to write code based on input descriptions and provide context-based suggestions. While not always perfect, AI-generated code offers developers a solid starting point, optimising the coding process with autocomplete predictions for repetitive coding patterns. This saves time and effort, reducing the need for extensive internet searches. By leveraging natural language processing and AI capabilities to detect bugs, AI code generation helps developers analyse code, identify issues, and suggest tests, enabling faster software delivery.

On one hand, it offers unprecedented opportunities for innovation, promising to speed up advancements across industries at an exponential rate. On the other, it presents us with myriads of unknown security risks — a vast, uncharted territory that defies our traditional approaches to risk management.

While AI promises increased efficiency, there are significant concerns about the quality and robustness of the code it produces. AI models, trained on generic datasets, may struggle to grasp nuanced requirements, potentially leading to code that lacks efficiency or contains hidden bugs.

In this paper, we are addressing both the risks and opportunities, without prescribing specific solutions, as these will naturally evolve as Quality Engineering companies navigate these challenges.

The rise of AI-generated code: Navigating risks and challenges

The potential for AI to generate unforeseen vulnerabilities raises serious concerns about the safety of AI-driven software development. Some of the risks that demand careful consideration are:

Code quality and performance concerns

While AI code generators offer consistency and time-saving benefits, they can be unreliable and don't always guarantee code quality or performance. The generated code, though functional, lacks the efficiency and optimisation crucial for high-performance applications. This deficiency can lead to substantial operational issues during runtime, potentially compromising overall system integrity.

According to a UC Davis research paper in 2023, titled, Large Language Models and Simple, Stupid Bugs[1], AI-generated code may overlook specific nuances in common situations. There's also a risk of inefficiency and hidden bugs because AI relies on generic data rather than a deep understanding of the specific development project.

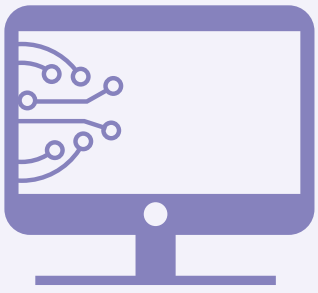

Increased risk of incorrect and insecure solutions

In industry sectors where code quality and security are paramount, the increased risk of introducing bugs and vulnerabilities through AI-generated code is a significant threat to system integrity and data security.

AI-generated code not only has the potential to give inaccurate output but also to inadvertently introduce flaws that could be exploited by malicious actors, leading to compromised systems and data breaches. The reliance on AI for coding, therefore, necessitates rigorous scrutiny to ensure both accuracy and security, emphasising the importance of human oversight and comprehensive testing protocols.

A study in 2024 by the National Institute of Standards and Technology (NIST)[2] has identified various types of adversarial attacks on AI systems, including 'evasion, poisoning, privacy, and abuse attacks'.

These vulnerabilities can result in AI systems malfunctioning, especially when manipulated by external actors.

A 2022 Stanford research paper titled Do Users Write More Insecure Code with AI Assistants?[3], highlighted a significant concern by demonstrating that 'participants who had access to Codex (Open AI's programming model) were more likely to write incorrect and 'insecure' solutions to programming problems compared to a control group'.

Organisations are actively addressing the key risks associated with generative AI, with inaccuracy and cyber security being the most frequently cited concerns. According to the 2023 McKinsey Global Survey on AI[4], these risks are not only recognised but are also prioritised for mitigation, reflecting a strategic commitment to harnessing AI's potential while safeguarding operational integrity.

Contextual understanding deficits

AI's limitations in contextual understanding can lead to significant security gaps. According to an article in Code Intelligence, titled The Risks of AI-Generated Code[5], tools generating AI code rely heavily on prompts and operate on algorithms with little contextual or project-specific understanding.

AI systems may generate code without fully understanding the broader context of the application, potentially leading to solutions that are technically correct but inappropriate for the specific use case.

A groundbreaking study in 2023, at Stanford University[6] has shed light on a significant limitation of Large Language Models (LLM) in code generation. The study reveals that 'LLMs can easily generate the syntax for a given programming language, but they struggle to use algorithmic reasoning to build complex programs with many parts'.

Skill degradation

The widespread adoption of AI coding tools raises concerns about the long-term implications for software development skills. Over-reliance on AI-generated code could potentially lead to a decline in developers' coding proficiency and their ability to critically evaluate and maintain codebases.

Reimagining quality engineering to test AI-generated code

Despite the risks, AI-generated code also presents transformative opportunities that could redefine the software development lifecycle. For Quality Engineering (QE) companies that test AI-generated code, this is a new opportunity. By integrating AI technologies, QE firms can enhance their testing processes and ensure high-quality code.

Code security and optimisation

An interesting article in GitLab published in 2024 titled AI Code Generation Explained: A Developer's Guide[7] highlights the use of AI code suggestions that are broader in scope, providing hints, improvements, and potential changes to existing code rather than merely completing the current line. These AI-powered code assistants offer refactoring options, performance enhancements, and best practice recommendations for secure code.

They base their suggestions on an analysis of the entire codebase, community standards, and the conventions specific to the programming language.

This paper also highlights AI's ability to analyse entire codebases in context, offering the potential for more thorough and efficient security testing, potentially uncovering vulnerabilities that human testers are likely to miss.

Contextual understanding in AI operates in both directions, presenting risks as well as opportunity. AI's capacity to comprehend and apply context can significantly enhance Quality Engineering.

> **"The future of AI testing is set to become more complex. As AI-generated code becomes more prevalent, we will need advanced tools and expertise to ensure both accuracy and usability."**
>
> **Jason McIvor,** Head of Transformation, Resillion

AI-powered defect prediction in quality assurance

According to the 2022 research paper on the Use of Deep Learning in Software Defect Prediction (SDP)[8], the traditional approach to identifying software defects through testing and reviews is both time-consuming and requires significant workforce resources. However, automatic prediction of defective software modules can reduce the cost of enhancing code quality when compared to a fully manual process. The main challenge of SDP is identifying the faulty parts of source code with superior fault prediction performance, an area where cutting-edge AI and deep learning techniques can be leveraged to gain a competitive advantage.

Pattern recognition for anomaly detection

AI's pattern recognition capabilities can be harnessed for advanced anomaly detection in defects and log monitoring, catching issues that traditional methods likely overlook. According to a 2023 article on Pattern Recognition in Medium titled Pattern recognition: Its use in IT system defect analysis[9] , the ability to analyse and fix system defects during the software development lifecycle, coupled with an understanding of system relationships, hinges on having the cognitive skills to identify patterns.

This is particularly pertinent for QE companies testing AI-generated code. Although few defects manifest in abstract patterns, code generally operates in a predictable manner, even when the outcome is a defect.

Ethical AI testing

As AI systems become more prevalent, there's also an opportunity to develop new frameworks for testing the ethical considerations of AI, including decision autonomy, bias, and regulatory compliance.

An interesting article in Medium written by 'Fx is AI' in 2024 titled What ethical considerations should be addressed when using AI code generators for software development?[10], highlights that AI ethics testing is a comprehensive framework designed to assess AI models for prejudice, bias, social and environmental impact, data privacy, transparency, explainability, safety, and security.

Incorporating AI ethics testing into the quality assurance process not only enhances the integrity and trustworthiness of AI systems but also positions QE firms at the forefront of ethical AI development. By ensuring AI models meet these rigorous standards, companies can offer a higher level of service, addressing critical concerns around AI fairness and transparency and significantly contributing to the development of responsible and ethical AI technologies.

This emerging field of ethical AI testing could be crucial in ensuring that AI-generated code aligns with regulatory and legal requirements.

Redefining cyber security strategies

Security's traditional role is undergoing a radical transformation, shifting from a mere protective measure to the very foundation upon which innovation thrives. In this new paradigm, security is no longer an afterthought but a strategic forethought, embedded into every product and service. As technology advances rapidly, the focus has shifted from simply mitigating risks to pre-empting them, ensuring that potential threats are anticipated and addressed from the outset. This approach is not just essential but transformative, setting the stage for a new era of secure innovation. Traditional security testing methods may be inadequate for AI-generated code, as they often fail to account for the unique vulnerabilities introduced by AI.

This code may require specialised security testing tools and methodologies, which are not yet widely available or adopted. As AI becomes more prevalent in software development, it increases the overall attack surface for potential cyber threats, making AI-generated code an attractive target for attackers due to its potential vulnerabilities.

These opportunities raise important questions. How can QA/QE companies best integrate AI into their testing processes to not just keep pace with AI-generated code but to transform QA and QE?

The recent amendment to the EU Cyber Security Act[11] introduces an EU-wide cyber security certification[12] framework for ICT products, services, and processes. This amendment aims to streamline and strengthen cyber security across the EU by allowing companies to certify their ICT products, processes, and services only once, with those certifications being recognised across all member states.

The regulation aims to enhance cyber security, cyber resilience, and trust within the European Union (EU) by strengthening the European Union Agency for Cyber Security (ENISA) with a permanent mandate and establishing a framework for voluntary cyber security certification schemes for ICT products, services, and processes.

ENISA's mandate includes achieving a high level of cyber security across the EU, supporting national and EU authorities, serving as a reference for technical advice, reducing market fragmentation, and developing its own resources while acting independently and leveraging national expertise.

Figure 2: Intriguing question for QE/QA companies to ponder

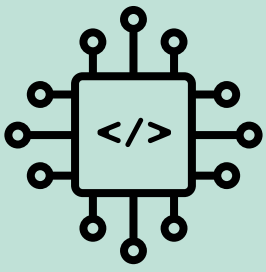

Should we allow AI to test its own code, even though it's traditionally been considered bad practice for developers to test their own work?

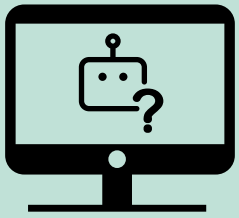

If AI's interpretation of requirements differs from human interpretation, whose version should be considered correct?

Can AI continuously self-learn and stay updated with new standards and innovations as effectively as humans?

The revolution of AI in quality engineering - Navigating new frontiers

At Resillion, our approach combines cutting-edge technology with deep expertise to address the unique challenges posed by AI-generated code and explore the opportunities it presents. These challenges and opportunities will drive transformation in the way QA and QE are approached by testing companies. Resillion is collaborating with leading academic partners, the Software Languages Lab at Vrije Universiteit Brussel and the DistriNet research unit at KU Leuven, Campus Group T to research complex defect prediction models that would allow the prediction of flaws in source code before test execution.

This proactive strategy could significantly shorten development cycles and increase software dependability.

Also, to manage the risks of inconsistent outputs from AI-generated code, we are integrating data science into QE processes to trace issues back to their root causes, even when they originate from legacy data.

Finally

Our commitment to innovation and excellence positions us at the forefront of this technological revolution, ready to guide you through the risks and complexities as you integrate AI tools into your business operations.

AI is evolving rapidly and so will its adoption. The question is, how can we harness the power of AI while ensuring the security, efficiency, and ethical implications of its outcomes?

Click to explore how Our AI Partnering Strategy will help you navigate the AI landscape.

"Generative AI has the potential to accelerate software development by automating mundane tasks, freeing up developers to focus on higher-level problem-solving and innovation. However, the speed of AI-powered development should not compromise the depth of testing. More than ever, a robust testing strategy is essential to prevent unforeseen issues"

Robby Putzeys,
Head of Software Testing Practice, Resillion

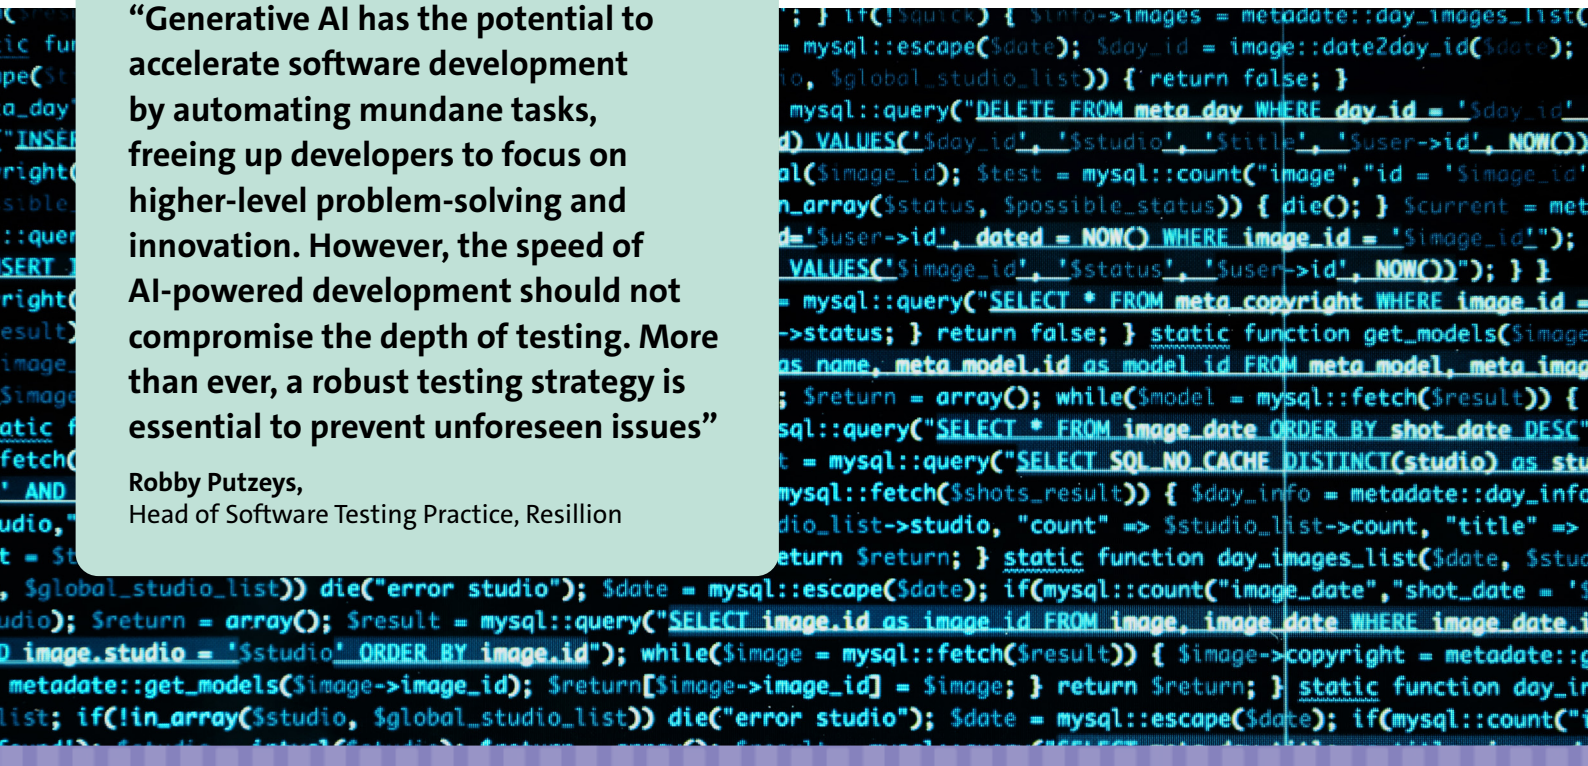

resillion

resillion

References

[1] Kevin Jesse, Toufique Ahmed, Premkumar T. Devanbu, Emily Morgan, Large Language Models and Simple, Stupid Bugs, arXiv, 2023.

[2] NIST Identifies Types of Cyberattacks That Manipulate Behavior of AI Systems, U.S. Department of Commerce, 2024

[3] Neil Perry, Megha Srivastava, Deepak Kumar, Dan Boneh, Do Users Write More Insecure Code with AI Assistants?, airXiv, 2023

[4] What's the future of generative AI?, Mckinsey & Company, 2023

[5] Sergej Dechand, The Risks of AI-Generated Code, code intelligence.

[6] Allison Whitten, New Tool Helps AI and Humans Learn To Code Better, Human Centered Artificial Intelligence,

[7] AI Code Generation Explained: A Developer's Guide, GitLab B.V., 2024

[8] Gorkem Giray, Kwabena Ebo Bennin, Omer Koksal, Onder Babur, Bedir Tekinardogan, Use of

[9] Evan Harbinson, Pattern recognition :: Its use in IT system defect analysis, Medium, 2023

[10] Fx is Ai, What ethical considerations should be addressed when using AI code generators for software development?, Medium, 2024

[11] The EU Cyber Security Act, EUR-LEX

[12] Shaping Europe's digital future, European Commission

Harnessing the combined power of human and artificial intelligence, Resillion is a leading end-to-end quality engineering company offering a comprehensive suite of cutting-edge services, including software testing, cyber security, conformance & interoperability, and digital content quality control.

Offering deep sector expertise and to partner with the world's leading companies in Media, Energy, Healthcare, Finance and Consumer from its operations in Europe, US, UK, India and China. Learn more at: www.resillion.com