

Towards Assurance of LLM Adversarial Robustness using Ontology-Driven Argumentation

Tomas Bueno Momcilovic^{1,*}, Beat Buesser², Giulio Zizzo³, Mark Purcell³ and Dian Balta¹

¹fortiss GmbH Research Institute, Munich, Germany

²IBM Research Europe, Zurich, Switzerland

³IBM Research Europe, Dublin, Ireland

Abstract

Despite the impressive adaptability of large language models (LLMs), challenges remain in ensuring their security, transparency, and interpretability. Given their susceptibility to adversarial attacks, LLMs need to be defended with an evolving combination of adversarial training and guardrails. However, managing the implicit and heterogeneous knowledge for continuously assuring robustness is difficult. We introduce a novel approach for assurance of the adversarial robustness of LLMs based on formal argumentation. Using ontologies for formalization, we structure state-of-the-art attacks and defenses, facilitating the creation of a human-readable assurance case, and a machine-readable representation. We demonstrate its application with examples in English language and code translation tasks, and provide implications for theory and practice, by targeting engineers, data scientists, users, and auditors.

Keywords

assurance, LLM, adversarial robustness, argumentation, ontologies

1. Introduction

Large language models (LLMs) have shown promise in various natural and domain-specific language tasks [1, 2], even without further training [3]. However, challenges hinder their trustworthiness [4], as LLMs have an inscrutable structure and dynamicity that make them a moving target for safety and security research [5]. In particular, they are brittle against adversarial attacks; slight perturbations in the input can cause a model to provide malicious output [6], and guardrails can often only be introduced post-incident [7].

Given the novelty and fast-paced evolution of LLMs, engineers need to rely on preprints and experiments (cf. [8]) to analyse the impact of novel attacks and envision suitable defenses. Unlike software security, for which maintained knowledge bases exist (e.g. Common Vulnerability Enumerations [9]), no such process is established for LLMs. Consequently, the required knowledge is captured in the data, code, documentation, and brains of individuals. This implicit knowledge base for assurance may not capture the entire picture of attacks and confidence in defenses over time. For instance, a very recent example by Microsoft shows extremely effective

Late-breaking work, Demos and Doctoral Consortium, colocated with The 2nd World Conference on eXplainable Artificial Intelligence: July 17–19, 2024, Valletta, Malta

*Corresponding author.

✉ momcilovic@fortiss.org (T. B. Momcilovic); beat.buesser@ibm.com (B. Buesser); giulio.zizzo2@ibm.com (G. Zizzo); markpurcell@ie.ibm.com (M. Purcell); balta@fortiss.org (D. Balta)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

“multi-turn jailbreaks” across LLMs, which would require engineers to redesign the existing defenses by combining heterogeneous knowledge: attack model and history analysis, prompt and response analysis, turn-pattern analysis, turn-by-turn and overall defenses [10].

Hence, the question we seek to address is: *How can one continuously assure that an LLM is robust enough against adversarial attacks in a particular domain?* In this research-in-progress work, we propose an assurance approach that allows for structuring the heterogeneous knowledge about LLM attacks and defenses (cf. [11]), as well as the application domain. We handle the knowledge involved in creating assurance arguments explicitly and comprehensively based on ontological models. The latter allow for a formal argumentation along human-readable assurance cases expressed with machine-readable ontologies, thus creating a shared understanding about training, guardrails, and implementation.

2. Background and Related Work

LLMs are neural network models that are pre-trained on a large amount of text data and have been shown to be capable of predicting, translating, or generating text for natural [2] and programming languages [1].

Traditional **adversarial attacks** add imperceptible perturbations δ to a given data point x so that a classifier f predicts $f(x) = c$ and $f(x + \delta) = c'$ where $c \neq c'$. Attacks on LLMs involve malicious prompts bypassing guardrails or model alignment to obtain harmful outputs [6]. Obtaining such prompts includes gradient-based optimizations of the input [6], persuasion patterns to bypass guardrails [12], and model inversions to generate vulnerable code in non-natural-language tasks [13]. **Robustness defenses** are similarly developing and highly heterogeneous; they include, for example, perplexity filters against gradient-based suffix-style attacks [14], estimation of the brittleness of jailbreaks [7], and instructions for LLMs to detect harmful prompts [15].

Assurance is the process of structuring an argument from claims about a system and its environment that are grounded by evidence [16]. An **assurance case** is a bundle of arguments, used to assess the level of confidence in a particular quality of a system [17] in a domain. Assurance cases have been shown to be suitable for complex and rapidly evolving AI technologies [18], and also usable for structuring claims about explainability and interpretability [11].

Assurance of AI security draws on traditional methods such as verification with test libraries [19], validation with human feedback [20], and manual [21] and automated [22] stress testing. However, the inscrutability of AI has motivated the proliferation of experimental interpretability [23], auditing [24], and forensic [25] methods to investigate the causes of problematic output. Research which makes use of both approaches includes the work of Kläs *et al.* [26] on risk-based assurance cases for autonomous vehicles, and Hawkins *et al.* [18] on a dynamic assurance framework for autonomous systems.

Since arguments may cover heterogeneous knowledge about the technology and its domain, **knowledge formalization** proves valuable for creating a common understanding. Knowledge representation and reasoning is a field of AI research [27], covering topics such as formalization based on ontologies to support explainable AI [28]. An **ontology** is “an explicit specification of a conceptualization” (p. 199, [29]) that allows machine-readable knowledge to be shared

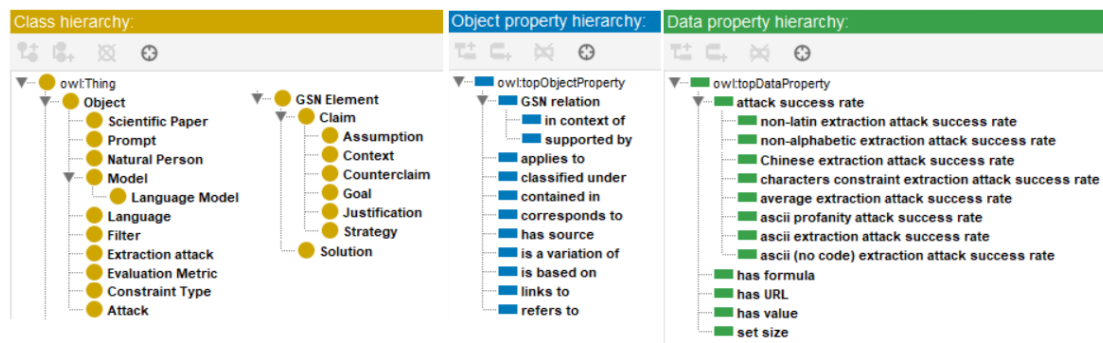


Figure 1: Structure of the ontology. Visualized using Stanford Protégé.

between humans in a common vocabulary.

While the combination of ontologies and assurance cases is not entirely novel - Gallina *et al.* [30] propose such a framework for assuring AI conformance with the EU Machinery Regulation - we note that continuously formalizing, assuring and reasoning about LLM security is a novel proposition. Our approach links two graph representations in the same ontology: a non-hierarchical, mixed-direction acyclic graph of attacks and defenses in the LLM’s application domain, and a hierarchical directed acyclic graph of corresponding claims and evidence about its robustness. The elements of both graphs are represented as subject-predicate-object semantic triples using the Resource Description Framework [31] and Web Ontology Language [32]. We additionally make use of the Goal Structuring Notation (GSN) metamodel [16] to structure assurance cases (cf. Figure 2) with goals (G), strategies (S), solutions (Sn), contexts (C) and justifications (J), and add attacks as counterclaims (CC) following community practice [33].

3. Assurance with Ontology-Driven Arguments

3.1. Robustness in Natural Language Tasks

Recent experiments show that simple attacks can have high success rates in the natural language application domain [6]. For example, Geiping *et al.* [8] demonstrate that in most tests, particular characters in seemingly benign prompts (e.g., Latin, Chinese, ASCII) can successfully induce a particular response from many pre-trained open-source LLMs (e.g., LLaMa-2 with 7 billion parameters). For example, an attack is deemed successful if an LLM responds with profanities (i.e., profanity attacks) or reveals its hidden system instruction (i.e., extraction attacks).

Several options can help reduce the vulnerability of an LLM to such attacks. Retraining the LLM to be robust to character-specific perturbations [34] is arguably more secure than simply filtering the input based on prompt properties [14], but also more resource and time intensive. Thus, an engineer may decide to combine defenses in stages: add a naive input filter to exclude prompts with reportedly “risky” character types in the short-term; perform experiments with benign and adversarial prompts, reconfiguring the filter to adjust the parameters according to results in the medium-term; and adversarially retrain the LLM to be deployed in the longer-term.

We develop an ontology that formalizes the relations between concepts (i.e., LLM, attack

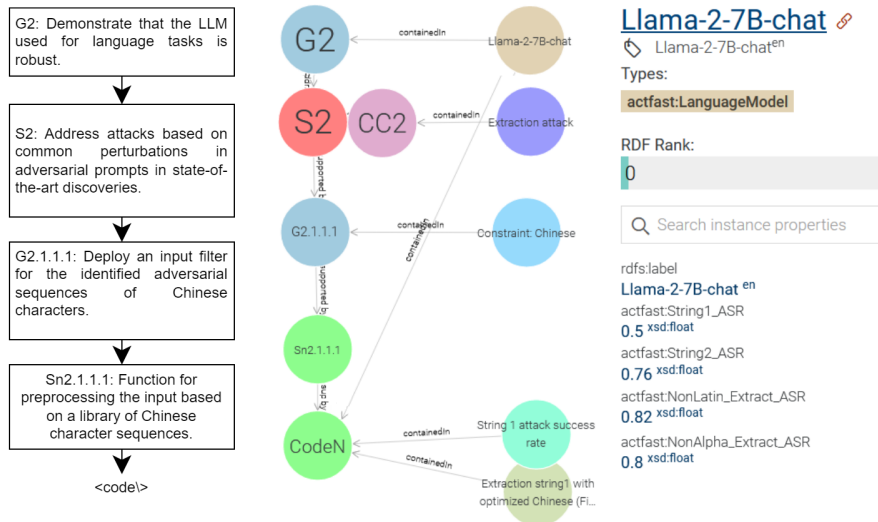


Figure 2: Assurance case in GSN (left), connected graphs (center), and the queriable values in concepts (right). The boxes are a manual illustration of the text in the argument.

and constraint type) and variables (i.e., attack success rate, character) as described in the paper [8]. The ontology is implemented with a trivial structure (cf. Figure 1), consisting of classes, object properties (i.e., relations) and data properties (i.e., values). In the example provides the attack success rate (e.g., `String1_ASR`: 0.5) of an individual attack (e.g., adversarial extraction-type prompt with Chinese-English characters) with the LLM (e.g., LLaMa-2-7B-chat) and the constraint under which the attack functions (e.g., Chinese language characters).

The ontology allows attack- and defense-relevant values to be retrieved, calculated, and inserted with complex queries, while showing the argument and architecture to readers. We posit that this setup and pipeline (cf. Figure 3) separates the following maintenance concerns while providing an explainable representation of robustness: (i) explication and structuring of the approaches to defend from adversarial attacks; (ii) continuous reasoning against changes by querying the parameter values from a central repository; (iii) inserting and maintaining values in the ontology based on experiments or external empirical data; and (iv) auditing the design and effectiveness of the operationalized robustness in the LLM (cf. Figure 2).

3.2. Robustness in Code Translation Tasks

LLMs used for domain-specific language tasks can similarly be susceptible to simple adversarial attacks [13]. We present a toy example where a function for calculating the factorial of a number is translated from C++ to Python. While users could attempt to jailbreak or translate intentionally harmful code, they may also be unaware of potential vulnerabilities in the input or output. These naive requests can happen with large codebases, imprecise mappings between languages, or users who lack security awareness or proficiency.

Regardless of the user's intent, the engineer could want to ensure that the LLM is not generating harmful code. Robustness would then include a sequence of specific claims (S1; G1.5)




Figure 3: Pipeline for representing data about attacks and defenses into an ontology and assurance case, for further reasoning by middleware implementing the LLM defenses. Original contributions emphasized.

about various defenses (cf. Figure 4). We show claims about three example mechanisms for the given context (C1.5): perplexity input filter (G1.5.1; [14]); and code analysis output filters to detect injections (G1.5.3; [9, 35]) or lack of input sanitization (G1.5.4; [36]).

Input filters could detect malicious requests. Perplexity (Sn1.5.1) filtering is a mechanism for determining if the prompt is an outlier (e.g., a gradient-based attack) by comparing it with the properties of data on which the LLM was trained. Randomization of input may not lead to comprehensible or functioning code in the prompt, but depending on the LLM training (e.g., helpfulness, correctness) and application (e.g., code autocompletion), the LLM may still generate executable output with vulnerable, malicious or toxic elements.

When an input filter fails to detect an attack, or the LLM generates problematic code from benign prompts, an engineer can rely on output filters. Code analysis, for example, could flag vulnerable or malicious elements with manually defined software tests [9, 36] or automatic queries from externally maintained tools [35]. Such flags can be treated differently. For vulnerable code, the LLM could provide three aspects in the same output: the translated function; a warning that end-users of the function could create problems with wrong or intentionally manipulated input (CC1.5.4), unless inputs are sanitized (Sn1.5.4); and error message patterns to fix this (Sn1.5.3.1; Sn1.5.4.1). For malicious code, such as a request to translate an injection that bypassed the input filter, the filter could prevent the translation from reaching the user without affecting the helpfulness of the LLM (Sn1.5.3).

4. Conclusion

In this research-in-progress paper, we explore assuring the robustness of LLMs using human-comprehensible assurance cases and machine-comprehensible semantic networks in ontologies. We show that our approach can be implemented alongside the LLM-based system, to make its robustness explainable by providing metadata for code variables, encoding the dependencies explicitly, and making the evidence transparent. Implications for researchers include studying different types of claim and evidence, as well as notations towards a shared knowledge for LLM assurance. Implications for practitioners include a novel idea for proactively engineering adversarially robust LLMs. Future work will center on exploring and evaluating this approach

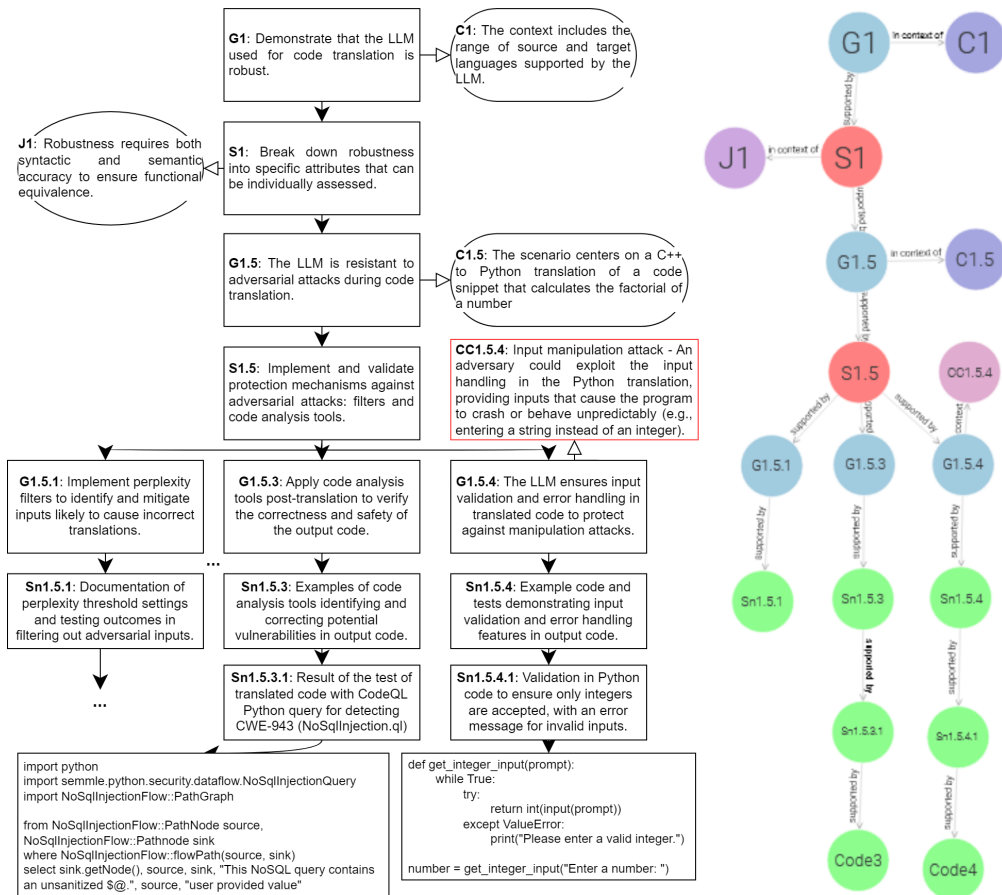


Figure 4: Sketched argument (left) and its representation in ontology (right).

with real-life implementations and industrial use cases, as well as addressing the limitation of manually formalizing arguments and ontologies, to cover various attacks and improve maintainability over time.

Acknowledgments

This research was partially funded by the Bavarian Ministry of Economic Affairs, Regional Development and Energy. We thank Kevin Eykholt for his valuable input. We also thank the reviewers for their valuable comments.

References

- [1] A. Fan, B. Gokkaya, M. Harman, M. Lyubarskiy, S. Sengupta, S. Yoo, J. M. Zhang, Large language models for software engineering: Survey and open problems, arXiv preprint arXiv:2310.03533 (2023).

- [2] F. Shi, M. Suzgun, M. Freitag, X. Wang, S. Srivats, S. Vosoughi, H. W. Chung, Y. Tay, S. Ruder, D. Zhou, et al., Language models are multilingual chain-of-thought reasoners, arXiv preprint arXiv:2210.03057 (2022).
- [3] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, Y. Iwasawa, Large language models are zero-shot reasoners, *Advances in neural information processing systems* 35 (2022) 22199–22213.
- [4] Expert Group on Liability and New Technologies, Liability for artificial intelligence and other emerging digital technologies, European Commission, 2019.
- [5] R. Bommasani, K. Klyman, S. Longpre, S. Kapoor, N. Maslej, B. Xiong, D. Zhang, P. Liang, The foundation model transparency index, arXiv:2310.12941 (2023).
- [6] A. Zou, Z. Wang, J. Z. Kolter, M. Fredrikson, Universal and transferable adversarial attacks on aligned language models, *CoRR* abs/2307.15043 (2023). arXiv:2307.15043.
- [7] A. Robey, E. Wong, H. Hassani, G. J. Pappas, SmoothLLM: Defending large language models against jailbreaking attacks, arXiv preprint arXiv:2310.03684 (2023).
- [8] J. Geiping, A. Stein, M. Shu, K. Saifullah, Y. Wen, T. Goldstein, Coercing llms to do and reveal (almost) anything, arXiv preprint arXiv:2402.14020 (2024).
- [9] MITRE, Common Vulnerability Enumeration., <https://www.cve.org/About/Overview>, 2023. Accessed: 2024/03/14.
- [10] M. Russinovich, A. Salem, R. Eldan, Great, now write an article about that: The crescendo multi-turn llm jailbreak attack, 2024. arXiv:2404.01833.
- [11] A. V. Silva Neto, J. B. Camargo, J. R. Almeida, P. S. Cugnasca, Safety assurance of artificial intelligence-based systems: A systematic literature review on the state of the art and guidelines for future work, *IEEE Access* 10 (2022) 130733–130770.
- [12] Y. Zeng, H. Lin, J. Zhang, D. Yang, R. Jia, W. Shi, How johnny can persuade LLMs to jailbreak them: Rethinking persuasion to challenge AI safety by humanizing LLMs, *CoRR* abs/2401.06373 (2024). arXiv:2401.06373.
- [13] H. Hajipour, T. Holz, L. Schönherr, M. Fritz, Systematically finding security vulnerabilities in black-box code generation models, arXiv:2302.04012 (2023).
- [14] G. Alon, M. Kamfonas, Detecting language model attacks with perplexity, arXiv preprint arXiv:2308.14132 (2023).
- [15] A. Helbling, M. Phute, M. Hull, D. H. Chau, LLM self defense: By self examination, LLMs know they are being tricked, arXiv preprint arXiv:2308.07308 (2023).
- [16] Assurance Case Working Group (ACWG), Goal Structuring Notation Community Standard, Version 3, <https://scsc.uk/scsc-141c>, 2021. Accessed: 2024/02/25.
- [17] F. A. Batarseh, L. Freeman, C.-H. Huang, A survey on artificial intelligence assurance, *Journal of Big Data* 8 (2021) 60.
- [18] R. Hawkins, C. Paterson, C. Picardi, Y. Jia, R. Calinescu, I. Habli, Guidance on the assurance of machine learning in autonomous systems (AMLAS), arXiv:2102.01564 (2021).
- [19] M.-I. Nicolae, M. Sinn, M. N. Tran, B. Buesser, A. Rawat, M. Wistuba, V. Zantedeschi, N. Baracaldo, B. Chen, H. Ludwig, et al., Adversarial robustness toolbox v1. 0.0, arXiv preprint arXiv:1807.01069 (2018).
- [20] J. MacGlashan, M. K. Ho, R. Loftin, B. Peng, G. Wang, D. L. Roberts, M. E. Taylor, M. L. Littman, Interactive learning from policy-dependent human feedback, in: *International conference on machine learning*, PMLR, 2017, pp. 2285–2294.
- [21] A. Wei, N. Haghtalab, J. Steinhardt, Jailbroken: How does LLM safety training fail?,

- Advances in Neural Information Processing Systems 36 (2024).
- [22] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, arXiv preprint arXiv:1706.06083 (2017).
 - [23] T. Räuher, A. Ho, S. Casper, D. Hadfield-Menell, Toward transparent AI: A survey on interpreting the inner structures of deep neural networks, in: 2023 IEEE Conference on Secure and Trustworthy Machine Learning, 2023, pp. 464–483.
 - [24] A. Koshiyama, E. Kazim, P. Treleaven, P. Rai, L. Szpruch, G. Pavey, G. Ahamat, F. Leutner, R. Goebel, A. Knight, et al., Towards algorithm auditing: a survey on managing legal, ethical and technological risks of ai, ml and associated algorithms, SSRN Preprint, 10.2139/ssrn.3778998 (2021).
 - [25] S. Shan, A. N. Bhagoji, H. Zheng, B. Y. Zhao, Poison forensics: Traceback of data poisoning attacks in neural networks, in: 31st USENIX Security Symposium (USENIX Security 22), 2022, pp. 3575–3592.
 - [26] M. Kläs, R. Adler, L. Jöckel, J. Groß, J. Reich, Using complementary risk acceptance criteria to structure assurance cases for safety-critical AI components., in: AISafety@ IJCAI, 2021, pp. 1–7.
 - [27] J. P. Delgrande, B. Glimm, T. Meyer, M. Truszczynski, F. Wolter, Current and future challenges in knowledge representation and reasoning, arXiv 2308.04161 (2023). arXiv:2308.04161.
 - [28] S. Chari, O. Seneviratne, D. M. Gruen, M. A. Foreman, A. K. Das, D. L. McGuinness, Explanation ontology: A model of explanations for user-centered AI, in: J. Z. Pan, et al. (Eds.), ISWC 2020, Springer, Cham, 2020, pp. 228–243.
 - [29] T. R. Gruber, A translation approach to portable ontology specifications, Knowledge Acquisition 5 (1993) 199–220.
 - [30] B. Gallina, T. Y. Olesen, E. Parajdi, M. Aarup, A knowledge management strategy for seamless compliance with the machinery regulation, in: European Conference on Software Process Improvement, Springer, 2023, pp. 220–234.
 - [31] World Wide Web Consortium (W3C), RDF/XML syntax specification (revised), <https://www.w3.org/TR/REC-rdf-syntax/>, 2023.
 - [32] World Wide Web Consortium (W3C), OWL 2 web ontology language, (second edition), <https://www.w3.org/TR/owl2-rdf-based-semantics/>, 2012.
 - [33] R. Bloomfield, J. Rushby, Assessing confidence with assurance 2.0, arXiv preprint arXiv:2205.04522 (2022).
 - [34] B. Cao, Y. Cao, L. Lin, J. Chen, Defending against alignment-breaking attacks via robustly aligned LLM, 2023. arXiv:2309.14348.
 - [35] Github, CodeQL, <https://codeql.github.com/>, 2024. Accessed: 2024/03/05.
 - [36] J. Liu, C. S. Xia, Y. Wang, L. ZHANG, Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation, in: A. Oh, et al. (Eds.), Advances in Neural Information Processing Systems, volume 36, 2023, pp. 21558–21572.