



Delft University of Technology

The (ab)use of Open Source Code to Train Large Language Models

Al-Kaswan, A.; Izadi, M.

DOI

[10.1109/NLBSE59153.2023.00008](https://doi.org/10.1109/NLBSE59153.2023.00008)

Publication date

2023

Document Version

Final published version

Published in

Proceedings of the 2nd International Workshop on NL-based Software Engineering

Citation (APA)

Al-Kaswan, A., & Izadi, M. (2023). The (ab)use of Open Source Code to Train Large Language Models. In *Proceedings of the 2nd International Workshop on NL-based Software Engineering* (pp. 9-10). (Proceedings - 2023 IEEE/ACM 2nd International Workshop on Natural Language-Based Software Engineering, NLBSE 2023). <https://doi.org/10.1109/NLBSE59153.2023.00008>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

The (ab)use of Open Source Code to Train Large Language Models

Ali Al-Kaswan
Delft University of Technology
Delft, The Netherlands
a.al-kaswan@tudelft.nl

Maliheh Izadi
Delft University of Technology
Delft, The Netherlands
m.izadi@tudelft.nl

Abstract—In recent years, Large Language Models (LLMs) have gained significant popularity due to their ability to generate human-like text and their potential applications in various fields, such as Software Engineering. LLMs for Code are commonly trained on large unsanitized corpora of source code scraped from the Internet. The content of these datasets is memorized and emitted by the models, often in a verbatim manner. In this work, we will discuss the security, privacy, and licensing implications of memorization. We argue why the use of copyleft code to train LLMs is a legal and ethical dilemma. Finally, we provide four actionable recommendations to address this issue.

I. LANGUAGE MODELS FOR CODE

Large Language Models (LLMs) have gained significant attention in the field of Natural Language Processing (NLP) in recent years due to their ability to perform a wide range of NLP tasks with impressive accuracy. These models, trained on massive amounts of data, improve in accuracy as they grow from millions to billions of parameters. LLMs for code are trained on massive amounts of data and can learn the structure and syntax of programming languages, making them well-suited for tasks such as code summarization, generation, and completion [1, 2]. LLMs are even making their way into commercial products like GitHub’s Copilot, Replit’s GhostWriter and Tabnine. Meanwhile, some have identified that LLMs can memorize large swaths of training data [3]. Memorization enables the extraction of the data using Data Extraction Attacks. Some attacks have even been able to extract addresses and other personal information from public models [3]. Memorization also impacts LLMs for code, with all its associated consequences. We will discuss these consequences in three categories: security, privacy and licensing.

II. SECURITY IMPLICATIONS

Text memorization has strong security implications. Firstly, massively mined code datasets are not sanitized or manually curated, the datasets could therefore contain many biases,¹ and instances of badly written or buggy and insecure code. A recent study found that around 40% of GitHub Copilot’s code generations for MITRE’s top 25 Common Weakness Enumerations, a list of the most dangerous software weaknesses were found to be vulnerable [4]. If these models become more prevalent and trusted, they can introduce more vulnerable code into software.

III. PRIVACY IMPLICATIONS

Memorization enables adversaries to access training data, and everything contained within, simply by accessing the model. This has major privacy implications since code can contain private information. Think of credentials, API keys, directory structures, logged info, or in-code discussions by developers. Code can also contain personal information like emails or contact information. If personal data is published on the Internet, the data could be retracted and deleted from the source. But once it is mined and used to train an LLM, the information is forever embedded in a compact representation, which is queryable at scale. With query access to these models, an adversary can potentially extract this data [3] and threaten Internet users’ privacy. There are many reasons why one could publicly share private information; (1) simply by accident, or (2) a malicious actor could share this information in a doxing campaign [3]. Even if the data is published willingly, the owner has a certain use and audience in mind and might not wish to share this information with the entire world. This is referred to as the re-purposed data problem.¹

IV. LICENSING

Publicly available source code is also subject to licences, some of which heavily regulate the use of the material. Initially, developers raised concerns about licensed code on social media. GitHub Copilot could be prompted to produce verbatim copies of copyrighted code, without providing the required attribution or licence terms.² Similarly, Copilot was producing copyrighted code while attributing the wrong author and providing the wrong license.³ Later, a lawsuit was filed against GitHub, Microsoft and OpenAI, claiming that Copilot is violating the licence of open-source code.⁴

Broadly, open-source code is licensed under two types of licences. **Permissive licences**, allow users to use, modify, and distribute the software for any purpose, without requiring that the user share their work. **Non-permissive licences**, also known as “copyleft” licences, require that users freely share their own software under the same licence if they distribute

¹Does GPT-2 Know Your Phone Number?: <http://archive.is/LxSyA>

²Matrix Transpose: <http://archive.is/YU5BI>

³Fast Inverse Square: <http://archive.is/HNiyg>

⁴GitHub Complaint (p.26): <http://archive.is/3PFAs>

the software or any *derivatives* of it. Creating closed or commercial software based on non-permissively licensed code is unethical and possibly even illegal [5]. But this does raise the following question: **Does training LLMs on copyleft code infringe on their license?**

Firstly, we must determine how many LLMs for code are trained on copyleft code. Looking at some of the most popular code models, we can observe that the vast majority are trained on open-source code. CodeBERT and CodeT5 are trained on CodeSearchNet, which contains copyleft code. We also found that CodeBERT, CodeGen, and CodeClippy make use of The Pile, a collection of 22 datasets, one of which is a GitHub dataset containing copyleft data.

We found that only InCoder makes an effort to prevent training on copyleft code. InCoder does however make use of a dataset of StackOverflow answers, which are licensed under varying CC-BY-SA licences, all of which require attribution.⁵

Despite the public attention, it is not completely clear whether Codex, the model behind Copilot, is trained on non-permissive code. Many imply that it does,⁶ citing the copying of copyleft code and the fact that the system has encountered a copy of the GPL licence many times during training. The training data for Codex is not publicly available, and neither OpenAI nor Github have provided any clarification.⁷

LLMs for code can be seen as derivatives of their training data. So unless the model is published under the same licence as the training data and includes the copyright notice, this would be a clear violation. Moreover, many licences are not inter-compatible, i.e., the inclusion of code licensed under them automatically warrants an infringement as the combined licence agreements contain irreconcilable conditions.

Some opponents,⁸ including OpenAI, argue that the use of public code is an instance of transformative fair use, which is a defence that allows the use of copyrighted works in new and unexpected ways and exists in many jurisdictions including the US.⁸ Yet, it is still unclear whether the fair use defence applies to ML-systems,⁴ as it has not yet been tested in court. Furthermore, the fair use argument is sometimes based on the assumption that models do not memorize and emit training data, which is false. Even if the fair use argument protects the use of the data, the verbatim outputting might not be protected.

A moral argument can also be made on this issue. Training LLMs on copyleft code goes against the will of some open-source developers, who share their code for the betterment of society and who believe in the principle of free and open software so profoundly, they're willing to add a full legal clause to their work to perpetuate this ideal. The use of their work, without attribution, especially by commercial parties, is not what they had in mind.

Finally, some researchers have also proposed a different approach to this issue by letting the authors of open-source code take matters into their own hands. Using data poisoning

techniques, the authors can reduce the performance and embed watermarks into the models [5].

V. DISCUSSION AND RECOMMENDATIONS

To conclude we recommend the following:

- The ML community should carefully consider the licence of their training material, from both a legal and an ethical point of view. The authors of published LLMs should be transparent about the licences of their training material.
- More research should be conducted on the nature and proportionality of text memorization in LLMs for code and LLMs in general. Other topics include memorized text extraction and prevention.
- Lawmakers should clarify whether the use of copyleft code (and copyrighted materials in general) and text to train LLMs constitutes fair use and under which conditions this clause applies.
- Finally, the software engineering community should clarify their stance on this issue. Developers could make informed decisions and clearly denote if their source code can be used to train AI models. LLMs for code are likely to stay and bring new tools that change the way software is engineered. So the community needs to answer important questions on this matter. For instance, should open-source code be allowed for training these models? If so, should the developers be credited and compensated, and under which license should the models be released? Alternatively, do we need to revise current code licenses to clarify the community's stance?⁹

REFERENCES

- [1] A. Al-Kaswan, T. Ahmed, M. Izadi, A. A. Sawant, P. Devanbu, and A. van Deursen, "Extending source code pre-trained language models to summarise decompiled binaries," in *Proceedings of the 30th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2023.
- [2] M. Izadi, R. Gismondi, and G. Gousios, "Codefill: Multi-token code completion by jointly learning from structure and naming sequences," in *Proceedings of the 44th International Conference on Software Engineering (ICSE)*. ACM, 2022, p. 401–412.
- [3] N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson *et al.*, "Extracting training data from large language models," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2633–2650.
- [4] H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, and R. Karri, "Asleep at the keyboard? assessing the security of github copilot's code contributions," in *IEEE Symposium on Security and Privacy (SP)*, 2022, pp. 754–768.
- [5] Z. Sun, X. Du, F. Song, M. Ni, and L. Li, "Coprotector: Protect open-source code against unauthorized training usage with data poisoning," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 652–660.

⁹Additional Reading Material: Link to our GitHub Repository

⁵StackOverflow license: <http://archive.is/obaoy>

⁶Comment on Copilot and OSS: <http://archive.is/6gE0U>

⁷If Software is My Copilot, Who Programmed My Software?: Link

⁸GitHub Copilot is not infringing your copyright: <http://archive.is/PY1m5>