

# Leveraging Generative AI for Enhancing Domain-Driven Software Design

Götz-Henrik Wiegand<sup>1</sup>, Filip Stepniak<sup>2</sup>, and Patrick Baier<sup>3</sup>

<sup>1</sup> Hochschule Karlsruhe University of Applied Sciences, Germany  
wigo1011@h-ka.de

<sup>2</sup> esentri AG, Ettlingen, Germany  
filip.stepniak@esentri.com

<sup>3</sup> Hochschule Karlsruhe University of Applied Sciences, Germany  
patrick.baier@h-ka.de

**Abstract.** Domain-Driven Design (DDD) is a key framework for developing customer-oriented software, focusing on the precise modeling of an application’s domain. Traditionally, metamodels that describe these domains are created manually by system designers, forming the basis for iterative software development. This paper explores the partial automation of metamodel generation using generative AI, particularly for producing domain-specific JSON objects. By training a model on real-world DDD project data, we demonstrate that generative AI can produce syntactically correct JSON objects based on simple prompts, offering significant potential for streamlining the design process. To address resource constraints, the AI model was fine-tuned on a consumer-grade GPU using a 4-bit quantized version of Code Llama and Low-Rank Adaptation (LoRA). Despite limited hardware, the model achieved high performance, generating accurate JSON objects with minimal post-processing. This research illustrates the viability of incorporating generative AI into the DDD process, improving efficiency and reducing resource requirements, while also laying the groundwork for further advancements in AI-driven software development.

**Keywords:** Generative AI, Domain-Driven Design, LoRA, QLoRA, Quantization, Consumer GPU, PEFT, Weighted Sum, Model Assessment

## 1 Introduction

Creating customer-oriented software demands efficient tools and methods. A promising approach is the Domain-Driven Design (DDD) pattern[1], a robust framework for software development emphasizing the understanding and modeling of the application’s domain. Initially, the software is described using Domain Specific Language (DSL) in JSON or UML, forming a Domain Model (DM) that underpins the iterative development process. From this DM, a code framework is derived, which is then endowed with logic to create a prototype. This prototype generates insights for refining the DM further.

The initial DM generation is typically a manual task performed by a system designer using a GUI tool. To enhance this process, we explore in this paper the possibility of partially automating it with the help of generative AI. We demonstrate how generative models can learn to create syntactically correct JSON objects for describing the DDD DM. Moreover, we show that being trained on real-world data from existing DDD projects, the AI model can automatically generate new parts of a DM through simple interactions with a system prompt. The AI model’s ability to produce syntactically correct JSON objects ensures machine readability, facilitating integration into existing DDD development tools.

Due to data confidentiality, the use of commercial Large Language Models (LLMs) are not an option, which led us to the constraint to develop the JSON code generator model on resource-restrictive hardware, specifically a single consumer-grade GPU.

The final results on the test dataset yielded impressively low loss on JSON generation and high BLEU[2] scores, underscoring the model’s proficiency. Most of the generated JSON objects exhibited syntactical correctness with minimal post-processing, and all JSON objects created from clear prompts were syntactically correct. The successful creation of a code generator for JSON objects in the DSL signifies a pivotal advancement towards incorporating generative AI into the DDD-based software development process, enhancing both efficiency and efficacy.

## 2 Related Work

The foundation for this work lies in the principles of DDD, as established by Eric Evans in his seminal works[1,3]. DDD provides a strategic approach to software development, emphasizing the modeling of complex systems based on their underlying business domains. The company internal framework used for this work builds upon these principles and knowledge regarding this were obtained from the internal documentation[4].

To address the challenges of efficient resource utilization in AI model training and deployment, techniques such as Parameter Efficient Fine-Tuning (PEFT) were employed, specifically the Low Rank Adaptation (LoRA) method introduced by Hu et al.[5]. Further refinements, including quantization methods like Quantized Low Rank Adaptation (QLoRA) by Dettmers et al.[6], played a crucial role in optimizing performance on resource-constrained hardware. The model in this work was quantized to 4-bit precision using the ‘BitsAndBytes’ library from Hugging Face[7]. This approach is supported by research on low-precision quantization, such as the work of Sun et al.[8] and Neshaei et al.[9].

The model used for the code generation component was *Code Llama* from *Meta*<sup>4</sup>, proposed by Rozière et al.[10]. Other models relevant to this field include StarCoder[11] and CodeT5[12]. Additionally, commercial AI code generation tools such as *GitHub Copilot*[13] and *Amazon CodeWhisperer*[14] provide further context and reference in evaluating the landscape of AI-assisted software development.

For the evaluation of the importance of the hyperparameters after hyperparameter Tuning we used a permutation importance analysis referencing the *Random Forest Regressor* from Louppe [15].

For performance evaluation, two key metrics, *BLEU*[16] and Loss, were used to assess the quality of the fine-tuned model. These metrics have been adopted in the evaluation of code generation models, as discussed in the works of Chen et al.[17] and Yetiştirin et al.[18], providing a foundation for assessing syntactic and semantic alignment in generated outputs.

## 3 Methods

In this section, a comprehensive outline of the methodological approach is presented, detailing the processes and techniques used for data handling, model development, and evaluation.

---

<sup>4</sup> <https://www.llama.com/code-llama/>

**Goals and Constraints:** This work explores integrating Generative AI into the software development process within a DDD framework, focusing on automating early-stage development by generating DMs from business requirements. The prototype uses causal language modeling to produce Unified Modeling Language (UML) representations in JSON format, aligning with the iterative nature of DDD.

Technical constraints include the use of open-weight models due to data privacy regulations, prohibiting commercial AI models and requiring local hosting and fine-tuning. Additionally, the project operates under a €1000 budget for external computational resources, demanding resource-efficient model selection and training. The limited dataset further challenges the generation of unbiased, generalizable results, requiring mitigation of data-induced biases. The research aims to evaluate the feasibility of AI-driven code generation within these constraints, emphasizing model performance, resource management, and compliance.

**Data Basis:** The dataset utilized for this study comprises 1,022 files, each containing a single JSON object. Of these, 821 files—accounting for 80% of the dataset—originate from a customer project, while the remaining 20% are derived from a test project. The data represent hierarchically structured DDD logic, encoded in JSON format.

Each JSON object consists of specific key-value pairs, which are defined within a specialized framework. These key-value structures are inherited from a metamodel, which serves as the basis for the framework’s logic. However, the metamodel itself is not included within the dataset, limiting direct access to the underlying inheritance structure.


**Data Pre-Processing:** The data pre-processing step is essential for developing a robust code generator, particularly given the dataset’s significant bias, with approximately 80% of the dataset sourced from a customer project and 20% from a test project. This dataset, comprising 1,022 files containing completed JSON[19] objects, necessitates careful handling to ensure effective model training.

The pre-processing process began with data import, followed by cleaning and abstraction, where high-variability keys were replaced with placeholder values to anonymize customer-specific information. This step not only protects sensitive data but also simplifies the dataset’s complexity, allowing for a clearer focus on the JSON structure. Subsequently, the data was chunked into non-overlapping segments of 2,048 tokens, which were shuffled to enhance randomness. The final step involved a double 80:20 split[20] of the data into training, evaluation, and test sets, resulting in 64% for training, 16% for evaluation, and 20% for testing. This structured approach to data pre-processing ensures that the dataset is well-prepared for effective model training while maintaining compliance with data privacy standards. After exportation, the datasets were versioned for future use, solidifying the pre-processing phase as a foundational element in the overall development process. A full process flow of data pre-processing is displayed in Figure 1.

**Training and Setup:** The foundation of the code generator utilizes the Code Llama 7B model, released by *Meta*<sup>5</sup> proposed by Rozière et al.[10]. With a VRAM size of approximately 25 GB, it necessitated adaptations for the limited hardware available, including a local PC with an RTX 2080 GPU (11 GB VRAM) and a *Lambda Cloud*<sup>6</sup> instance with an RTX A6000 GPU (48 GB VRAM). Due to financial constraints, the

<sup>5</sup> <https://www.llama.com/code-llama/>

<sup>6</sup> <https://lambdalabs.com/service/gpu-cloud>



**Fig. 1.** Abstracted visualization of the various steps of data pre-processing with data cleaning, chunking and splitting to the various data sets for training.

cloud instance was primarily used for hyperparameter tuning. To facilitate fine-tuning on the local setup, a 4-bit quantization was applied, reducing the model’s size to around 4 GB VRAM. The LoRA method, part of PEFT, was selected due to its proven effectiveness when combined with model quantization. The *Hugging Face Transformers*<sup>7</sup> framework’s Trainer[21] was employed alongside a LoRA adapter to optimize training on constrained hardware. This was used to fine-tune *Code Llama 7B* to generate JSON with the help of Next Token Prediction from the dataset. Key training arguments were established to manage resource use, including batch size, gradient accumulation steps, and mixed-precision training. Evaluation metrics were critical for assessing model performance; while the built-in loss function was utilized, metrics such as *BLEU*[22] and *ROUGE-L-F1*[23] were employed to guide the training process. Memory overflow issues were addressed by implementing a custom function for pre-processing logits, ensuring efficient metric evaluation. Overall, the training utilized both local and cloud resources.

**Hyperparameter Tuning:** Hyperparameter tuning is crucial for optimizing model performance by selecting the most effective values for key hyperparameters. In this process, two categories of hyperparameters were identified for tuning: basic training parameters—learning rate, number of training epochs, and warm-up steps—and adapter-specific parameters such as the rank (R-value) and alpha value of the LoRA adapter[21,24,25]. These parameters were chosen due to their significant impact on the model’s performance, particularly in hardware-constrained environments.

To guide the tuning process, initial ranges were defined: a learning rate of 1e-5 to 5e-5, 1 to 5 training epochs, 200 to 1200 warm-up steps, an R-value of 4 to 16, and an alpha value of 4 to 16[24,25]. These ranges were informed by engineering practices and recommendations from existing documentation. After 100 trials, adjustments were made to refine the search, particularly for the number of training epochs and the LoRA rank. In the second phase of tuning, the R-value range was expanded to 4 to 32, and the number of training epochs increased to 5 to 12, allowing further exploration of these critical parameters while leaving the other ranges unchanged. This iterative approach helped maximize model efficiency under the available hardware constraints.

To determine the optimal hyperparameters for final training, a multi-objective weighted sum approach was used, following Bazgan et al.[26]. The weighted sum function  $f(x)$  was initially defined for three evaluation metrics (Equation (1)).

However, hyperparameter tuning results showed that the *ROUGE-L-F1* metric was outside the expected range. As a result, the weight for *ROUGE-L-F1* was set to zero in Equation (2), removing its influence on  $f(x)$ .

Finally, introducing the Inverse Loss  $\tilde{L}(x)$  as  $1-L(x)$ , the weighted sum was simplified to  $f(x)$  in Equation (3).

<sup>7</sup> <https://huggingface.co/docs/transformers/index>

$$f(x) = w_{\text{Loss}} \cdot (1 - L(x)) + w_{\text{BLEU}} \cdot B(x) + w_{\text{ROUGE-LF1}} \cdot R(x) \quad (1)$$

$$f(x) = w_{\text{Loss}} \cdot (1 - L(x)) + w_{\text{BLEU}} \cdot B(x) + 0 \cdot R(x) \quad (2)$$

$$f(x) = w_{\text{Loss}} \cdot \tilde{L}(x) + w_{\text{BLEU}} \cdot B(x) \quad (3)$$

**Model Assessment** To comprehensively evaluate the model’s performance for a generative DDD system, a three-phase assessment approach was used, as traditional metrics alone offer limited insight. In the first phase, the evaluation metrics Loss and *BLEU* [16] from both the training and test datasets were reviewed. The second phase assessed the syntactic correctness and machine-readability of the generated JSON objects. Here, 100 JSON samples were generated from 10 *clear* and 10 *experimental* prompts. *Clear* prompts specify a distinct DDD class object, guiding the model to create a corresponding JSON object, while *experimental* prompts progressively reduce detail, giving the model more room for errors and issues. If any of the generated samples exceeded the token length limit of 4,000, post-processing was applied to ensure completeness, followed by verification through a JSON parser. The final phase involved a qualitative review of the generated JSON objects to identify potential errors and issues. This multi-step evaluation offers a more detailed understanding of the model’s quality and its suitability for real-world applications.

## 4 Results and Discussion

In this section, the results of the Hyperparameter Tuning, Final Model Training, and Model Assessment are summarized and discussed.

**Hyperparameter Tuning:** Table 1 displays the results of the top three values for the different evaluation metrics (objectives) from the hyperparameter tuning process.

The results from the hyperparameter tuning also allow for the derivation of the importance of individual hyperparameters. A Permutation Importance analysis was conducted using a *Random Forest Regressor* [15] to assess the influence of each parameter on the evaluation metrics (objectives) shown in Figure 2. This method helps to quantify how changes in specific hyperparameters affect the model’s performance.


In the analysis of the results from hyperparameter tuning, a noticeable discrepancy was observed between the expected and actual values of the ROUGE-*L-F1* score. The ROUGE-*L-F1* score was anticipated to approach 1. During hyperparameter tuning, it reached a maximum of only approximately 0.062 in the second trial (see Table 1). Due to this significant deviation, the ROUGE-*L-F1* metric was excluded from the determination of the optimal hyperparameters.

Subsequently, the weighted sum method described in Section 3 was applied and calculated for each trial. Assuming that all evaluation metrics converge towards 1 (using the inverse loss as  $1 - \text{Loss}$ ), it can be inferred that the trial with the maximum weighted sum defines the optimal hyperparameters, denoted as  $\theta^*$ . Table 2 presents the top five trials, ranked by their weighted sum along with their respective objectives. Figure 3 illustrates the convergence of the weighted sum towards 1, in relation to BLEU and inverse loss, providing a visual representation of this progression.

The maximum of the weighted sum  $f(x)$  with  $\theta^* = \max(f(x)) \mid x \in \text{Trials}$  is reached at trial 116. Concluding to  $\theta^* = \theta_{f(\text{Trial}_{116})}$ . Therefore, trial 116 defines the optimal

Trial Number	Ranking	Loss↓	BLEU↑	ROUGE-L-F1 ↑	Learning Rate	LoRA Alpha	LoRA R	Train Epochs	Warmup Steps
116	1. Loss	<b>0.031224</b>	0.991329	0.046125	3.4e-05	30	10	6	448
128	2. Loss	0.03168	0.990963	0.04753	3.5e-05	29	13	6	419
127	3. Loss	0.031686	0.991554	0.046529	4.2e-05	22	5	6	1044
110	1. BLEU	0.034611	<b>0.991905</b>	0.047163	3.8e-05	17	10	12	1194
124	2. BLEU	0.032367	0.991696	0.047422	3.4e-05	24	10	9	974
125	3. BLEU	0.033012	0.991665	0.04615	3.3e-05	27	10	9	968
2	1. ROUGE	0.054714	0.987621	<b>0.062322</b>	1.3e-05	16	4	2	1033
130	2. ROUGE	0.038864	0.99047	0.061189	3e-05	11	9	6	675
42	3. ROUGE	0.04477	0.989132	0.05983	3.1e-05	12	11	2	852

**Table 1.** List of the top three Trials of hyperparameter tuning for each evaluation metric (objective) along with the marking of the best values for each objective.



**Fig. 2.** Parameter importance for multiple evaluation metrics (*Objectives*) with importance calculated using permutation importance with *Random Forest Regressor*[15].

hyperparameters  $\theta^*$  for the final training as follows:  $\theta_{\text{learning\_rate}}^* = 3.4e - 5 = 0.000035$ ,  $\theta_{\text{num\_train\_epochs}}^* = 6$ ,  $\theta_{\text{warmup\_steps}}^* = 448$ ,  $\theta_{\text{lora\_r\_value}}^* = 10$  and  $\theta_{\text{lora\_alpha\_value}}^* = 30$ .

**Final Model Training:** The final model training was conducted using the optimal hyperparameters, denoted as  $\theta^*$ , which were determined during the hyperparameter optimization phase. This training was performed on an NVIDIA RTX 2080 GPU with 11 GB of VRAM. Key statistics related to training times and memory usage are summarized in Table 3.


Figure 4 illustrates the progression of the loss function for both the evaluation and test datasets. The training loss shows a high degree of fluctuation, while the evaluation loss remains stable throughout the process. Given that the evaluation loss follows a similar trend to the training loss without significant deviation, it can be inferred that overfitting did not occur during the training.

**Model Assessment** In addition to the final training evaluation metrics (Table 3), the results on the test dataset are shown in Table 4. A slight improvement in loss by 0.0028 and a minor decrease in BLEU by 0.0006 were observed, both considered negligible, indicating stable model performance.

During model assessment, hardware limitations restricted generation to 4,000 tokens per sample. Of 100 JSON samples from 20 prompts, only one terminated correctly within this limit. After post-processing to remove incomplete key-value pairs and close JSON objects, 81 out of 100 samples were successfully parsed.

Trial Number	Loss↓	Inverse Loss↑	BLEU↑	ROUGE-L-F1 ↑	Weighted Sum $f(x)$ ↑
116	0.0312	0.9688	0.9913	0.0461	<b>0.9801</b>
127	0.0317	0.9683	0.9916	0.0465	0.9799
124	0.0324	0.9676	0.9917	0.0474	0.9797
108	0.0319	0.9681	0.9912	0.0488	0.9797
102	0.032	0.968	0.9913	0.0482	0.9797

**Table 2.** List of the top five trials with the highest results for the weighted sum  $f(x)$  sorted in descending order.



**Fig. 3.** Weighted sum in reference to inverse Loss  $\tilde{L}(x)$  and BLEU  $B(x)$ .

Two prompt types were used: *experimental* and *clear* (see Section 3). All 50 JSON samples from the *clear* prompts were parsed without errors after post-processing. However, 19 parsing errors occurred in samples from the *experimental* prompts.

The third phase involved qualitative analysis. Most samples were content-wise comparable to the original dataset, but limitations emerged. In some cases, the model repeated certain sections (e.g., Field Model) until reaching the 4k token limit. This repetitive behavior is permitted in JSON structure but unrealistic for real-world applications, requiring further investigation with more computational resources.


The analysis of the erroneously generated samples revealed two main types of issues:

1. Some generated JSON objects began within another JSON object, such as within a key-value pair. This led to parsing problems due to violations of JSON syntax and structure. This error is likely caused by data chunking during preprocessing.
2. The generation of unwanted characters, such as the zero-width space Unicode symbol (Unicode U+200B [27]), was observed. Since these characters were not present in the training data, it is assumed that they are artifacts originating from the Code Llama 7B model.

**Results Conclusion:** We can conclude that the model is capable of generating machine-readable JSON objects when given appropriate prompts. Fine-tuning on a GPU with sufficient memory produced strong results, and the abstraction of datasets reduced complexity, allowing for the use of customer data. This first version is already capable of being integrated into possible applications with respect to the described bias.

Training Duration	Training Steps	Loss↓ (Train Data)	Loss↓ (Eval Data)	BLEU↑ (Eval Data)
36.43hr	11.54k	0.0337	0.0393	0.9924
Model Size (Quantized)	LoRA File Size	Train Dataset Size	Eval Dataset Size	Available VRAM
4.0046GB	21MB	30.3MB	7.4MB	11GB

**Table 3.** Summary of the results of the final training.



**Fig. 4.** The development of the Training Loss (blue) and Evaluation Loss (red) are plotted over the training steps of the final training, with an update every 50 steps.

Loss ↓	BLEU↑	ROUGE-1-F1 ↑	ROUGE-2-F1 ↑	ROUGE-L-F1 ↑	ROUGE-L-F1 ↑
0.0309	0.9918	0.0565	0	0.0565	0.0565

**Table 4.** Results of evaluation of the model from final training using the test dataset. In this case, the ROUGE values were not within the expected range and were included only for completeness.

## 5 Conclusions

The rapid development of AI technology and the increasing prevalence of LLMs have created opportunities for new applications and tools. With continuous improvements in the efficiency of LLMs, they are increasingly being used in the consumer and Small and Medium-Sized Enterprises (SME) sectors. This work demonstrates the potential and capabilities of open LLMs, the extent to which the development of efficient model trainings with approaches such as LoRA and quantization has already progressed and how these can be combined under restrictive resources.

The work of this paper mark an initial step towards a generative AI assistance system showed the potential for further development into an assistant system for DDD software development framework. In particular, the ability to generate machine-readable JSON objects enabled the use of the final model of this paper in potential tool chains and systems. The results and findings as well as the limitations and challenges, form a broad basis for further development.

Thus, the work of this thesis is a functional code generation model prototype that offers further possibilities and learning, paving the way for the development of an “*Artificial Intelligent*” assistant system that meets the requirements and needs of DDD software development.

The code we used to train and evaluate our models is available at <https://github.com/Tr33Bug/DomainlifecyclesCodeGenerator>.



## References

1. Evans, E.: Domain-Driven Design: Tackling Complexity in the Heart of Software. Addison-Wesley (2004)
2. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In Isabelle, P., Charniak, E., Lin, D., eds.: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, Association for Computational Linguistics (July 2002)
3. Evans, E.: Domain-Driven Design Reference: Definitions and Pattern Summaries. Dog Ear Publishing (2014)
4. esentri, c.: Internal Company Documentation (2024)
5. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models (2021)
6. Dettmers, T., Pagnoni, A., Holtzman, A., Zettlemoyer, L.: Qlora: Efficient finetuning of quantized llms (2023)
7. Hugging Face: Bits and bytes documentation v0.44.1. <https://huggingface.co/docs/bitsandbytes/main/en/index> (2024)
8. Sun, X., Wang, N., Chen, C.y., Ni, J.m., Agrawal, A., Cui, X., Venkataramani, S., El Maghraoui, K., Srinivasan, V., Gopalakrishnan, K.: Ultra-low precision 4-bit training of deep neural networks. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS '20, Red Hook, NY, USA, Curran Associates Inc. (2020)
9. Neshaei, S.P., Boreshban, Y., Ghassem-Sani, G., Mirroshandel, S.A.: The impact of quantization on the robustness of transformer-based text classifiers (2024)
10. Rozière, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Sauvestre, R., Remez, T., Rapin, J., Kozhevnikov, A., Evtimov, I., Bitton, J., Bhatt, M., Ferrer, C.C., Grattafiori, A., Xiong, W., Défossez, A., Copet, J., Azhar, F., Touvron, H., Martin, L., Usunier, N., Scialom, T., Synnaeve, G.: Code llama: Open foundation models for code (2024)
11. Li, R., Allal, L.B., Zi, Y., Muennighoff, N., Kocetkov, D., Mou, C., Marone, M., Akiki, C., Li, J., Chim, J., Liu, Q., Zheltonozhskii, E., Zhuo, T.Y., Wang, T., Dehaene, O., Davaadorj, M., Lamy-Poirier, J., Monteiro, J., Shliazhko, O., Gontier, N., Meade, N., Zebaze, A., Yee, M.H., Umapathi, L.K., Zhu, J., Lipkin, B., Oblokulov, M., Wang, Z., Murthy, R., Stillerman, J., Patel, S.S., Abulkhanov, D., Zocca, M., Dey, M., Zhang, Z., Fahmy, N., Bhattacharyya, U., Yu, W., Singh, S., Luccioni, S., Villegas, P., Kunakov, M., Zhdanov, F., Romero, M., Lee, T., Timor, N., Ding, J., Schlesinger, C., Schoelkopf, H., Ebert, J., Dao, T., Mishra, M., Gu, A., Robinson, J., Anderson, C.J., Dolan-Gavitt, B., Contractor, D., Reddy, S., Fried, D., Bahdanau, D., Jernite, Y., Ferrandis, C.M., Hughes, S., Wolf, T., Guha, A., von Werra, L., de Vries, H.: Starcoder: may the source be with you! (2023)
12. Wang, Y., Wang, W., Joty, S., Hoi, S.C.H.: Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation (2021)
13. GitHub Inc.: Github copilot. <https://github.com/features/copilot> (2024)
14. Amazon Web Services (AWS): AWS codewhisperer. <https://aws.amazon.com/de/codewhisperer/> (2024)
15. Louppe, G.: Understanding random forests: From theory to practice (2015)
16. Tran, N., Tran, H., Nguyen, S., Nguyen, H., Nguyen, T.: Does bleu score work for code migration? In: 2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC). Volume 10., IEEE (May 2019) 165–176
17. Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto, H.P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F.P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W.H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A.N., Leike, J., Achiam,

- J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., Zaremba, W.: Evaluating large language models trained on code (2021)
18. Yetiştirilen, B., Özsoy, I., Ayerdem, M., Tüziin, E.: Evaluating the code quality of ai-assisted code generation tools: An empirical study on github copilot, amazon codewhisperer, and chatgpt (2023)
  19. Bourhis, P., Reutter, J.L., Suárez, F., Vrgoč, D.: Json: data model, query languages and schema specification (2017)
  20. Joseph, V.R.: Optimal ratio for data splitting. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **15**(4) (April 2022) 531–538
  21. Hugging Face: Trainer class for fine-tuning models (2023)
  22. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In Isabelle, P., Charniak, E., Lin, D., eds.: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Philadelphia, Pennsylvania, USA, Association for Computational Linguistics (July 2002) 311–318
  23. Lin, C.Y.: ROUGE: A package for automatic evaluation of summaries. In: *Text Summarization Branches Out*, Barcelona, Spain, Association for Computational Linguistics (July 2004) 74–81
  24. Hugging Face: Performance considerations for training on gpus. [https://huggingface.co/docs/transformers/main/en/perf\\_train\\_gpu\\_one](https://huggingface.co/docs/transformers/main/en/perf_train_gpu_one) (2024)
  25. Hugging Face: Lora (low-rank adaptation) for training diffusion models. <https://huggingface.co/docs/diffusers/training/lora> (n.d.)
  26. Bazgan, C., Ruzika, S., Thielen, C., Vanderpooten, D.: The power of the weighted sum scalarization for approximating multiobjective optimization problems. *Theory of Computing Systems* **66**(1) (November 2021) 395–415
  27. The Unicode Consortium: The unicode standard, version 15.0. <https://www.unicode.org/versions/Unicode15.0.0/UnicodeStandard-15.0.pdf> (2022)