

---

# Planning in Natural Language Improves LLM Search for Code Generation

---

Evan Wang<sup>† 2</sup> Federico Cassano<sup>† 3,4</sup> Catherine Wu<sup>† 5</sup> Yunfeng Bai<sup>1</sup> William Song<sup>1</sup>  
Vaskar Nath<sup>1</sup> Ziwen Han<sup>1</sup> Sean Hendryx<sup>1</sup> Summer Yue<sup>1</sup> Hugh Zhang<sup>1</sup>  
<sup>1</sup>Scale AI <sup>2</sup>California Institute of Technology <sup>3</sup>Anysphere  
<sup>4</sup>Northeastern University <sup>5</sup>Anthropic, <sup>†</sup> *work done while at Scale AI*  
Correspondence to [evan.wang@scale.com](mailto:evan.wang@scale.com) and [hugh.zhang@scale.com](mailto:hugh.zhang@scale.com)

## Abstract

While scaling training compute has led to remarkable improvements in large language models (LLMs), scaling inference compute has not yet yielded analogous gains. We hypothesize that a core missing component is a lack of diverse LLM outputs, leading to inefficient search due to models repeatedly sampling highly similar, yet incorrect generations. We empirically demonstrate that this lack of diversity can be mitigated by searching over candidate plans for solving a problem in natural language. Based on this insight, we propose PLANSEARCH, a novel search algorithm which shows strong results across HumanEval+, MBPP+, and LiveCodeBench (a contamination-free benchmark for competitive coding). PLANSEARCH generates a diverse set of observations about the problem and uses these observations to construct plans for solving the problem. By searching over plans in natural language rather than directly over code solutions, PLANSEARCH explores a significantly more diverse range of potential solutions compared to baseline search methods. Using PLANSEARCH on top of Claude 3.5 Sonnet achieves a pass@200 of 77.0% on LiveCodeBench, outperforming both the best pass-rate achieved without any search (pass@1 = 41.4%) and using standard repeated sampling on top of existing non-search models (pass@200 = 60.6%). Finally, we show that, across all models, search algorithms, and benchmarks analyzed, we can accurately predict performance gains from search as a function of the diversity over generated ideas. Code can be found at <https://github.com/scaleapi/plansearch>.

## 1 Introduction

The bitter lesson [42] famously posits that two forms of scaling trump everything else: learning and search. While recent advances in large language models (LLMs) have shown that learning is extremely effective, search has not yet proven its value for LLMs, despite its success with classical machine learning techniques [10, 39, 40, 7, 8, 4, 18].

Here, we refer to search as any method of spending compute at inference time to improve overall performance [29]. In this work, we focus our efforts on improving LLM search for code generation, one of the most important current applications of LLMs. We hypothesize the major bottleneck preventing widespread use of search for code is a lack of high-level diversity in model outputs, since repeated sampling with little to no diversity returns highly similar programs, resulting in minimal gain from additional inference-time compute. We empirically demonstrate that this is the case for many open-source language models which have undergone significant post-training. Specifically, we show that in many cases, despite instruction tuned models outperforming base models by large margins on a single sample regime (pass@1), this trend disappears—sometimes even reversing—when generating many samples. We refer to Figure 29 for examples of this phenomenon.

This diversity problem is also *not* reflected in many public leaderboards (e.g. LMSYS Chatbot Arena [15], LiveCodeBench [23], OpenLLMLLeaderboard [1]), which often report only the pass rate from a single sample of the model, ignoring an entire dimension along which to compare models. While the performance of one sample is the primary metric of relevance for applications such as chatbots, as users typically are sensitive to latency, this scalar is insufficient to fully capture the quality of a model with greater inference-time compute budget.

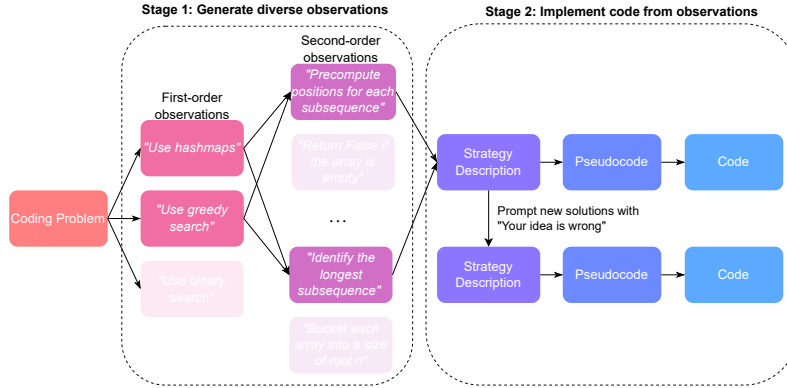


Figure 1: PLANSEARCH first generates observations, then combinatorially samples subsets of these observations to generate the next step in the search process. To generate the next layer of observations, the combinations derived from the first observations are used as stepping stones to generate the next observations, and the process repeats. After generating both the first and second order observations, PLANSEARCH then generates a natural language description of a strategy to solve the problem. For additional diversity, the model is prompted to regenerate its strategy as an additional sample before generating code. See Section 4.3 for additional information.

In this paper, we explore several directions for improving the diversity of LLMs at inference time. We hypothesize that the right level of diversity to search over is the natural language conceptual/idea space, and we validate our hypothesis across several experiments. First, we show that models can produce the correct final program when fed the correct solution sketches, where these sketches have been “backtranslated” from passing solution code into sketches in idea space (Section 3.2). Second, we show that when models are asked to generate their own ideas before implementing them on LiveCodeBench (IDEASEARCH), their accuracy conditioned on a particular sketch trends towards either 0% or 100%, suggesting that most of the variance in passing a particular problem is captured by whether the sketch is correct. These two experiments suggest a natural method to improving LLM search for code generation: by searching for the correct idea to implement.

Guided by this principle of *maximizing exploration of ideas*, we propose PLANSEARCH. In contrast to many existing search methods that search over individual tokens [50, 53], lines of code [25], or even entire programs [26], PLANSEARCH searches over possible *plans* for solving the problem at hand, where a plan is defined as a collection of high level observations and sketches helpful to solve a particular problem (Figure 1). We find that searching over plans outperforms both standard repeated sampling and directly searching over ideas (IDEASEARCH, introduced in Section 4.2) in terms of effectively using compute at inference time.

Applying PLANSEARCH on top of Claude 3.5 Sonnet achieves a state-of-the-art pass@200 of 77.0% on LiveCodeBench, outperforming both the best score achieved without search (pass@1 = 41.4%) and the standard best-of-n sampling score (pass@200 = 60.6%). Furthermore, consistent with recent findings on the effectiveness of search on top of small models [13, 6, 5, 44], allowing PLANSEARCH based on a small model (GPT-4o-mini) outperforms larger models not augmented with search after merely 4 attempts. Evaluations of PLANSEARCH across two other coding benchmarks, HumanEval+ and MBPP+ [27], suggest similar improvements.

Finally, we measure the idea diversity of output codes via an LLM-as-a-judge procedure (Section 6.1) and show that the resulting diversity score is highly correlated with the performance gains generated by a given search method. This provides further support for our hypothesis that the effective exploration of plans in idea space is key to LLM search for code generation (Figure 3).

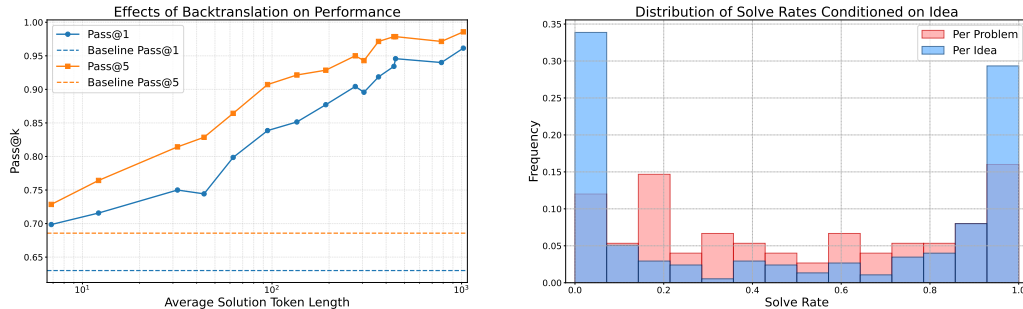
## 2 Related Work

We reiterate that search as defined in the context of our paper refers to any method which expends inference-time compute to improve performance. We further specify planning as any form of high level observation or abstract thought that assists a model in generating a final solution. Our work builds off a long history of work in scaling search and planning. For information on relevant work in classical AI, general search, and filtering, see Appendix T.

Regarding searching over plans in natural language, several approaches have proposed generalizing chain-of-thought [43] reasoning into a search-like process, such as Tree of Thoughts [45] and Reasoning via Planning [21]. However, prior methods have largely demonstrated effectiveness on somewhat contrived problems designed to highlight the power of search, such as the game of 24, or classic planning benchmarks such as Blocksworld [28], where both benchmarks are easier to solve by explicitly considering many options, and where the ‘steps’ over which to search over are fairly obvious. By contrast, most real-world planning is used to assist in domains that are complex enough to benefit from, but not require, the additional exploration of plans. We demonstrate that PLANSEARCH, which plans in natural language, outperforms baseline search methods in one such domain: code generation. Moreover, our analysis reveals the underlying reason that such search is effective: it increases the diversity of the generated ideas, allowing more efficient search relative to other methods which repeatedly submit highly similar, incorrect solutions. This is consistent with prior work suggesting the importance of diversity in natural language generation [22, 51]. Other directions for search include decomposing programs down into smaller parts before solving each one individually [48, 54, 20].

PLANSEARCH is also distinct from other methods which explicitly train a model to search or on reasoning traces sampled from the model [49, 52, 47, 31] in that PLANSEARCH induces diversity at inference-time and converts an LLM API not designed for search into one that is capable of showing strong gains from search. Separately, there is a large family of work in the agent space, in which outputs from terminal commands or other tools are fed back into the model before the agent is queried for the next step [38, 46, 35, 11].

## 3 Motivation



(a) Performance of GPT-4o-mini on LiveCodeBench when provided with backtranslated solutions of varying lengths. The baselines plot performance without backtranslated solutions. Providing the model with a compressed solution in natural language, even as short as 10 tokens, significantly increases performance.

(b) We plot the distribution of solve rates conditioned on being given a solution sketch and without. When conditioning on a given sketch, we notice that downstream solve rates polarize towards either 0% or 100%. Most of the variance in performance is predicted by whether a given idea is correct or not.

Figure 2: Backtranslation shows the promise of providing good sketches, and conditioning on idea shows the presence of a solution sketch polarizes performance.

### 3.1 Defining the Search Space

Perhaps the most important question for eliciting strong search capacities is determining which space to search over, as finding the proper layer of abstraction is critical to progress in the field. Prior approaches have varied, with many people searching over individual tokens [50, 53], lines

of code [25], or even entire programs [26]. We hypothesize that the key factor is obtaining *the correct solution sketch*, which we define as a description of the correct program in natural language space. Intuitively, conducting the reasoning process in natural language space allows us to effectively harness the training process of LLMs, which have observed many human reasoning traces in both pre- and post-training. Prior work [43] has observed strong positive effects from being allowed to conduct such reasoning in natural language, making it a natural place to search over. We describe two experiments providing evidence for this hypothesis below.

### 3.2 Backtranslation

To investigate the hypothesis whether the idea space, instantiated as solution sketches, is the right area of exploration, a natural question is whether LLMs can correctly implement a correct code solution given a correct sketch. Inspired by approaches to backtranslation in machine learning [36, 33, 17], we experiment with “backtranslating” passing code solutions back into idea space.

First, we generate code solutions using GPT-4o to generate 1000 attempts to solve the problem and filter out problems without any passing solutions. As we also do not have a dataset of correct solution sketches associated with each solution, we generate a candidate correct idea via backtranslation. We do this by feeding an LLM both the problem and code solution and asking the LLM to convert said solution into a natural language description of the solution. Additionally, we vary the detail of the backtranslated idea via instructions to the LLM in the prompt (e.g. ‘in  $w$  words’). A full description of the prompts can be found in Appendix O.1, alongside several example backtranslated solutions of various lengths.

We observe that prompting a model with a backtranslated idea significantly improves accuracy, increasing with the length of the translated idea (Figure 2a), which suggests that having a correct sketch is sufficient to produce the correct final solution with relatively high accuracy, even only after 10 tokens of backtranslated solution.

### 3.3 Conditioning on Idea Quality

Next, we prompt an LLM to generate its own sketches to solve LiveCodeBench problems instead of providing it with golden ones via backtranslation. First, we generate 5 ideas per problem using IDEASEARCH, defined in Section 4.2. For each idea, we then sample 25 candidate solutions and measure their pass rate. For this experiment, we filter out any problem that GPT-4o-mini solves with either a 100% or a 0% solve rate, since such problems are either too easy or too hard for the model and would not be informative for this experiment. We end with 75 problems and 375 sketches.

To test our hypothesis that generating a correct sketch is a critical factor for solving problems, we compare the distribution of solve rates for generating correct code solutions *conditioned* on a given sketch to the distribution over solve rates given a sketch drawn at random, i.e., just the distribution over solve rates.

While verifying whether a sketch is correct or incorrect is difficult without access to external labels, a key insight is that if generating the correct idea is a critical factor in solving the problem, then conditioning on a particular sketch should polarize the distribution of solve rates towards  $\{0, 1\}$ . If the model is given a correct sketch, it should consistently generate correct solutions, while if given a bad sketch, it should consistently generate incorrect solutions.

Our results confirm this to be the case. Figure 2b shows the distribution of solve rates across problems, both unconditionally (in red) and conditioned on each sketch (in blue). We notice that when grouping by sketches, the solve rates indeed become polarized towards  $\{0, 1\}$ . This result has important implications for improving code generation, suggesting that a large portion of variance in performance may come from whether the model is given a correct idea or not. Therefore, a natural path for improvement is to focus on the sketch generation step and search for correct sketches and observations in idea space before generating solution code.

## 4 Methods

We provide a description of the various methods of search we explore in our work. If additional background on competitive programming and related notation is desired, we provide more (optional) information in Appendix S.

### 4.1 REPEATED SAMPLING

We consider the basic prompting approach as a baseline, in which we use few-shot prompting by providing the LLM with a number of problem-solution pairs before asking it to solve the desired question [9]. A full example of the prompt is given in Appendix O.2. In code generation, the most common variant of search utilized is repeated sampling, where models are repeatedly sampled from until they generate an output that passes the test or the maximum number of samples is reached. Refer to the Related Work for more information (Section T.2).

### 4.2 IDEASEARCH

A natural extension of the REPEATED SAMPLING approach discussed in Section 4.1 is to avoid prompting the LLM for the solution code immediately. This can be viewed as an application of the commonly used “chain-of-thought” prompting to programming problems [43], although we find that IdeaSearch shows non-negligible performance boosts over standard “chain-of-thought” prompting (see Appendix H).

In IDEASEARCH, the LLM is given the problem  $P$  and is asked to output a natural language solution  $S$  of the problem. Then, a separate instance of the LLM is given  $P$  and  $S$ , and tasked to follow the proposed solution  $S$  to solve the problem  $P$ . The purpose of IDEASEARCH is to isolate the effectiveness of having the correct “idea/sketch” for solving the problem. Empirically, we find that explicitly forcing the search algorithm to articulate an idea for solving the problem increases diversity. See Appendix O.3 for detailed prompts.

### 4.3 PLANSEARCH

While both REPEATED SAMPLING and IDEASEARCH are successful and lead to improvement in the results on benchmark results, we observe that in many of the cases, prompting multiple times (pass@k) (even at high temperatures) will only lead to small, narrow changes in the output code that change minor aspects but fail to improve upon pitfalls in idea.

Ablations for many of the choices in the subsequent description of PLANSEARCH can be found in Appendix K.

#### 4.3.1 Prompting for Observations

Starting from the problem statement  $P$ , we prompt an LLM for “observations”/hints to the problem.

We denote these observations as  $O_i^1$ , where,  $i \in \{1, \dots, n_1\}$  due to the fact that they are first-order observations. Typically,  $n_1$  is on the order of 3 to 6. The exact number depends on the LLM output. To use these observations to inspire future idea generation, we create all subsets with size at most  $S = 2$  of  $s^1 = \{O_1^1, \dots, O_{n_1}^1\}$ . Each of these subsets is a combination of observations, and for clarity we denote each subset as  $C_i^1, i \in \{1, \dots, l_1\}$ , where  $l_1 = 1 + n_1 + \binom{n_1}{2}$ .

#### 4.3.2 Deriving New Observations

The set of all observations can be thus defined as a directed tree with depth 1, where the root node is  $P$ , and an edge exists for each  $C_i^1$  pointing from  $P$  to  $C_i^1$ . We then repeat this procedure from Section 4.3.1 on each leaf node  $C_i^1$  to generate a set of second order observations,  $s_i^2 = \{O_{i,1}^2, \dots, O_{i,n_{i,2}}^2\}$ . To obtain second order observations, we prompt the model with both the original problem  $P$  and all observations contained in  $C_i^1$ , framed as primitive observations that are necessary in order to solve  $P$ . The LLM is then prompted to use/merge the observations found in  $C_i^1$  in order to derive new ones.

The same procedure as Section 4.3.1 is used to create all subsets  $C_{i,j}^2$ , for all  $i \in \{1, \dots, l_1\}$ . This process may be arbitrarily repeated, but we truncate the tree at depth  $L = 2$  for computational constraints.

Note that there is no assumption any of the observations generated are correct. In fact, it is critical to note that many of them may be incorrect. The observations merely serve to elicit the model to search over a more diverse set of ideas.

### 4.3.3 Observations to Code

After the observations have been made, they must be implemented as ideas before being translated into code. For each leaf node, we prompt the model with all observations, along with the original problem  $P$ , in order to generate a natural language solution to the problem  $P$ . To add more diversity, for each generated idea, we generate an additional idea by supposing the idea is wrong, and asking an LLM to give criticisms/feedback, thus increasing our proposed ideas by a factor of 2.

These natural language solutions are then translated into pseudocode, which are subsequently translated into actual Python code. We take a more granular approach to reduce the translation error (which may cause the model to revert to its original mode, disregarding the reasoned-through observations). We provide all prompts for all sections in Appendix O.4.

## 5 Experimental Results

Model	Benchmark	Pass@1	Pass@200	PLANSEARCH@200
GPT-4o-mini	LiveCodeBench	39.0	53.3	64.9
GPT-4o	LiveCodeBench	41.3	60.6	73.0
DeepSeek-Coder-V2	LiveCodeBench	41.4	53.2	70.3
Claude-Sonnet-3.5	LiveCodeBench	40.3	55.6	77.0
GPT-4o-mini	HumanEval+	83.7	95.0	98.2
GPT-4o	HumanEval+	86.4	98.2	99.5
DeepSeek-Coder-V2	HumanEval+	82.8	91.4	99.3
Claude-Sonnet-3.5	HumanEval+	81.6	88.9	98.5
GPT-4o-mini	MBPP+	73.5	83.8	91.0
GPT-4o	MBPP+	77.2	87.4	92.2
DeepSeek-Coder-V2	MBPP+	76.3	81.9	92.6
Claude-Sonnet-3.5	MBPP+	77.1	83.0	93.7

Table 1: We find that PLANSEARCH and IDEASEARCH improve upon search baselines across all models, with PLANSEARCH achieving the best results across all models and benchmarks considered. Notably, using PLANSEARCH on top of Claude 3.5 Sonnet [2] has a pass@200 of 77.0 on LiveCodeBench, which is nearly double the performance of the top model without using search (41.4). We highly encourage readers to check Appendix D for complete results and pass@k curves.

Our summarized results for REPEATED SAMPLING, IDEASEARCH, and PLANSEARCH can be found in Table 1 and Figure 9. Additionally, we plot our full pass@k curves for all methods, models, and datasets in Appendix D. For sake of easy comparison, we also plot all relative gains compared to REPEATED SAMPLING@1 averaged over all models in Appendix F. For a compute-normalized comparison between REPEATED SAMPLING and PLANSEARCH, see Figure 16.

See Appendix C for experiment and dataset details.

## 6 Analysis

Our results suggest that both PLANSEARCH and IDEASEARCH outperform basic sampling by a wide margin (Figures 10, 11, 12), with PLANSEARCH achieving the best score across all methods and models considered. We show the detailed pass@k results for each dataset in Figures 4, 5 and 6. We

also compare with Chain-of-Thought [43] in Appendix H. Interestingly, we find that IDEASEARCH performs somewhat better, which we speculate comes from differences in splitting solution sketch into *two* model responses, instead of doing both chain-of-thought and code solution in one model response.

Investigating the differences in specific models, we notice that trends exhibited by the pass@k curves are not uniform across all models; in fact, each curve seems unique. We hypothesize that these differences are in part due to changes in idea diversity, as investigated in Figures 3, 24, 25. From the figures, we can see that our approximate diversity score accounts for much of the variance we see in the relative improvement that arrives from scaling-up inference-time compute. This correlation holds across all methods and models on the same dataset, thus suggesting that diversity score can be used as a proxy to predict for relative pass@k improvement. For further discussion on the specifics of the diversity score, see Section 6.1.

One interesting point of observation is that PLANSEARCH often hurts pass@1 for several models, including most notably Sonnet 3.5 on LiveCodeBench, our best performing combination. Intuitively, this is because increasing the diversity across ideas likely dilutes the probability that any *particular* idea is generated, while simultaneously increasing the chance of having *at least one* correct idea within said pool. Therefore, pass@1 may be slightly lower than usual, yet pass@k will likely surpass “pools” of ideas lacking diversity for this reason. See Figure 39 for a graphical intuition.

Finally, in Table 1, we present our main results normalized across attempts/completion, where each search method is allowed *k* attempts to solve each problem. Although we did not design for compute-optimality, we show in Appendix G that PLANSEARCH also outperforms other methods if sufficient compute is expended.

### 6.1 Measuring Diversity

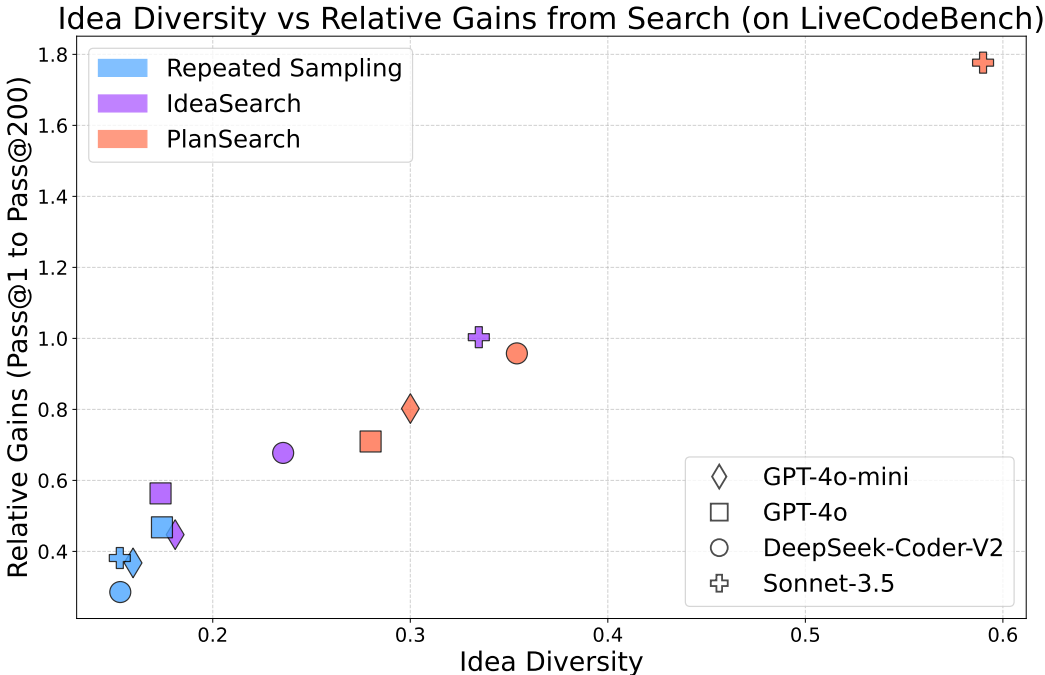


Figure 3: We observe a strong positive correlation between the measured amount of idea diversity in a search algorithm and the resulting improvements due to search (Section 6.1). Diversity score is the probability that GPT-4o-mini believes two randomly selected output codes implement different ideas (higher is more diverse). Our findings suggest that diversity in idea space is essential for effective LLM search.

We find that diversity as measured in idea space is highly predictive of search performance, as measured by the relative improvement between a model/method’s pass@1 and its pass@200 (Fig-

ure 3). While the most common measure of diversity is entropy [37], entropy is insufficient for a number of reasons for the precise setting of LLMs [22, 51]. As a simple example, consider two different language models, one of which generates minor variations of the same program while another generates a variety of programs with different underlying ideas. Even if both models have the same entropy, the latter model will be significantly better when augmented with search capabilities.

In our setting, we measure diversity by grounding it in *idea* space using a simple pair-matching strategy across all generated programs. Formally, suppose we have a pool of  $n$  code generations,  $\{c_1, \dots, c_n\}$ .

To compute the diversity of such a given generation pool, we ask an LLM to judge the similarity of two ideas in the following manner. First, we construct each of the  $\binom{n}{2}$  pairs. For each pair  $(c_i, c_j)$ , we judge (using an LLM) whether both  $c_i$  and  $c_j$  implement the same idea. We define this as the function  $S(c_i, c_j) \in \{0, 1\}$ , which evaluates to 1 if  $c_i$  and  $c_j$  implement the same idea and 0 otherwise. Our overall diversity score for a particular problem is then defined as:

$$D = 1 - \frac{\sum_{i < j} S(c_i, c_j)}{\binom{n}{2}} \quad (1)$$

Models that output programs implementing the same idea will have a score of  $D = 0$ , while models that output completely unique programs will have a score of  $D = 1$ . Overall, a score of  $D$  implies that if two codes are chosen at random, the probability that they implement similar ideas (as measured by the LLM) is  $D$ . In Appendix U, we describe this measure in additional mathematical depth.

We detail the prompts used in Appendix R.1 and use OpenAI’s GPT-4o-mini as the supporting LLM.

## 7 Conclusion

In this work, we find that diversity in idea space is incredibly useful to unlock significant achievements in the effectiveness of inference-time compute—otherwise referred to as search—particularly in code generation tasks. We propose PLANSEARCH, which obtains SOTA performance on all datasets tested, almost doubling baseline performance at pass@200. Additionally, we find strong correlation between our diversity metric and resulting performance gains from evaluating at pass@k instead of pass@1, which underscores the importance of idea diversity in effective search.

Looking forward, we believe that these insights and techniques can be applied to many other domains and unlock similar significant performance gains as seen here. Embracing idea diversity during inference will be crucial in realizing the full potential of LLMs.

## Acknowledgments and Disclosure of Funding

We would like to thank Jason Wei, Charlie Snell, Miles Turpin, Sail Wang, Horace He, Kenneth Li, Celia Chen, Rahul Chalamala, Alan Wu, Kevin Chang, and Jerry Wei for their helpful comments, suggestions and discussion over the course of this project.

## References

- [1] Zhiqiang Shen Aidar Myrzakhan, Sondos Mahmoud Bsharat. Open-llm-leaderboard: From multi-choice to open-style questions for llms evaluation, benchmark, and arena. *arXiv preprint arXiv:2406.07545*, 2024.
- [2] Anthropic. Claude 3.5 sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>, June 2024.
- [3] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models, 2021.
- [4] Anton Bakhtin, David J. Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H. Miller, and Noam Brown. Mastering the Game of No-Press Diplomacy via Human-Regularized Reinforcement Learning and Planning, October 2022. arXiv:2210.05492 [cs].



- [5] Hritik Bansal, Arian Hosseini, Rishabh Agarwal, Vinh Q. Tran, and Mehran Kazemi. Smaller, weaker, yet better: Training llm reasoners via compute-optimal sampling, 2024.
- [6] Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling, 2024.
- [7] Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018. Publisher: American Association for the Advancement of Science.
- [8] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456):885–890, 2019. Publisher: American Association for the Advancement of Science.
- [9] 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. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [10] Murray Campbell, A. Joseph Hoane, and Feng hsiung Hsu. Deep blue. *Artificial Intelligence*, 134(1):57–83, 2002.
- [11] Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. Codet: Code generation with generated tests, 2022.
- [12] Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. Codet: Code generation with generated tests. *arXiv preprint arXiv:2207.10397*, 2022.
- [13] Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei Zaharia, and James Zou. Are more llm calls all you need? towards scaling laws of compound inference systems, 2024.
- [14] 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. Evaluating Large Language Models Trained on Code, July 2021. *arXiv:2107.03374* [cs].
- [15] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating llms by human preference, 2024.
- [16] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [17] Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. Understanding back-translation at scale, 2018.

- [18] FAIR, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyang Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, December 2022. Publisher: American Association for the Advancement of Science.
- [19] Markus Freitag and Yaser Al-Onaizan. Beam search strategies for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, 2017.
- [20] Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models, 2023.
- [21] Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. Reasoning with language model is planning with world model, 2023.
- [22] Tatsunori B. Hashimoto, Hugh Zhang, and Percy Liang. Unifying Human and Statistical Evaluation for Natural Language Generation. *North American Association for Computational Linguistics (NAACL)*, April 2019. arXiv: 1904.02792.
- [23] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
- [24] Andy L. Jones. Scaling scaling laws with board games, 2021.
- [25] Sumith Kulal, Panupong Pasupat, Kartik Chandra, Mina Lee, Oded Padon, Alex Aiken, and Percy Liang. Spoc: Search-based pseudocode to code, 2019.
- [26] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, 2022.
- [27] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [28] Drew M. McDermott. The 1998 ai planning systems competition. *AI Magazine*, 21(2):35, Jun. 2000.
- [29] Aidan McLaughlin. AI Search: The Bitter-er Lesson, 2024. Accessed on September 3, 2024.
- [30] Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, Jilin Chen, Alex Beutel, and Ahmad Beirami. Controlled decoding from language models, 2024.
- [31] OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, Sep 2024. Accessed: October 1, 2024.
- [32] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad 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, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, 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, Reiichiro Nakano, Rajeev 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 H. 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 B. 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, C. J. Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, 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. GPT-4 Technical Report, March 2024. arXiv:2303.08774 [cs].

- [33] Hieu Pham, Xinyi Wang, Yiming Yang, and Graham Neubig. Meta back-translation, 2021.
- [34] Stuart Russell and Peter Norvig. Artificial intelligence: a modern approach. 2002.
- [35] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools, 2023.
- [36] Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models with monolingual data, 2016.
- [37] Claude E Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 623–656, 1948.
- [38] Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023.
- [39] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, 529(7587):484–489, January 2016.

- [40] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, and others. Mastering the Game of Go Without Human Knowledge. *Nature*, 550(7676):354–359, 2017. Publisher: Nature Publishing Group.
- [41] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024.
- [42] Richard S Sutton. The bitter lesson. *Incomplete Ideas*, 2019.
- [43] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv, 2022. arXiv:2201.11903 [cs].
- [44] Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. An empirical analysis of compute-optimal inference for problem-solving with language models, 2024.
- [45] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of Thoughts: Deliberate Problem Solving with Large Language Models, May 2023. arXiv:2305.10601 [cs].
- [46] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models, 2023.
- [47] Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. Quiet-star: Language models can teach themselves to think before speaking, 2024.
- [48] Eric Zelikman, Qian Huang, Gabriel Poesia, Noah D. Goodman, and Nick Haber. Parsel: Algorithmic reasoning with language models by composing decompositions, 2023.
- [49] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. STaR: Bootstrapping Reasoning With Reasoning, May 2022. arXiv:2203.14465 [cs].
- [50] Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. Rest-mcts\*: Llm self-training via process reward guided tree search, 2024.
- [51] Hugh Zhang, Daniel Duckworth, Daphne Ippolito, and Arvind Neelakantan. Trading off diversity and quality in natural language generation. In *Proceedings of the workshop on human evaluation of NLP systems (HumEval)*, pages 25–33, Online, April 2021. Association for Computational Linguistics.
- [52] Hugh Zhang and David C. Parkes. Chain-of-thought reasoning is a policy improvement operator, 2023.
- [53] Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B. Tenenbaum, and Chuang Gan. Planning with large language models for code generation. In *The Eleventh International Conference on Learning Representations*, 2023.
- [54] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables complex reasoning in large language models, 2023.

## A Limitations and Future Work

While PLANSEARCH substantially improves diversity over idea space at inference-time, fundamentally, improvements in diversity should come at the post-training stage, like with methods such as o1 [31]. This likely requires re-imagining the post-training pipeline for LLMs around search, instead of the current paradigm optimized for a single correct response. We are optimistic about future work in designing improved post-training objectives to maximize both quality and diversity, while specifically optimized to use inference-time compute to maximum effectiveness.

PLANSEARCH and IDEASEARCH tradeoff a slight deterioration of pass@1 performance for a large improvement in pass@k performance. However, in many such cases outside of code generation, it is infeasible to run an LLM-based model for more than a few attempts at most. For example, in Figure 6, PLANSEARCH does not significantly outperform REPEATED SAMPLING until  $k \geq 4$ .

Fortunately, many filtering algorithms exist, which mitigates this tradeoff by implicitly bringing pass@k (for high  $k$ ) to pass@1 (or lower  $k$ ), i.e. shifting the original pass@k curve leftward. Even the simplest filtering—public test filtering—improves PLANSEARCH’s pass@1 significantly above REPEATED SAMPLING’s pass@1, which continues as  $k$  increases. Moreover, *most to almost all* base models with public test filtering outperform their instruct model variants at pass@1, no matter the dataset (see Appendix M). Since base models’ pass@1 is known to be worse than instruct models to trade off for higher diversity, we suggest a new paradigm—developing search algorithms which tradeoff pass@1 performance for much stronger pass@k performance, then filtering the generated solutions to extract the pass@k *back into* pass@1.

With good filtering methods, which we demonstrate can be simple in nature, pass@k, for medium  $k$ , can be effectively brought down to pass@1, emphasizing a similar paradigm of increasing diversity, then strengthening existing filtering methods, even for domains outside of code generation that are out of scope of this paper.

Finally, a natural extension of this work is training the underlying model itself on successful plans and code solutions obtained from PLANSEARCH. This has the potential to distill the pass@k into the pass@1—without inference-time methods like filtering—by reducing the likelihood of the model going down unfavorable branches of the search tree. We believe that such training is likely to significantly improve the model and look forward to future work in this direction.

## B Public Test Filtering

Public test filtering is a method which only chooses samples out of the original pool  $n$  which pass the public tests. This is particularly useful in settings such as code deployment where executing the full suite of tests may be computationally costly or otherwise undesirable (e.g. in a coding contest where every incorrect submission is penalized). Thus, instead of submitting all  $n$  codes, after public test filtering, only codes  $c_i$  would be submitted such that  $c_i(x_j) = y_j$  for all  $j \in \{1, \dots, u\}$ , where  $c_i(x)$  refers to the output from running the code on some input  $x$ . The primary effect of public test filtering is to shift the pass@k curve leftward, since public test filtering will discard low quality candidate solutions that either fail to compile or fail elementary test cases for the problem.

All problems in MBPP+, HumanEval+, and LiveCodeBench come with a few public tests which are usually used to sanity check any submissions. We can further improve performance by filtering on these public tests before a final submission, as described. Applying public test filtering reduces the number of samples to achieve the same accuracy by tenfold: PLANSEARCH to achieve a 77.1% accuracy on LiveCodeBench after just 20 submissions (pass@20) compared to a pass@200 of 77.0% without using public filtering (see Figure 9). We provide full results for the other datasets in Appendix E.

## C Result Details

### C.1 Datasets

We evaluate our search methods on three benchmarks: MBPP+, HumanEval+ [27], and LiveCodeBench [23]. MBPP [3] and HumanEval [14] are some of the most widely used code benchmarks

in the field. However, since both benchmarks provide only a few test cases, [27] updates both benchmarks with additional test cases that increase the benchmarks’ robustness to reward hacking. LiveCodeBench is a benchmark for coding that consists of competitive programming problems which typically require advanced reasoning capabilities. Given the reality that coding data is often highly upsampled during pre-training [32, 16], LiveCodeBench differentiates itself from other benchmarks by taking care to segregate problems by date to avoid data contamination concerns. For this paper, we use only the subset of problems between May 2024 and September 2024 to avoid possibilities of contamination. We choose May 2024 as the cutoff date to ensure that our results with our best performing model (Claude 3.5 Sonnet) are not due to contamination, because Claude 3.5 Sonnet has a knowledge cutoff of April 2024. To ensure fair comparison, we use the same cutoff for all models evaluated, even though the precise cutoff dates for other models may vary slightly from May 2024.

## C.2 Experiment Details

For all search algorithms, we require that all output code be in the correct format specified, and we mark a solution as incorrect if it does not follow the intended formatting. The extracted code is then run through all tests of the program and marked as correct if and only if it passes all tests.

All models are run with temperature 0.9 and top- $p$  of 0.95. Temperature was determined through a coarse hyperparameter sweep on REPEATED SAMPLING and IDEASEARCH from  $T \in \{0.0, 0.1, 0.2, \dots, 1.2\}$ , which we describe in Appendix I.

The OpenAI, Claude, and DeepSeek endpoints were used for their respective model inference. For other models not available on those APIs, a single node of 8 NVIDIA H100s were used for inference. For test execution, a cluster on the order of 100 CPU nodes were used.

Both REPEATED SAMPLING and IDEASEARCH generate exactly  $n$  codes, whereas PLANSEARCH generates a variable number of codes, usually ranging on the order of 300 to 400. To compute pass@ $k$ , we use the unbiased estimator in Equation 4 [14]<sup>1</sup>. If  $k > n$ , we assume the remaining generations did not pass. To compute pass@ $k$  for filtering, we limit the pool of codes to those that are filtered, meaning that both  $n$  and  $c$  may shrink in size. This can be thought of as a conditional probability, where the condition is that the code passes public tests.

## D Full Pass@K curves for All Models and All Benchmarks

See Figures 4, 5, 6. We plot all models and methods on HumanEval+, MBPP+ [27], and LiveCodeBench [23], respectively.

We include error bars as derived in Appendix Q.

---

<sup>1</sup>Note that the estimator in Equation 4 theoretically requires that the number of successes follows a binomial distribution. REPEATED SAMPLING and IDEASEARCH obey this, but PLANSEARCH generations may not be independent. See Appendix R for more discussion.

Pass@k vs k for Methods on HumanEval+

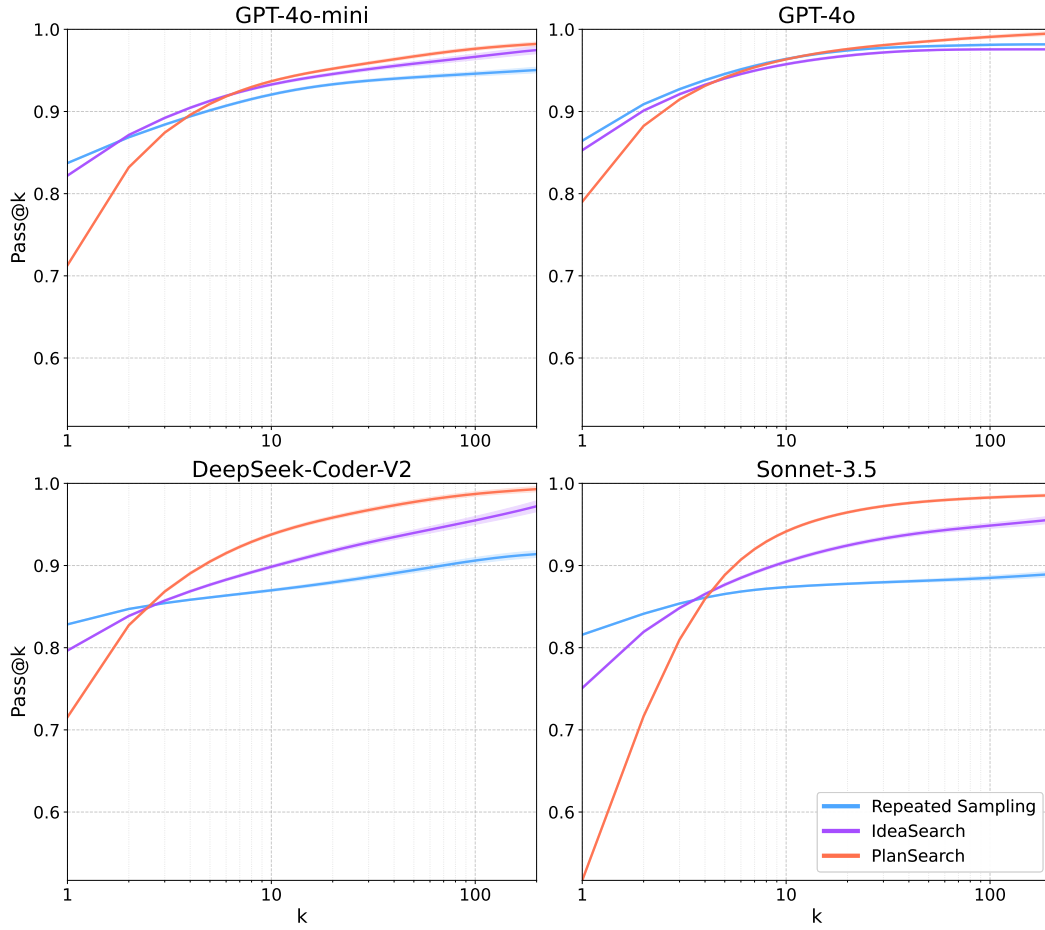


Figure 4: Pass@k performance of all models and methods on HumanEval+, plotted over  $k \in \{1, \dots, 200\}$ .

Pass@k vs k for Methods on MBPP+

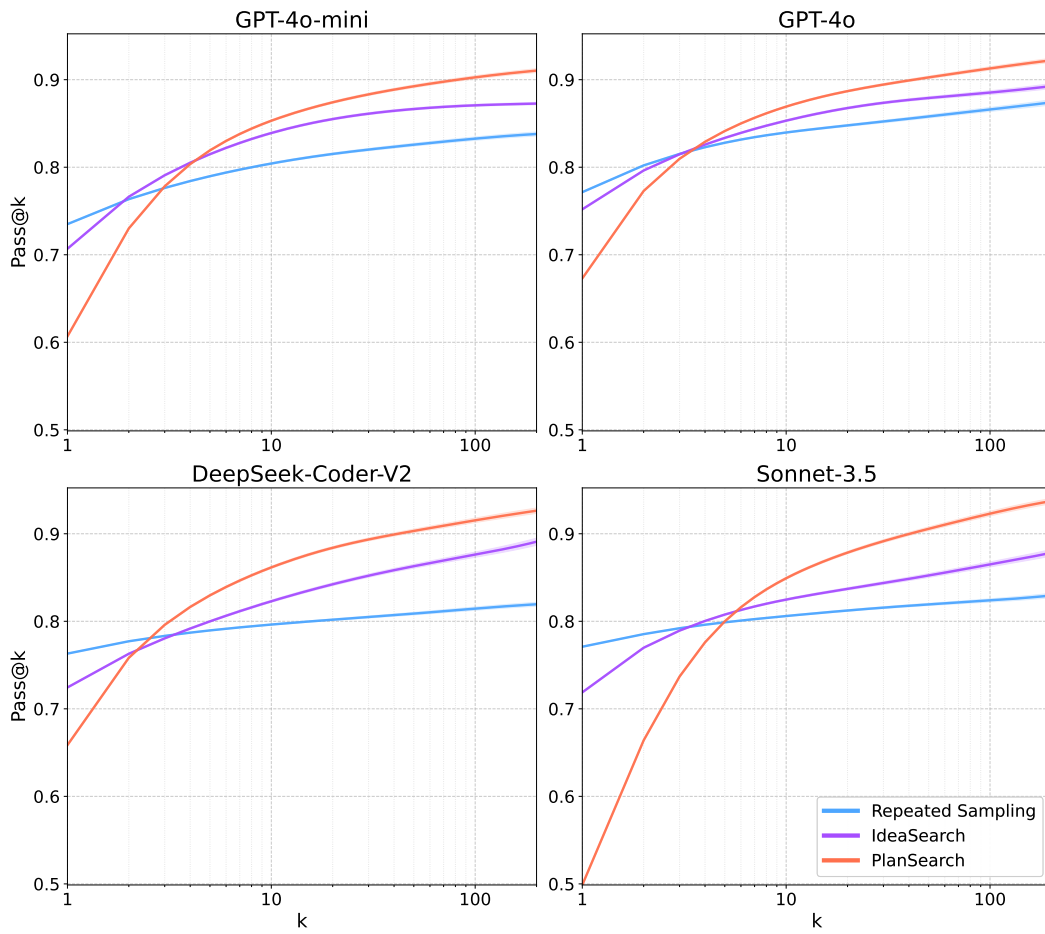


Figure 5: Pass@k performance of all models and methods on MBPP+, plotted over  $k \in \{1, \dots, 200\}$ .



### Pass@k vs k for Methods on LiveCodeBench

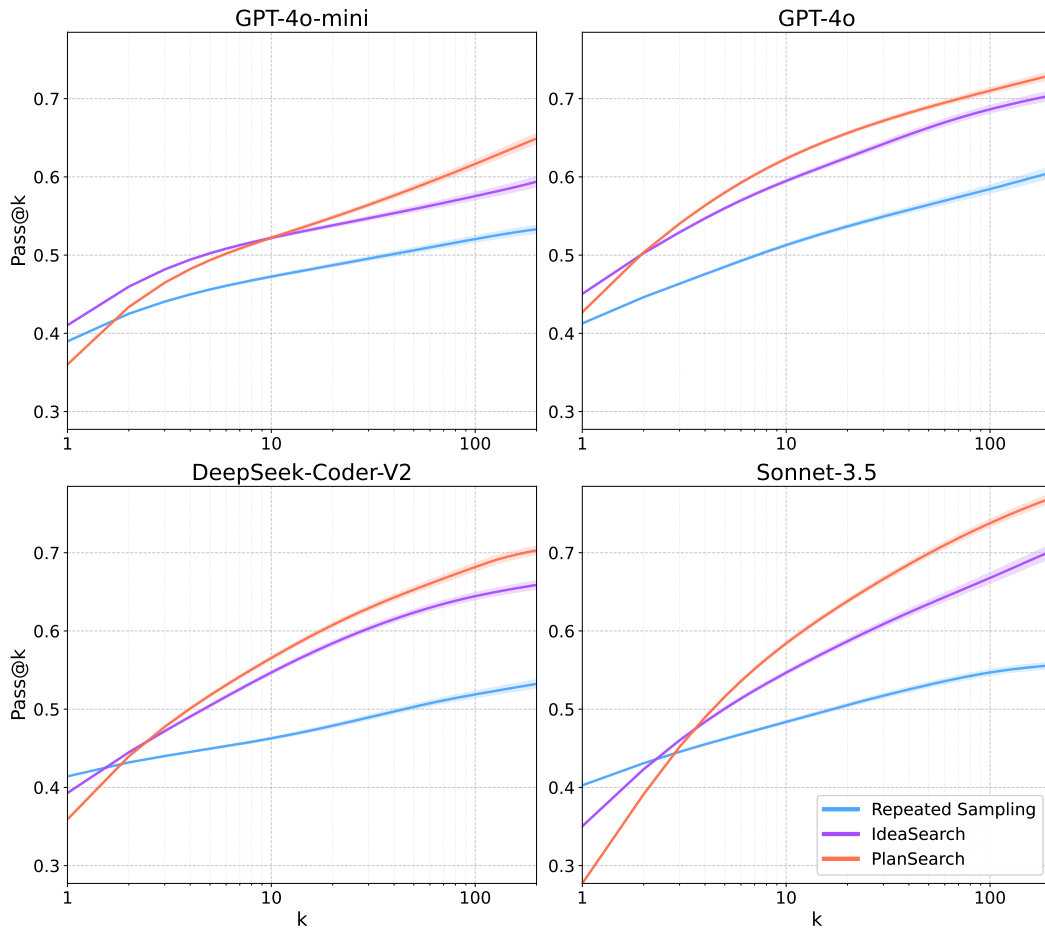


Figure 6: Pass@k performance of all models and methods on LiveCodeBench, plotted over  $k \in \{1, \dots, 200\}$ .

## E Full Pass@k Curves with Public Filtering

The idea of public test filtering is to shift pass@k curves leftward (i.e., bringing high  $k$  to low  $k$ ), so we plot curves in detail over  $k \in \{1, \dots, 20\}$ . Dotted lines are provided for reference of the base method pass@k before filtering. For LiveCodeBench, even at 10 completions, PLANSEARCH outperforms *filtered* REPEATED SAMPLING by a flat 30 to 40%.

See Figures 7, 8, 9. We plot all models and methods with public test filtering on HumanEval+, MBPP+ [27], and LiveCodeBench [23], respectively.

Pass@k vs k for Methods with Public Filtering on HumanEval+

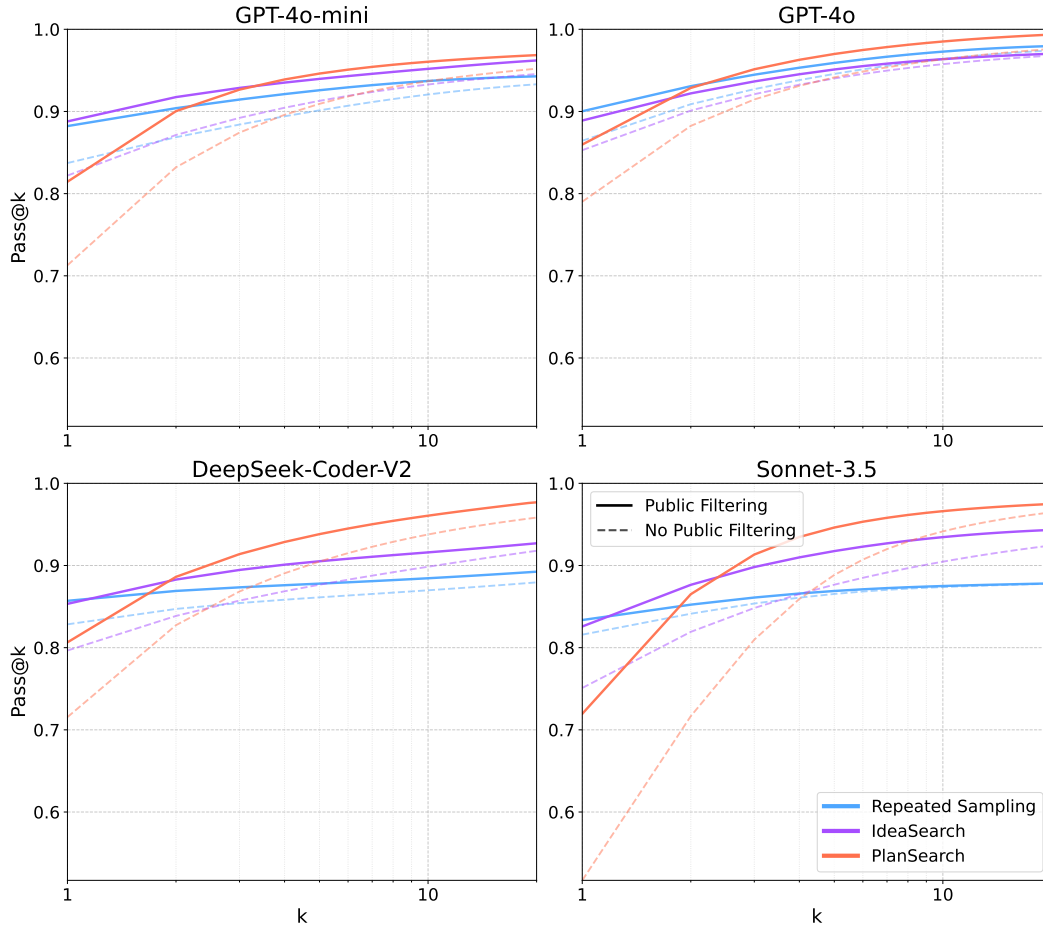


Figure 7: Pass@k performance of all models and methods on HumanEval+, with public test filtering, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k before filtering.

### Pass@k vs k for Methods with Public Filtering on MBPP+

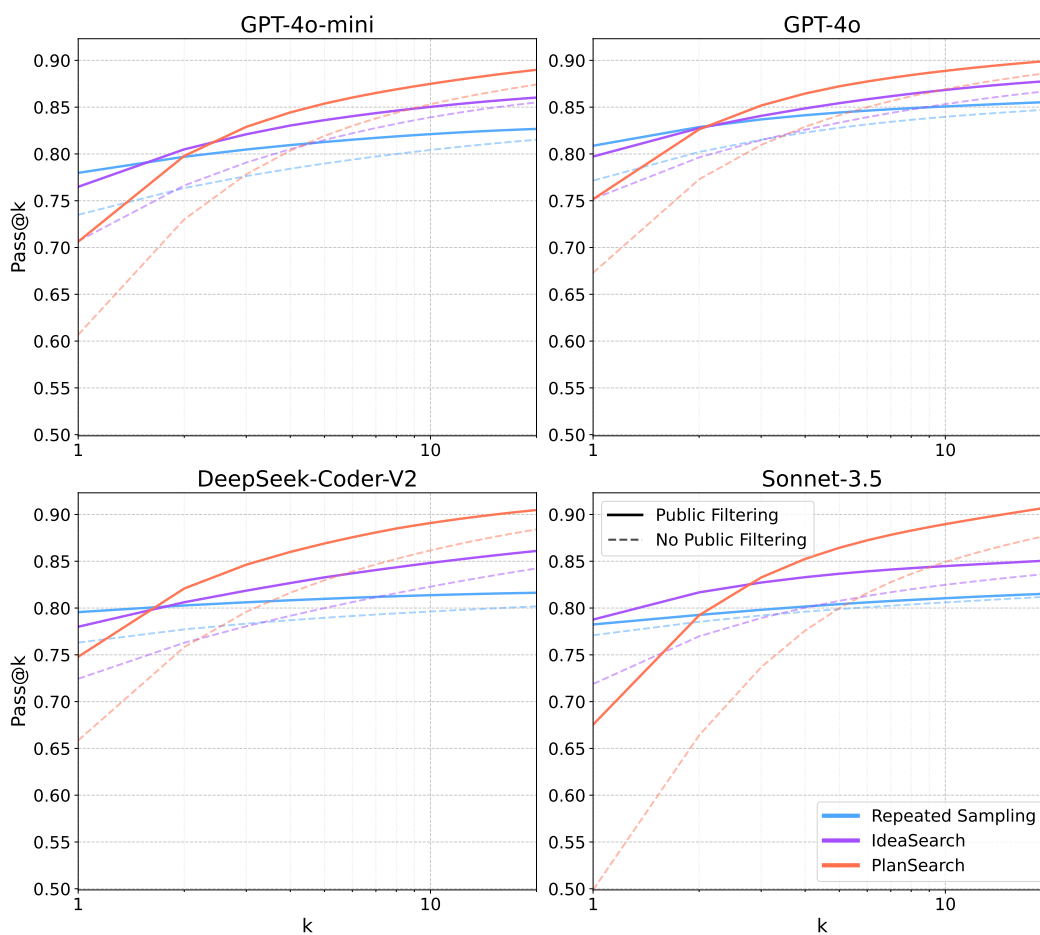


Figure 8: Pass@k performance of all models and methods on MBPP+, with public test filtering, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k before filtering.

Pass@k vs k for Methods with Public Filtering on LiveCodeBench

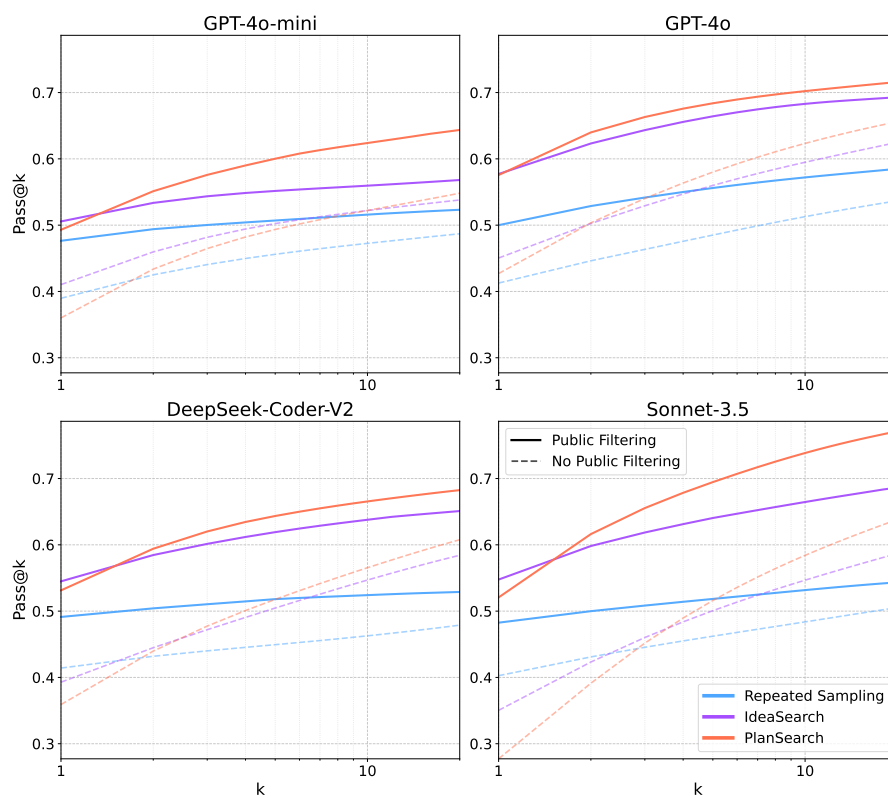


Figure 9: Pass@k performance of all models and methods on LiveCodeBench, with public test filtering, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k before filtering.

## F Average Relative Improvements

See Figures 10, 11, 12. To create these graphs, the relative improvements of each point on all pass@k curves are computed and compared to the respective pass@1 of REPEATED SAMPLING. Then these values are averaged over all models, so that there is one curve per method per dataset. The datasets are HumanEval+, MBPP+ [27], and LiveCodeBench [23], respectively. For the public test filtered versions, see Figures 13, 14, 15.

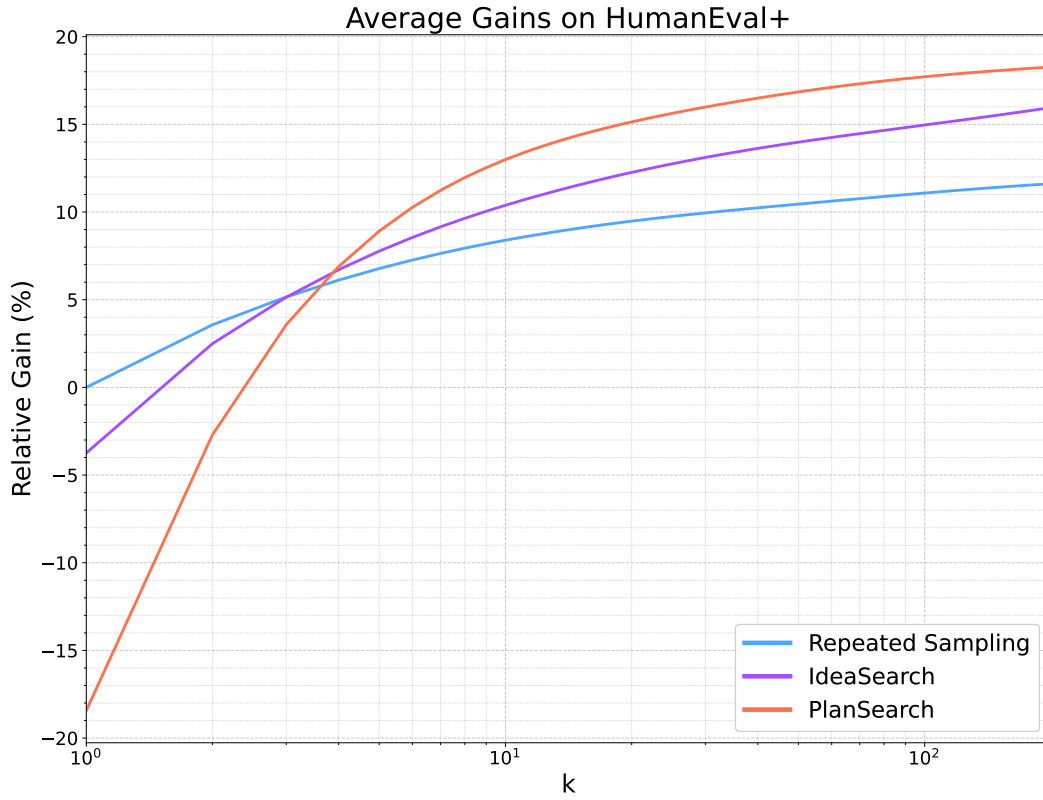


Figure 10: Performance gain over REPEATED SAMPLING@1 averaged over all models on HumanEval+, plotted over  $k \in \{1, \dots, 200\}$ .

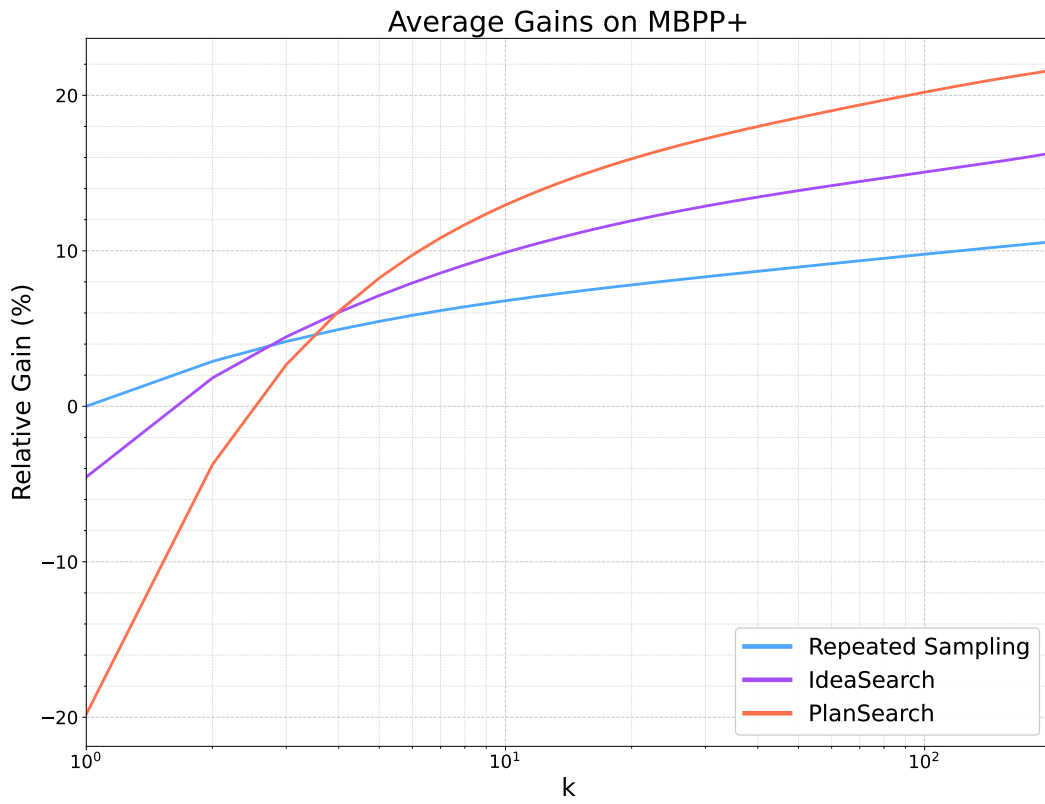


Figure 11: Performance gain over REPEATED SAMPLING@1 averaged over all models on MBPP+, plotted over  $k \in \{1, \dots, 200\}$ .

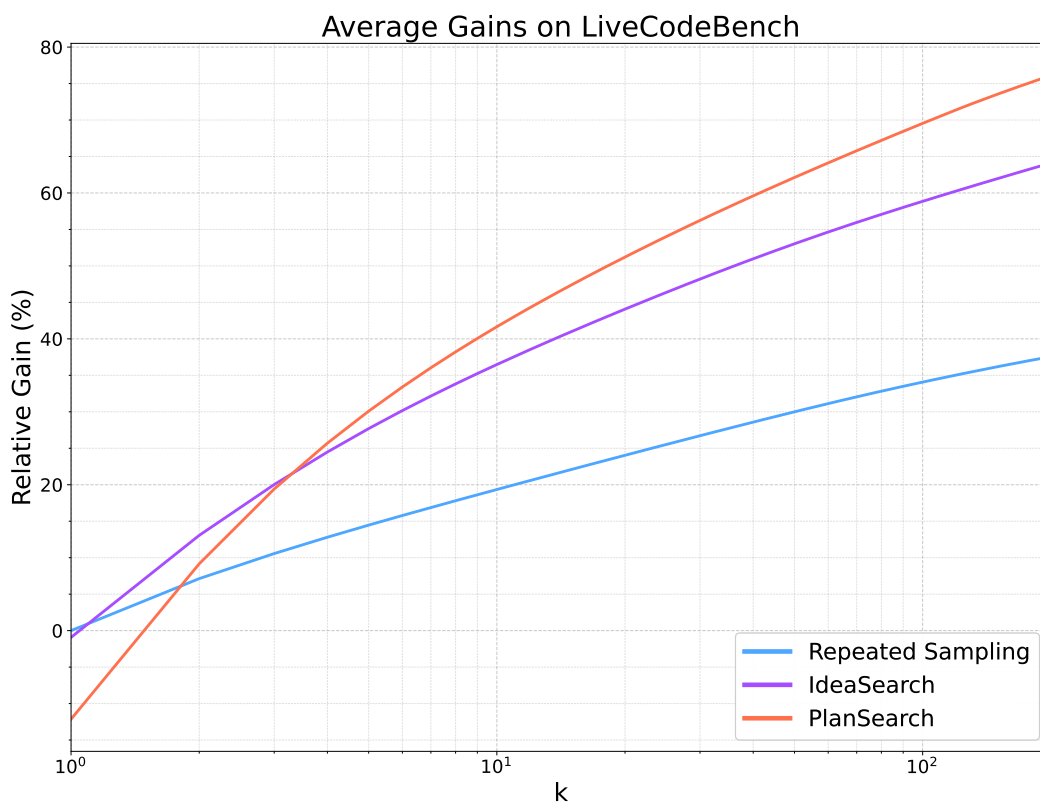


Figure 12: Performance gain over REPEATED SAMPLING@1 averaged over all models on LiveCodeBench, plotted over  $k \in \{1, \dots, 200\}$ .

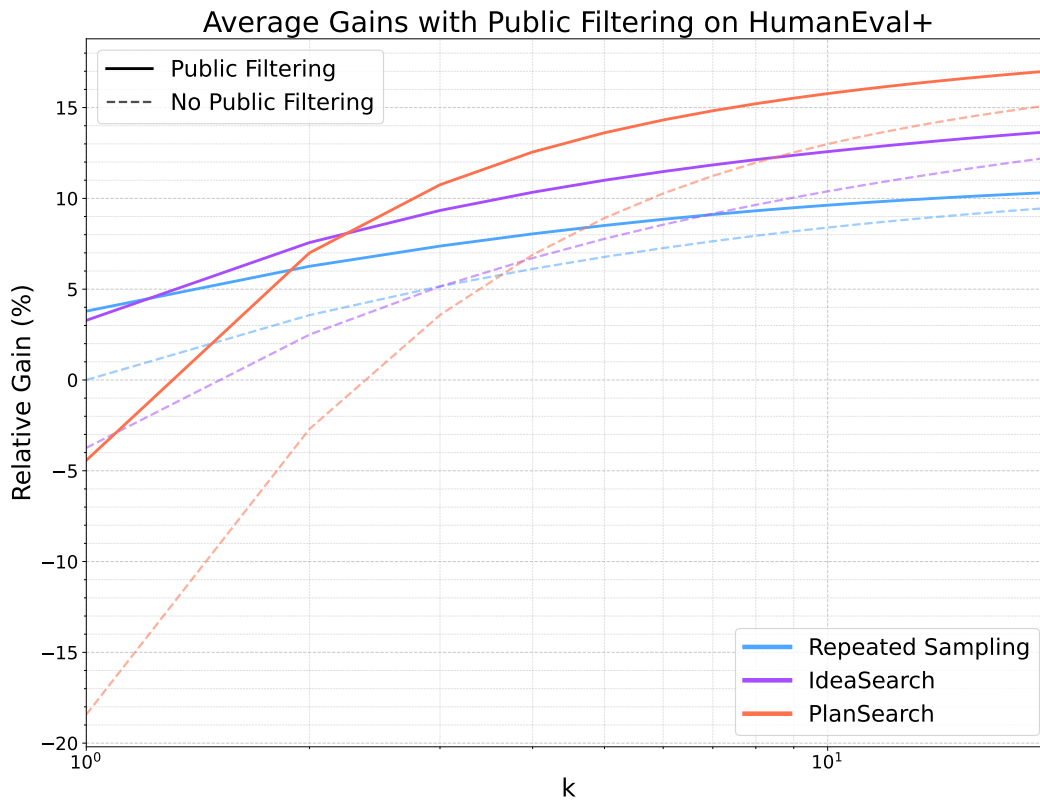


Figure 13: Average performance gain over all models of methods with public test filtering compared to REPEATED SAMPLING@1, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k (before filtering).



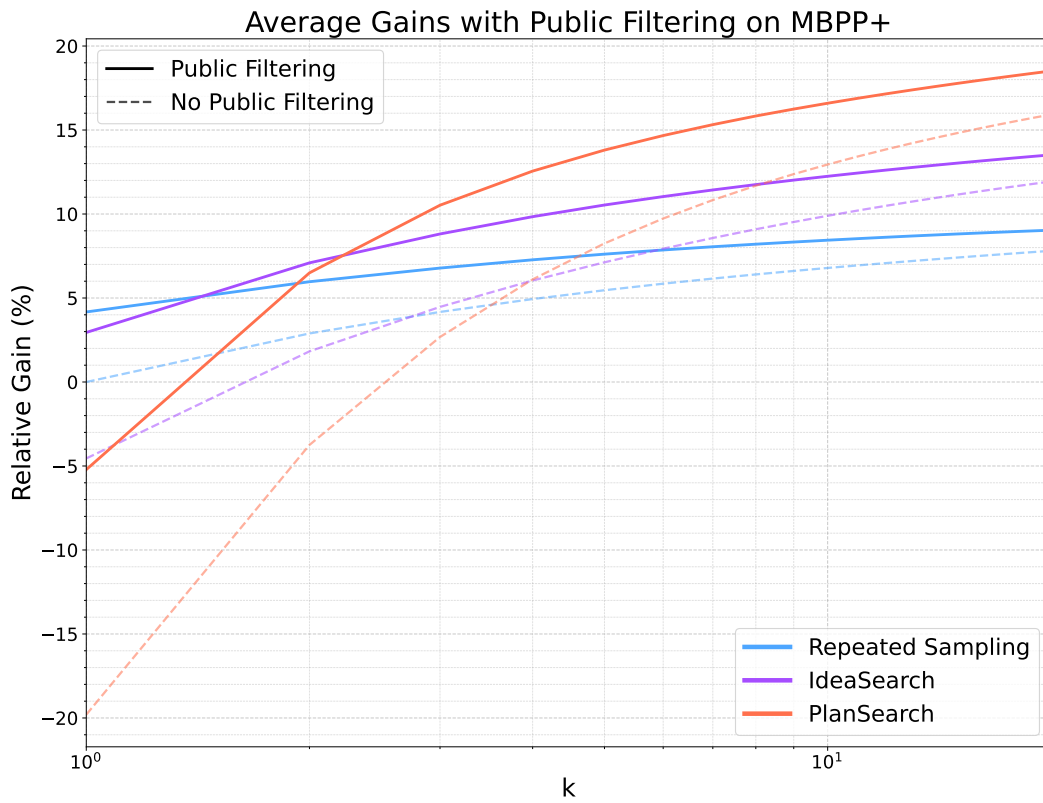


Figure 14: Average performance gain over all models of methods with public test filtering compared to REPEATED SAMPLING@1, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k (before filtering).

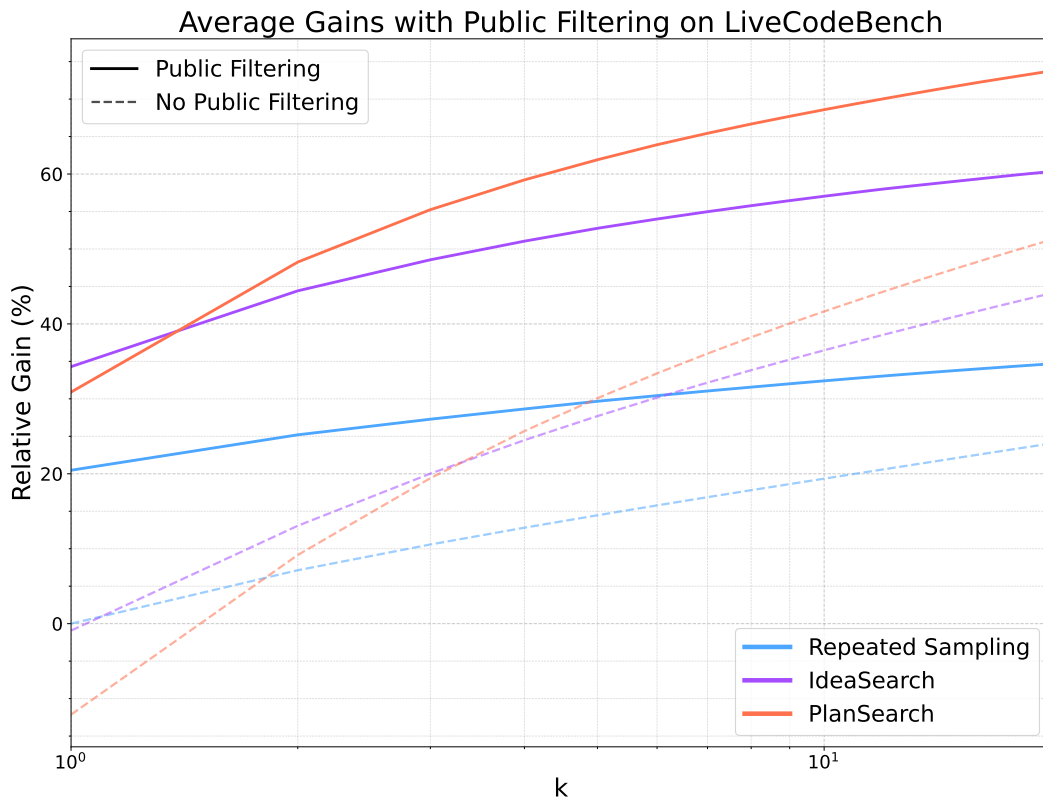


Figure 15: Average performance gain over all models of methods with public test filtering compared to REPEATED SAMPLING@1, plotted over  $k \in \{1, \dots, 20\}$ . Note that dotted lines are provided for reference of the base method pass@k (before filtering).

## G Compute Normalized Pass@K Graphs

See Figure 16. For each run of a method in Appendix D, we compute the number of generated tokens needed per completion, per problem, independently on each dataset. Then, we average across all datasets to obtain 244 generated tokens per completion per problem for REPEATED SAMPLING, and 1,428 generated tokens per completion per problem for PLANSEARCH.

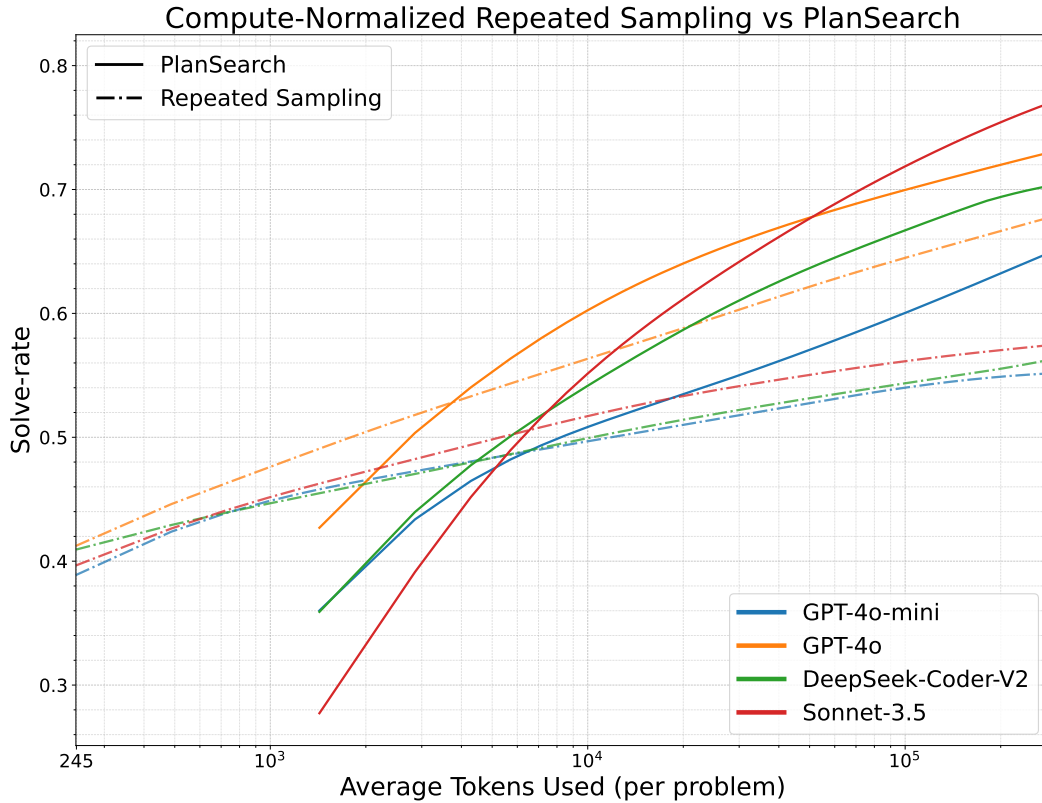


Figure 16: Normalized pass@k by average tokens used per problem. REPEATED SAMPLING uses roughly 244 tokens per completion per problem, and PLANSEARCH uses roughly 1428 tokens per completion per problem. When we normalize compute across methods, we find that PLANSEARCH begins to be more effective than repeated sampling if the user is willing to sample at least 10,000 tokens per problem.

## H Comparison with Chain-of-Thought

See Figures 17, 18, 19, which are run on LiveCodeBench [23], MBPP+, and HumanEval+ [27], respectively. These are the same plots as Appendix D, with CoT [43]. See Figures 20, 21, 22 for the public test filtered versions.

Pass@k vs k with CoT on LiveCodeBench

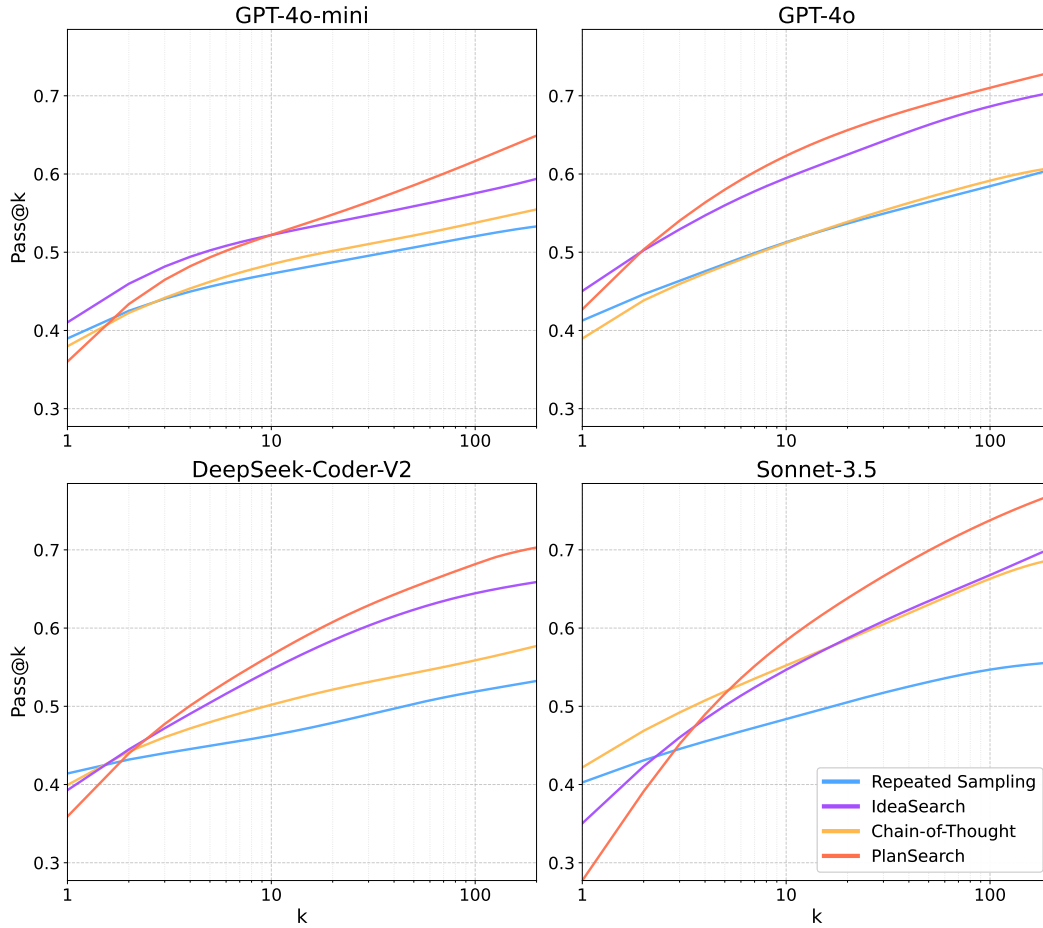


Figure 17: Pass@k graphs on LiveCodeBench, with the Chain-of-Thought baseline.

Pass@k vs k with CoT on MBPP+

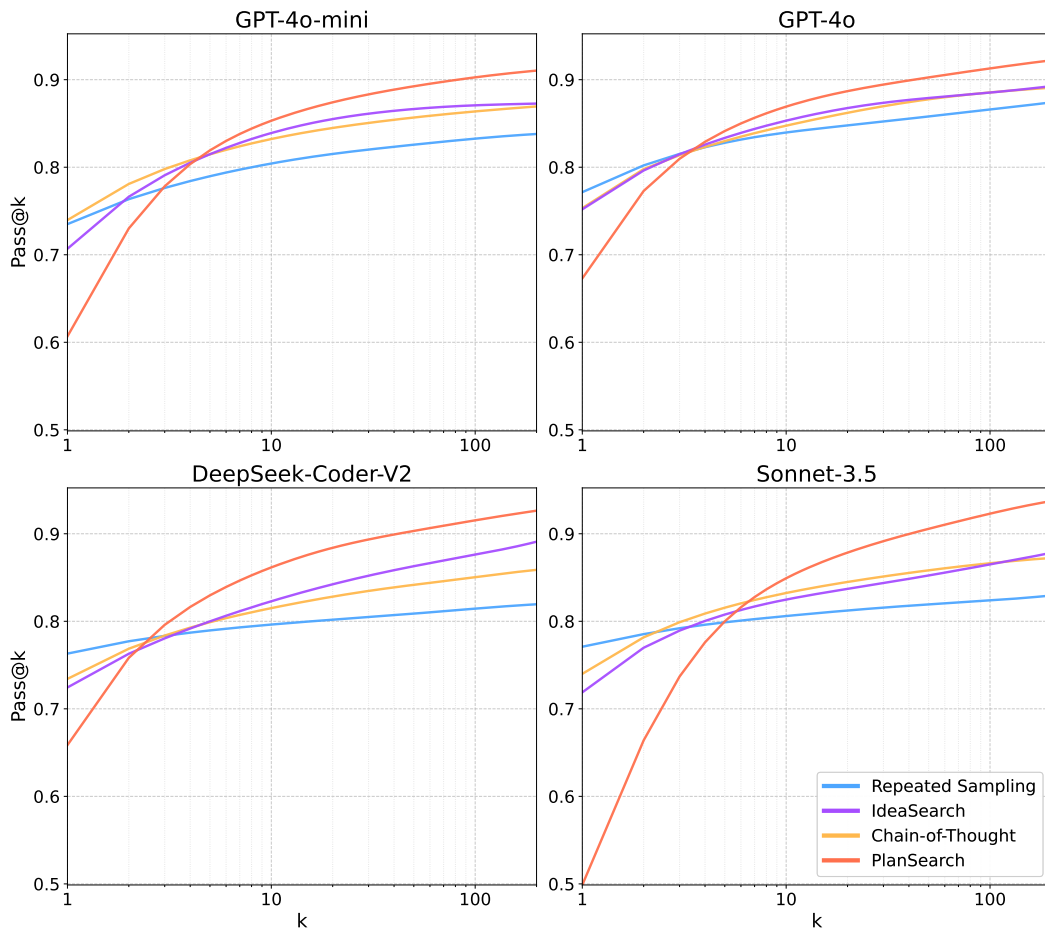


Figure 18: Pass@k graphs on MBPP+, with the Chain-of-Thought baseline.

Pass@k vs k with CoT on HumanEval+

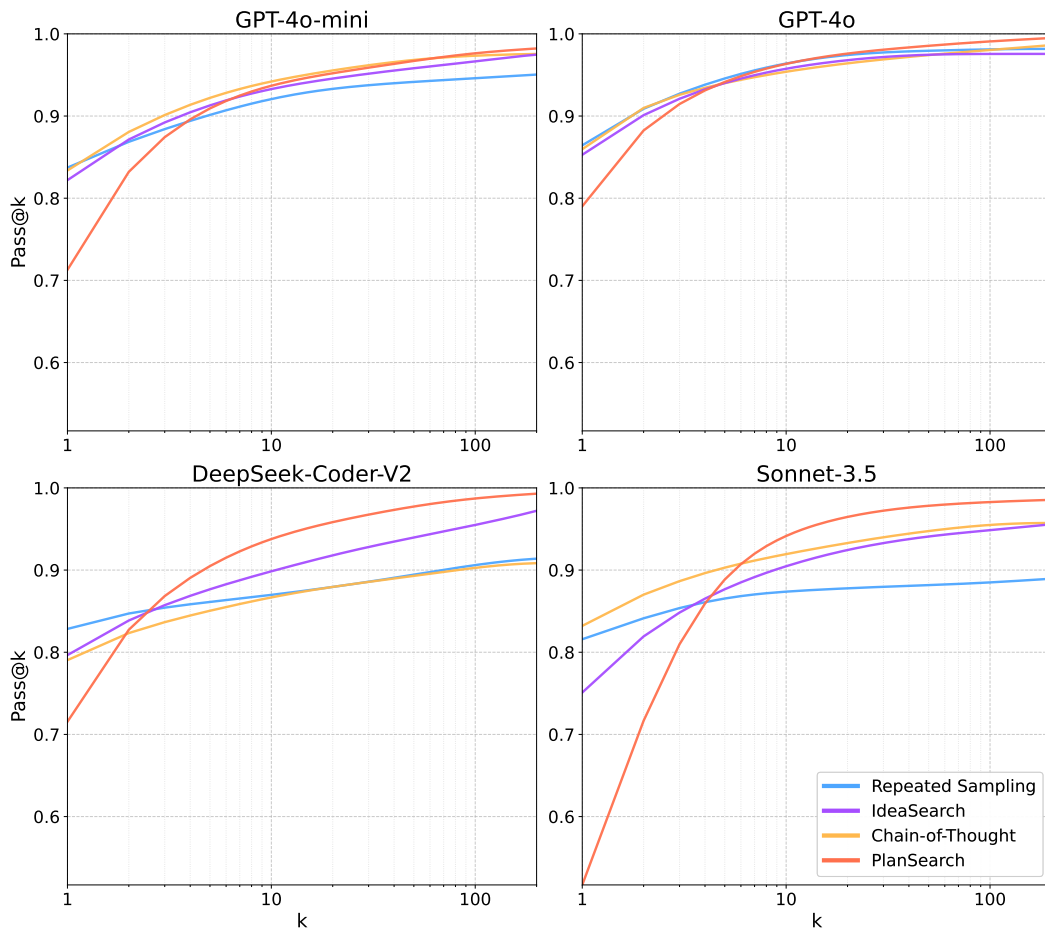


Figure 19: Pass@k graphs on HumanEval+, with the Chain-of-Thought baseline.

### Pass@k vs k with CoT (Public Filtering) on LiveCodeBench

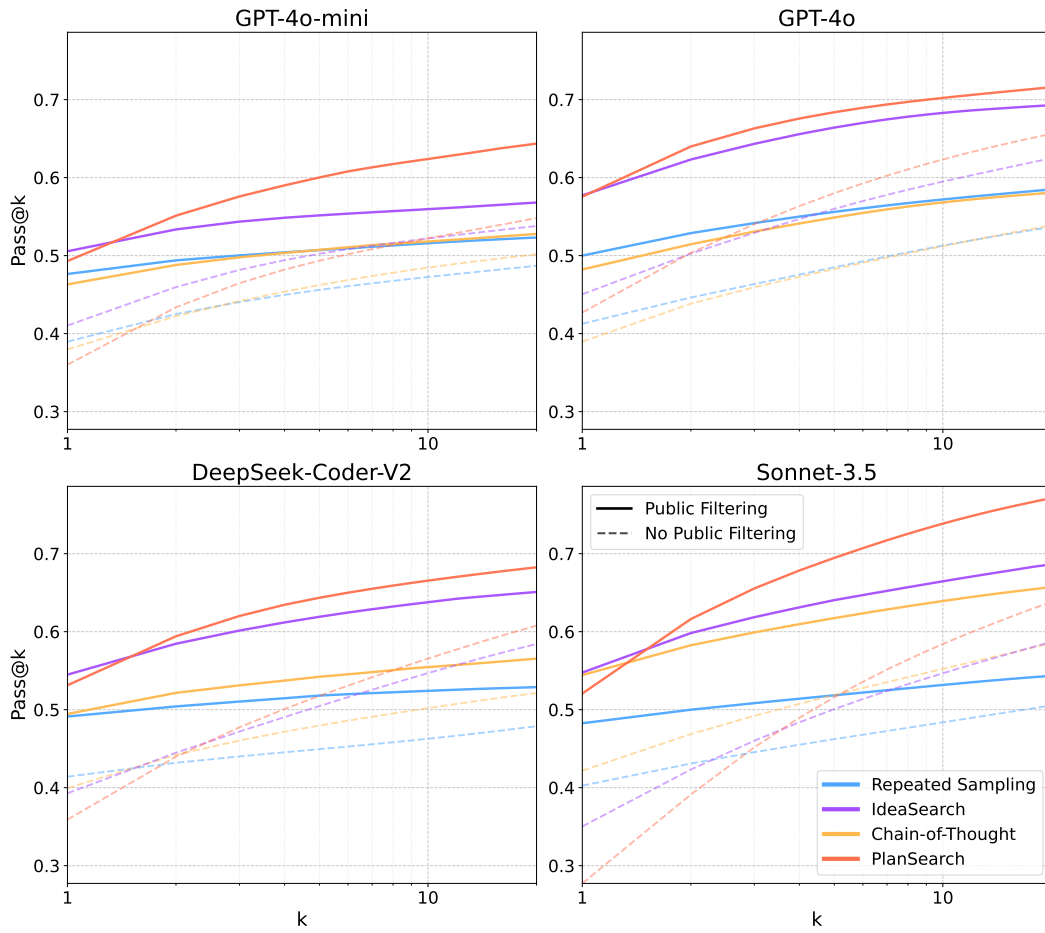


Figure 20: Pass@k graphs on LiveCodeBench, with the Chain-of-Thought baseline and public filtering.

Pass@k vs k with CoT (Public Filtering) on MBPP+

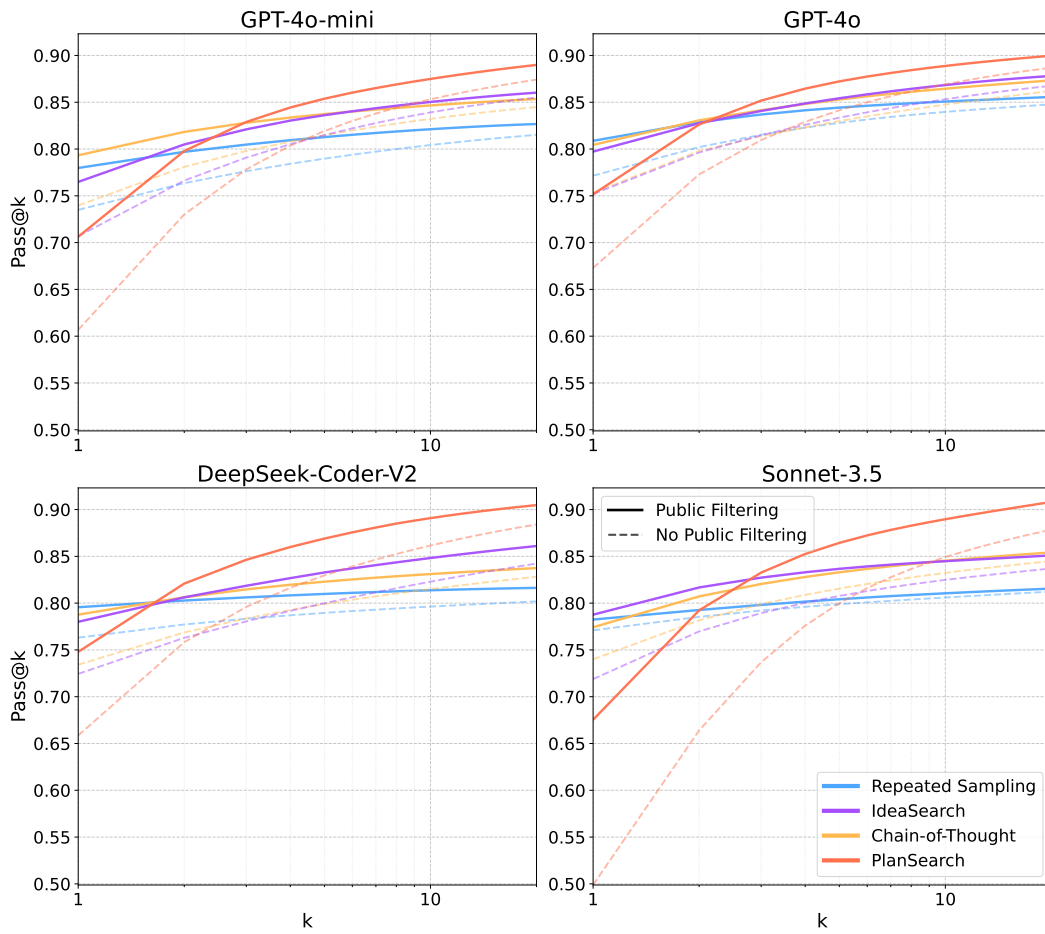


Figure 21: Pass@k graphs on MBPP+, with the Chain-of-Thought baseline and public filtering.



Pass@k vs k with CoT (Public Filtering) on HumanEval+

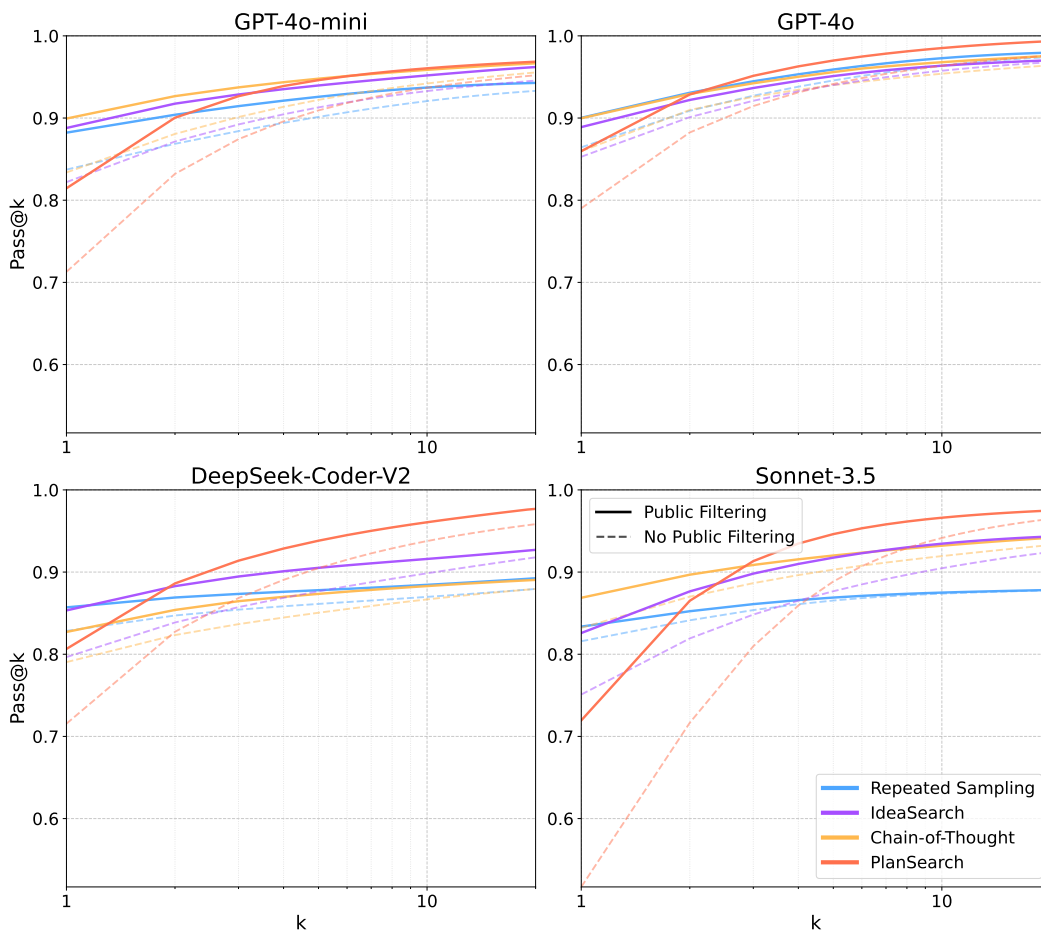


Figure 22: Pass@k graphs on HumanEval+, with the Chain-of-Thought baseline and public filtering.

## I Ablation on Temperature for REPEATED SAMPLING and IDEASEARCH

See Figure 23. We sweep over temperature increments of 0.1 from 0.0 to 1.2, inclusive, with top- $p$  of 0.95, on REPEATED SAMPLING and IDEASEARCH.

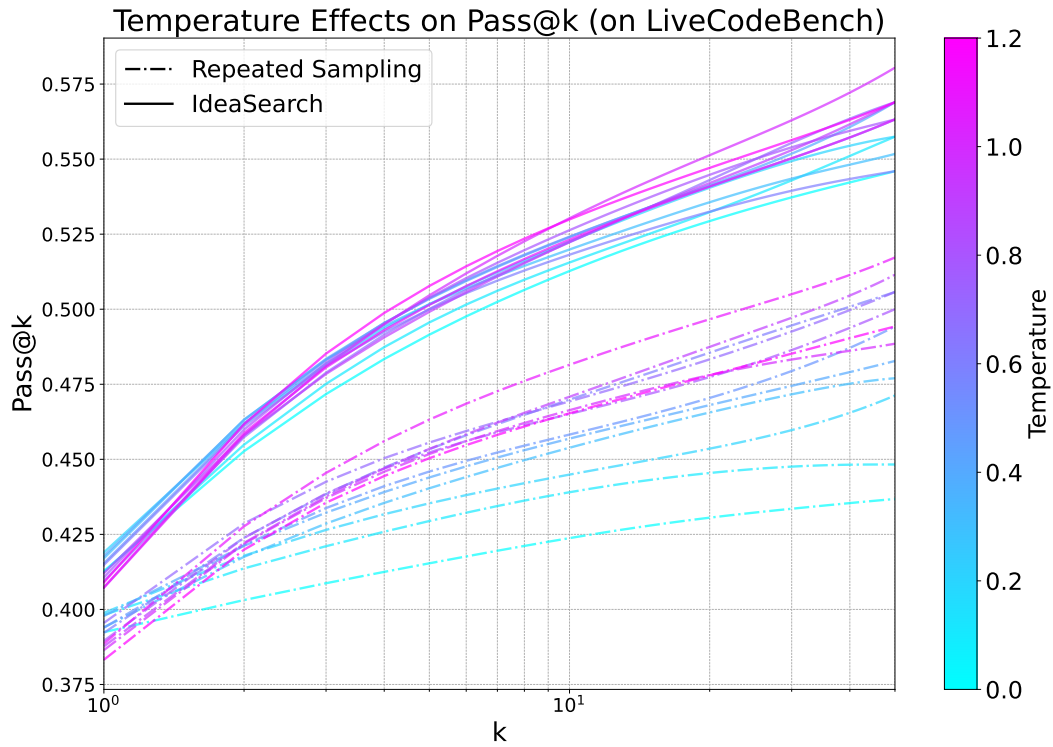


Figure 23: Sweep over temperature in 0.1 increments from 0.0 to 1.2. REPEATED SAMPLING and IDEASEARCH both exhibit pass@k improvements at higher temperature, although it seems that higher temperatures may begin to plateau.

## J Diversity Score vs Search Improvement Plots for MBPP+ and HumanEval+

See Figures 24, 25, 3. Each figure is made through running the diversity measure as described in Section 6.1 on the generated codes of each run, then compared with the relative gain from pass@k compared to pass@1.

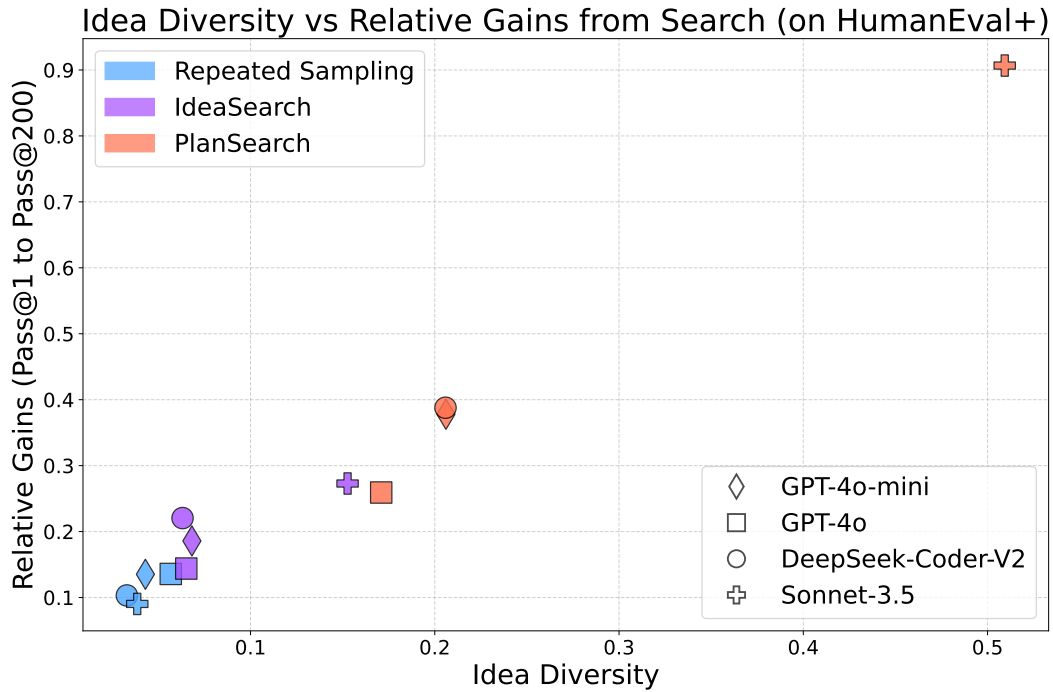


Figure 24: Relationship between the measured diversity score as described in Section 6.1 (where higher is more diverse) and relative improvement from the pass@1 of the method to the pass@200 of the method.

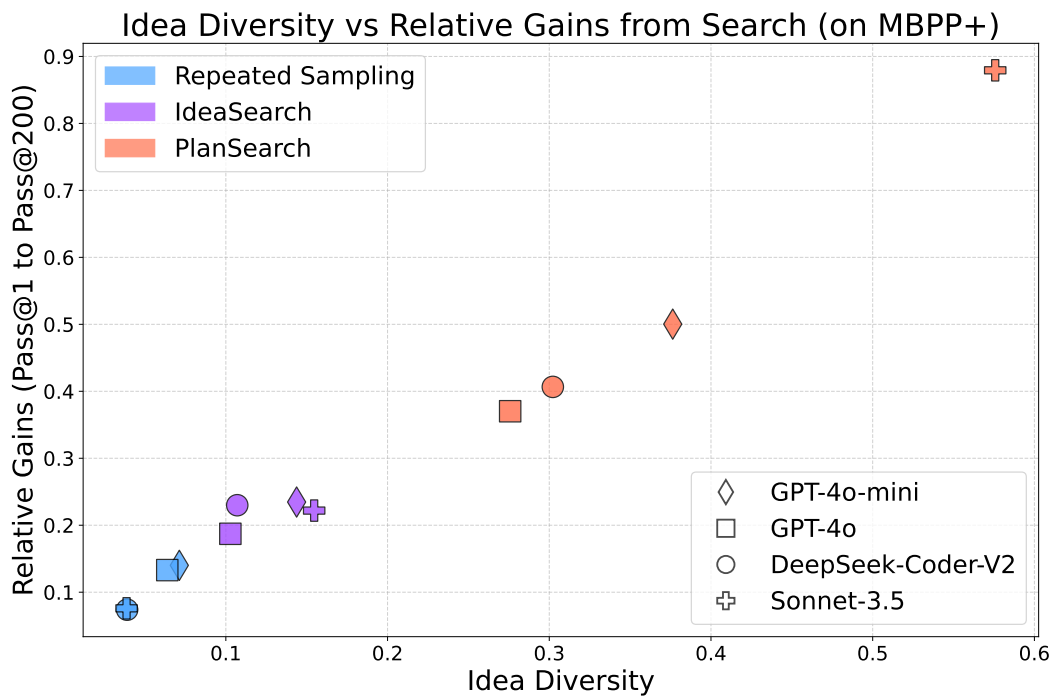


Figure 25: Relationship between the measured diversity score as described in Section 6.1 (where higher is more diverse) and relative improvement from the pass@1 of the method to the pass@200 of the method on MBPP+.

## K Ablations

We run ablations on the core parts of PLANSEARCH, using GPT-4o-mini on LiveCodeBench [23]. On the diverse observation generation side, we first verify our choice of  $S = 2$ —the maximum subset size to sample from out of a given pool of observations—and also compare performance across varying the number of observation layers used (Figures 26, 27). On the implementation side, we compare different sections of the proposed pipeline (see Figure 1) to translate combinations of observations to code in Figure 28. We compare each method’s pass@ $k$  from  $k = 1$  to  $k = 200$ .

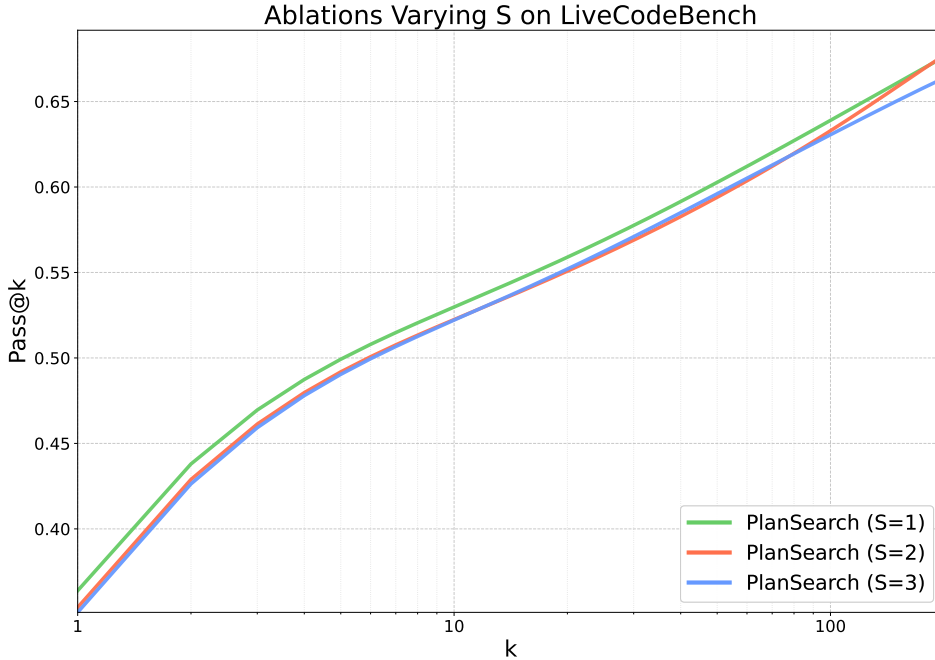


Figure 26: We run ablations of PLANSEARCH with different  $S$ —controlling the maximum subset size from a given pool of observations. There is not a large difference between different  $S$ , but we find that 1 or 2 is slightly more optimal, although if more completions are desired,  $S$  can be increased to 3 as well.

The effect  $S$ , the maximum observation subset size to build upon, has on performance is not overly significant; there are small degradations as  $S$  is increased to 3, but not noticeable. We choose  $S = 2$  to obtain more code completions. See Figure 26.

Increasing  $L$ , the maximum number of layers of the observation tree, increases pass@ $k$  at large enough  $k$  (above 50). We choose  $L = 2$  to strike a balance between extracting a large pass@ $k$  gain while keeping compute costs reasonable. See Figure 27.

From Figure 28, we see that our overall translation step adds minor pass@ $k$  gains. We deconstruct the translation step into parts: the pseudocode step, the fix step (i.e., asking the model to fix its proposed solution sketch), and creating the solution sketch at all (which includes the fix step).

- “No solution sketch, no pseudocode” implies using a given observation combination to directly prompt for the solution code.
- “No pseudocode” implies skipping the pseudocode step. In other words, given a solution sketch, the sketch is directly translated into code.
- “No fix step, no pseudocode” implies the fix step is skipped, as well as the pseudocode. In order to have the same number of completions, the whole PLANSEARCH pipeline is run twice.

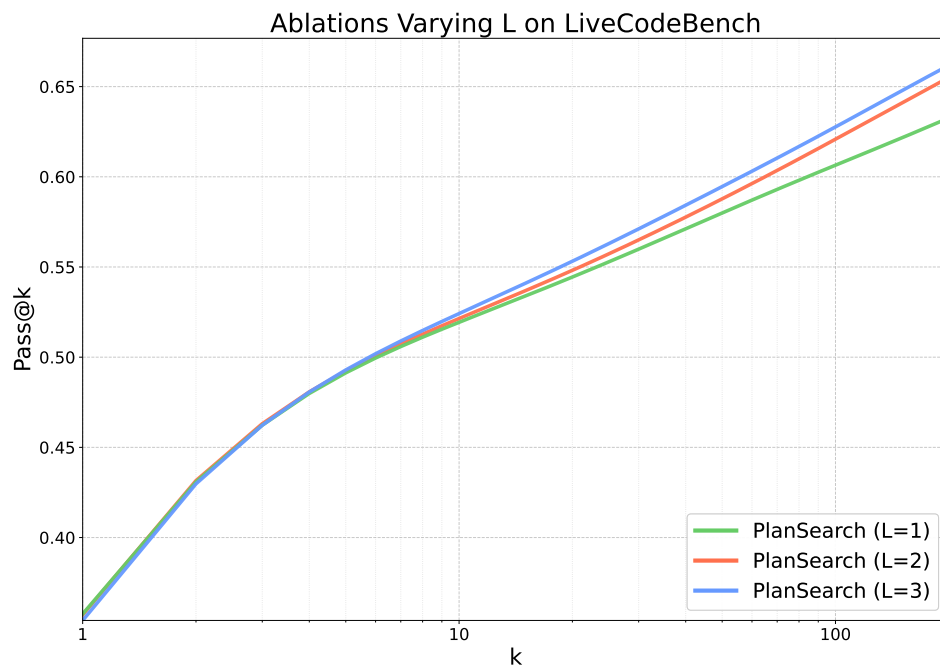


Figure 27: We run ablations of PLANSEARCH with different  $L$ —controlling the maximum order of observation used, i.e., how many layers the observation tree will search. We find pass@k scales with respect to increasing  $L$

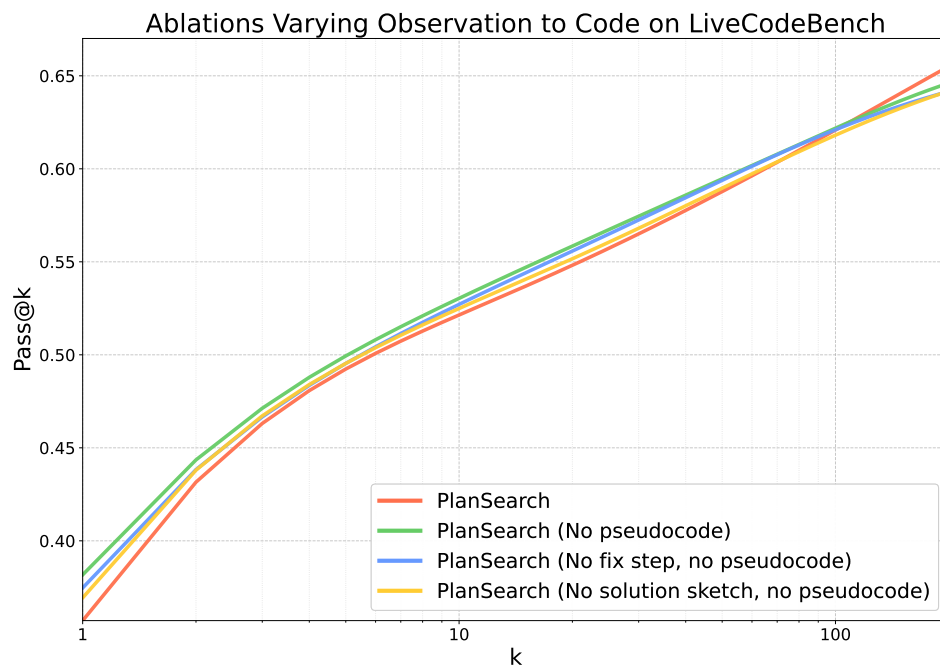


Figure 28: We run ablations of PLANSEARCH with different methods of translating a combination of observations to code.

## L Base Models vs. Instruct Models for Large Samples

We find that base models, despite performing poorly relative to their instruct counterparts for evaluated with pass@1, will frequently match or even exceed performance on pass@k for sufficiently high  $k$ . This is likely due to higher amounts of diversity in base models, which have not undergone post-training designed to elicit a single strong response from the model.

We see this effect across all models for HumanEval+ and MBPP+, but only the DeepSeek-Coder-V2 family for LiveCodeBench.

See Figure 29 for all base and instruct model comparisons between DeepSeek-Coder-V2-Lite, Llama-3.1-8B, and Llama-3.1-70B on all three datasets.

We also provide Llama-3.1-8b and DeepSeek-Coder-V2-Lite pass@k comparisons for  $k$  up to 10,000; see Figures 31, 30.

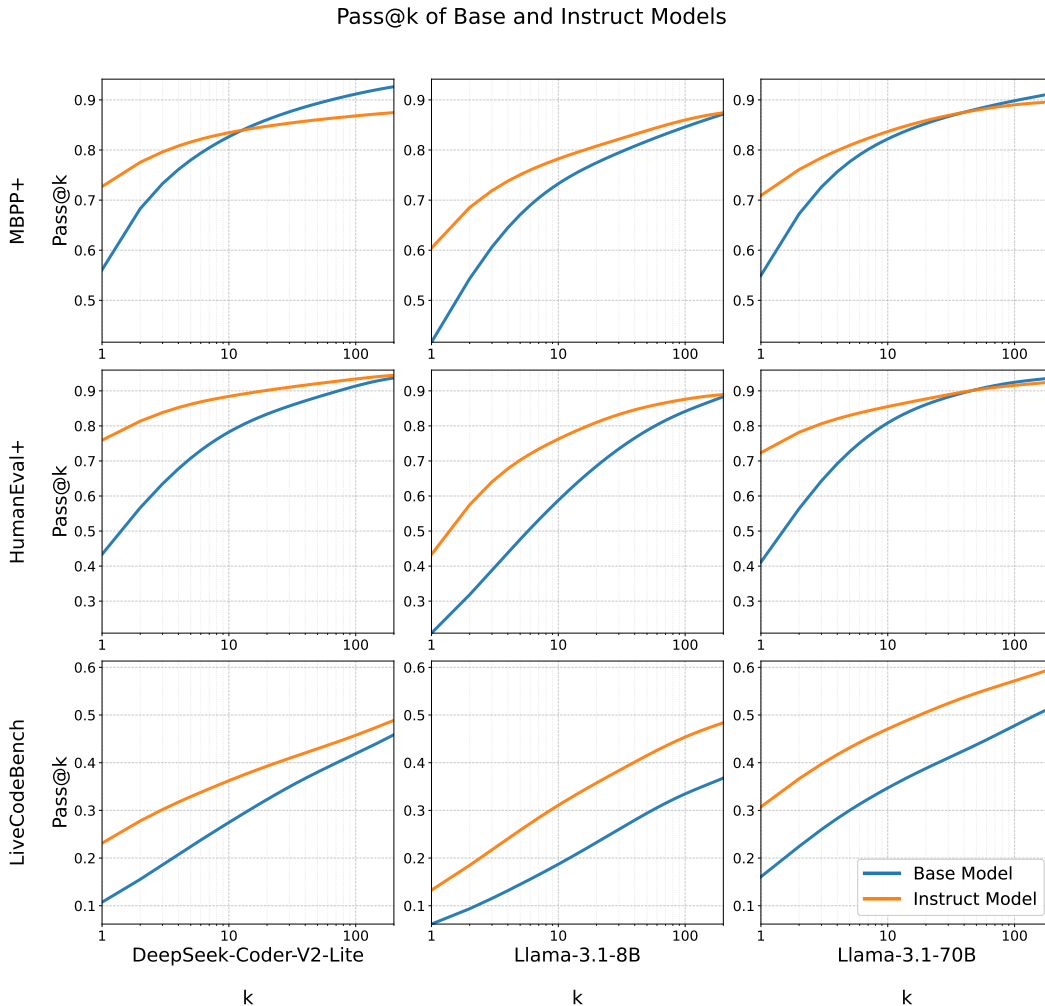


Figure 29: Pass@k curves comparing DeepSeek-Coder-V2-Lite, Llama-3.1-8B, and Llama-3.1-70B base and instruct performance.

## M Base Models vs. Instruct Models with Public Test Filtering

We repeat the graphs from Appendix L, but with public test filtering. We find that base models with public test filtering almost always exceed the pass@1 of their instruct model variants.

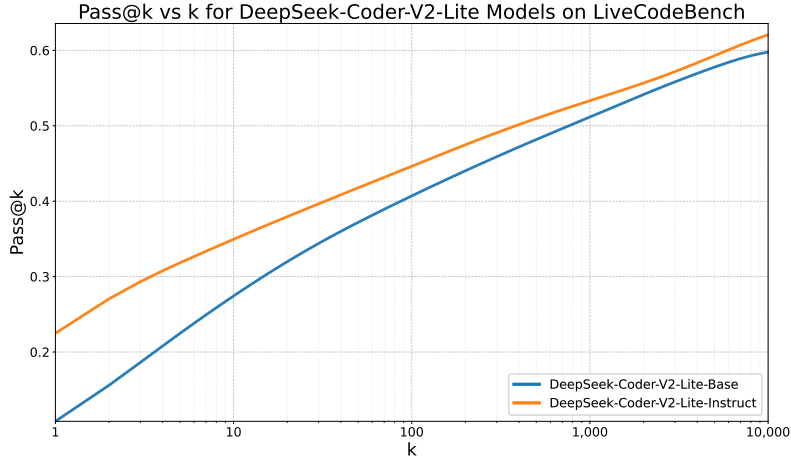


Figure 30: Pass@k curves comparing DeepSeek-Coder-V2-Lite’s base and instruct versions on LiveCodeBench with up to 10,000 completions.

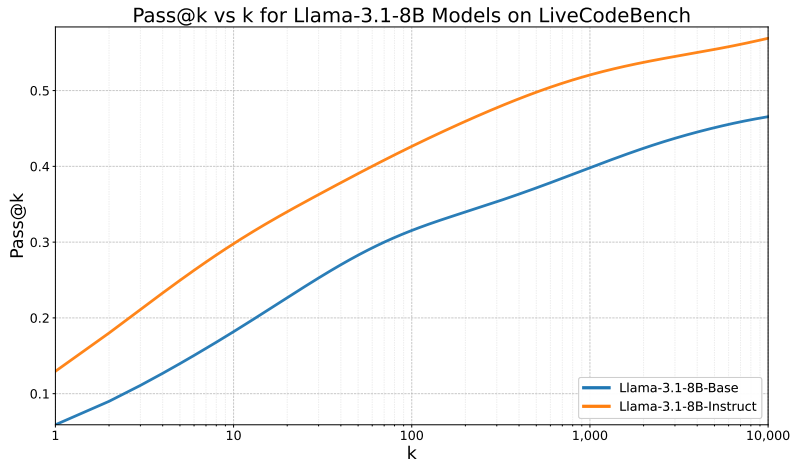


Figure 31: Pass@k curves comparing Llama-3.1-8B’s base and instruct versions on LiveCodeBench with up to 10,000 completions.

See Figure 32 for all base and instruct model comparisons between DeepSeek-Coder-V2-Lite, Llama-3.1-8B, and Llama-3.1-70B on all three datasets with public test filtering.

We also report Llama-3.1-8b and DeepSeek-Coder-V2-Lite pass@k comparisons with public test filtering for  $k$  up to 10,000; see Figures 33, 34.



Pass@k of Base and Instruct Models

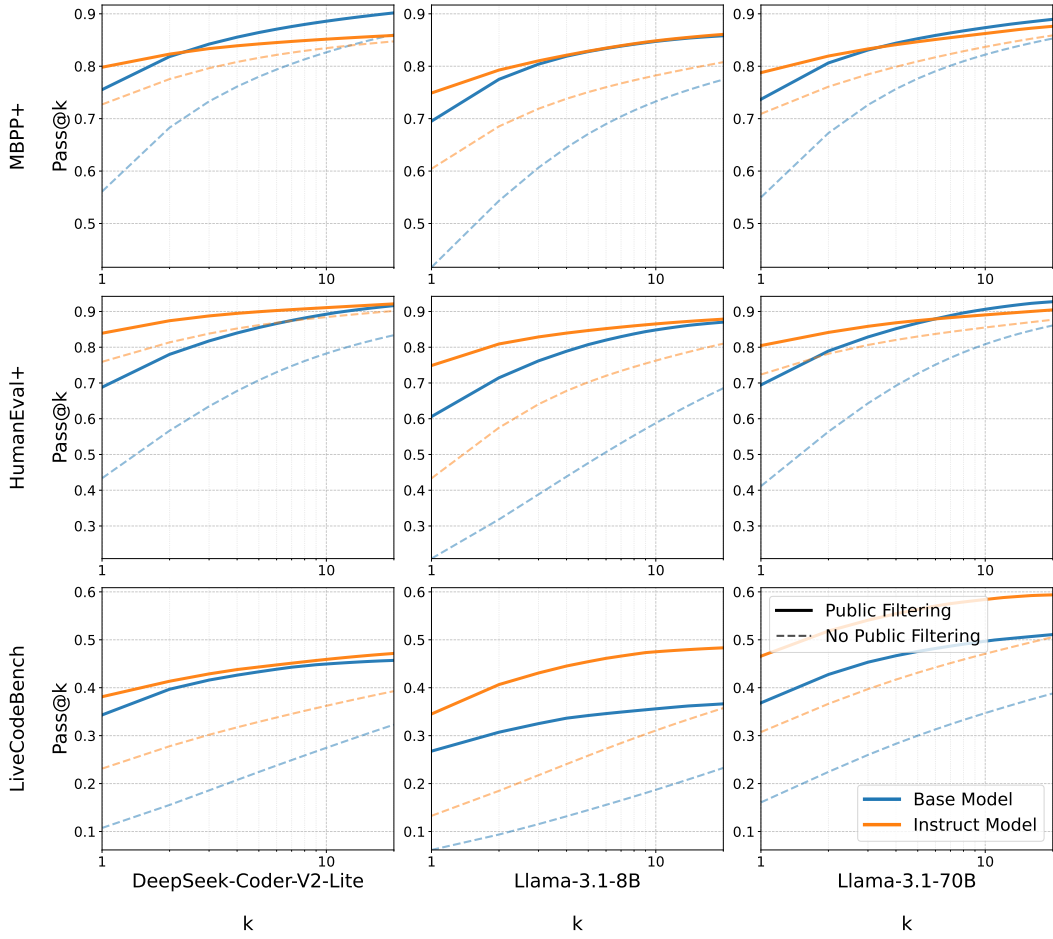


Figure 32: Pass@k curves comparing DeepSeek-Coder-V2-Lite, Llama-3.1-8B, and Llama-3.1-70B base and instruct performance with public test filtering.

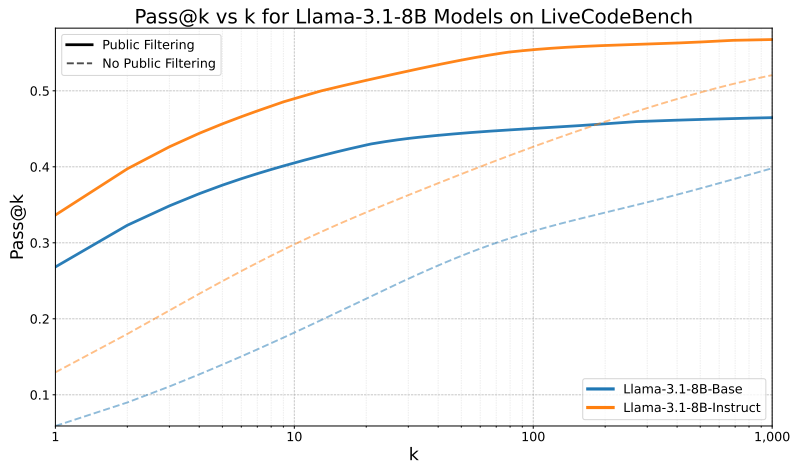


Figure 33: Pass@k curves comparing Llama-3.1-8B's base and instruct versions on LiveCodeBench with up to 1,000 completions and with public test filtering.

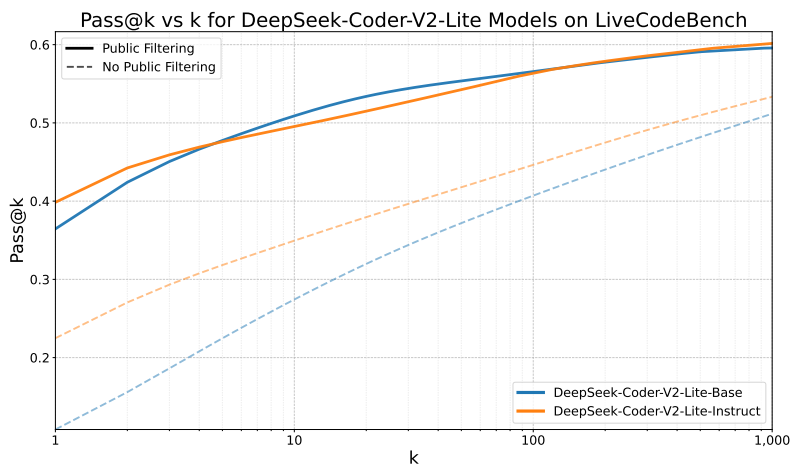


Figure 34: Pass@k curves comparing DeepSeek-Coder-V2-Lite’s base and instruct versions on LiveCodeBench with up to 1,000 completions and with public test filtering.

## N Bar Charts

See Figures 35, 36, 37. These plot pass@1 and pass@200 of select methods between REPEATED SAMPLING and PLANSEARCH, on datasets MBPP+, HumanEval+ [27], and LiveCodeBench [23].

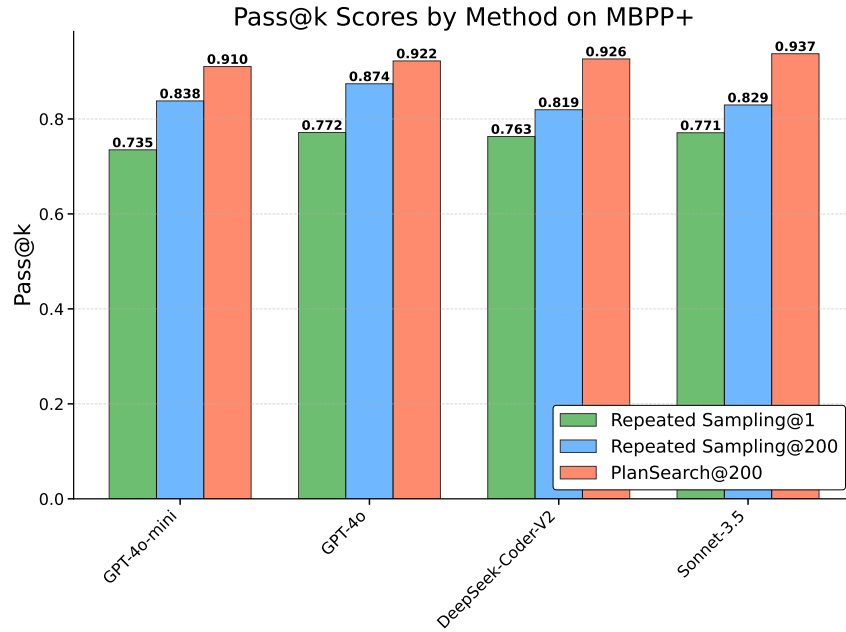


Figure 35: Bar chart with REPEATED SAMPLING@1, REPEATED SAMPLING@200, and PLANSEARCH@200, on MBPP+.

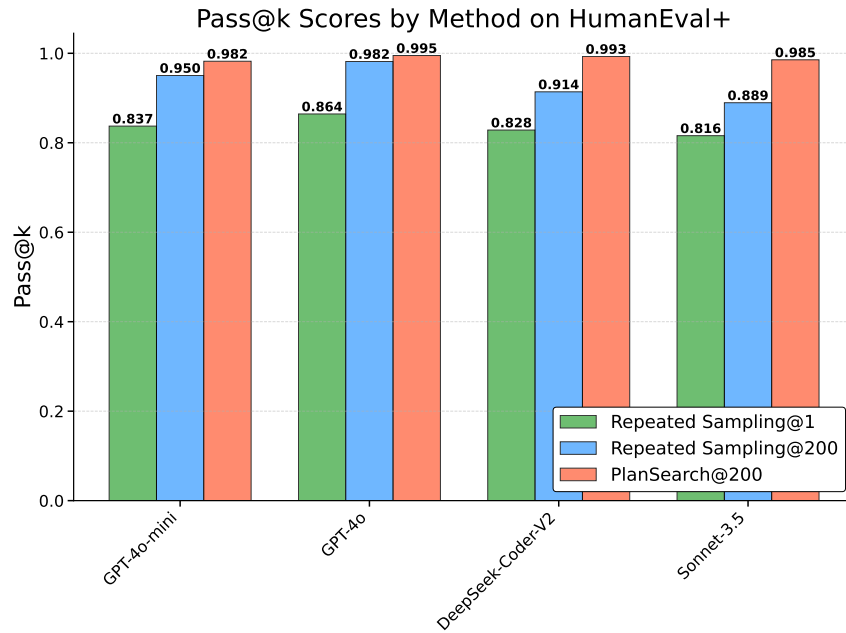


Figure 36: Bar chart with REPEATED SAMPLING@1, REPEATED SAMPLING@200, and PLANSEARCH@200, on HumanEval+.

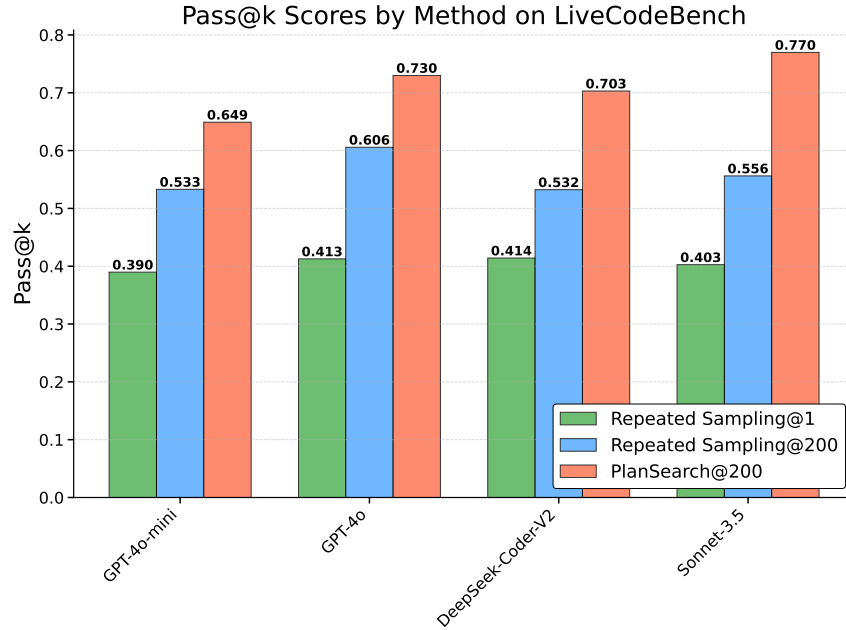


Figure 37: Bar chart with REPEATED SAMPLING@1, REPEATED SAMPLING@200, and PLANSEARCH@200, on LiveCodeBench.

## O Prompts

### O.1 Backtranslation

#### O.1.1 Backtranslate System Prompt

You are an expert Python programmer. You will be given an algorithmic question (problem specification). You will return a high-level, natural language solution to the question, like an editorial. You will NOT return any code. Be as creative as possible, going beyond what you think is intuitively correct.

#### O.1.2 Implement Backtranslation Idea

You are an expert Python programmer. You will be given a question (problem specification) and a natural language solution/tutorial that describes how to solve the problem. You will generate a correct Python program that matches said specification and tutorial and passes all tests. You will NOT return anything except for the program inside markdown codeblocks.

### O.2 Repeated Sampling

You are an expert Python programmer. You will be given a question (problem specification) and will generate a correct Python program that matches the specification and passes all tests. You will NOT return anything except for the program inside Markdown codeblocks.

### O.3 Simple Idea

You will given a competitive programming problem; please output a high-level description of how to solve the problem in natural language. Below are examples:

Example input: PROBLEM DESCRIPTION HERE  
Example output: EXAMPLE OUTPUT HERE  
Here is the competitive programming problem: PROBLEM TO SOLVE

Brainstorm a high-level, natural language solution to the problem above. Note that your intuition may lead you astray, so come up with simple, creative ideas that go beyond what you would usually come up with and go beyond your narrow intuition. Brainstorming solutions that do not seem intuitively correct IS CRUCIAL.

## O.4 PLANSEARCH

### O.4.1 Prompt for Observation Part 1

You are an expert Python programmer. You will be given an competitive programming question (problem specification). You will return several useful, non-obvious, and correct observations about the problem, like hints to solve the problem. You will NOT return any code. Be as creative as possible, going beyond what you think is intuitively correct.

### O.4.2 Prompt for Observation Part 2

You are an expert Python programmer. You will be given an competitive programming question (problem specification) and several correct observations about the problem.

You will brainstorm several new, useful, and correct observations about the problem, derived from the given observations. You will NOT return any code. Be as creative as possible, going beyond what you think is intuitively correct.

### O.4.3 Combining Observations

Here is a sample prompt from the function with placeholders:

Here is the competitive programming problem:

Problem statement placeholder

Here are the intelligent observations to help solve the problem:

Observation 1 placeholder

Observation 2 placeholder

Observation 3 placeholder

Use these observations above to brainstorm a natural language solution to the problem above. Note that your intuition may lead you astray, so come up with simple, creative ideas that go beyond what you would usually come up with and exceeds your narrow intuition.

Quote relevant parts of the observations EXACTLY before each step of the solution. QUOTING IS CRUCIAL.

## P A Model of Repeated Sampling: Pass@k

Consider a simplified model of repeated sampling for code generation. Suppose we have a dataset  $D = \{P_1, \dots, P_l\}$  with  $l$  problems. For some problem  $P_i$ , define the probability  $p_i$  as the probability that our code generation model solves the problem  $P_i$  in one submission. The pass@k [14, 25] metric (for problem  $P_i$ ) is defined as the probability that our code generation model solves the problem  $P_i$  at least once out of  $k$  submissions. Thus, if we know the true  $p_i$  of our model, we may compute our

pass@k simply:

$$\text{pass@k}_i = 1 - (1 - p_i)^k \quad (2)$$

$$\text{pass@k} = \sum_i \text{pass@k}_i / l \quad (3)$$

However, it turns out that for  $k > 1$ , the naïve estimator as seen in Equation 2 is biased, if we sample  $n_i \geq k$  from our code model to solve  $P_i$ ,  $c_i \leq n_i$  are correct, and compute  $p_i = c_i/n_i$  [14]. Instead,  $\text{pass@k}_i$  is typically computed using the unbiased estimator:

$$\text{pass@k}_i = 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \quad (4)$$

Note that reporting  $\text{pass@k}$  on a dataset where  $l = 1$  is rather pointless, since  $\text{pass@k}$  can be derived using only  $\text{pass@1}_1$  and  $n_1$ . Every curve, over a suitable range of  $k$  values, will look like the *S-curve* seen in Figure 38 (as  $k$  is plotted on a log scale).

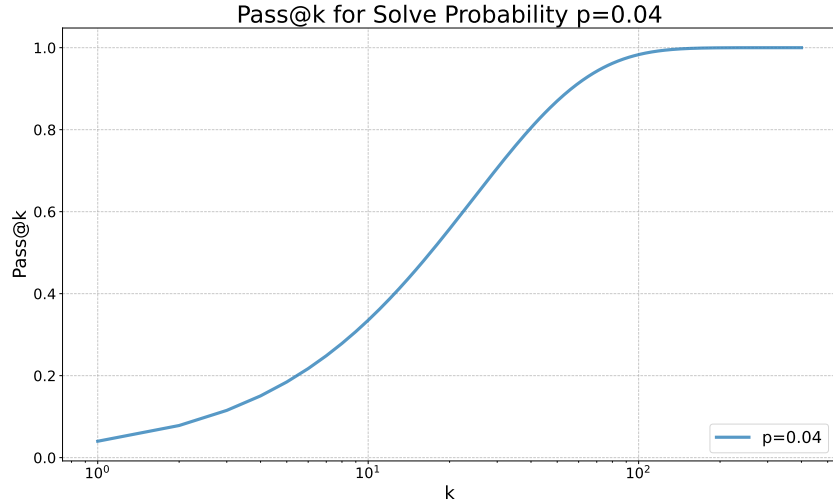


Figure 38: A simple  $\text{pass@k}$  ‘*S-curve*’ plotted with  $1 - (1 - p)^k$ , where  $p = 0.04$ .

However, with datasets where  $l > 1$ , models are able to differentiate themselves through larger  $k$ , since the overall  $\text{pass@k}$  is an average of these  $l$  curves. For example, for  $l = 3$ , it is less optimal to have solved probabilities of Set1 = {0.001, 0.7, 0.9} versus Set2 = {0.05, 0.1, 0.25}, in the regime of roughly  $k = 20$  to  $k = 2,000$  (in which both converge to 1), even though Set1 has a  $\text{pass@1}$  of 53% and Set2 has a  $\text{pass@1}$  of 13%. See Figure 39.

Although not shown in the graph, Set2 converges close to 1 at roughly  $k = 400$ , several orders of magnitude below Set1. In addition, note that the slight notch seen in Set1’s curve at large  $k$  is due to the presence of low, but non-zero solve-rates, which can be seen in empirical  $\text{pass@k}$  curves later on. (These can be thought as the beginning of the ‘ramping-up’ regime of the typical *S-curves* in Figure 38.)

## Q Computing a Standard Error Estimate of the Pass@k Estimator

For some approximation of the overall variance in our estimates of  $\text{pass@k}$ , we compute the theoretical variance of the estimator (assuming each sample solves the problem with an independent probability  $p$  and there are  $n$  samples):

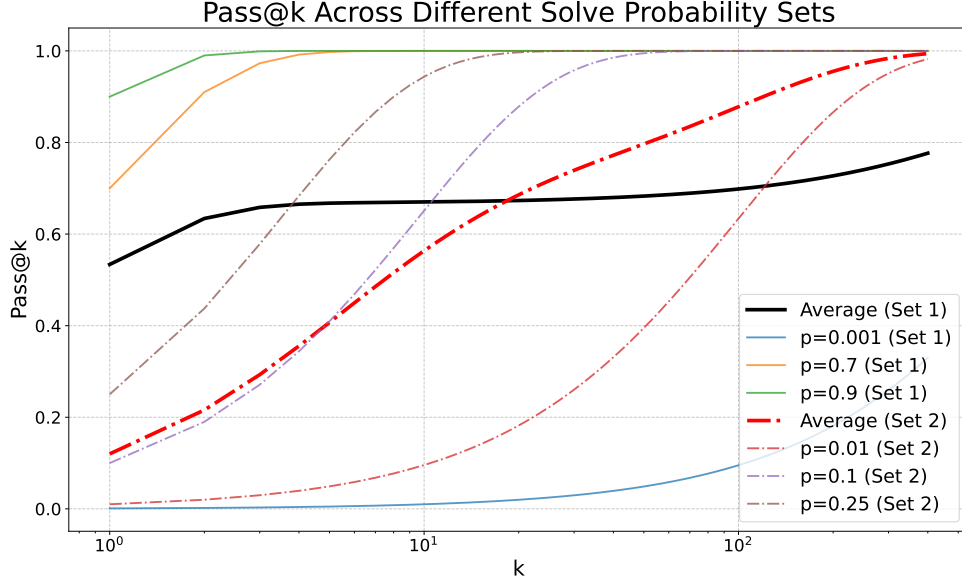


Figure 39: Two pass@k curves on a hypothetical dataset of length  $l = 3$ , and the solve probabilities of Set 1 are  $\{0.001, 0.7, 0.9\}$  and Set 2 are  $\{0.05, 0.1, 0.25\}$ . Note that the pass@1 is 53% and 13%, respectively. However, at roughly  $k = 20$ , Set 2 surpasses Set 1 and within an order of magnitude, achieves pass@k of roughly 1.0.

$$\begin{aligned}
 \text{Var} \left( 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right) &= \mathbb{E} \left[ \left( 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right)^2 \right] - (1 - (1-p)^k)^2 \\
 &= 1 - 2(1 - (1-p)^k) + \mathbb{E} \left[ \frac{\binom{n-c}{k}^2}{\binom{n}{k}^2} \right] - 1 + 2(1-p)^k - (1-p)^{2k} \\
 &= \mathbb{E} \left[ \frac{\binom{n-c}{k}^2}{\binom{n}{k}^2} \right] - (1-p)^{2k} \\
 &= \left( \sum_{c=0}^n \binom{n}{c} p^c (1-p)^{n-c} \right) \cdot \frac{\binom{n-c}{k}^2}{\binom{n}{k}^2} - (1-p)^{2k} \\
 &= \left( \sum_{c=0}^n \frac{\binom{n-k}{c}}{\binom{n}{c}} p^c (1-p)^{n-c} \right) - (1-p)^{2k} \tag{5}
 \end{aligned}$$

Note that Equation 5 depends on  $p$ . We naïvely substitute  $p = c/n$ , which introduces extra variance which is unaccounted for, thus making this estimate an underestimate.

Empirically applying this equation and comparing to known  $p$ , the unaccounted variance increases as  $k$  increases. With  $k$  and  $n$  on the order of 200, simulations suggest this standard error estimate to be roughly a factor of 3 off the theoretical standard deviation, if ground-truth  $p$  of all problems were known.

## R Biased Estimator for Pass@K Due to Non-Independence of PLANSEARCH

From a pure theoretical standpoint, the expression is biased (if using the same interpretation), but it still leads to a similar interpretation—computing the probability that a subset of size  $k$  drawn from the set of samples we already generated contains at least one success. (These given samples were

generated by one run of PLANSEARCH.) As such, in theory, the estimator may be slightly biased in the PLANSEARCH case when computing its true pass@k. In practice, we do not believe this to be a large concern, especially as our primary results feature a relatively large  $k = 200$ .

## R.1 Measuring Diversity

You are an expert Python programmer. You will be given a competitive programming problem and two pieces of code which are attempts to solve the problem. For your convenience, you will also be given the idea for each code, summarized in natural language. You will be asked to answer whether the ideas behind the code are the same. You must ONLY output 'Yes.' or 'No.'

## S Competitive Programming

Competitive programming is a popular subset of programming tasks that involve solving complex algorithmic reasoning. Typically, problems consist of a problem statement (written in natural language)  $P$ , with associated tests:  $(x_i, y_i), i \in \{1, \dots, m\}$ , for which any solution must pass all of them.

The number of tests  $m$  depends on the problem, but typically ranges on the order of 25 to 100. A small subset of the tests are typically given to the solver (we call these public tests) to use as validation that their program passes simple cases. The rest of the tests are hidden. Solutions to the problems must generally pass all the tests to be considered correct. Formally, we let  $f(x)$  denote the output of said code ran on input  $x$ . The solution code is considered correct (passing) if and only if  $f(x_i) = y_i$  for all  $i \in \{1, \dots, m\}$ .

Each dataset consists of many (on the order of low-hundreds) independent problems, and models are evaluated on each of these problems independently.

## T Other Related Work

### T.1 Search in Classical AI

Classical search algorithms like breadth-first search, depth-first search, and A\* search have been widely used for pathfinding, planning, and optimization [34]. More advanced search techniques like Monte Carlo Tree Search (MCTS) have achieved remarkable success in domains like game playing, enabling superhuman performance in Go [39, 40], Poker [7, 8] and Diplomacy [18]. More recently, scaling laws have been found for the performance of AI systems in board games, where ELO improves logarithmically with the amount of compute spent at inference [24].

### T.2 Search with Language Models

Applying search on top of LLMs has been a topic of much interest, especially with an eye towards code generation [14, 26]. Historically, methods such as beam search significantly improved performance for translation systems [19]. Closer to the present day, several recent works have explored repeated sampling [13, 6, 5, 44] as a search method for improving performance. Repeated sampling is a method which directly generates candidate code solutions from the model many times at moderate to high temperatures in hopes that one of the resulting generations will be correct. However, although these works address the roughly linear increase in pass@k with respect to  $\log k$ , they only focus on the most basic version of repeated sampling, without searching in idea space.

When combined with a verifier, reward model, or other filtering algorithm to select the best generation (in cases where pass@k is not a viable metric due to lack of test cases), it is also known under the name of best-of-n sampling [30]. Many works show somewhat good results under intelligent selection of such a filtering algorithm [12, 13]. Recently, several approaches have demonstrated the power of repeated sampling. For example, repeated sampling from a small model can sometimes outperform taking a single sample from a large model on an equalized compute bases [41]. Unlike algorithms such as repeated sampling, which search over the output space, the key insight of PLANSEARCH is



that it is far more effective to instead search plans over the *latent idea space*. By explicitly searching over different natural language plans before generating the code, we significantly increase the diversity of the final code outputs and thus, the resulting pass@k scores for sufficiently large k.

## U Mathematics and Examples of the Diversity Measure

While our choice of a diversity metric is intuitive, one should note that there are a number of intriguing details that result from our definition.

For example, *it is* the case that with  $k$  unique ideas and  $n$  samples of each idea, respectively (for a total of  $kn$  total generated codes), we achieve a diversity score approaching  $(k - 1)/k$ .

For a quick proof, suppose that there are  $k$  cliques, each of  $n$  size. Each clique represents a unique idea. We wish to capture the number of unfilled edges over the number of possible edges as our diversity score:

$$\frac{\binom{k}{2}n^2}{\binom{kn}{2}} = \frac{(k-1)n}{kn-1} \quad (6)$$

which converges to  $\frac{k-1}{k}$  as  $n$  grows large.

Our formulation seems more intuitive than other proposals, such as one that simply counts how many unique ideas  $k$  lie within a pool of size  $n$  to compute  $k/n$  as the diversity score.

For instance, suppose there are only two unique ideas, one clique of which is of size  $2n - 1$ , and the other only output once. The simple proposal would compute both the diversity score of this uneven group *and* that of an even group (where both ideas are output  $n$  times) to be  $\frac{2}{2n} = \frac{1}{n}$ .

However, our score would compute the even group to be  $\frac{n}{2n-1}$ , and the uneven group as:

$$\frac{1}{2n} \cdot 1 + \frac{2n-1}{2n} \cdot \frac{1}{2n-1} = \frac{1}{n} \quad (7)$$

Instead of counting edges, we compute the probability that two randomly selected outputs have similar idea to each other, which is another interpretation of our diversity score.

It seems clear that a case with  $2n - 1$  instances of idea 1 and 1 instance of idea 2 is ‘less diverse’ than a case with  $n$  instances of both idea 1 and idea 2. A naïve proposal may score these two as being the same diversity, whereas our score scores them as  $1/n$  and roughly  $1/2$ , respectively.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and the introduction only claim what is supported by experiments, most notably that idea diversity is strongly correlated with relative gains in inference-time-compute, which is in turn unlocked by searching in idea space.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See Appendix A.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: See Section 4, Appendix D, and Appendix C. In addition, all prompts are provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We do not provide open access to the code, although all data used is open-source. We plan to open-source the code sometime around acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Same justification as the “Experimental Result Reproducibility Section.”

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: See Appendix D, which contains the main pass@k figures used for the critical claim that PLANSEARCHimproves drastically at pass@k. Standard error estimates are somewhat of an underestimate since the estimator is not a naïve mean, as discussed in Appendix Q, and thus is not reported in other graphs. It was also somewhat infeasible to run experiments for much longer for compute concerns.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification: There are no foreseeable points in the Code of Ethics which do not conform with the research conducted. The nature of the research revolves around a somewhat contrived, academic problem of competitive programming, with very little societal impact. However, there is the case in which inference-time compute gains may cause safety concerns in the future, but that is beyond obvious foresight.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: As discussed above, since the work revolves around the setting of competitive programming and inference-time compute gains, there are no foreseeable societal impacts in the near-future.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: There are no released models or data.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.