

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

4-2024

Creative and correct: Requesting diverse code solutions from AI

Scott BLYTH

Monash University

Markus WAGNER

Monash University

Christoph TREUDE

Singapore Management University, ctreude@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Software Engineering Commons](#)

Citation

BLYTH, Scott; WAGNER, Markus; and TREUDE, Christoph. Creative and correct: Requesting diverse code solutions from AI. (2024). *FORGE '24: Proceedings of the 2024 IEEE/ACM First International Conference on AI Foundation Models and Software Engineering, Lisbon, Portugal, April 14*. 119-123.

Available at: https://ink.library.smu.edu.sg/sis_research/8959

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.



Creative and Correct: Requesting Diverse Code Solutions from AI Foundation Models

Scott Blyth
sbly0002@student.moansh.edu
Monash University
Clayton, Australia

Markus Wagner
markus.wagner@monash.edu
Monash University
Clayton, Australia

Christoph Treude
ctreude@smu.edu.sg
Singapore Management University
Singapore, Singapore

ABSTRACT

AI foundation models have the capability to produce a wide array of responses to a single prompt, a feature that is highly beneficial in software engineering to generate diverse code solutions. However, this advantage introduces a significant trade-off between diversity and correctness. In software engineering tasks, diversity is key to exploring design spaces and fostering creativity, but the practical value of these solutions is heavily dependent on their correctness. Our study systematically investigates this trade-off using experiments with HumanEval tasks, exploring various parameter settings and prompting strategies. We assess the diversity of code solutions using similarity metrics from the code clone community. The study identifies combinations of parameters and strategies that strike an optimal balance between diversity and correctness, situated on the Pareto front of this trade-off space. These findings offer valuable insights for software engineers on how to effectively use AI foundation models to generate code solutions that are diverse and accurate.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Software and its engineering** → **Software libraries and repositories**; **Automatic programming**.

KEYWORDS

Foundation models; correctness; creativity

ACM Reference Format:

Scott Blyth, Markus Wagner, and Christoph Treude. 2024. Creative and Correct: Requesting Diverse Code Solutions from AI Foundation Models. In *AI Foundation Models and Software Engineering (FORGE '24)*, April 14, 2024, Lisbon, Portugal. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3650105.3652302>

1 INTRODUCTION

The role of diversity in programming solutions is increasingly recognised as essential for innovation and problem solving in software engineering [2]. The need for diversity in programming solutions is rooted in two main reasons. First, the selection of the right solution for a programming task depends on the context, which involves

factors such as project requirements, the skills of the development team, the environment for deployment, and key aspects such as security, performance, and maintainability [16]. For example, in high-performance scenarios, such as real-time systems, the focus might shift from code that is easy to read and maintain to one that delivers on performance [18].

Second, being aware of diverse solutions is important in encouraging creativity and new ideas in software engineering [17]. Presenting different solutions helps to address ambiguous challenges and promotes creative thinking and serendipity, while also avoiding repetitive or limited approaches [19].

However, it is important to maintain a balance between this diversity and the correctness of the solutions, as too much diversity might lead to impractical or incorrect outcomes. The specific challenge addressed in this paper is how best to utilise AI foundation models to generate a wide array of programming solutions that not only exhibit diversity but are also correct. Our exploration includes adjusting parameters and modifying the prompts in these models to broaden the spectrum of viable solutions while ensuring that quality is not compromised.

Our approach involves systematic experimentation with HumanEval tasks, using similarity metrics from the code clone community [15] to evaluate the diversity of code solutions generated. We find that employing multiple rounds of prompts is an effective method for generating code solutions that are both diverse and accurate. In addition, integrating various strategies, such as adjusting frequency penalties, moves the Pareto front toward more optimal trade-offs between diversity and accuracy, demonstrating the potential of nuanced model configuration to achieve superior results in code generation. These findings aim to help software engineers effectively use AI foundation models to create code solutions that are diverse, accurate, and practical. Our scripts and data are available at https://github.com/scottb341/diverse_code_generator.

2 HUMANEVAL TASKS

We experiment with the HumanEval task dataset [11], which is a set of 164 tasks used to test the performance of foundation models on programming tasks.

Each task has a description and a set of test cases (average of 8.08 per task). The description of a task contains the function signature (name, type hinting) and some example test cases. As an example, here is the description for Task 4 in the HumanEval dataset:

```
from typing import List
```

```
def mean_absolute_deviation(numbers: List[float]) -> float:
    """ For a given list of input numbers, calculate Mean
    Absolute Deviation around the mean of this dataset.
    Mean Absolute Deviation is the average absolute
```



This work licensed under Creative Commons Attribution International 4.0 License.

FORGE '24, April 14, 2024, Lisbon, Portugal
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0609-7/24/04
<https://doi.org/10.1145/3650105.3652302>

```

difference between each element and a centerpoint
(mean in this case):
MAD = average | x - x_mean |
>>> mean_absolute_deviation([1.0, 2.0, 3.0, 4.0])
1.0
"""

```

For this study, we use different strategies and parameter settings to create ten code solutions for each HumanEval task, with the goal of exploring the trade-off between diversity and accuracy.

We call the response of a foundation model to a task a “(solution) candidate”. We do this to distinguish it from an actual solution that could be considered functionally correct with respect to the specified task.

3 DIVERSITY ASSESSMENT

To assess how diverse the solutions are, we employ clone detectors. Code clones refer to code snippets that have similar functionality [8]. In our study, we assess diversity using the SourcererCC tool [15] and using the cosine similarity of the vector embeddings.

SourcererCC is a token-based, Type 3 [8] code clone detector: it can detect “near miss clones”, which are syntactically similar except for the addition or removal of statements and the changes in comments, identifiers, types, literals, and layouts; therefore, they tend to be functionally similar. When presented with a set of candidates to analyse, SourcererCC reports the IDs of the pairs of clones. The granularity of this clone detector is functional, that is, method by method. To measure the similarity score (which we call *sccSim*) for a set of solutions for a task, we calculate the ratio of the number of detected clones to the number of possible clones; this means that if all code solutions are identical clones, then *sccSim*=1.0, and if all code solutions are pairwise different, then *sccSim*=0.0.

As an alternative to a token-based tool, we measure the cosine similarity of candidates in the vector embedding of a model trained on the CodeSearchNet [9] dataset, CodeBERTaPy [14]. This is similar to the approaches taken by Chung et al. [7] and Cevoli et al. [5], who used vector embeddings to measure the linguistic similarity/diversity of generated texts. In our case, to calculate the similarity of two candidates x and y , using the embedding from CodeBERTaPy, we compute the cosine of the vectors \vec{u} and \vec{v} , where \vec{u} and \vec{v} are the vector embeddings for x and y , respectively. When given a set of two or more candidates, we calculate all pairwise similarities and then return as the embedding-based similarity (which we call *cosineSim*) the mean over all pairwise results.

4 CORRECTNESS ASSESSMENT

To determine the correctness of the code generated by foundation models, we employ the *pass@k* metric as defined by Chen et al. [6]. They measure performance by generating k solution candidate samples, and if one candidate passes all test cases, then the foundation model is considered to have “passed” the task for that value of k ; this means that, (1) in the case of $k = 1$, the model gets one shot at passing the task and (2) in the case of $k > 1$, a single functionally correct candidate (out of multiple) suffices to pass the task.

However, computing *pass@k* in this way can have a high variance, depending on which subset of k samples from a total of n returned code solutions is considered. Instead, just like Chen et al. [6], to evaluate *pass@k*, we generate $n \geq k$ candidate solutions,

count the number of correct samples $c \leq n$ which pass the unit tests, and then calculate the unbiased estimator (an estimate for the ratio passed problem to total problems):

$$pass@k = \mathbb{E}_{Task} \left[1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right]$$

In our analyses, we focus almost exclusively on *pass@1* scores for two reasons: (1) no approach consistently gives us the required 10 solutions, thus preventing reliable comparisons at the *pass@10* level and also below; (2) it is difficult to interpret *pass@k* scores in our context when k is neither 1 (which shows how many correct code solutions are produced) nor 10 (which is our target number).

5 CONFIGURING FOUNDATION MODELS

While a plethora of foundation models are available, we consider (due to practical reasons) only gpt-3.5 and gpt-4 from OpenAI [12].

The operation of both foundation models can be influenced by a number of configuration parameters. In the following, we describe the parameters, list their default values, and outline their intended effects.

temperature. Temperature is used to control the randomness of the output of the foundation model. It is a value from 0 to 2, where 0 is the most deterministic and 2 is the most random. The default value is 1.

top_p. The *top_p* parameter is a value from 0 (exclusive) to 1 (inclusive, default) that controls what the next generated token is pooled from. A *top_p* value of 0.5 means that only the top 50% of the tokens is considered and the bottom 50% is discarded.

frequency_penalty. Changes the probability that a token appears on the basis of its frequency. It ranges from -2 to 2, where a positive value reduces (i.e., penalises) the probability and a negative value increases the probability. The default value is 0.

presence_penalty. Similar to *frequency_penalty*, the *presence_penalty* modifies the probability of each token based on its prior appearance in the text. This means that a token already used is “penalised” or “encouraged” based on its individual occurrence. The values for this penalty range from -2 to 2, with positive values decreasing the probability of a token reappearing and negative values increasing it. The default value is 0, which implies that there is no adjustment for token repetition.

logit bias. The *logit bias* parameter is used to decrease the probability that the foundation model generates certain tokens. Chung et al. [7] found that the *logit bias* parameter was successful in increasing the linguistic diversity of the responses. Our process of mapping solutions to a *logit bias* dictionary follows Chung et al.: first, the top 100 tokens are taken from the previously generated solutions (the concatenation of all of the solutions), and then, based on their counts and the total number of tokens, the bias is computed for those tokens. Specifically, the bias for the token “TOKEN” is:

$$Bias[TOKEN] = -max_bias \cdot \frac{COUNT[TOKEN]}{num_tokens}$$

where `num_tokens` is the total number of tokens in the text that are in the top 100. We employ a maximum bias of 7.5 as seen in Chung et al.

6 PROMPTING TECHNIQUES

As an alternative to changing the operations of a foundation model by changing its operational configuration, we investigate different approaches to prompting it. Depending on the method, this can give the model extra context for its statistical inference, allowing it to find different solutions.

Regeneration. Our regeneration approach (*regen*) is a simple process of giving the description of the task to the model and then repeating this n times to produce n solution candidates. Note that it suffices to hand over the HumanEval tasks (see Section 2 for an example), since they provide Python code as context request to the task. The prompts for generating two solutions for a task would be:

```
- establish a new connection to the model -
Prompt: <task description>
Response: <candidate_1>
- establish a new connection to the model -
Prompt: <task description>
Response: <candidate_2>
```

n_different. The `n_different` prompt asks the foundation model for n different solutions:

```
<task description> + "\n Give me n different solutions for
this problem in Python.
```

Due to the maximum number of tokens that can be generated, this prompt has an upper bound on the number of candidates that can be returned. The maximum number of tokens is specified by the parameter `max_tokens` which is capped by the model’s context length, which for gpt-3.5 is 4096 tokens. A value of $n = 10$ is used for this paper.

n_k_different. The `n_k_different` prompt is similar to the `n_different` prompt, but instead of asking the model to generate all n solutions at once, k solutions are asked to be generated ($n = 10$ and $k = 3$ are used). This is repeated until all n solutions have been generated, where the last iteration asks for the remaining number of solutions. The advantage of this technique over `n_different` is the ability to use the `logit bias` parameter as seen in Chung et al. [7] between API calls (see Figure 1).

Prompt Classification. As mentioned earlier, each problem description comes with some example test cases, making all prompts few-shot/ n -shot prompts [4] with regard to the goal of achieving correctness. However, in terms of the task of increasing diversity, we consider all prompts to be classified as zero-shot, as no examples relating to diversity are given in the prompts.

7 EXPERIMENTAL STUDY

In the following, we report our experimental results. In the first set of experiments, where we request diverse code solutions from foundation models, we change one aspect at a time. In the second set of experiments, we combine those aspects that represented the best diversity-correctness trade-offs in an attempt to create even further trade-offs.

Implementation Details. For security reasons, since we execute code generated by a foundation model, we follow the approach of the HumanEval study [6] and execute tests in a sand-boxed Ubuntu virtual machine.

Because models may output more than one function as part of the solution, this raises the problem of determining which function is the “solution”

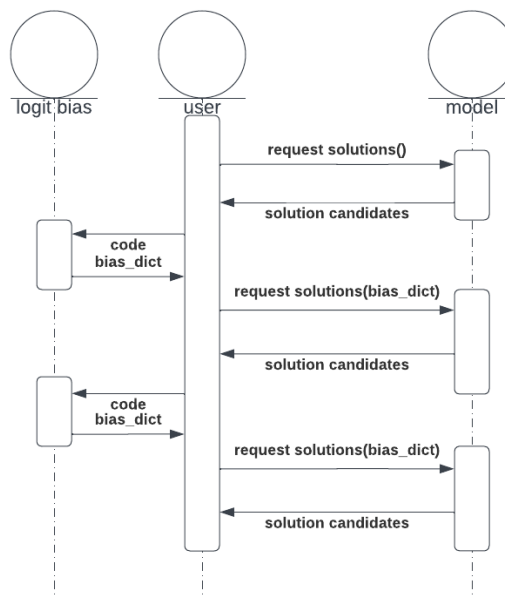


Figure 1: `n_k_different` prompt with Logit Bias; UML sequence diagram.

function. We solve this by checking each function until the test cases are passed or all functions are exhausted (in which case the candidate is flagged as having failed).

For practical reasons, we give each model only one attempt at producing a correct set of solutions for a task, even though foundation models can be non-deterministic in their behaviour. To mitigate this, we aggregate results over all 164 tasks and test each diversification approach in isolation.

7.1 Results: One Technique at a Time

In this part of the investigation, we consider gpt-3.5 as the starting point of our investigations and then vary one aspect at a time. In total, this results in 20 different configurations (see Table 1).

We show the results in Figure 2; the data is also available online [3] (anonymized). We highlight the Pareto front in either violet or green, i.e., the solutions that represent the best-possible trade-offs between solution correctness (as expressed by the `pass@1` scores, where larger values are better) and the diversity metrics `cosineSim` and `sccSim`, where smaller values are better.

Interestingly, in both cases, A15 and A20 form the extreme solutions on the Pareto front: A15 is the solution with the best diversity among its outputs (0.69 mean `cosineSim`, 0.04 mean `sccSim`) while sporting a very low `pass@1` rate (0.11); A20 has the highest `pass@1` rate (0.69) while sporting a relatively low diversity (0.88 mean `cosineSim`, 0.20 mean `sccSim`). In contrast to these two solutions, the default configuration A0 (red star) does not represent a trade-off that should be considered, as it is worse than both approaches in terms of diversity.

Remarkably, our `pass@1` scores are much higher than those observed in the HumanEval study [6]: theirs range (across a range of approaches) from 0.01 to 0.28 (while focussing only on correctness), and our scores range from 0.08 to 0.69 (while considering diversity at the same time). However, we cannot claim that this is due to our diversity focus, as our scores are aligned with those listed by Achiam et al. [10], who report `pass@1` scores of 0.48 for gpt-3.5 and 0.67 for gpt-4. For custom approaches to HumanEval, [13] lists approaches that achieve `pass@1` scores of up to 0.94.

Regarding the operational parameters, we make the following observations. First, the diversity increases (and `pass@1` decreases) consistently as the temperature increases. Second, for `presence_penalty`, we observe

Figure 2: Code similarity and correctness for all 20 approaches from Section 7.1 and for the 2 approaches from Section 7.2. Mean similarity scores are reported across all 164 HumanEval tasks. The red star shows the starting point of our investigations (A0). Sub-figures (a) and (b): the purple/green triangles represent the Pareto fronts (from left to right): purple A15, A14, A1, A20, green A15, A1, A20. The light-green diamonds represent A21 and A22 in Section 7.2. Sub-figure (c): we show the correlation of the two clone detection approaches (Spearman correlation coefficient 0.993).

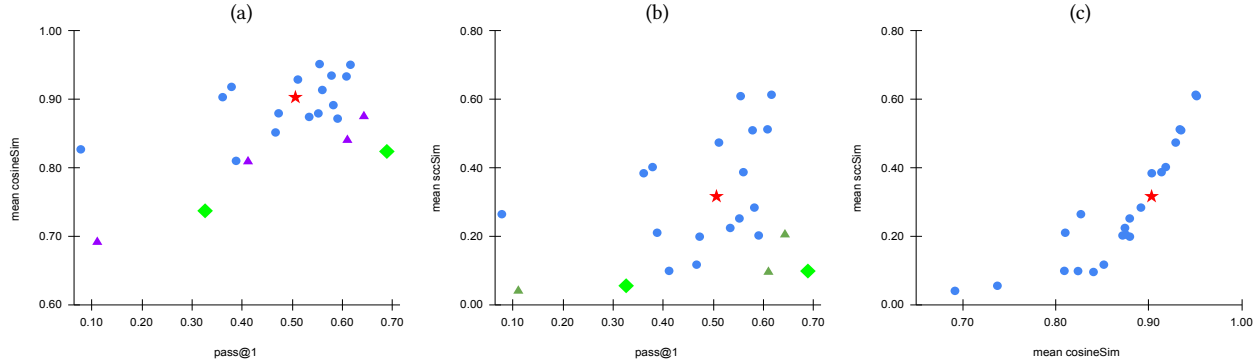


Table 1: Approaches investigated to request diverse code solutions. A0 represents the starting point of our investigations, as it uses gpt-3.5 with default parameters.

model	prompt	temperature	top_p	frequency penalty	presence penalty	logit bias
A0	gpt-3.5 n_different	1	1	0	0	-
A1	gpt-4 n_different	1	1	0	0	-
A2	gpt-3.5 regen	1	1	0	0	-
A3	gpt-3.5 n_k_different	1	1	0	0	-
A4	gpt-3.5 n_different	0.3	1	0	0	-
A5	gpt-3.5 n_different	0.7	1	0	0	-
A6	gpt-3.5 n_different	0.9	1	0	0	-
A7	gpt-3.5 n_different	1.3	1	0	0	-
A8	gpt-3.5 n_different	1	0.2	0	0	-
A9	gpt-3.5 n_different	1	0.4	0	0	-
A10	gpt-3.5 n_different	1	0.6	0	0	-
A11	gpt-3.5 n_different	1	0.8	0	0	-
A12	gpt-3.5 n_different	1	1	-2.0	0	-
A13	gpt-3.5 n_different	1	1	-0.5	0	-
A14	gpt-3.5 n_different	1	1	0.5	0	-
A15	gpt-3.5 n_different	1	1	2.0	0	-
A16	gpt-3.5 n_different	1	1	0	-2.0	-
A17	gpt-3.5 n_different	1	1	0	-0.5	-
A18	gpt-3.5 n_different	1	1	0	0.5	-
A19	gpt-3.5 n_different	1	1	0	2.0	-
A20	gpt-3.5 n_k_different	1	1	0	0	Chung et al. [7]

that the diversity increases with increasing penalty value; however, the pass@1 score does not follow an easily identifiable trend. Third, for frequency_penalty, the default settings appear to be those with the worst diversity and best pass@1 scores; towards either extreme (-2.0 and 2.0), diversity increases and pass@1 scores worsen. Lastly, for top_p, we did not observe notable trends.

On the topic of diversity assessments, we note that both metrics — although structurally very different — are highly correlated in our study, with a Spearman correlation coefficient of 0.993 (see Figure 2(c)).

Table 2: Approaches investigated based on combining aspects of Pareto optimal trade-offs (see Figure 2).

model	prompt	temperature	top_p	frequency penalty	presence penalty	logit bias
A21	gpt-4 n_k_different	1	1	0.5	0	Chung et al. [7]
A22	gpt-4 n_k_different	1	1	2.0	0	Chung et al. [7]

7.2 Results: Combination of Techniques

Next, we combine aspects of the solutions that represent the Pareto front with the goal to create approaches that represent further Pareto-optimal trade-offs between diversity and correctness.

The approaches A21 and A22 are listed in Table 2: both employ frequency penalty (values 0.5 (A21) and 2.0 (A22), due to A14/A15), gpt-4 (due to A1) and logit bias according to Chung et al. [7] (due to A20).

The results are also shown in Figure 2, as light-green diamonds. We observe that both recombinations result in new, Pareto-optimal trade-offs. A22 even achieves the highest pass@1 score across all our approaches (0.69).

8 SUMMARY AND FUTURE WORK

The ability of AI foundation models to produce a large number of responses for a single prompt is a significant asset, especially in software engineering, where diverse code solutions are essential for addressing different contexts and fostering creativity. In this paper, we have explored strategies to leverage foundation models to generate code solutions that are both diverse and correct. We achieved this by (1) altering the operational parameters of the models and (2) experimenting with various prompting strategies. Our research demonstrates that there is a trade-off between creativity and correctness, and we have shown that by recombining aspects of Pareto-optimal trade-offs, we can create new, effective trade-offs.

Moving forward, our focus will be on two main areas. First, our aim is to inform practitioners about effective prompting strategies that can yield diverse and correct solutions, such as iterative prompting and adjusting the frequency penalty. Second, we plan to automate this process. Applying the DUO principle [1], we will use optimisation to fine-tune these strategies to achieve the best trade-offs. This approach will enable us to make the most of existing models for creating code solutions that are both creative and correct. Our initial efforts in manual recombination show the potential of this approach.

REFERENCES

- [1] Amritanshu Agrawal, Tim Menzies, Leandro L. Minku, Markus Wagner, and Zhe Yu. 2020. Better software analytics via "DUO": Data mining algorithms using/used-by optimizers. *Empirical Software Engineering* 25, 3 (2020), 2099–2136. <https://doi.org/10.1007/s10664-020-09808-9>
- [2] Benoit Baudry and Martin Monperrus. 2015. The multiple facets of software diversity: Recent developments in year 2000 and beyond. *ACM Computing Surveys (CSUR)* 48, 1 (2015), 1–26.
- [3] Scott Blyth, Markus Wagner, and Christoph Treude. [n. d.]. Creative and Correct: Requesting Diverse Code Solutions from AI Foundation Models (Dataset). ([n. d.]). <https://doi.org/10.5281/zenodo.10515808>
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- [5] Benedetta Cevoli, Chris Watkins, and Kathleen Rastle. 2021. What is semantic diversity and why does it facilitate visual word recognition? *Behavior research methods* 53 (2021), 247–263.
- [6] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. (2021). [arXiv:2107.03374 \[cs.LG\]](https://arxiv.org/abs/2107.03374)
- [7] John Joon Young Chung, Ece Kamar, and Saleema Amershi. 2023. Increasing Diversity While Maintaining Accuracy: Text Data Generation with Large Language Models and Human Interventions. *arXiv preprint arXiv:2306.04140* (2023).
- [8] Shilpa Dang. 2015. Survey Based Analysis of Effect of Code Clones on Software Quality. *International Journal of Engineering Research & Technology (IJERT) ISSN: 2278-0181* 4 (03 2015), 371–379. <https://doi.org/10.17577/IJERTV4IS030495>
- [9] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. CodeSearchNet Challenge: Evaluating the State of Semantic Code Search. *arXiv:1909.09436 [cs, stat]* (Sept. 2019). <http://arxiv.org/abs/1909.09436> [arXiv: 1909.09436](https://arxiv.org/abs/1909.09436).
- [10] OpenAI, :, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Lukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Lukasz Kondrasiuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reichihiro Nakano, Rajeew Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Wong, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. GPT-4 Technical Report. [arXiv:2303.08774 \[cs.CL\]](https://arxiv.org/abs/2303.08774)
- [11] OpenAI. 2021. HumanEval: Hand-Written Evaluation Set. (2021). <https://github.com/openai/human-eval>
- [12] OpenAI. 2023. Request body. (2023). <https://platform.openai.com/docs/api-reference/chat/create>
- [13] Paperswithcode. 2023. Code Generation on HumanEval. (2023). <https://paperswithcode.com/sota/code-generation-on-humaneval>
- [14] Manuel Romero. [n. d.]. mrm8488/CodeBERTaPy. ([n. d.]). <https://huggingface.co/mrm8488/CodeBERTaPy>
- [15] Hitesh Sajani, Vaibhav Saini, Jeffrey Svajlenko, Chanchal K Roy, and Cristina V Lopes. 2016. Sourcerercc: Scaling code clone detection to big-code. In *Proceedings of the 38th International Conference on Software Engineering*, 1157–1168.
- [16] Luca Traini, Daniele Di Pompeo, Michele Tucci, Bin Lin, Simone Scalabrino, Gabriele Bavota, Michele Lanza, Rocco Oliveto, and Vittorio Cortellessa. 2021. How software refactoring impacts execution time. *ACM Transactions on Software Engineering and Methodology* 31, 2 (2021), 1–23.
- [17] Christoph Treude. 2022. Taming Multi-Output Recommenders for Software Engineering. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 1–5.
- [18] Christoph Treude. 2023. Navigating Complexity in Software Engineering: A Prototype for Comparing GPT-n Solutions. In *Proceedings of the 5th International Workshop on Bots in Software Engineering*.
- [19] Reza Jafari Ziarani and Reza Ravanmehr. 2021. Serendipity in recommender systems: a systematic literature review. *Journal of Computer Science and Technology* 36 (2021), 375–396.