

Applying RLAIIF for Code Generation with API-usage in Lightweight LLMs

Sujan Dutta^{1,†}, Sayantan Mahinder², Raviteja Anantha², and Bortik Bandyopadhyay²

¹Rochester Institute of Technology, sd2516@rit.edu

²Apple, {smahinder, raviteja_anantha, bbandyopadhyay}@apple.com

Abstract

Reinforcement Learning from AI Feedback (RLAIIF) has demonstrated significant potential across various domains, including mitigating harm in LLM outputs, enhancing text summarization, and mathematical reasoning. This paper introduces an RLAIIF framework for improving the code generation abilities of lightweight (<1B parameters) LLMs. We specifically focus on code generation tasks that require writing appropriate API calls, which is challenging due to the well-known issue of hallucination in LLMs. Our framework extracts AI feedback from a larger LLM (e.g., GPT-3.5) through a specialized prompting strategy and uses this data to train a reward model towards better alignment from smaller LLMs. We run our experiments on the Gorilla dataset and meticulously assess the quality of the model-generated code across various metrics, including AST, ROUGE, and Code-BLEU, and develop a pipeline to compute its executability rate accurately. Our approach significantly enhances the fine-tuned LLM baseline’s performance, achieving a 4.5% improvement in executability rate. Notably, a smaller LLM model (780M parameters) trained with RLAIIF surpasses a much larger fine-tuned baseline with 7B parameters, achieving a 1.0% higher code executability rate.

1 Introduction

LLMs have demonstrated unprecedented natural language understanding and generation capabilities in recent times (Brown et al., 2020; Chowdhery et al., 2022; OpenAI, 2023; Anil et al., 2023; Jiang et al., 2023; Touvron et al., 2023). Reinforcement Learning with Human Feedback (RLHF) is a key contributor to this success. RLHF is a fine-tuning approach that uses human feedback to train models by incorporating human evaluations into the reward signal. This method improves model performance on complex tasks by aligning the model’s

behavior with human preferences. However, this technique is expensive due to the requirement for high-quality human feedback. RLAIIF (Bai et al., 2022; Lee et al., 2023) has emerged as a promising alternative to replace human feedback with AI feedback, making the fine-tuning more scalable. Concurrently, there is growing research interest in teaching LLMs how to use external tools (APIs) (Schick et al., 2024; Nakano et al., 2021; Patil et al., 2023; Qin et al., 2023; Li et al., 2023; Zhuang et al., 2024; Hao et al., 2024). However, the focus on lightweight models (<1B parameters) is limited. In this work, we propose an RLAIIF framework to enhance lightweight LLMs’ capability to generate code and effectively integrate API calls. Following Patil et al. (2023), we consider the task of generating Python codes that include suitable API calls given instructions across a wide array of applications. The authors published the Gorilla dataset and showed that fine-tuned LLaMA-7B (Touvron et al., 2023) on this dataset outperforms non-finetuned LLMs like GPT-4 (OpenAI, 2023) in terms of understanding a natural language request and mapping it to API calls. Using our RLAIIF framework, we fine-tune GPT-2-large (Radford et al., 2019) (780M parameters), which not only demonstrates comparable API call correctness to (Patil et al., 2023) but also surpasses its code generation performance.

Code Generation. Although extensively studied since the early days of AI research, code generation (Waldinger and Lee, 1969; Budinsky et al., 1996; Svyatkovskiy et al., 2020; Li et al., 2022) remains a challenging problem. In recent years, the community has explored ways to apply RL in training machine learning models for code generation tasks. For instance, *Seq2SQL* (Zhong et al., 2017) proposed a neural network trained through RL for generating SQL queries given a text description. During training, a generated query is executed

[†] Work done as a part of an internship at Apple.

against a database, and the result is utilized as the reward in the RL algorithm. Le *et al.* (2022) developed *CodeRL*, a sequence-to-sequence language model fine-tuned through an actor-critic RL approach for program synthesis. The code-generator LM is treated as the actor during the training, and the critic model, which is trained to predict unit test results, provides the reward for a generated code. Another work (Shojaee *et al.*, 2023) similar to the above, proposed using feedback from code execution and a ground truth target code to compute the reward. While these approaches may perform well on classical programming tasks (e.g., writing SQL queries, solving competitive/interview-level coding problems, etc.), they are inapplicable on Gorilla-like (Patil *et al.*, 2023) code generation where the program is required to load and execute ML models using the correct API. The bottleneck comes from the fact that the above-mentioned techniques require execution of the generated code to either compute the reward directly or train the critic model, but running thousands of such programs is prohibitively expensive.

Reinforcement Learning with AI Feedback.

Bai *et al.* (2022) introduced the concept of Reinforcement Learning with AI Feedback (RLAIF), which combines preferences labeled by LLMs with human-labeled preferences to optimize for helpfulness and harmlessness. Since then, many studies have explored the usefulness of AI-generated feedback as an alternative to expensive human annotations in various tasks. For instance, Luo *et al.* (2023) proposed *WizardMath*, which enhances the mathematical reasoning abilities of LLaMA-2 using AI feedback in the training process. In another work (Zhang *et al.*, 2023), researchers used real-world data along with RLAIF to improve LLMs as medical consultants. Prior research has also explored AI evaluation for improving factual correctness in LLM-generated medical summaries (Mishra *et al.*, 2023). Kwon *et al.* (2023) explored the usefulness of LLMs in the reward design for RL agents in Ultimatum Game, matrix games, and the DealOrNoDeal negotiation task. Recently, Lee *et al.* (2023) demonstrated that RLAIF can achieve human-level performance in summarization and helpful and harmless text generation. However, the possibility of using RLAIF to improve the code generation and API usage ability in small models (<1B parameters) is under-explored. We demonstrate that even with a few model parameters, AI feed-

back significantly improves code generation quality over simple fine-tuning baselines. Moreover, we found RLAIF applied on smaller 780M parameter GPT-2-large model outperforms LLaMA-7B fine-tuned models, which has nine times more parameters.

2 Dataset

We applied our proposed method to the Gorilla dataset published by Patil *et al.* (2023). The Gorilla dataset consists of three parts - HuggingFace, TensorFlow, and PyTorch. In this work, we only focus on the HuggingFace dataset, which is the most extensive among the three, featuring over 925 unique APIs. These APIs belong to 37 different domains (e.g., Multimodal Text-to-Image, Computer Vision Image Classification, Audio Text-to-Speech, etc.), and for each API, there exist ten unique instructions. Each instance of the data contains an *instruction* (task description), *domain*, *API call* (a single code line), *explanation* (how to solve the task using the API), and a complete *code* (Python script) to accomplish the task. Here, we highlight some key differences between Gorilla and the traditional code generation datasets. Most of the problem statements and corresponding code snippets present in the benchmark datasets including CodeSearchNet (Husain *et al.*, 2019), XLCOST (Zhu *et al.*, 2022), APPS (Hendrycks *et al.*, 2021) and MBPP (Austin *et al.*, 2021) are related to traditional software engineering tasks, representative of common interview questions, require minimal computational resources to execute and do not require internet connection. On the contrary, the Python scripts in the Gorilla dataset focus on AI-related tasks and require an internet connection and significant computing resources (storage and processing power) to execute. The scripts are expected to download ML models hosted on HuggingFace, load them in memory, and run inference. So, techniques where code execution or unit test outcomes are treated as feedback (Zhong *et al.*, 2017; Le *et al.*, 2022; Shojaee *et al.*, 2023) become inapplicable.

While Patil *et al.* (2023) focused only on generating the API call, we demonstrate the effectiveness of our approach both on API call correctness and the ability to use that API in a complete code.

3 Methodology

Our framework follows a similar pipeline to RLHF (Ouyang *et al.*, 2022). However, instead of asking

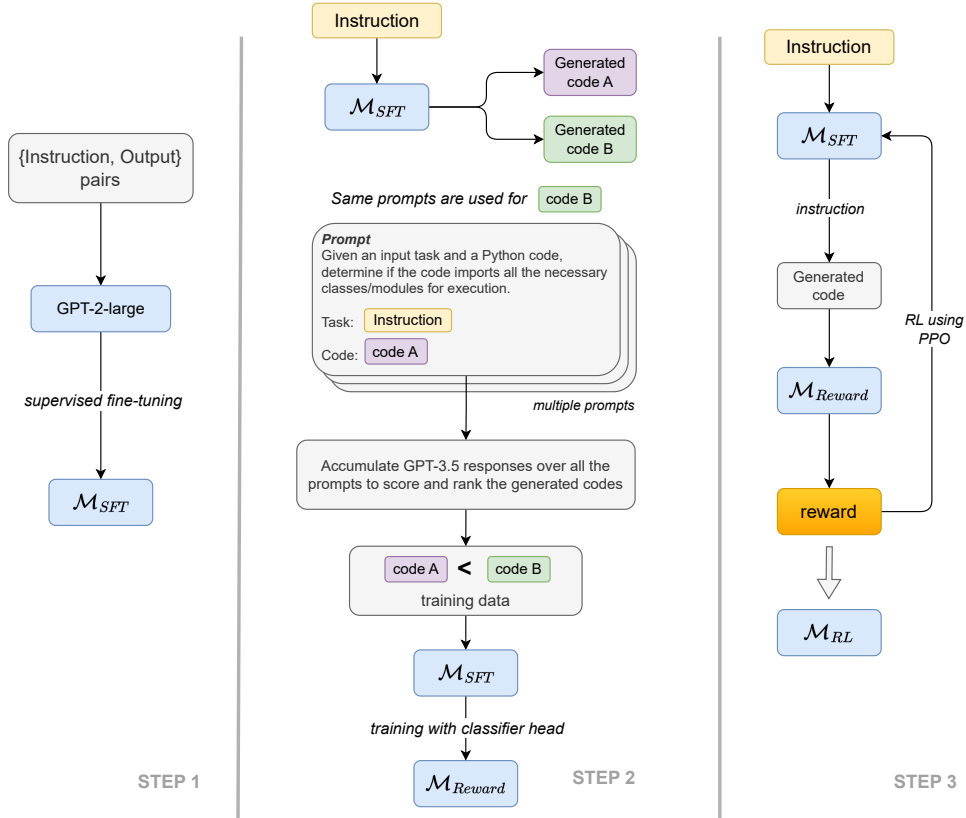


Figure 1: Schematic diagram of the proposed framework. Step 1 is to fine-tune a base model on the dataset. In step 2, we score the \mathcal{M}_{SFT} generated outputs based on the GPT-3.5 feedback using the technique described in section 3. Using this score, we prepare preference data and train a reward model. Finally, in step 3, we use RL to fine-tune \mathcal{M}_{SFT} where \mathcal{M}_{reward} provides the reward.

human annotators to rank the generated responses, we employ a bigger LLM by using a novel prompting strategy. More specifically, for a given instruction and generated code (containing an API call), we ask multiple binary (*yes/no*) questions that capture different aspects of the generated code (and API call) to determine its quality. Our intuition is, that while generating code from natural language might be still challenging for LLMs, providing binary (*yes/no*) answers guided by few-shot exemplars is a much easier task. These feedbacks in turn could be aggregated as a preference ground truth to train the reward model in the RLHF (Ouyang et al., 2022) process. Thus our approach eliminates the need for expensive human annotation cost. We describe the proposed framework (Figure 1) in detail.

- **Step 1: Training a base model**

The first step in the pipeline is to fine-tune a language model on the dataset to get a base model. We choose GPT-2-large and train it on the Gorilla dataset using the supervised fine-tuning technique for causal language models. We denote the fine-

tuned model by \mathcal{M}_{SFT} .

- **Step 2: Training a reward model using LLM feedback**

Instead of human feedback from annotators, we employed a bigger LLM to generate the labels for the reward model. We realized that human graders while judging the correctness of a response, consider different aspects of the generated output. Based on this intuition, we created multiple prompts (P_i) that ask different questions (Q_i) for the same input-output pair. More specifically, we created a set of 8 questions which we feed as prompts to a state-of-the-art language model (GPT-3.5) to get a binary response. Each of these questions addresses a different desired quality (free of bugs, correct imports, no undefined variables, correct syntax, etc.) of the output relevant to the task. Step 2 in Figure 1 presents a sample prompt made using one of the questions. The appendix contains the complete list of prompts. As the questions are binary (*yes/no*) in nature, we simply count the number of *yes* replies by $\mathcal{M}_{GPT3.5}$ to score each input-

output pair. More formally, given a task t , generated output o , and question set $\{Q_i\}$ the prompt set is defined as $P(t, o) = \{P_i \mid P_i = [Q_i, t, o]\}$. The corresponding score (S) is given as:

$$S(t, o) = \frac{\sum_{P_i \in P(t, o)} \mathbb{I}(\mathcal{M}_{GPT 3.5}(P_i) = y_{\epsilon S})}{|P(t, o)|}$$

where \mathbb{I} is the indicator function and $\mathcal{M}_{GPT 3.5}(P_i)$ is the reply from $\mathcal{M}_{GPT 3.5}$ for the prompt P_i . We use this score to prepare the training data for \mathcal{M}_{reward} in the following way. For each instruction in the training data, we generate two outputs from \mathcal{M}_{SFT} by varying the generation parameters (topk, temperature, etc.). Then they are scored using the method described above and labeled (accept or reject) based on this score. These tuples of $\{input\ instruction, accepted\ output, rejected\ output\}$ are then combined to form the dataset for \mathcal{M}_{reward} . In the training phase, \mathcal{M}_{reward} learns to classify whether a machine-generated code is acceptable (or not) for a given input instruction. We append a classifier head on top of \mathcal{M}_{SFT} and use this as the starting point of \mathcal{M}_{reward} and train for three epochs.

- **Step 3: Reinforcement Learning**

Finally, in the RL step, we fine-tune \mathcal{M}_{SFT} using the proximal policy optimization (PPO) algorithm (Schulman et al., 2017). The reward in this step is given by \mathcal{M}_{reward} 's logit scores. We denote our final fine-tuned model by \mathcal{M}_{RL} .

4 Results and Discussions

Model Name (Size)	Executability Rate (%)	ROUGE ($\times 100$)	CodeBLEU ($\times 100$)	AST (%)
$\mathcal{M}_{Gorilla}$ (2023) (7B)	26.9	41.2	36.8	71.68
\mathcal{M}_{SFT} (780M)	23.4	47.2	40.6	72.96
\mathcal{M}_{RL} (780M)	27.9	47.5	42.2	73.62


Table 1: Performance comparison of different models on the Gorilla dataset.

We compute the code generation quality using multiple metrics by comparing the generated output with the ground truth. The reported *ROUGE* is the average of *ROUGE-1*, *ROUGE-2*, *ROUGE-L*, and *ROUGE-sum* metrics introduced in (Lin, 2004). *CodeBLEU* (Ren et al., 2020) was specifically designed for evaluating code synthesis. Ren et al. (2020) defined *CodeBLEU* as the weighted average of standard *BLEU* (Papineni et al., 2002), the weighted n-gram match ($BLEU_{weight}$), the syntactic AST match ($Match_{ast}$), and the semantic dataflow match ($Match_{df}$). $CodeBLEU = \alpha \cdot BLEU + \beta \cdot$

Instruction: We need to extract useful features from Russian text for downstream tasks.


$\mathcal{M}_{Gorilla}$

```
from transformers import AutoModel, AutoTokenizer
model = AutoModel.from_pretrained('DeepPavlov/rubert-base-cased')
tokenizer = AutoTokenizer.from_pretrained('DeepPavlov/rubert-base-cased')
encoded_text = tokenizer(russian_text, return_tensors='pt')
features = model(**encoded_text).last_hidden_state
```



\mathcal{M}_{SFT}

```
from transformers import AutoTokenizer, AutoModel
tokenizer = AutoTokenizer.from_pretrained('DeepPavlov/rubert-base-cased')
model = AutoModel.from_pretrained('DeepPavlov/rubert-base-cased')
tokens = tokenizer.encode(russian_text, return_tensors='pt')
features = model(tokens)
```



\mathcal{M}_{RL}

```
from transformers import AutoTokenizer, AutoModel
tokenizer = AutoTokenizer.from_pretrained('DeepPavlov/rubert-base-cased')
model = AutoModel.from_pretrained('DeepPavlov/rubert-base-cased')
text = "Russian text here..."
tokens = tokenizer(text, return_tensors='pt')
outputs = model(**tokens)
```




Figure 2: Example code generated by different models for the same instruction. In the generations of $\mathcal{M}_{Gorilla}$ and \mathcal{M}_{SFT} the variable `russian_text` is undefined and hence will result in an error. Whereas \mathcal{M}_{RL} defines the variable `text` before using it.

$BLEU_{weight} + \gamma \cdot Match_{ast} + \delta \cdot Match_{df}$. We set $\alpha = \beta = \gamma = \delta = 0.25$ to give equal importance to all the components. The AST sub-tree-matching metric was proposed in (Patil et al., 2023) to capture the correctness of the API calls. In addition to that, we also report the successful execution rate of the generated code (*Executability Rate*). It is worth noting that running this amount of machine-generated programs that download and use large AI models is challenging. We created a pipeline to automatically run the machine-generated codes in an isolated environment.

Table 1 compares the performance of the proposed model with $\mathcal{M}_{Gorilla}$ (finetuned LLaMA-7B) (2023). The results clearly show that the proposed \mathcal{M}_{RL} boosts the performance of the supervised fine-tuned \mathcal{M}_{SFT} in terms of *CodeBLEU* (1.6 points abs), *AST* (0.66% abs) and *Executability Rate* (4.5% abs). We also note that \mathcal{M}_{RL} outperforms the $\mathcal{M}_{Gorilla}$ despite having only 1/9-th of the parameters. It is also reflected in the *Executability Rate* of the generated code. Figure 2 shows an instance where our framework helps in fixing a common error present in $\mathcal{M}_{Gorilla}$ and \mathcal{M}_{SFT} generations.

5 Ethics statement

This work adheres to the ethical guidelines and principles set out in the ACM Code of Ethics and followed by the broader research community. The dataset used in this paper was originally collected from public repositories hosted on HuggingFace. The authors are aware of the growing literature on jailbreaking language models to generate unsafe content. We hope the community will use the proposed models responsibly and only for the intended use cases.

6 Limitations

One of the common limitations faced by similar fine-tuned models is the presence of biases inherited from the pre-trained model. We anticipate that the biases present in the chosen base model (GPT-2-large) also exist in the final model \mathcal{M}_{RL} which might lead to the generation of biased code comments.

Another limitation of this work is the lack of diversity in programming language. The public dataset we considered contains only Python code. Future work should consider expanding this approach to encompass additional programming languages such as C++, Java, JavaScript, etc. Besides, we have not analyzed the performance between more frequent APIs (head) and infrequent APIs (tail). There might be some scope for improvements by focusing on tail APIs more.

Lastly, the learning methodology applied in this study is offline. Given the rapid evolution and proliferation of machine learning models and the corresponding APIs for specific tasks, the model may not leverage more suitable APIs that emerge post-training. To address this, periodic updates to the model are necessary. Our framework's reliance on machine-generated feedback significantly reduces the resource intensity associated with the RLHF process, making these updates more feasible and less costly than a human feedback-based approach.

References

Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen

Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Frank J. Budinsky, Marilyn A. Finnie, John M. Vlisides, and Patsy S. Yu. 1996. Automatic code generation from design patterns. *IBM systems Journal*, 35(2):151–171.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Shibo Hao, Tianyang Liu, Zhen Wang, and Zhiting Hu. 2024. Toolkengpt: Augmenting frozen language models with massive tools via tool embeddings. *Advances in neural information processing systems*, 36.

Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. 2021. Measuring coding challenge competence with apps. *NeurIPS*.

Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. CodeSearchNet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436*.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. 2023. Reward design with language models. *arXiv preprint arXiv:2303.00001*.

Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. 2022. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:21314–21328.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbone, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*.

- Minghao Li, Feifan Song, Bowen Yu, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. Apibank: A benchmark for tool-augmented llms. *arXiv preprint arXiv:2304.08244*.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*.
- Prakanya Mishra, Zonghai Yao, Shuwei Chen, Beining Wang, Rohan Mittal, and Hong Yu. 2023. Synthetic imitation edit feedback for factual alignment in clinical summarization. *arXiv preprint arXiv:2310.20033*.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. 2023. Gorilla: Large language model connected with massive apis. *arXiv preprint arXiv:2305.15334*.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2024. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Parshin Shojaee, Aneesh Jain, Sindhu Tipirneni, and Chandan K Reddy. 2023. Execution-based code generation using deep reinforcement learning. *arXiv preprint arXiv:2301.13816*.
- Alexey Svyatkovskiy, Shao Kun Deng, Shengyu Fu, and Neel Sundaresan. 2020. Intellicode compose: Code generation using transformer. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1433–1443.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, and Shengyi Huang. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Richard J Waldinger and Richard CT Lee. 1969. Prow: A step toward automatic program writing. In *Proceedings of the 1st international joint conference on Artificial intelligence*, pages 241–252.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Hongbo Zhang, Junying Chen, Feng Jiang, Fei Yu, Zhihong Chen, Jianquan Li, Guiming Chen, Xiangbo Wu, Zhiyi Zhang, Qingying Xiao, et al. 2023. HuatuoGPT, towards taming language model to be a doctor. *arXiv preprint arXiv:2305.15075*.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.

Ming Zhu, Aneesh Jain, Karthik Suresh, Roshan Ravindran, Sindhu Tipirneni, and Chandan K. Reddy. 2022. [Xlcost: A benchmark dataset for cross-lingual code intelligence](#).

Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2024. Toolqa: A dataset for llm question answering with external tools. *Advances in Neural Information Processing Systems*, 36.

A Prompts

Table 2 lists all the prompts used in accessing the quality of the generated codes.

B Experimental details

B.1 Dataset

The HuggingFace part of the Gorilla dataset (Patil et al., 2023) consists of over 9k instruction-output pairs. We trained our model on 90% of the data and kept the rest for evaluation.

B.2 Model and implementation details

\mathcal{M}_{SFT} , \mathcal{M}_{reward} and \mathcal{M}_{RL} all have 780M parameters. While training \mathcal{M}_{SFT} and \mathcal{M}_{reward} we used a learning rate of 5×10^{-4} and 5×10^{-5} respectively. In the RL step (PPO algorithm), we set the learning rate to 6×10^{-6} . We did not perform any hyperparameter search. The results are reported by taking the mean of three inference runs. We implemented the training pipeline using the following Python libraries: transformers (Wolf et al., 2020) and TRL (von Werra et al., 2020).

B.3 Computational cost

We used a cluster of NVIDIA A100 40GB GPUs for our experiments. We spent in total ~ 60 GPU hours for all of the experiments.

Prompt
Given an input task and a Python code, determine if the code is functional. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code imports all the necessary classes/modules for execution. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code uses the correct functions/APIs. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code is free of bugs and code smells. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code is sufficient to accomplish the task. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code uses indentations correctly. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code uses quotes in string literals correctly. TASK: [instruction] CODE: [code]
Given an input task and a Python code, determine if the code uses duplicate parameters in a function. TASK: [instruction] CODE: [code]

Table 2: Complete set of prompts. The tokens [instruction] and [code] are used to denote an instruction from the dataset and the corresponding generated code respectively.