

Summary Review: On the Naturalness of Software

Eimaan Saqib
e2saqib@uwaterloo.ca
University of Waterloo
Waterloo, Ontario, Canada

ACM Reference Format:

Eimaan Saqib. 2018. Summary Review: On the Naturalness of Software. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 2 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 PAPER SUMMARY

The paper explores the use of statistical models trained on large code corpora to improve various aspects of software engineering. The idea is that like natural language, code has patterns and regularities that can be captured with probabilistic techniques. The paper uses n-gram language models to capture regularities by predicting the probability of a token based on the preceding context. The authors use the models to show that programming languages have predictable patterns like human language.

A major challenge is data sparsity. The paper suggests using smoothing techniques to handle the issue. The authors compare cross-entropy loss generated by training models on corpora with English language vs Java projects, between different Java projects by training on one project and testing on others, and across Ubuntu application domains. They find that code has more predictable patterns than natural language, the models train to patterns specific to the project instead of just learning patterns in the language used, and that models also train to recognize patterns specific to the respective domains.

The authors also create a code suggestion heuristic by merging suggestions from the Eclipse IDE and their trained n-gram model. They find that the merged heuristic consistently performs better than just using the heuristic-based Eclipse suggestion engine.

2 POSITIVE POINTS

- (1) The paper explores statistical language models for code, which laid the groundwork for tools like GitHub Copilot, CodeT5, and Codex. This early insight into "naturalness" in code could have helped inspire deep learning-based approaches that now dominate AI-assisted programming.
- (2) The paper suggests that statistical models can capture different aspects of code structure (syntax, type, scope, and semantics). Modern approaches, such as transformer-based

models, have indeed shown this to be true, validating the paper's hypothesis.

- (3) Instead of just focusing on autocompletion, the paper explores applications in error detection, accessibility, and machine translation of code, which are still relevant research directions.
- (4) The idea of using statistical methods to approximate expensive static analysis tools is insightful. Even today, researchers are trying to balance soundness with efficiency in static analysis tools.

3 NEGATIVE POINTS

- (1) The paper briefly mentions deep learning as a possible future direction, but by 2016, neural networks were already proving effective for natural language tasks. The paper could have explored neural approaches (such as LSTMs, which were popular then) in more depth.
- (2) The primary approach discussed is based on n-gram language models, which are effective for local patterns but struggle with long-range dependencies. Modern tools use transformers but even at the time, the paper could have used neural-network based techniques.
- (3) The paper does not address issues like bias in training data, security risks, or ethical concerns about AI-generated code.
- (4) The paper does not address the computational cost of training and using such models in real-world software engineering workflows.

4 FUTURE WORK

- (1) Future work could examine how transformer-based models like Codex, Copilot, or StarCoder outperform statistical models and how hybrid approaches (e.g., combining statistical heuristics with deep learning) could enhance reliability.
- (2) Future research could explore combining textual descriptions, diagrams, and execution traces to build better software understanding tools.

5 RATING

4.5/5

The paper was ahead of its time in applying statistical and probabilistic models to software engineering tasks like code completion, bug detection, and summarization. There could have been a discussion on scalability and ethical/security concerns. But overall this is still a great paper for its time.

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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

6 DISCUSSION POINTS

- (1) **The Role of Statistical Models vs. Deep Learning in Software Engineering:** The discussion began with an exploration of whether simple statistical models still have relevance in an era dominated by deep learning (DL). While DL has become the mainstream approach, participants agreed that this should not lead to the outright dismissal of statistical methods. Key points raised include:

- Deep learning models are resource-intensive, requiring significant computational power and large datasets, which may not always be practical. A statistical model that achieves 90% effectiveness at a fraction of the cost still holds value.
- Some tasks, such as syntax-based code completion, may be better suited for traditional models like syntax trees rather than DL-based approaches. While deep learning excels in broader suggestion tasks, explainability and efficiency remain strengths of statistical models.
- The professor noted that with the increasing cost of running large-scale DL models, practical alternatives should not be overlooked, especially when explainability is a concern.

- (2) **Trade-off Between AI-Assisted Coding Speed and Code Quality:** The class debated whether tools like GitHub Copilot and ChatGPT improve productivity at the cost of software quality. The following perspectives emerged:

- AI tools are often overly confident in their responses, which can mislead developers. One student suggested using them for idea generation rather than direct implementation, as AI-generated code may work in one context but fail in another.
- The professor expressed concerns that companies might prioritize cost-cutting over quality by replacing developers with AI, raising ethical and workforce-related issues.
- Some students argued that programmers would still be needed for oversight. However, they emphasized that blindly trusting AI-generated code without careful validation could be harmful.

- (3) **AI Coding Tools and Reinforcement of Biases:** The discussion also touched on how AI models, trained on large datasets, might perpetuate biases present in human-written code. Key concerns included:

- If training data contains insecure, inefficient, or biased code, AI models might reinforce bad programming practices.
- There is an open question about whether AI should actively correct such biases, even if that contradicts human developers' preferences.

- (4) **AI-Driven Software Maintenance: Who Maintains AI-Written Code?:** A major concern was the long-term maintenance of AI-generated software. The discussion revolved around the role of human engineers in ensuring AI-driven code remains sustainable:

- Some students argued that instead of reducing the workforce, companies should invest in maintaining AI systems while ensuring compliance with ethical and technical standards.

- The professor expressed skepticism, citing past trends where companies prioritized short-term profits over long-term technical needs. He referenced the shift in corporate strategies during the 1990s, where profitability began taking precedence over technological innovation.

- One student countered that industries such as gaming have faced backlash for prioritizing shareholder profits at the expense of quality, suggesting that a similar pushback may occur in AI-driven software development.

- (5) **The Concept of "Naturalness" in Code:** The final discussion centered on whether the notion of "natural" programming should always be the goal:

- Some students argued that predictability in code is beneficial, as programming is fundamentally about communication between developers. Readable and structured code improves maintainability.
- However, others pointed out that certain domains, such as security, require unpredictability. For instance, passwords and cryptographic implementations should not follow easily guessable patterns.
- The discussion concluded that while naturalness is useful for general software development, there are cases where deviation from common patterns is necessary.