

# A Path Less Traveled: Reimagining Software Engineering Automation via a Neurosymbolic Paradigm

Antonio Mastropaolo  
William & Mary  
Williamsburg, Virginia, USA  
amastropaolo@wm.edu

Denys Poshyvanyk  
William & Mary  
Williamsburg, Virginia, USA  
dposhyvanyk@wm.edu

## Abstract

The emergence of Large Code Models (LCMs) has transformed software engineering (SE) automation, driving significant advancements in tasks such as code generation, source code documentation, code review, and bug fixing. However, these advancements come with trade-offs: achieving high performance often entails exponential computational costs, reduced interpretability, and an increasing dependence on data-intensive models with hundreds of billions of parameters. In this paper, we propose *Neurosymbolic Software Engineering*, in short NSE, as a promising paradigm combining neural learning with symbolic (rule-based) reasoning, while strategically introducing a controlled source of *chaos* to simulate the complex dynamics of real-world software systems. This hybrid methodology aims to enhance efficiency, reliability, and transparency in AI-driven software engineering, while introducing controlled randomness to adapt to evolving requirements, unpredictable system behaviors, and non-deterministic execution environments. By redefining the core principles of AI-driven software engineering automation, NSE lays the groundwork for solutions that are more adaptable, transparent, and closely aligned with the evolving demands of modern software development practices.

## CCS Concepts

• **Software and its engineering** → **Software development methods**.

## Keywords

AI, Large Code Models, Software Engineering, Neurosymbolic

### ACM Reference Format:

Antonio Mastropaolo and Denys Poshyvanyk. 2025. A Path Less Traveled: Reimagining Software Engineering Automation via a Neurosymbolic Paradigm. In *33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)*, June 23–28, 2025, Trondheim, Norway. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3696630.3728720>

## 1 Introduction

During the past decades, software engineering automation has undergone a transformative evolution, driven by the rise of deep neural networks (DNN) and the development of large-scale deep

learning (DL) models with hundreds of billions of parameters, such as DeepSeek-CoderV2 [8]. The scalability of these models, achieved through extensive pre-training on diverse datasets and increasingly larger architectures, has been a cornerstone of their success. This has enabled large language models (LLMs) to develop a deep understanding of software engineering principles, automating tasks such as code generation, documentation, optimization, and bug fixing [12, 23]. These advancements have revolutionized development workflows, significantly boosting the efficiency and accuracy of software creation and maintenance.

The impact of LLMs is further evidenced by the rapid adoption of commercial AI-powered tools such as GitHub Copilot [3] and ChatGPT [2]. For example, GitHub Copilot, Microsoft's flagship AI-driven solution for software engineering, surpassed one million active users by June 2023<sup>1</sup>, underscoring the increasing dependence on AI-assisted tools within the developer community. What was once a distant vision: the deep integration of AI into software engineering has now become a tangible reality, signaling a transformative shift in the field. AI-powered tools have transitioned from experimental prototypes to essential components of modern software engineering workflows.

This rapid progress has been fueled by two key factors: (i) the unprecedented accessibility of data, particularly from public code repositories like GitHub, and (ii) significant technological advancements in hardware acceleration, such as GPUs and TPUs. These developments have enabled the scaling of LLMs and large code models (LCMs) to hundreds of billions of parameters, empowering AI-driven coding assistants like GitHub Copilot [3] and AWS CodeWhisperer [1] to deliver highly accurate, efficient, and context-aware solutions. As a result, AI-enabled software engineering has become not only feasible but also a cornerstone of modern development practices, reshaping the field in profound ways.

However, as these models become increasingly complex, the increasing demand for computational power results in higher energy consumption, higher hardware costs, and a greater environmental impact [12]. The relentless drive toward larger architectures and expanding resource requirements not only strains existing infrastructure but also raises a fundamental question: *Is indefinite scaling truly a sustainable and effective path for advancing software engineering automation?*

A possible answer lies in the work of Villalobos *et al.* [22], anticipating the beginning of a new phase, where the volume of human-generated data can no longer sustain the exponential growth of LLMs. This imminent change, which we refer to with the term **singularity of automation** marks a turning point where AI's demand for training data exceeds the production rate of human knowledge.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

FSE Companion '25, Trondheim, Norway

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1276-0/25/06

<https://doi.org/10.1145/3696630.3728720>

<sup>1</sup><https://aibusiness.com/companies/one-year-on-github-copilot-adoption-soars>

In the context of software engineering, this phenomenon was first introduced and explored in the pioneering work of Velasco *et al.* [21], highlighting the urgent need for alternative approaches beyond the mere scaling of models. In particular, Velasco’s work *et al.* challenged the conventional belief that scaling deep learning models alone is the most effective way to push the boundaries of AI-driven software engineering. In this context, the authors introduced the first Neurosymbolic framework specifically tailored for program comprehension tasks, leveraging the complementary strengths of deep learning techniques (e.g., LLMs) and traditional symbolic methods to enhance program comprehension activities.

Building on this foundation, we introduce Neurosymbolic Software Engineering (NSE), a generalized extension of the Neurosymbolic Program Comprehension (NsPC) framework, expanding its scope beyond program comprehension to encompass the full spectrum of SE-related tasks. NSE offers a compelling alternative to the indiscriminate scaling of deep learning models, while preserving the key benefits that deep learning-based techniques have introduced to the field.

A key improvement of the novel paradigm is the integration of chaos, a critical yet previously unexplored component in the original framework by Velasco *et al.* [21]. This addition allows us to address a fundamental challenge in software engineering: modeling and managing uncertainty in complex, evolving systems. To achieve this, we leverage the mathematical foundation of chaos theory [11], a powerful framework that provides a structured approach to understanding unpredictability, shifting requirements, and dynamic behaviors in software development. Unlike pure randomness, which is entirely unpredictable, a chaotic system operates under deterministic rules but is highly sensitive to initial conditions. This means that even small changes in input can lead to drastically different outcomes over time, a phenomenon widely known as the “butterfly effect”—where a minor variation in one part of a system can cascade into significant and unexpected consequences elsewhere.

In this position paper, we advocate for the integration of chaos theory into software engineering automation pipelines as an effective and efficient proxy to model “true randomness”. With that in mind, we introduce a unified vision that merges symbolic reasoning and probabilistic learning with chaos, naturally bridging between these two domains. This hybrid approach seeks to improve transparency, efficiency, trustworthiness, and sustainability in AI-driven software engineering, ensuring that automation remains scalable, interpretable, and practically applicable while capturing the dynamic aspects of real-world software systems. The remainder of this paper explores the key components of the NSE paradigm we envision, its advantages, and its broader implications for the future of AI in software engineering.

## 2 Background and Related Work

In this section, we first provide an overview of recent advances in LCMs and their impact on various software engineering tasks. We then highlight first attempts of Neurosymbolic approaches in SE.

### 2.1 Large Code Models

Recent advancements in LCMs have substantially contributed to the automation of SE tasks, allowing greater support and greater

accuracy in various SE-related processes [12, 23]. Although general-purpose LLMs, such as GPT-4 [4] and Claude [6], exhibit broad applicability across multiple domains, they often lack the domain-specific granularity required for code-related tasks. To address this limitation, researchers have developed specialized LCMs, explicitly designed to enhance SE applications. Notable examples include CodeLlama [19], StarCoder [16], Incoder [10] and DeepSeek-Coder [15]. These domain-adapted models leverage code-centric training data and employ task-aware optimizations, thereby improving the precision and automation of SE activities, including various code-related tasks. A complete analysis is provided in the Systematic Literature Review by Watson *et al.* [24].

### 2.2 Neurosymbolic AI in SE

Princis *et al.* [18] introduce a hybrid approach that combines symbolic reasoning with LCMs to improve SQL query generation. Their system uses symbolic checks for query validation and repair, along with partial query evaluation and early elimination of invalid queries, resulting in improved runtime efficiency and accuracy.

Arakelyan *et al.* [7] enhance semantic code search (SCS) systems by integrating neural and symbolic methods. Their approach uses rule-based parsing to match natural language queries with code snippets, though the manually created rules may lack generalizability across languages.

In different research, Jana *et al.* [14] introduces CoTran, a neurosymbolic system powered by large language models to translate code between programming languages. CoTran incorporates symbolic execution feedback to ensure the functional equivalence of translated code, specifically supporting translations between Java and Python. The inclusion of the symbolic component enhances the system’s ability to preserve the original code’s logic, resulting in more robust and reliable translations.

Parisotto *et al.* [17] propose a Neuro-Symbolic Program Synthesis approach to improve program induction, addressing key challenges such as high computational costs, limited interpretability, and task-specific training requirements. Their method constructs executable programs in a domain-specific language using input-output examples, demonstrating a strong generalization to unseen tasks.

In a separate line of research, Hu *et al.* [13] introduces NSEdit, a Transformer-based code repair method that predicts editing sequences to fix buggy source code across varied program structures and code repair benchmarks.

Further applications of neurosymbolic AI techniques can be found in the realms of representation learning [5], error correction [25] and semantic code repair [9].

## 3 Conceptualizing Neurosymbolic Software Engineering

The rise of LCMs has closely driven the evolution of software engineering automation, enabling breakthroughs in tasks like code generation, bug fixing, and vulnerability detection—long viewed as beyond the reach of automation (Section 1). However, this progress has also revealed new challenges that now constrain further advances in the field.

This section introduces the core components of the proposed NSE paradigm and explains how their interaction supports AI-driven

software automation (Section 3.1). We then highlight opportunities and open challenges for adopting NSE, outlining key directions for future research and practice (Section 3.2).

### 3.1 Establishing a New Paradigm

We conceptualize Neurosymbolic Software Engineering as an innovative paradigm to advance software engineering automation, harnessing the combined strengths of neural probabilistic learning and symbolic reasoning, unified through the presence of controlled chaos. The latter serves as the connective element between these two domains, enabling a synergistic collaboration that integrates the advantages of data-driven probabilistic learning, facilitating pattern recognition and deep semantic understanding, with the precision, interpretability, and correctness of symbolic reasoning, which has long been the foundation of rule-based frameworks and verification methods.

Figure 1 illustrates the three fundamental components of our envisioned NSE paradigm. At its core, this paradigm integrates the following elements:

- **Probabilistic Method:** This component capitalizes on the latest advancement in LLMs and their ability to distill knowledge from vast amounts of code and natural language data. By identifying patterns and relationships within the data, LLMs and LMCs in particular enable systems to generalize and adapt to new unseen problems. These models excel at tasks such as code completion, bug detection, and natural language understanding, providing a data-driven foundation for automation in software engineering.
- **Symbolic Method:** This component encompasses logical reasoning frameworks, constraint solvers, and rule-based inference systems. Symbolic methods ensure that software engineering processes are interpretable, correct, and aligned with formal specifications. By combining symbolic reasoning with probabilistic learning, the NSE paradigm achieves a balance between flexibility and precision.
- **Chaos-driven Component:** This component injects controlled randomness into the system to simulate real-world uncertainty. Its primary function is to bridge neural probabilistic learning with symbolic reasoning, enabling the exploration of alternative solutions when deterministic logic or statistical inference alone falls short. During training, this component can be leveraged to introduce structured noise to simulate out-of-distribution data, situations where the input deviates from what the model has previously encountered. This ability to remain robust and effective despite such variations is a significant advantage, especially in dynamic software environments. Crucially, because the chaos-driven component employs structured randomness *i.e.*, variability governed by deterministic rules, it provides a principled approach to modeling uncertainty, allowing NSE to handle unpredictability with both flexibility and precision.

Together, these three components form a cohesive framework that advances software engineering automation by combining the

strengths of probabilistic learning, symbolic reasoning, and chaos-driven adaptability. This integration ensures that the resulting systems are not only powerful and efficient but also interpretable, trustworthy, and capable of addressing the complexities of real-world software engineering challenges.

In practice, one possible instantiation of NSE intentionally introduces controlled perturbations or noise into inputs and execution conditions during testing to observe how small changes propagate through the software pipeline. These chaos-driven trials can reveal hidden vulnerabilities and edge case behaviors that purely probabilistic or deterministic methods might overlook. Unlike random testing, which can lack focus, chaos in NSE is carefully managed and targeted, guided by theoretical principles to explore critical variations without introducing irrelevant or meaningless noise. This approach not only improves testing activities, but also provides valuable benefits for program repair, code generation, and security vulnerability detection in a broader context.




Figure 1: Neurosymbolic Software Engineering Paradigm.

### 3.2 Opportunities and Open Challenges

By bridging neural learning and symbolic reasoning while incorporating deliberately controlled *chaos*, NSE introduces a new class of automation solutions with unique advantages. The following sections provide an overview of the opportunities (💡) that this paradigm offers in addressing key software engineering challenges.

Furthermore, we explore the open challenges (⚠️) that the SE community may face when transitioning to this approach. These challenges highlight critical areas that require further research, development, and innovation to fully exploit the potential of NSE and establish it as a transformative actor in the field, with the goal of fully integrating neurosymbolic principles into real-world software engineering workflows.

#### 3.2.1 Opportunities for a better, faster, stronger AI-driven software engineering automation 💡

**Improved interpretability and scalability:** Unlike black-box deep learning models, NSE incorporates rule-based logic that allows developers to trace and comprehend the reasoning behind automated suggestions and predictions. Furthermore, the presence of the symbolic reasoner facilitates performance optimization via structured constraints, avoiding the necessity to scale up the size of the training dataset and DL architecture, to improve the performance of the system on the given task.



**Trustworthiness and robustness against bias:** With the possibility to integrate formal verification thanks to the symbolic methods, NSE guarantees that the generated software engineering artifacts comply with well-defined correctness standards. Additionally, while traditional deep learning models are susceptible to adversarial perturbations, often resulting in misleading predictions, the integration of a symbolic reasoning component helps mitigate biases and robustness issues by enforcing logical consistency and rule-based validation mechanisms.

**Energy-Efficient AI for Sustainable Automation:** Scaling deep learning models requires substantial computational resources, raising concerns about their environmental sustainability [20]. NSE presents a more energy-efficient alternative by combining symbolic reasoning, structured constraints, and compact, specialized neural models. Specifically, NSE reduces the dependence on massive datasets by using structured symbolic knowledge, minimizing the need for continuous data expansion. In addition, as an extension of its scalability benefits, NSE enhances model sustainability by integrating predefined formal rules, logical constraints, and domain-specific knowledge, ultimately reducing the reliance on large-scale datasets for training while maintaining high performance and accuracy.

**Context-Aware and Adaptive Automation:** A defining strength of NSE, as envisioned, is its ability to incorporate adaptive and context-sensitive intelligence into software engineering automation. Unlike purely deep learning-based approaches, which, while efficient in executing specific tasks, often require extensive retraining to adapt to new contexts, NSE leverages symbolic reasoning combined with controlled chaos to dynamically adjust to evolving requirements, environments, and software development paradigms.

### 3.2.2 Challenges and limitations

**Handling randomness effectively:** Controlling and managing randomness to achieve the desired level of variability without compromising reliability is a challenge. The predictions of neural models can lead to unwanted behaviors that can impact the consistency and dependability of automated software engineering tasks.

**Encoding knowledge in symbolic learning systems:** Symbolic AI depends on explicitly encoded rules and logic, which require domain expertise to define. Manually encoding symbolic constraints can be labor intensive and may not scale efficiently across diverse SE tasks. Ensuring that symbolic representations remain up-to-date as software engineering practices evolve is a nontrivial challenge.

**Generalization across SE-related practices:** Despite NSE born as a generalized paradigm based on NsPC by Velasco *et al.* [21], ensuring that models generalize effectively across different tasks and software development environments is critical and time-consuming. In this regard, domain-specific optimizations may be required, making it harder to develop a one-size-fits-all NSE solution.

**Deployment challenges:** The absence of well-established frameworks, tools, and benchmarks for NSE poses significant challenges for practitioners seeking to adopt and integrate it into existing software engineering workflows. Unlike traditional DL-based systems or symbolic approaches, NSE requires a hybrid infrastructure that also contemplates chaos, via deep learning non-determinism.

This practical challenge can hinder the practical deployment of NSE-based solutions that would otherwise require interoperability with existing SE-tools, frameworks, and workflows. Bridging this gap is vital, as a lack of standardized support and tools would slow adoption, increase integration complexity, and limit the scalability of NSE-based systems.

**Keeping Pace with Rapidly Evolving Software Engineering Practices:** Software engineering is a rapidly evolving discipline, and automation approaches designed for current practices risk becoming obsolete as new methodologies, frameworks, and technologies emerge. Ensuring that NSE remains adaptable and resilient to these continuous advancements is crucial for its long-term viability and requires joint effort provided by academics and industrials.

**Ensuring Explainability Without Reducing Performance:** Increasing explainability often introduces additional computational complexity that can negatively impact performance. For example, while neural models rely on fast probabilistic approximations, symbolic reasoning methods on the other hand, may require explicit rule processing, constraint solving, or logic-based inference, which can slow down execution.

## 4 Conclusion and Future Directions

Traditional probabilistic methods or chaos theory alone, while capable of modeling uncertainty, struggle to keep pace with the scale and adaptability of LLMs—that have fundamentally reshaped the way automation is approached in software engineering. For example, LLMs’ capability to generate context-aware suggestions in real-time requires handling the nondeterministic nature of the model, which exceeds the limitations of earlier approaches based solely on statistical inference or chaotic modeling. The NSE framework addresses these challenges by integrating symbolic reasoning with neural learning and chaos-driven adaptability. The symbolic component ensures that the results of LLMs are interpretable and verifiable, enforcing correctness and providing structured inference. Meanwhile, the chaos-driven component simulates real-world uncertainty, helping LLMs adapt to changing environments and unseen data. This hybrid approach preserves the flexibility and scalability of LLMs while also ensuring their reliability and transparency in the face of dynamic, unpredictable software engineering tasks that increasingly rely on seamless collaboration between humans and intelligent systems to achieve the desired goals.

Looking ahead, we envision NSE as a key component of the future of software engineering automation, thanks to its ability to integrate controlled uncertainty. By adopting this hybrid approach, NSE facilitates more dependable, sustainable, and flexible automation solutions to address the complexities underlying modern software engineering practices that are blending with LLMs in increasingly dynamic, data-rich, and non-deterministic development environments. This integration enables systems not only to perform traditional tasks such as code generation and bug fixing, but also to adapt intelligently to changing requirements, ambiguous input, and evolving software ecosystems.

## Acknowledgment

Dr. Poshyvanyk acknowledges the support of the National Science Foundation (NSF) under grant CCF-2311469.

## References

- [1] [n. d.]. Amazon CodeWhisperer. <https://aws.amazon.com/codewhisperer/>. Accessed: 2024-03-10.
- [2] [n. d.]. ChatGPT <https://openai.com/blog/chatgpt>.
- [3] [n. d.]. GitHub Copilot – Your AI pair programmer. <https://github.com/features/copilot/>. Accessed: 2024-03-10.
- [4] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [5] Miltiadis Allamanis, Pankajan Chanthirasegaran, Pushmeet Kohli, and Charles Sutton. 2017. Learning continuous semantic representations of symbolic expressions. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) (ICML '17). JMLR.org, 80–88.
- [6] Anthropic. n.d.. The Claude 3 Model Family: Opus, Sonnet, Haiku. [https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model\\_Card\\_Claude\\_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf). Accessed: 2025-02-26.
- [7] Shushan Arakelyan, Anna Hakhverdyan, Miltiadis Allamanis, Luis Garcia, Christophe Hauser, and Xiang Ren. 2024. NS3: neuro-symbolic semantic code search. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '22). Curran Associates Inc., Red Hook, NY, USA, Article 761, 16 pages.
- [8] DeepSeek-AI, Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y. Wu, Yukun Li, Huazuo Gao, Shirong Ma, Wangding Zeng, Xiao Bi, Zihui Gu, Hanwei Xu, Damai Dai, Kai Dong, Liyue Zhang, Yishi Piao, Zhibin Gou, Zhenda Xie, Zhewen Hao, Bingxuan Wang, Junxiao Song, Deli Chen, Xin Xie, Kang Guan, Yuxiang You, Aixin Liu, Qiushi Du, Wenjun Gao, Xuan Lu, Qinyu Chen, Yaohui Wang, Chengqi Deng, Jiashi Li, Chenggang Zhao, Chong Ruan, Fuli Luo, and Wenfeng Liang. 2024. DeepSeek-Coder-V2: Breaking the Barrier of Closed-Source Models in Code Intelligence. *arXiv:2406.11931* [cs.SE] <https://arxiv.org/abs/2406.11931>
- [9] Jacob Devlin, Jonathan Uesato, Rishabh Singh, and Pushmeet Kohli. 2017. Semantic Code Repair using Neuro-Symbolic Transformation Networks. <http://arxiv.org/abs/1710.11054> arXiv:1710.11054.
- [10] Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Wen-tau Yih, Luke Zettlemoyer, and Mike Lewis. 2022. InCoder: A Generative Model for Code Infilling and Synthesis. *CoRR* abs/2204.05999 (2022). doi:10.48550/arXiv.2204.05999
- [11] James Gleick. 2008. *Chaos: Making a new science*. Penguin.
- [12] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2024. Large language models for software engineering: A systematic literature review. *ACM Transactions on Software Engineering and Methodology* 33, 8 (2024), 1–79.
- [13] Yaojie Hu, Xingjian Shi, Qiang Zhou, and Lee Pike. 2022. Fix Bugs with Transformer through a Neural-Symbolic Edit Grammar. doi:10.48550/arXiv.2204.06643 arXiv:2204.06643.
- [14] Prithwish Jana, Piyush Jha, Haoyang Ju, Gautham Kishore, Aryan Mahajan, and Vijay Ganesh. 2024. CoTran: An LLM-based Code Translator using Reinforcement Learning with Feedback from Compiler and Symbolic Execution. doi:10.48550/arXiv.2306.06755 arXiv:2306.06755.
- [15] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437* (2024).
- [16] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. Starcoder 2 and the stack v2: The next generation. *arXiv preprint arXiv:2402.19173* (2024).
- [17] Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. 2016. Neuro-Symbolic Program Synthesis. doi:10.48550/arXiv.1611.01855 arXiv:1611.01855.
- [18] Henrijs Princis, Cristina David, and Alan Mycroft. 2024. Enhancing SQL Query Generation with Neurosymbolic Reasoning. doi:10.48550/arXiv.2408.13888 arXiv:2408.13888.
- [19] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).
- [20] Jieke Shi, Zhou Yang, and David Lo. 2024. Efficient and green large language models for software engineering: Vision and the road ahead. *ACM Transactions on Software Engineering and Methodology* (2024).
- [21] Alejandro Velasco, Aya Garryeva, David N Palacio, Antonio Mastropaolo, and Denys Poshyvanyk. 2025. Toward Neurosymbolic Program Comprehension. *arXiv preprint arXiv:2502.01806* (2025).
- [22] Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbhahn. 2024. Will we run out of data? Limits of LLM scaling based on human-generated data. doi:10.48550/arXiv.2211.04325
- [23] Cody Watson, Nathan Cooper, David Nader Palacio, Kevin Moran, and Denys Poshyvanyk. 2022. A Systematic Literature Review on the Use of Deep Learning in Software Engineering Research. *ACM Trans. Softw. Eng. Methodol.* 31, 2, Article 32 (March 2022), 58 pages. doi:10.1145/3485275
- [24] Cody Watson, Nathan Cooper, David Nader Palacio, Kevin Moran, and Denys Poshyvanyk. 2022. A Systematic Literature Review on the Use of Deep Learning in Software Engineering Research. *ACM Transactions on Software Engineering and Methodology*, TOSEM 31, 2 (2022), 1–58.
- [25] Min Xue, Artur Andrzejak, and Marla Leuther. 2024. An interpretable error correction method for enhancing code-to-code translation. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=fVxiEHGnVT>