

Flaming-hot Initiation with Regular Execution Sampling for Large Language Models

Weizhe Chen

University of Southern California
weizhech@usc.edu

Zhicheng Zhang

Carnegie Mellon University
zhichen3@cs.cmu.edu

Guanlin Liu, Renjie Zheng, Wenlei Shi, Chen Dun, Zheng Wu, Xing Jin, Lin Yan
ByteDance Inc.

{guanlin.liu, renjie.zheng, wenlei.shi}@bytedance.com
{chen.dun, zheng.wu1, jinxing.9, neil}@bytedance.com

Abstract

Since the release of ChatGPT, large language models (LLMs) have demonstrated remarkable capabilities across various domains. A key challenge in developing these general capabilities is efficiently sourcing diverse, high-quality data. This becomes especially critical in reasoning-related tasks with sandbox checkers, such as math or code, where the goal is to generate correct solutions to specific problems with higher probability. In this work, we introduce Flaming-hot Initiation with Regular Execution (FIRE) sampling, a simple yet highly effective method to efficiently find good responses. Our empirical findings show that FIRE sampling enhances inference-time generation quality and also benefits training in the alignment stage. Furthermore, we explore how FIRE sampling improves performance by promoting diversity and analyze the impact of employing FIRE at different positions within a response.

1 Introduction

Large language models (LLMs) have achieved remarkable success in a wide range of tasks since the release of ChatGPT (OpenAI, 2022). In addition to traditional natural language processing tasks such as summarization and sentiment analysis, LLMs have demonstrated effectiveness in many new domains, including code generation (Chen et al., 2023; Roziere et al., 2023), human-computer interaction (Li et al., 2023), and math problem-solving (Wei et al., 2022; Yu et al., 2024). Although standalone LLMs have limited reasoning capabilities (Sun et al., 2023; Valmeekam et al., 2023; Chen et al., 2024b), researchers have tried to enhance them by incorporating tool-use and developing integrated systems known as LLM agents (Xi et al., 2023; Wang et al., 2024), which expands the applications of LLMs to more general domains like robot control (Wang et al., 2023a) and autonomous driving (Mao* et al., 2023).

To develop general capabilities, LLMs are typically trained through a three-stage process: pre-training, supervised fine-tuning (SFT), and alignment (Bai et al., 2022; Ouyang et al., 2022). During pretraining, the model learns from a vast array of data gathered from publicly available sources. Then, in the SFT and alignment stages, the model’s abilities are further refined, allowing it to increase reasoning abilities and better follow users’ instructions. In order to enhance reasoning tasks, a sandbox checker — a tool used to verify the correctness of solutions — is often used during training (Liu et al., 2023b). Therefore, one of the key challenges in achieving effective and efficient training is determining how to obtain more successful samples within a fixed number of trials, particularly when addressing complex problems.

In this paper, we introduce Flaming-hot Initiation with Regular Execution (FIRE), a simple yet effective sampling method for training large language models. Inspired by recent findings on attention sink (Xiao et al., 2023), our approach begins by sampling the initial token at a very high temperature and proceeds with the regular sampling process for the remaining sequence. Our algorithm can be viewed as a simplified and more general version of CoT-decoding (Wang and Zhou, 2024), especially with a focus on training in math and coding domains where a sandbox checker is available at a relatively cheap cost.

We first show that our method, at inference time, can improve the pass rate within N trials (pass@ n), also known as the best-of- N (BoN) when only the correctness of the final answer is considered. To demonstrate its effectiveness in training, we show that it can be directly integrated into the reinforcement learning process of large language models. Our approach proves to be effective across multiple open-source models and various LLM capabilities, including mathematical reasoning and coding. We highlight how our method promotes diversity in

generated samples, a key factor linked to performance improvements in pass rate. Importantly, this diversity is maintained even after training with our sampling method, indicating room for further enhancement. We also discuss the effects of simple variations of our method, where the temperature change occurs mid-process rather than at the start, on performance outcomes.

2 Related Works

Researchers have been exploring two primary directions to efficiently improve response quality under a frozen pre-trained LLM. The first direction focuses on prompting techniques such as Chain-of-Thought (Wei et al., 2022) and Tree-of-Thought (Yao et al., 2023a). The second direction involves letting LLMs fix their own mistakes (Wang et al., 2023b; Yao et al., 2023b; Shinn et al., 2023; Madaan et al., 2023; Chen et al., 2024a). In line with these two directions, there has been increasing focus on controlled decoding in LLMs to enhance reasoning capabilities during inference, ranging from search-based approaches applied to policy models (Mudgal et al., 2023; Huang et al., 2024) to utilizing value models trained in the alignment phase (Liu et al., 2023a; Feng et al., 2023).

In this paper, we also focus on inference time; however, our approach extends to the sampling processes used during the training of large language models, as commonly practiced in Instruct-GPT (Ouyang et al., 2022). This process consists of three key stages: pretraining, supervised fine-tuning (SFT), and alignment, also known as reinforcement learning with human feedback (RLHF). For large language models trained in this paradigm, there could be some helpful properties that, without strong theoretical guarantees, are empirically true and thus helpful for LLMs. Our work is related to attention sink (Xiao et al., 2023). An attention sink refers to a token or set of tokens that disproportionately receive attention from other tokens during the attention mechanism within transformer architectures. In their study, they found that one of the most identifiable tokens was shown to be the initial token. While there are no theoretical guarantees, they propose an intuition that initial tokens are visible and used in all later token generations, making them more readily trained to be attention sinks.

Our work is closely related to CoT-decoding (Wang and Zhou, 2024), which uncovers the CoT-paths by enumerating over

the top- k alternative tokens and aggregating the responses by scoring the decoded responses with confidence on the final answer. However, our approach differs in three key aspects: (1) we introduce a differentiable sampling method that can be directly integrated with existing inference and training frameworks, (2) we focus on improving model performance in scenarios with a sandbox checker, where aggregating responses is less data-efficient, and (3) our method operates without assumptions about the prompts, even when a chain of thought (CoT) is included, extending beyond the scope of CoT-decoding. Prior work has explored dynamic temperature allocation during processing (Chang et al., 2023; Zhu et al., 2023; Zhang et al., 2024). However, these approaches restrict the temperature values to the range of $[0, 1]$, which contrasts sharply with the extreme temperature settings employed in our algorithm.

Orthogonal to our work, numerous efforts in the NLP domain have proposed prompt-based algorithms to enhance the diversity of generated responses (Naik et al., 2023; Huang et al., 2023; Wu et al., 2023; Zhou et al., 2024). In contrast, our method focuses on scenarios where the prompt remains fixed, improving generation diversity by controlling the sampling process without relying on domain-specific modifications to the prompt. Notably, our approach is complementary to these works and can be combined with prompt-based methods to further enhance diversity, particularly in cases where generating a larger number of diverse samples for the same question is required.

3 Flaming-hot Initiation Regular Execution

3.1 Method

In this work, we propose a sampling method, Flaming-hot Initiation with Regular Execution (FIRE), inspired by the attention sink phenomenon (Xiao et al., 2023) that demonstrates the importance of initial tokens.

FIRE first samples the initial token at a very high temperature $p \gg 1$, combined with top- k filtering to make the candidate tokens more controllable. At higher temperatures, the candidate tokens are sampled from a probability distribution that approaches uniform sampling. After the initial token is sampled, FIRE proceeds with the decoding stage using a regular temperature setting.

Our approach FIRE is similar to CoT-

Model	Regular		FIRE		
	Pass%	#EA	Pass%	#EA	
GSM8K	DeepSeek	97.57	2.26	98.71	2.76
	Gemma-2	86.81	3.87	87.57	4.01
	Qwen2	95.90	2.58	98.25	3.17
	Qwen2-RL	96.90	2.63	97.90	3.26
MATH	DeepSeek	76.16	5.63	78.16	7.89
	Gemma-2	49.20	9.24	51.48	10.39
	Qwen2	76.60	7.44	79.08	9.03
	Qwen2.5-72B	79.30	2.39	80.40	2.60

Table 1: Inference results for different models on different datasets with best hyperparameters combinations. Specifically, Qwen2-RL is a fine-tuned model trained by ourselves. We show the pass rate (%) with 40 samples, and the effective answers (EA) among the 40 samples.

decoding (Wang and Zhou, 2024) that enumerates the top- k candidates of the initial token. However, while CoT-decoding focuses more on the decoding stage and extracting Chain-of-Thought without prompt, our approach FIRE serves as a general differentiable sampling method, which can be combined with existing sampling frameworks and can be more efficient in the training stage where a sandbox checker that judges whether a specific answer is correct or not is available with a cheap cost.

While FIRE can be applied to any token in the decoding stage, we restrict its application to the initial token to prevent the generation of random tokens that are wrong in the context. For example, if we apply FIRE after the prefix "1+2=", it would sample, in addition to the token "3", other tokens like "4" or "5", which are very likely to be wrong. In contrast, since FIRE is only applied to the initial token, it would unlikely lead to broken sentences or code with syntax errors. In our empirical experiments, we found that the initial token frequently consists of words like "Let's", "Sure", "So", and "The", which do not directly convey any information. But what these initial tokens affect is the reasoning steps afterward, with the same intuition as StreamingLLM (Xiao et al., 2023).

3.2 Experiments

In this section, we evaluate our algorithm, FIRE, by addressing several key research questions that guide our experiments.

How effective is FIRE during inference? We first showcase the effectiveness of FIRE sampling in inference-only scenarios. We tested four open-source models: Qwen2-7B-Instruct (Qwen2) (Yang et al., 2024), Qwen2.5-72B-Instruct (Qwen2.5-

	Regular		FIRE	
	Pass@1	Pass@10	Pass@1	Pass@10
MBPP	61.2	82.8	50.6	86.6
MBPP+	52.7	74.2	44.1	77.0

Table 2: Pass rate (%) with different number of samples from Qwen2-7B-Instruct on MBPP and MBPP+.

p	k	min-p	Regular		FIRE	
			n=10	n=40	n=10	n=40
0.7	16	0.01	66.4	75.8	70.0	78.9
		0	66.4	75.8	70.0	78.9
	32	0.01	66.2	75.3	70.1	78.9
		0	66.2	75.2	70.1	78.9
0.9	16	0.01	66.1	76.6	69.5	78.9
		0	66.2	76.6	69.5	78.9
	32	0.01	66.7	76.4	69.5	79.1
		0	66.8	74.4	69.1	79.0

Table 3: Pass rate (%) for Qwen2-7B-Instruct on MATH dataset with different hyperparameter combinations. p: nucleus sampling parameter, k: top-k sampling parameter, min-p: minimum probability threshold (0 indicates min-p is not used). n=10 and n=40 represent the number of samples for calculating the pass rate.

72B) (Yang et al., 2024), DeepSeek-coder-v2-Instruct (DeepSeek)(Zhu et al., 2024), and Gemma-2-2b-it (Gemma-2)(Team et al., 2024), on a diverse set of datasets, including GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), and MBPP(+) (Austin et al., 2021; Liu et al., 2023c). In GSM8K and MATH, we extend the prompts with phrase "Please reason step-by-step" to ensure CoT reasoning in models' responses, a setting where the original motivation of CoT-decoding becomes less meaningful as CoT paths would naturally occur. For the regular sampling settings, we use a combination of nucleus sampling and top-k sampling. To ensure a fair comparison, we conducted a thorough enumeration over hyperparameters, including p , k , and min-p (Huggingface, 2023). Table 1 and Table 2 present the aggregated results, where the reported numbers represent the best outcomes from the enumeration. We observe that FIRE consistently improves the pass rate compared to regular settings across all models on different benchmarks. To further demonstrate the consistent improvement over different hyperparameters, we provide an example result of Qwen2-7B-Instruct on the MATH dataset in Table 3. Full results for all models and datasets are provided in the appendix. Table 3 reveals that although FIRE may alter the hyperparam-

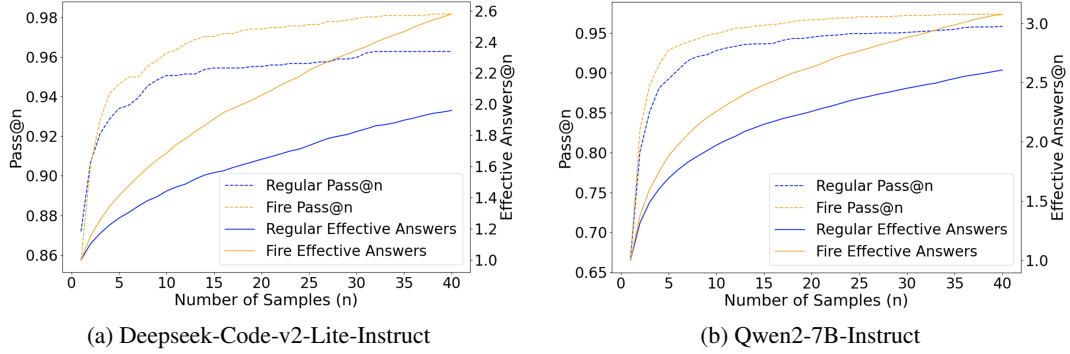


Figure 1: Curves for pass rate and number of effective answers with different numbers of samples on GSM8K.

Dataset	Model	PPO	PPO+FIRE
GSM8K	Deepseek	80.64	82.16
	Qwen2	80.16	82.02
	Gemma	40.39	42.91
	Gemma-2	58.07	61.20
MATH	Qwen2	53.50	55.07

Table 4: Pass@1 on GSM8K and Math for Different models trained with PPO with different sampling.

eter combination that yields optimal performance, it consistently outperforms regular sampling across all hyperparameter combinations.

Why is FIRE effective? FIRE introduces more diversity to the initial token that is generated at a high temperature, and due to the strong attention scores towards initial tokens (Xiao et al., 2023), this diversity benefits the entire subsequent generation. To measure diversity quantitatively, we use the number of unique answers (effective answers) within a set of responses as our metric. We choose not to use some popular metrics like n-grams since we only control the initial token, and in tasks with long reasoning paths, such as math and coding, similar n-grams will likely always appear, making it unsuitable for measuring diversity. As shown in Figure 1, Table 1 (#EA), FIRE demonstrates increased diversity across various models and datasets, which contributes to enhanced pass@n performance. As anticipated, FIRE does not improve Pass@1 performance due to its focus on promoting diversity. However, it consistently delivers improvements when more samples are considered.

Is FIRE helpful when integrated into training? Having established that our method improves pass@n by improving diversity, we directly apply FIRE to boost language model training. To

	1st-line	2nd-line	3rd-line	PRM-line
Regular	66.40	74.36	74.77	75.73
FIRE	69.98	74.96	75.92	78.21

Table 5: Pass@10 Results from Qwen2-7B-Instruct on the training set of MATH dataset for FIRE variants with different sampling points, compared to regular sampling method that does not change the temperature.

test this, we use Proximal Policy Optimization (PPO) (Schulman et al., 2017) to finetune several models using the GSM8K and MATH datasets, and assess their performance through the final pass rate for single samples (Pass@1). As shown in Table 4, integrating FIRE into the training process leads to an improvement in Pass@1. Notably, even though each data point is sampled only once during PPO training following common practice (Ouyang et al., 2022; Sheng et al., 2024), our method still yields improvements. The results also show that the improvements are consistent for different models. Furthermore, after our RL training, the model still exhibits diversity and continues to benefit from inference-time pass rate improvements, as evidenced by Qwen2-RL in Table 1. Consequently, FIRE can be applied iteratively to refine the model, leading to an even bigger improvement margin.

Can FIRE sampling work in mid-sequence?

Finally, we explore the effect of applying FIRE sampling midway through a response. We first construct a dataset that ensures the correctness of the initial sentences, by utilizing a Process Reward Model (PRM) to identify the first sentences at which the response becomes incorrect. We then evaluate the effect of applying FIRE sampling at the beginning of different sentences (1st, 2nd, and 3rd-line) or at the first token deemed incorrect by the PRM ("PRM-line"). We refer the reader to the

Dataset	Model	Sampling	Pass@1	Pass@5	Pass@10	Pass@20
GSM8K	Qwen	Reg	66.95 ± 0.09	89.33 ± 0.06	92.58 ± 0.04	94.48 ± 0.03
		FIRE	65.65 ± 0.12	92.92 ± 0.04	95.58 ± 0.03	97.20 ± 0.02
	Deepseek	Reg	87.02 ± 0.06	93.64 ± 0.03	95.06 ± 0.03	96.02 ± 0.02
		FIRE	85.85 ± 0.07	94.86 ± 0.03	96.76 ± 0.02	97.60 ± 0.02
MATH	Qwen	Reg	35.94 ± 0.05	59.32 ± 0.03	66.04 ± 0.03	71.19 ± 0.02
		FIRE	39.64 ± 0.05	63.34 ± 0.03	69.75 ± 0.03	74.77 ± 0.02
	Deepseek	Reg	51.03 ± 0.04	64.45 ± 0.03	68.58 ± 0.02	71.80 ± 0.01
		FIRE	48.79 ± 0.05	65.39 ± 0.02	70.25 ± 0.02	73.99 ± 0.02

Table 6: Average and standard deviation of pass rate results for different models on different datasets.

appendix for a more detailed description of the construction of this dataset. As shown in Table 5, while FIRE sampling offers benefits throughout different settings, its advantages diminish for tokens beyond the initial ones, despite an overall increase in accuracy due to the prefix guaranteed to be correct.

Is the improvement stable? In previous sections and the appendix, we demonstrated that FIRE consistently enhances performance across various models, datasets, and hyperparameter configurations. Here, we further substantiate these findings by evaluating their statistical significance. Specifically, for each question, we leverage a pool of 40 pre-collected samples and randomly select different subsets of n samples, repeating this process 10 times. The results, summarized in Table 6, reveal minimal standard deviations across the datasets, underscoring the robustness and statistical stability of the improvements.

4 Conclusion

In this paper, we introduced a novel sampling method called Flaming-hot Initiation with Regular Execution (FIRE). Through empirical analysis, we demonstrated that FIRE enhances both inference-time performance and reinforcement learning, particularly when a chain of thought is integrated into the prompt. We showed that FIRE improves generation diversity, and we believe that this diversity contributes to its overall effectiveness. Additionally, we explored several variants of FIRE that modify the sampling process not only immediately after the question but also during the middle of the generation, further showcasing its versatility.

5 Limitations

While this work focuses on improving the efficiency of LLM training through better sampling methods, there are two limitations. First, our approach lacks a strong theoretical guarantee, mean-

ing that there is a possibility that future models, especially ones that are with different model architectures, may not benefit from it. Second, although our method is designed for training LLMs, the inference-time algorithm could potentially bypass safety measures by sampling out-of-distribution data. However, we argue that this concern can be inherently mitigated in models trained with our proposed sampling technique.

References

- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Chung-Ching Chang, David Reitter, Renat Aksitov, and Yun-Hsuan Sung. 2023. K1-divergence guided temperature sampling. *arXiv preprint arXiv:2306.01286*.
- Weizhe Chen, Sven Koenig, and Bistra Dilkina. 2024a. Reprompt: Planning by automatic prompt engineering for large language models agents. *arXiv preprint arXiv:2406.11132*.
- Weizhe Chen, Sven Koenig, and Bistra Dilkina. 2024b. Why solving multi-agent path finding with large language model has not succeeded yet. *arXiv preprint arXiv:2401.03630*.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

- Xidong Feng, Ziyu Wan, Muning Wen, Ying Wen, Weinan Zhang, and Jun Wang. 2023. Alphazero-like tree-search can guide large language model decoding and training. *arXiv preprint arXiv:2309.17179*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *NeurIPS*.
- Baizhou Huang, Shuai Lu, Weizhu Chen, Xiaojun Wan, and Nan Duan. 2023. Enhancing large language models in coding through multi-perspective self-consistency. *arXiv preprint arXiv:2309.17272*.
- James Y Huang, Sailik Sengupta, Daniele Bonadiman, Yi-an Lai, Arshit Gupta, Nikolaos Pappas, Saab Mansour, Katrin Kirchoff, and Dan Roth. 2024. Deal: Decoding-time alignment for large language models. *arXiv preprint arXiv:2402.06147*.
- Huggingface. 2023. New sampling strategy dropped in transformers – min p sampling. <https://huggingface.co/posts/joaogante/319451541682734>. Accessed: 2024-10-15.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. **CAMEL: communicative agents for "mind" exploration of large language model society**. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jiacheng Liu, Andrew Cohen, Ramakanth Pasunuru, Yejin Choi, Hannaneh Hajishirzi, and Asli Celikyilmaz. 2023a. Making ppo even better: Value-guided monte-carlo tree search decoding. *arXiv preprint arXiv:2309.15028*.
- Jiate Liu, Yiqin Zhu, Kaiwen Xiao, Qiang Fu, Xiao Han, Wei Yang, and Deheng Ye. 2023b. Rlhf: Reinforcement learning from unit test feedback. *arXiv preprint arXiv:2307.04349*.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023c. **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*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jiageng Mao*, Junjie Ye*, Yuxi Qian, Marco Pavone, and Yue Wang. 2023. A language agent for autonomous driving. *arXiv preprint arXiv:2311.10813*.
- Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, et al. 2023. Controlled decoding from language models. *arXiv preprint arXiv:2310.17022*.
- Ranjita Naik, Varun Chandrasekaran, Mert Yuksekogonul, Hamid Palangi, and Besmira Nushi. 2023. Diversity of thought improves reasoning abilities of large language models. *arXiv preprint arXiv:2310.07088*.
- OpenAI. 2022. **Introducing chatgpt**.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. **Reflection: language agents with verbal reinforcement learning**. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jiankai Sun, Chuanyang Zheng, Enze Xie, Zhengying Liu, Ruihang Chu, Jianing Qiu, Jiaqi Xu, Mingyu Ding, Hongyang Li, Mengzhe Geng, et al. 2023. A survey of reasoning with foundation models. *arXiv preprint arXiv:2312.11562*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Planbench: An extensible benchmark for evaluating large language models on planning and reasoning about change. *Advances in Neural Information Processing Systems (NeurIPS)*.

- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.
- Xuezhi Wang and Denny Zhou. 2024. Chain-of-thought reasoning without prompting. *arXiv preprint arXiv:2402.10200*.
- Yen-Jen Wang, Bike Zhang, Jianyu Chen, and Koushil Sreenath. 2023a. Prompt a robot to walk with large language models. *arXiv preprint arXiv:2309.09969*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023b. **Self-instruct: Aligning language models with self-generated instructions**. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 13484–13508. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*.
- Ning Wu, Ming Gong, Linjun Shou, Shining Liang, and Daxin Jiang. 2023. Large language models are diverse role-players for summarization evaluation. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 695–707. Springer.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. 2023b. **ReAct: Synergizing reasoning and acting in language models**. In *International Conference on Learning Representations (ICLR)*.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2024. **Meta-math: Bootstrap your own mathematical questions for large language models**. In *International Conference on Learning Representations, (ICLR)*.
- Shimao Zhang, Yu Bao, and Shujian Huang. 2024. Edt: Improving large language models’ generation by entropy-based dynamic temperature sampling. *arXiv preprint arXiv:2403.14541*.
- Yue Zhou, Yada Zhu, Diego Antognini, Yoon Kim, and Yang Zhang. 2024. Paraphrase and solve: Exploring and exploiting the impact of surface form on mathematical reasoning in large language models. *arXiv preprint arXiv:2404.11500*.
- Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y Wu, Yukun Li, Huazuo Gao, Shirong Ma, et al. 2024. Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence. *arXiv preprint arXiv:2406.11931*.
- Yuqi Zhu, Jia Allen Li, Ge Li, YunFei Zhao, Jia Li, Zhi Jin, and Hong Mei. 2023. Improving code generation by dynamic temperature sampling. *arXiv preprint arXiv:2309.02772*.

A Implementation Details

In the paper, we proposed FIRE sampling, which is similar to CoT-decoding, and removed the need to calculate the confidence score. One of the biggest benefits of simplifying the method is getting an extremely easy implementation. For inference, we use vLLM (Kwon et al., 2023) and do a two-stage sampling, with the first stage sampling only one token with high temperature and the second stage continuing the sampling with regular sampling. For training, we implement based on Hybrid-Flow(Sheng et al., 2024), a newly released RLHF code base, which supports sampling with vLLM. Thus, we only changed the sampling part of the code in the RLHF framework. As shared in all experiments, the temperature used for the initial token is set at 30.

In our experiment, we enumerate the parameters of top-p sampling, top-k sampling, and min-p sampling. We list all the parameters we have tried in the next section. Due to computation costs, some of the models are not enumerated in the same number as others. However, our conclusion that FIRE outperforms regular sampling is consistent, as we will show later. Specifically for MBPP(+), and for Qwen-RL, the model after our fine-tuning, we test on a single hyperparameter combination of $top - p = 0.9$, $top - k = 16$, which follows the best configuration from previous trials. For Qwen2.5-72b-Instruct, we follow the

recommended hyperparameters of $top - p = 0.8$, $top - k = 16$. For reinforcement learning problems, we use the default parameters in HybridFlow, specifically, $top - k = 16$, $top - p = 1.0$. For training with FIRE sampling, to enable PPO to accept the relatively out-of-distribution samples, we change the clipping ratio for PPO from 0.2 to 0.5. We observe that for PPO+FIRE to use the original clip rate, it will generally match the original performance, while pure PPO with a higher clip ratio will lead the training to a failure and converge to a pass rate close to 0.

In the paper, we use three different datasets: GSM8K, MATH, and MBPP(+). GSM8K is a dataset with 8.5K total instances of math problem, of which 7.5K is in the training set and 1.3K is in the test set. MATH is a math dataset that is slightly more difficult and more comprehensive than GSM8K, with 7.5K training data and 5K test data. MBPP is a benchmark consisting of around 1,000 crowd-sourced Python programming problems, and MBPP+ is a benchmark that enlarges MBPP with some harder problems, reaching around 35K total test problems. While MBPP+ is still under regular update, we use version 0.1.0 in our paper.

For the final part of the experiment about generation in the middle sequence, we use a dataset that guarantees a certain number of sentences of prefixes to be correct. Here, the sentences are defined based on '.' in the answer. This dataset is generated on the training set of the MATH dataset, for which we first use Qwen2-7b-Instruct to sample 10 responses for each question. Then, for each response, we enumerate the sentences and sample 20 times using different numbers of sentences as the prefix. Thus, we obtained an approximation of the point at which the original samples became wrong. Specifically, if one response is wrong before the number of lines we enumerate in Table 5, we use all the prefix up to the point that is still correct for that response, i.e., if for a specific sample, the correct sentences are less than 2, 3rd-line pass rate will be calculated in the same way as PRM-line.

B Extra Experiment Results

We provide our full inference experiment table in Table 7, Table 8, and Table 9. We observe that among all hyperparameter combinations, FIRE stably outperforms regular sampling, starting at Pass@10 to Pass@20 and Pass@40. In most set-

tings, FIRE is superior to regular sampling at Pass@5, and for certain settings in the MATH dataset, FIRE could even show an advantage in Pass@1.

Dataset	top-p	top-k	min-p	Sampling	Pass@1	Pass@5	Pass@10	Pass@20	Pass@30	Pass@40	EA@40
GSM8K	0.7	16	0.01	Reg	87.19	93.40	95.07	95.53	95.98	96.29	1.96
	0.7	16	0.01	FIRE	85.75	94.62	96.21	97.42	97.95	98.18	2.58
	0.7	16	0.0	Reg	87.19	93.40	95.07	95.53	95.98	96.29	1.96
	0.7	16	0.0	FIRE	85.75	94.62	96.21	97.42	97.95	98.18	2.58
	0.7	32	0.01	Reg	87.72	93.63	94.77	95.45	96.13	96.36	1.97
	0.7	32	0.01	FIRE	85.06	95.07	96.74	97.73	98.18	98.26	2.64
	0.7	32	0.0	Reg	86.58	93.40	94.84	96.06	96.36	96.82	1.97
	0.7	32	0.0	FIRE	85.67	94.84	96.66	97.57	98.10	98.26	2.63
	0.9	16	0.0	Reg	87.41	94.16	95.68	96.66	97.35	97.57	2.26
	0.9	16	0.0	FIRE	84.46	95.83	96.89	97.88	98.26	98.71	2.76
	0.9	32	0.01	Reg	86.05	94.31	95.83	96.74	97.04	97.27	2.28
	0.9	32	0.01	FIRE	84.76	94.84	96.59	97.88	98.33	98.41	2.86
	0.9	32	0.0	Reg	86.43	94.01	95.53	96.66	97.19	97.35	2.29
	0.9	32	0.0	FIRE	84.76	94.84	96.59	97.88	98.33	98.41	2.86
MATH	0.7	16	0.01	Reg	51.04	64.66	69.20	72.36	73.86	74.82	5.08
	0.7	16	0.01	FIRE	49.68	64.94	70.16	73.52	75.58	76.68	6.33
	0.7	16	0.0	Reg	51.04	64.56	68.56	71.84	73.44	74.62	5.07
	0.7	16	0.0	FIRE	49.58	65.48	70.22	74.34	76.02	77.16	6.34
	0.7	32	0.01	Reg	51.00	64.36	68.42	71.74	73.46	74.38	5.06
	0.7	32	0.01	FIRE	49.08	65.64	70.22	73.92	76.10	77.06	6.90
	0.7	32	0.0	Reg	51.00	64.36	68.42	71.74	73.46	74.38	5.06
	0.7	32	0.0	FIRE	49.08	65.64	70.22	73.92	76.10	77.06	6.90
	0.9	16	0.01	Reg	50.42	64.98	69.44	72.98	75.00	76.08	5.66
	0.9	16	0.01	FIRE	48.96	65.34	70.36	74.12	76.16	77.64	7.29
	0.9	16	0.0	Reg	50.82	65.36	69.62	73.12	75.06	76.16	5.64
	0.9	16	0.0	FIRE	48.26	65.00	69.98	74.42	76.18	77.64	7.26
	0.9	32	0.01	Reg	50.00	65.40	69.32	72.88	74.72	75.98	5.65
	0.9	32	0.01	FIRE	47.66	65.48	70.48	74.58	76.86	78.16	7.90
0.9	32	0.0	Reg	50.00	65.40	69.32	72.88	74.72	75.98	5.65	
0.9	32	0.0	FIRE	47.66	65.48	70.48	74.58	76.86	78.16	7.90	

Table 7: Deepseek-coder-v2-Instruct on different datasets with regular sampling (Reg) and FIRE (ours). We show the pass rate with different number of samples (Pass@n), and the effective answers (EA) of the total 40 samples.

Dataset	top-p	top-k	min-p	Sampling	Pass@1	Pass@5	Pass@10	Pass@20	Pass@30	Pass@40	EA@40
MATH	0.7	16	0.01	Reg	15.90	29.94	36.40	42.36	45.74	48.14	8.44
	0.7	16	0.01	FIRE	17.20	32.28	39.22	45.30	48.52	51.18	9.82
	0.7	16	0.0	Reg	15.90	29.94	36.40	42.36	45.74	48.14	8.44
	0.7	16	0.0	FIRE	17.20	32.28	39.22	45.30	48.52	51.18	9.82
	0.7	32	0.01	Reg	15.78	29.84	36.16	41.70	45.20	47.70	8.40
	0.7	32	0.01	FIRE	16.68	32.40	38.80	45.32	48.90	51.26	9.76
	0.7	32	0.0	Reg	15.78	29.84	36.16	41.70	45.20	47.70	8.40
	0.7	32	0.0	FIRE	16.68	32.40	38.80	45.32	48.90	51.26	9.76
	0.9	16	0.01	Reg	14.74	30.46	37.02	43.30	46.98	49.20	9.23
	0.9	16	0.01	FIRE	15.12	31.48	38.30	45.48	48.90	51.48	10.39
	0.9	16	0.0	Reg	14.74	30.46	37.02	43.30	46.98	49.20	9.23
	0.9	16	0.0	FIRE	15.12	31.48	38.30	45.48	48.90	51.48	10.39
	0.9	32	0.01	Reg	14.58	30.16	36.20	42.28	45.98	48.34	9.17
	0.9	32	0.01	FIRE	15.04	31.24	37.60	44.26	47.84	50.54	10.34
0.9	32	0.0	Reg	15.02	30.06	36.48	43.12	46.52	49.08	9.15	
0.9	32	0.0	FIRE	14.58	31.40	38.36	44.92	48.48	51.06	10.35	
GSM8K	0.7	16	0.01	Reg	36.54	66.41	75.66	82.34	84.46	86.81	3.86
	0.7	16	0.01	FIRE	32.45	66.57	76.57	83.32	85.97	87.26	3.97
	0.7	16	0.0	Reg	36.54	66.41	75.66	82.34	84.46	86.81	3.86
	0.7	16	0.0	FIRE	32.45	66.57	76.57	83.32	85.97	87.26	3.97
	0.7	32	0.01	Reg	36.92	67.25	75.66	82.11	84.08	85.52	3.91
	0.7	32	0.01	FIRE	31.24	66.79	76.27	82.87	85.97	87.57	4.01
	0.7	32	0.0	Reg	36.92	67.25	75.66	82.11	84.08	85.52	3.91
	0.7	32	0.0	FIRE	31.24	66.79	76.27	82.87	85.97	87.57	4.01

Table 8: Gemma-2-2b-it on different datasets with regular sampling (Reg) and FIRE (ours). We show the pass rate with different number of samples (Pass@n), and the effective answers (EA) of the total 40 samples.

Dataset	top-p	top-k	min-p	Sampling	Pass@1	Pass@5	Pass@10	Pass@20	Pass@30	Pass@40	EA@40
GSM8K	0.7	16	0.01	Reg	66.72	89.23	92.80	94.47	95.07	95.83	2.61
	0.7	16	0.01	FIRE	66.49	92.87	94.92	96.66	97.19	97.35	3.08
	0.7	16	0.0	Reg	66.72	89.23	92.80	94.47	95.07	95.83	2.61
	0.7	16	0.0	FIRE	66.49	92.87	94.92	96.66	97.19	97.35	3.08
	0.7	32	0.0	Reg	67.02	89.16	92.27	94.31	95.30	95.91	2.58
	0.7	32	0.0	FIRE	66.34	92.95	95.75	97.19	97.88	98.26	3.17
	0.9	16	0.0	Reg	64.52	90.83	94.16	95.75	96.89	97.42	2.96
	0.9	16	0.0	FIRE	64.22	92.27	95.07	97.04	97.65	97.95	3.33
MATH	0.7	16	0.01	Reg	35.80	59.40	66.40	71.68	74.22	75.76	6.47
	0.7	16	0.01	FIRE	40.26	63.74	69.98	75.08	77.42	78.90	7.86
	0.7	16	0.0	Reg	35.80	59.40	66.40	71.68	74.22	75.76	6.47
	0.7	16	0.0	FIRE	40.26	63.74	69.98	75.08	77.42	78.90	7.86
	0.7	32	0.01	Reg	36.42	59.42	66.22	71.10	73.52	75.26	6.47
	0.7	32	0.01	FIRE	39.52	63.54	70.10	75.06	77.42	78.92	8.11
	0.7	32	0.0	Reg	36.42	59.42	66.22	71.10	73.52	75.26	6.47
	0.7	32	0.0	FIRE	39.52	63.54	70.10	75.06	77.42	78.92	8.11
	0.9	16	0.01	Reg	35.30	59.48	66.16	71.86	74.68	76.60	7.44
	0.9	16	0.01	FIRE	38.70	62.44	69.50	74.64	77.36	78.86	8.76
	0.9	16	0.0	Reg	35.30	59.48	66.16	71.86	74.68	76.60	7.44
	0.9	16	0.0	FIRE	38.70	62.44	69.50	74.64	77.36	78.86	8.76
	0.9	32	0.01	Reg	35.82	59.84	66.70	72.02	74.64	76.38	7.40
	0.9	32	0.01	FIRE	37.44	62.72	69.50	75.08	77.50	79.08	9.03
	0.9	32	0.0	Reg	35.14	59.84	66.80	72.34	74.72	76.40	7.43
	0.9	32	0.0	FIRE	36.70	62.52	69.12	74.54	77.10	79.04	9.04

Table 9: Qwen2-7B-Instruct on different datasets with regular sampling (Reg) and FIRE (ours). We show the pass rate with different number of samples (Pass@n), and the effective answers (EA) of the total 40 samples.