

Tag-Evol: Achieving Efficient Instruction Evolving via Tag Injection

Yixuan Wang[†], Shiqi Zhou[†], Chuanzhe Guo, Qingfu Zhu*
Research Center for Social Computing and Interactive Robotics,
Harbin Institute of Technology, China
{wyx, sqzhou, czguo, qfzhu}@ir.hit.edu.cn

Abstract

Evol-Instruct has made significant improvements as a data synthesis method in several areas. Existing methods typically rely on a fixed set of strategies to evolve, which require manual design and are monolithic in form. In addition, iterative evolution also makes the acquisition of hard samples expensive. In view of this, we propose the Tag-Evol framework, a more diverse and efficient instruction evolving method. Specifically, Tag-Evol uses diverse and specific knowledge tags as strategies to achieve controlled evolution by injecting different combinations of tags into the original instructions. Experiments with multiple backbones in diverse domain benchmarks show that the proposed method generates significantly better evolved data than other methods. Furthermore, we conduct a thorough analysis of the evolved data, demonstrating that Tag-Evol is not only efficient but also generates more diverse and challenging data.

1 Introduction

As a key stage in the training of large language models (LLMs), supervised fine-tuning (SFT) greatly improves the performance of the models on various natural language processing tasks (Ouyang et al., 2022; Wang et al., 2023). To achieve better alignment results in the SFT phase, many high-quality instruction datasets (e.g., ShareGPT (Chiang et al., 2023), LIMA (Zhou et al., 2024)) have been collected and proposed. However, human-annotated high-quality datasets are usually expensive and limited, which cannot support larger scale SFT of the model.

In light of this, recent researches (Tan et al., 2024; Dubey et al., 2024; Yang et al., 2024) has focused on how to synthesize instruction data for training using LLMs themselves, which is relatively cheap and scalable. Evol-Instruct (Xu et al.,

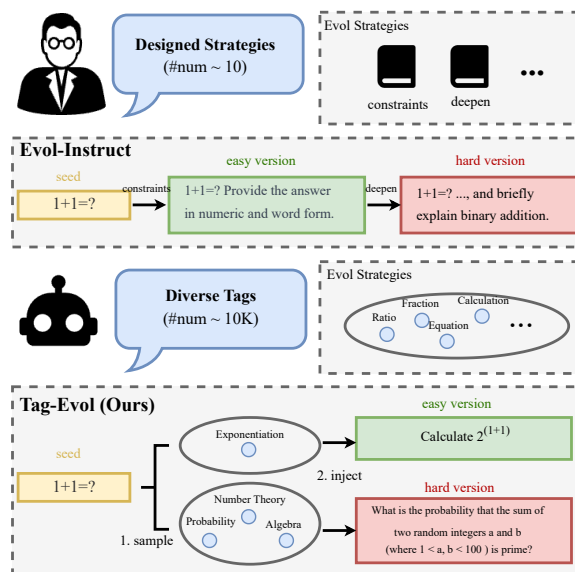


Figure 1: An illustration of the differences between the Evol-Instruct and Tag-Evol. Tag-Evol uses diverse tags as evolution strategies and can generate difficult sample at once by controlling the number of tags.

2023) achieves high-quality instruction synthesis by utilizing rewriting ability of LLMs. This method iteratively improves the difficulty and diversity of the samples from an initial seed dataset through heuristic prompts, which are ultimately merged as the instruction dataset. Benefiting from well-designed evolution prompts and the power of LLMs, Evol-Instruct has achieved a wide range of applications in several domains (Luo et al., 2023a,b). Besides, Zeng et al. (2024) leverage LLMs to automatically optimize the evolution prompts, which further improves the quality of the evolution process.

While Evol-Instruct has shown excellent results on a variety of tasks and models, there are still some issues that need to be addressed: (1) **Fixed Evolution Strategies**. Existing methods obtain a fixed set of evolution strategies through heuristic human annotation and machine optimization, which is costly

[†]Equal contribution.

*Corresponding author.

and hardly adaptable to other domains. Simultaneously, fixed strategies also limit the diversity of evolved data due to preference bias of the model itself. (2) **Inefficient Synthesis Mode.** In order to obtain the hard version of instructions, Evol-Instruct requires multiple iterations of evolution, which not only take extra time, but also introduce cumulative errors due to hallucinations. An ideal instruction evolving algorithm should have diverse and specific evolution strategies and be able to efficiently generate synthetic data of varying difficulty.

To address the above challenges, we propose **Tag-Evol** framework, a more efficient approach for evolved instruction data synthesis through knowledge tag information injection. Inspired by Lu et al. (2023), who evaluates the difficulty of instructions based on the number of knowledge tags typed by LLMs, we redesign the instruction evolving method from the perspective of knowledge tags. These tags represent the attributes contained in the instruction. As shown in Figure 1, we use diverse and specific knowledge tags as enriched evolution strategies, and use different combinations and numbers of knowledge tags to explicitly control the degree of evolution. In this way we can avoid the fixed strategy by combining tags and control the number of injected tags to directly generate hard samples without iteration.

Specifically, our Tag-Evol framework can be divided into two parts, tag pool construction and tag sampling evolution. Firstly, we obtain a domain-specific tag pool by tagging and merging the samples in the seed dataset. Unlike tagging methods used for evaluation, we want to acquire tags that are as diverse and specific as possible, and for this reason we propose a multi-step fine-grained knowledge tagging approach (see details in §3.1) that can significantly improve the quantity and quality of the tag pool. After obtaining the tag pool, we guide the synthesis of the instruction data by sampling different numbers of knowledge tags as the evolution strategy (see details in §3.2). By prompting the LLMs to inject the knowledge of different tag combinations into the seed dataset, we are able to achieve more accurate instruction evolution.

We have conducted experiments using several mainstream LLMs as backbones on multiple tasks including multi-turn conversation, instruction following, mathematical reasoning, and code generation. For a fair comparison, we replicate the Evol-Instruct (Xu et al., 2023) and Auto Evol-Instruct (Zeng et al., 2024) as our baseline methods un-

der the same open source setup. The results show that the proposed Tag-Evol significantly outperforms the baseline model under different backbones, achieving better downstream task performance on multiple domains. In addition, we have also conducted exhaustive analytical experiments on the proposed method in terms of the size of the required evolution model, the number of iterative convergence, demonstrating the great potential of the Tag-Evol method in the field of synthetic data. Our main contributions can be summarised as follows:

- We propose the Tag-Evol framework¹, which achieves efficient instruction evolving by injecting different combinations of knowledge tags.
- We propose a multi-step fine-grained knowledge tagging approach that can obtain more diverse and specific tags for evolution.
- Experiments show that the proposed method outperforms existing Evol-Instruct methods on diverse tasks under multiple base models.

2 Related Work

Synthesis with LLMs. Compared to human annotated datasets, synthetic data has a wide range of applications (Tan et al., 2024; Dubey et al., 2024; Yang et al., 2024; Abdin et al., 2024) in synthetic instruction datasets due to its lower cost. Wang et al. (2022) first propose the Self-Instruct method, which utilizes the LLM itself to efficiently synthesize instruction data for alignment. Subsequently, synthetic data (e.g., Alpaca (Taori et al., 2023), CodeAlpaca (Chaudhary, 2023), etc.) based on Self-Instruct targeting various domains are released for training in the SFT phase. The Evol-Instruct proposed by the Wizard series (Xu et al., 2023; Luo et al., 2023a,b) dramatically increases the upper bound of the available instruction data through well-designed evolution strategies. To address the high cost of manually designed strategies, Zeng et al. (2024) propose an automated strategy optimization method to achieve a wider range of Evol-Instruct.

While most synthetic data methods are implemented using closed-source models (e.g., ChatGPT (Achiam et al., 2023)), there has been some recent works focus on how to use open-source models for

¹<https://github.com/fghccv/TagEvol>

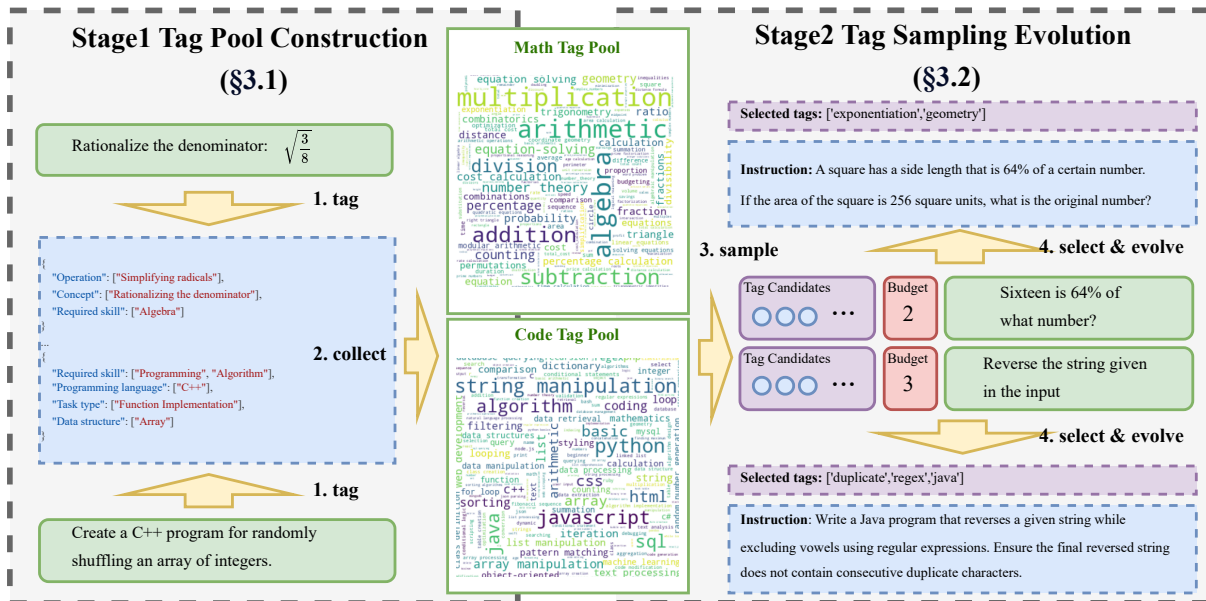


Figure 2: An illustration of the proposed Tag-Evol framework. First, we utilize LLMs to tag seed datasets to construct diverse and fine-grained tag pools. In the evolution phase, we sample a batch of candidate tags and let the model select the appropriate tags to inject into the original samples according to the budget, achieving an increase in difficulty.

data synthesis. Xu et al. (2024) exploit the design of open-source model templates and utilize special tokens to distill high quality instruction data from the SFT phase of the model. Ding et al. (2024) further propose the question fine-tuning (QFT) task, which takes instruction data synthesis as a task for the model to learn. It enables the fine-tuned QFT model to extract more high-quality instruction data. Hui et al. (2024) compare the performance of different scale models in the Evol-Instruct method and demonstrate that small models are more suitable for generating diverse evolved data.

Evaluation with LLMs. In addition to the generation of synthetic data, how to automatically evaluate and filter synthetic data using LLMs is also an important issue. Zheng et al. (2024) propose the MT-Bench dataset to evaluate the performance of LLM-based chat assistants by scoring them with a powerful LLM judge. Liu et al. (2023b) effectively evaluate and filter the SFT phase training data in terms of three dimensions: complexity, quality, and diversity. Unlike directly scoring, Lu et al. (2023) propose the InsTag method to assess the difficulty and diversity of instruction data from the perspective of knowledge tags. By tagging each sample using LLMs, InsTag can estimate the difficulty of different samples by the number of tags and the diversity of the dataset by the total number of tags.

InsTag proposes a systematic approach to data evaluation, but existing synthetic data approaches do not take full advantage of the benefits of tags in knowledge representation. In this paper, we propose the Tag-Evol framework for efficient and high-quality instruction evolution using the content and number of tags as evolution strategies.

3 Tag-Evol Framework

As shown in Figure 2, our Tag-Evol framework can be divided into two parts as a whole: tag pool construction (§3.1) and tag sampling evolution (§3.2). In the tag pool construction phase, our goal is to extract tags containing extensive knowledge from the seed dataset as strategies for subsequent evolution. After obtaining the tags, we implement the evolution of instructions with various levels of difficulty by randomly sampling a certain number of tags and injecting them into the original dataset utilizing the rewriting ability of LLMs.

3.1 Tag Pool Construction

InsTag (Lu et al., 2023) implement a straightforward tagging method that aims to obtain coarse-grained tags, which can be used to estimate the difficulty and diversity of the dataset. However, for evolution strategies, coarse-grained tags do not provide rich and specific guidance goals. To this end, we propose a multi-step fine-grained tagging

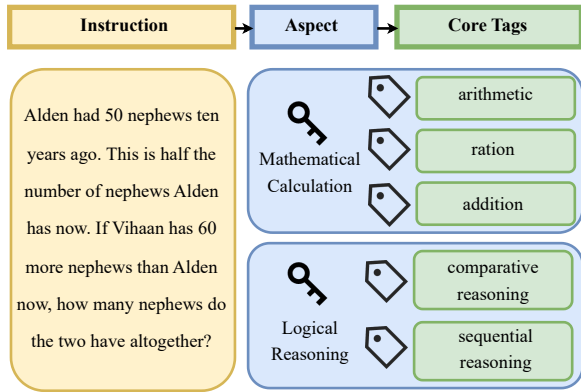


Figure 3: An illustration of the proposed multi-step fine-grained tagging method.

method that is capable of yielding a larger number of more specific tags for subsequent use.

Specifically, we specify the structure of the original tags from lists to dictionaries, allowing the model to generate more comprehensive knowledge tags from keys to values step by step. As shown in Figure 3, the whole process is divided into two parts: (1) **Aspect Generation**. For a more comprehensive tagging, we first prompt the model to summarize in abstract terms, describing the macroscopic characteristics of the sample. For example, the type of task, the skills required, the type of arithmetic, and so on. These abstract aspects ensure the diversity of evolution strategies, with the ability to change the task categories of the samples, add new constraints. (2) **Tag Generation**. Subsequently, we ask the model to generate concrete knowledge tags based on the abstract aspects, which will be sampled in combination as final evolution strategies. It is these specific tags that distinguish Tag-Evol from previous Evol-Instruct method. Tag-Evol can give more specific details of constraints (e.g., exponential correlation) than a generalized strategy of adding constraints, thus making the goal of evolution more clear and the content of evolution more diverse.

The model generates both content together in chain of thought format (see Appendix B for prompt details). The analytical experiment in Section 5.1 demonstrate the advantages of our multi-step tagging over single-step, significantly improving the performance of Tag-Evol.

3.2 Tag Sampling Evolution

The motivation for tag sampling evolution is to inject samples with the proper tag combination so

that they can contain more tags and thus evolve in difficulty. Specifically, after obtaining the tag pool \mathcal{P} , for each sample x , we assign it a difficulty budget b indicating the number of tags to be injected into it.

Budget controlled injection. In order to minimize the insertion of irrelevant tags that lead to hallucinations, we choose to provide the evolution model with a batch collection of tags, so that it can choose the appropriate subset for evolution according to the budget. For the evolution model M_θ , the entire tag injection process can be formulated as:

$$\hat{x}, t = M_\theta(x, b, cand) \quad (1)$$

where x represents the original data, b is the difficulty budget, and $cand$ is a subset sampled from \mathcal{P} . For the output of the model, \hat{x} represents the evolved instruction, and t is the tag combination chosen by M_θ with $t \subset cand$ and $|t| = b$.

Similar to Auto Evol-Instruct (Zeng et al., 2024), we use designed prompt to make M_θ complete the complex operations in Equation 1. The whole tag injection process can be divided into four steps (see Appendix B for details of prompt): (1) Select proper tag combination from candidate tag batch based on the instruction and budget. (2) Generate an injection plan based on original instruction and selected tags. (3) Execute the plan to complete the generation of evolution instruction. (4) Rewrite the instruction to remove the hallucinations. The proposed process allows the model to choose a more appropriate combination of tags based on different instructions. Together with the budget hyper-parameter, it can provide a unique evolution strategy for each sample. In addition, we have conducted experiments for the scale of the evolution model M_θ in Section 5.3. Due to the provision of specific evolution strategies in the form of tags, we find that even 7B level models can execute the Tag-Evol well.

Evolution in multiple rounds. Although Tag-Evol can directly generate samples of varying difficulty by choosing different budgets, we choose to perform multiple rounds of evolution on the seed dataset in order to make a fair comparison with the Evol-Instruct method. Considering that the Evol-Instruct method usually performs three rounds of evolution and merge them as the final dataset, we also used three different budgets to generate three rounds of synthesized data with different levels of difficulty.

Domain	Source	Num#
General	Databricks-Dolly	15,000
Math	GSM8K, MATH	7,473 + 7,500
Code	Code Alpaca	20,022

Table 1: Statistical information on seed datasets used.

4 Experiment

4.1 Experiment Setup

Dataset setup. To evaluate the proposed Tag-Evol framework effectively, we choose three commonly used domains for validation, general-purpose task, mathematical reasoning and code generation, all of which require the model to have strong reasoning capabilities. The multi-turn conversation dataset MT-Bench (Zheng et al., 2023) and the instruction-following evaluation benchmark IFEval (Zhou et al., 2023) are selected to represent general-purpose tasks. For the math domain, we use the GSM8K (Cobbe et al., 2021) and the MATH-500 (Lightman et al., 2023) as benchmarks, while for the code domain, we use the HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021) benchmarks. We construct seed datasets across three domains: for general-purpose tasks, we use the Databricks-Dolly dataset (Conover et al., 2023); for the math domain, the training sets of GSM8K and MATH are adopted; and for the code domain, we follow prior work by incorporating CodeAlpaca (Chaudhary, 2023). Detailed statistical information is shown in Table 1.

Model details. Existing datasets vary in their seed sources and the evolution models used, making comparisons difficult. To fairly evaluate existing Evol-Instruct methods, we construct detailed experiments using Qwen2.5-72B-Instruct (Yang et al., 2024) as an evolution model. The model is also used to generate all the responses of the instruction data. In addition, we choose to train on three different pre-trained models Mistral-7b-v0.1 (Jiang et al., 2023), Llama3-8B (Dubey et al., 2024) and Qwen2.5-7B (Yang et al., 2024) which represent backbones with different abilities. See Appendix A for specific training setups and hyper-parameter details.

Baseline methods. In this paper, we choose two representative Evol-Instruct methods as baselines: (1) **Evol-Instruct** (Xu et al., 2023) manually design a number of evolution prompt as strategies,

and complete each round of evolution by randomly selecting different strategies. Specifically, we follow Hui et al.’s (2024) setup² and use four depth instructions and one width instruction as evolutionary strategies. (2) **Auto Evol-Instruct** (Zeng et al., 2024) leverage LLMs to reflect on generated evolutionary trajectories in order to optimize the original evolution strategy automatically, enabling it to generate higher quality data. In our experiments, we directly choose the optimal prompt reported in Auto Evol-Instruct as the evolution strategy, which has been automatically optimized 12 times.

In our experiments on downstream tasks, we adapt the setup from Zeng et al. (2024) and conduct three rounds of evolution merging to establish the baseline dataset. For the proposed Tag-Evol, we choose budgets of 1, 3, and 5 tags for math and 3, 5, and 7 tags for code to generate the same number of datasets for comparison, respectively.

4.2 Evaluation Results

Main results. After constructing the evolved dataset, we conduct experiments on three different base models as backbones to evaluate their performance across three domains: general-purpose, mathematical reasoning, and code generation. The results are shown in Table 2, Tag-Evol significantly outperforms the existing Evol-Instruct method at all backbone settings, delivering an improvement of 2-3 points on average. For some models released early such as Mistral, whose overall performance is relatively low, the diverse and complex evolved data generated by Tag-Evol can well enhance its performance. For more advanced models like Qwen2.5, the power of its own capabilities results in a decrease for improvements of all Evol-Instruct methods. But this case Tag-Evol still gives a boost of about 1.5 points compared to other evolutionary methods.

In particular, we also observe a trend in the performance difference between the evolved data and the seed data under different backbone models. In both the Mistral and Llama3 settings, the evolutionary approach shows a 8-9 points improvement compared to the seed dataset, suggesting that by further increasing the difficulty of the knowledge involved in the seed dataset can help to improve the model’s capabilities. When it comes to the Qwen2.5 setting, however, the performance of the existing evolved dataset is very close to that of the seed dataset, de-

²<https://github.com/HypherX/Evolution-Analysis/>

Method	General-purpose			Math Reasoning		Code Generation		Average
	MTBench	IF-L	IF-S	GSM	MATH	HEval	MBPP	
<i>Mistral-7B-base</i>								
Seed	5.4	31.6	27.3	58.1	25.0	39.6	50.5	33.9
Evol-Ins	5.2	36.2	32.9	63.3	33.4	53.0	56.3	40.0
Auto Evol-Ins	5.6	40.6	34.3	65.3	35.4	53.7	55.0	41.4
Tag-Evol (Ours)	5.7	43.8	39.5	67.6	35.4	55.5	58.2	43.7
<i>Llama3-8B-base</i>								
Seed	6.0	40.4	35.4	68.4	36.6	45.7	59.0	41.6
Evol-Ins	6.2	46.0	40.4	70.6	42.3	59.8	66.4	47.4
Auto Evol-Ins	6.6	48.9	44.7	70.7	41.0	58.5	65.3	48.0
Tag-Evol (Ours)	6.8	51.2	47.1	72.6	46.2	62.8	66.4	50.4
<i>Qwen2.5-7B-base</i>								
Seed	6.7	48.4	44.1	87.4	68.2	78.0	79.9	59.0
Evol-Ins	6.8	48.4	44.5	87.1	69.0	78.7	78.8	59.0
Auto Evol-Ins	6.9	52.3	46.9	88.4	69.2	78.7	80.2	60.4
Tag-Evol (Ours)	7.2	54.0	49.0	89.4	71.6	80.5	80.2	61.7

Table 2: The main experimental results cover the domains of general-purpose tasks, mathematical reasoning, and code generation. IF-L and IF-S denote the IFEval-loose and IFEval-strictly evaluation benchmarks, respectively.

spite the fact that they have three times as much data as the seed dataset. We analyze that it may be that the Qwen2.5 model has already mastered the knowledge involved in the seed dataset, and the SFT phase serves more as alignment than learning new skills. Existing Evol-Instruct methods evolve only based on one single sample and do not bring sufficiently diverse new combinations to stimulate new knowledge in LLMs. The proposed Tag-Evol can be seen as a weighted combination of different samples in the seed dataset by injecting knowledge tags, leading to a more diverse evolved samples, and also a 2-point improvement in the Qwen2.5 strong backbone setting.

5 Analysis

5.1 Effect of Multi-step Tagging

Setup. As mentioned in Section 3.1, the diversity and specificity of tags directly affects the effectiveness of the Tag-Evol method. In order to validate the effectiveness of the proposed fine-grained tagging method, we conduct analytical experiments on the tagging method. We choose two baseline methods: the original method reported in the Lu et al.’s (2023) paper, and the method of modifying its content for one-step fine-grained tagging. We compare the performance of Tag-Evol method under different tagging pools on the math domain setting using Llama3-8B as the base model.




Figure 4: Impact of different tagging methods on the quality of the final evolved dataset. The y-axis represents task performance and the number of tag pools, respectively

Results. The experiment results are shown in Figure 4. The most obvious difference between these methods is the number of tags in the tag pool, and with the two-phase tagging method, we can obtain tags about more aspects (20 times for the original, 2 times for the single step). The number of tags represents the diversity of evolutionary strategies. As shown in the table, the mathematical performance of the model increases gradually as the number of tags increases, from 67.0 to 69.3 on GSM8K. For the more difficult MATH-500 dataset, the improvement from complex tags is even more significant, from 34.6 to 38.0.




Figure 5: Multi-round performance at different data sizes using Llama3-8B as backbone.

5.2 Analysis of Data Scale

Setup. In the main experiment, we use a 3-iteration evolutionary setup for all methods. In this section we analyze the data at different scales for each iteration in more detail. We compare the performance on Llama3-8B with medium capacity under three different amounts of evolved data. It is worth noting that Evol-Instruct methods use iterative evolution, where the input to each round of evolution is the output of the previous round. By contrast, all three rounds of the proposed Tag-Evol are evolved on the same seed dataset with different budgets.

Results. The performance of each method for different number of rounds is shown in Figure 5. As the number of evolution rounds increases, the size of the merged dataset keeps growing and all the evolutionary methods show an upward trend compared to the seed dataset. Overall, Tag-Evol demonstrate strengths at almost all data scales across various benchmarks. It suggests that injecting different numbers of knowledge tags directly to the original instructions generates equally or even more challenging samples compared to iterative evolution, resulting in improved performance for training on evolved data.

In addition to this, Tag-Evol shows a more stable performance improvement compared to the Evol-Instruct method. Our analysis suggests that this is due to the fact that strategy of Evol-Instruct is selected randomly, independent of the original instruction. This may produce some inappropriate evolutions leading to hallucinations. Differently, our approach mitigates this effectively by sampling a batch of tags for the model to actively select, enabling more stable training.

Method	Evolution LLM	
	7B-Instruct	72B-Instruct
<i>Evol-Instruct</i>		
Evol-Ins	68.3	66.1
Auto Evol-Ins	68.1	68.6
<i>Tag-Evol</i>		
7B-Tagger	68.7	68.2
72B-Tagger	69.2	69.3

Table 3: Analysis experiments on the scale of tagging and evolution models using the Qwen2.5 series of models. 7B-Tagger represents the use of Qwen2.5-7B-Instruct to build the mentioned tag pool.

5.3 Analysis of Model Scale

Setup. In order to evaluate in detail the requirements of the proposed method on the capability of instruct model in both the tag pool construction and tag sampling evolution phases, we conduct experiments using different model scales. We choose the Qwen2.5 series 7B-Instruct and 72B-Instruct as the evolutionary LLMs to evaluate the performance of three different Evol-Instruct approaches. In addition to this, we also used these two scales of models to construct tag pools (7B-Tagger and 72B-Tagger in Tabel 3) separately to explore the relationship between the quality of tags and the ability of the tagging model. Regardless of the way the instruction datasets are constructed, we use the Qwen2.5-72B-Instruct model to generate all the responses of the instruction data for a fair comparison.

Results. We conduct experiments on the GSM8K dataset and the results are shown in Table 3. For Evol-Instruct and Tag-Instruct, where the strategy has been given, we find that the evolved instruc-

Method	GSM8K		MATH-500		HumanEval		MBPP	
	8-gram	13-gram	8-gram	13-gram	8-gram	13-gram	8-gram	13-gram
Seed	60	4	528	76	38	0	83	2
Evol-Ins	31	1	341	31	33	0	46	2
Auto Evol-Ins	38	0	388	24	26	0	5	1
Tag-Evol	30	3	357	36	15	0	50	1

Table 4: Leakage testing of different evolved datasets on benchmark data.

tions generated by the smaller model achieve higher performance, consistent with Hui et al.’s (2024) findings. It is shown that given an evolutionary strategy, evolution is essentially a rewriting task that does not require much modeling capability. The uncertainty of small models may bring diversity to evolved data. However, for the Auto Evol-Instruct approach, it requires the LLM to generate concrete strategies based on abstract prompt, resulting in its evolution being more dependent on the capabilities of the evolution model.

For the models used for tagging, the experiments illustrate that we still need some high-quality tags generated by large models to better guide the evolutionary process. Fortunately, high-quality tags are not disposable and are easy to collect and maintain. Once we have maintained a collection of high-quality tags in a specific domain, using Tag-Evol to evolve existing data will be an efficient process.

5.4 Data Leakage Detection

Setup. We follow the Liu et al.’s (2023a) setup and use gram level matching to detect data leakage. Specifically, we choose 8-gram matching and 13-gram matching on four benchmarks in the math and code domains.

Results. As shown in Table 4, all the evolved datasets have a smaller number of matches to the benchmark than the seed dataset, having a lower risk of data leakage. The matching number in the math domain is significantly higher than that in the code domain, which we analyze may be due to the fact that the seed set uses a homogeneous training set in math. In addition, the high number of matches in the MATH dataset may be because this dataset has more formulas, which are not tokenized efficiently and are prone to multi-gram matches.

5.5 Discussion of Complexity and Diversity

Setup. In addition to the performance on downstream tasks, we also use the automated metric

Method	Seed	R1	R2	R3
<i>Difficulty</i>				
Evol-Ins	2.12	2.38	2.30	2.44
Auto Evol-Ins	2.12	2.32	2.76	2.89
Tag-Evol	2.12	2.44	2.96	3.28
<i>Diversity</i>				
Evol-Ins	45	61	66	62
Auto Evol-Ins	45	51	69	71
Tag-Evol	45	60	69	103

Table 5: Automated evaluation metrics defined by InsTag for evolved data, where R1 represents Round 1.

used in InsTag to evaluate our evolved data, which is divided into two dimensions: diversity and difficulty. Lu et al. (2023) use the average number of tags to indicate the difficulty of the dataset, and the full set of tags in the overall dataset to indicate diversity. We follow his setting and experiment using the prompt reported in the paper. Specifically, we use the latest GPT-4o (gpt-4o-2024-08-06) for our experiments. We randomly select 50 samples and evaluate different versions of these data evolved different rounds using different methods.

Results. The results are shown in Table 5. The Tag-Evol evolved dataset has significant advantages in both difficulty and diversity. For the difficulty metric, the experiment shows that our multi-tag direct injection strategy is feasible. Compared to Evol-Instruct methods that iteratively evolve 3 times, we achieve the goal of efficiently constructing difficult samples by injecting multiple knowledge tags into the original instructions at once. For the diversity metric, the samples evolved by Tag-Evol contain significantly more tags than the other methods in the third round by injecting multiple tags. It shows that the knowledge tags we inject via Tag-Evol are effective and can be validated by re-tagging the evolved data.

6 Conclusion

In this paper, we propose a novel instruction evolving method, TagEvol, that can generate diverse and challenging instruction data more efficiently. Experiments in multiple domains illustrate the advantages of the Tag-Evol framework. Its efficient generation with flexible strategies makes it a very promising method for data synthesis.

Limitations

Considering the cost factor, the main experiments in this paper all use the current powerful open source model Qwen2.5-72B-Instruct for tag pool construction and tag sampling generation. The Tag-Evol evolved dataset may have higher performance if experiments are conducted using more powerful closed-source models, allowing the trained model to be compared with existing open-source SOTA models. Moreover, for a fair comparison, only the tags extracted from the seed dataset are chosen as the evolutionary strategy in this paper. Whether by mixing high-quality tags from different sources can bring unexpected enhancement in the general domain is a worthwhile thing to try, which will also serve as our future research direction.

Ethics Statement

The source data for the proposed methods is derived entirely from publicly available project resources on legitimate websites, ensuring that no sensitive information is involved. Furthermore, all baselines and datasets utilized in our experiments are also publicly accessible, and we have credited the respective authors by citing their work.

Acknowledgements

We gratefully acknowledge the support of the National Natural Science Foundation of China (NSFC) via grant 62236004, 62206078 and 62476073.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman,

Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

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*.

Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world's first truly open instruction-tuned llm](#).

Yuyang Ding, Xinyu Shi, Xiaobo Liang, Juntao Li, Qiaoming Zhu, and Min Zhang. 2024. Unleashing reasoning capability of llms via scalable question synthesis from scratch. *arXiv preprint arXiv:2410.18693*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Tingfeng Hui, Lulu Zhao, Guanting Dong, Yaqi Zhang, Hua Zhou, and Sen Su. 2024. Smaller language models are better instruction evolvers. *arXiv preprint arXiv:2412.11231*.

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

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let's verify step by step. *arXiv preprint arXiv:2305.20050*.

- Bingbin Liu, Sebastien Bubeck, Ronen Eldan, Janardhan Kulkarni, Yuanzhi Li, Anh Nguyen, Rachel Ward, and Yi Zhang. 2023a. Tinygsm: achieving > 80% on gsm8k with small language models. *arXiv preprint arXiv:2312.09241*.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2023b. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning. *arXiv preprint arXiv:2312.15685*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023. # instag: Instruction tagging for analyzing supervised fine-tuning of large language models. In *The Twelfth International Conference on Learning Representations*.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023a. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023b. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*.
- 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.
- Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattarjee, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. Large language models for data annotation and synthesis: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 930–957.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.
- Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. 2024. Magpie: Alignment data synthesis from scratch by prompting aligned llms with nothing. *arXiv preprint arXiv:2406.08464*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Weihao Zeng, Can Xu, Yingxiu Zhao, Jian-Guang Lou, and Weizhu Chen. 2024. Automatic instruction evolving for large language models. *arXiv preprint arXiv:2406.00770*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

A Training Hyper-parameter

B Prompt Details

We have designed a simple and effective prompt to guide LLM for tag pool construction (Table 7) and tag sampling evolution (Table 8).

C Case Study

We analyze the examples in code (Table 9) and math (Table 10) for three methods, including Evol-Instrut, Auto Evol-Instrut and Tag-Evol (Ours). It can be found that the other two methods yield similar outcomes over three rounds of evolution. However, our method stands out in terms of diversity.

Hyperparameter	Value
General	
Number of Epochs	3
Number of Devices	8
Per-device Batch Size	4
Learning Rate	2×10^{-5}
Learning Rate Scheduler	WarmupDecayLR
Warmup step	15
Max Sequence Length	2048
Code Dataset	
Gradient Accumulation Steps	16
Math Dataset	
Gradient Accumulation Steps	8

Table 6: Hyperparameters and their values

Moreover, by sampling a greater number of tags in just one round, we are able to generate instructions that are more challenging compared to those produced by the other methods.

Prompt for Tagging

You are a tagging system that provides useful tags for task intentions to distinguish tasks for a helpful AI assistant.

#Task {Instruction}

Please follow the steps below to assign tags to the given task.

Step 1: Consider from which aspects that tags can be assigned to cover main features of the task and provide a brief explanation. Aspects need to be summarizing, such as 'Required skill'. etc. Please summarize this task with as few aspects as possible.

Step 2: Based on the #Aspect List# obtained in Step 1, assign core tags to the given task from each aspect.

Please reply strictly in the following format:

Step 1 #Aspect List and Explanation#:

Step 2 #Aspect2Tags#:

#Aspect2Tags#

{"xxx": [tag1, tag2, ...], "xxx": [tag1, tag2, ...], ...}

where xxx means aspect you get in step 1.

Table 7: Prompt for Tagging. [Instruction](#) stands for the instruction used to extract the tags.

Prompt for Evolution

You are an Instruction Rewriter that rewrites the given #Instruction# into a more challenging version based on the given tags.

Here is the #Instruction#: {Instruction}

Here is the #Tag List#: {TagList}

Please follow the steps below to rewrite the given #Instruction# into a more intricate and demanding version.

Step 1: Carefully read the #Instruction# and #Tag List#. Select a subset from the #Tag List# that will allow the #Instruction# to evolve and become more complex. The chosen subset should provide richer, more nuanced information that enhances the original prompt, ultimately increasing its difficulty and quality. Aim to challenge advanced AI assistants like ChatGPT and GPT-4. The subset should contain {Budget} tags and must exclude any tags already present in the #Instruction#.

Step 2: Develop a comprehensive plan based on the #Tag subset# generated in Step 1 to make the #Instruction# more challenging. This plan should focus on seamlessly integrating multiple tags from the #Tag subset# into the original #Instruction#.

Step 3: Execute the plan step by step and provide the #Rewritten Instruction#. The #Rewritten Instruction# may only add between 10 and 20*{Budget} words to the original #Instruction#.

Step 4: Thoroughly review the #Rewritten Instruction# to identify any inconsistencies. Ensure that the #Rewritten Instruction# is solely a more challenging version of the original #Instruction#. Provide only the #Final Rewritten Instruction# without any additional explanation.

Please reply strictly in the following format:

Step 1 #Tag subset#:

Step 2 #Plan#:

Step 3 #Rewritten Instruction#:

Step 4 #Finally Rewritten Instruction#:

Table 8: Prompt for Evolution. [Instruction](#) represents based instruction for evolution; [TagList](#) represents tags batch for selection; [Budget](#) represents the number of tags selected each time.

<p>Original:Reverse the string given in the input Reverse the string given in the input</p>
<p>Evol-Instruct</p> <p>Round1:Reverse the string given in the input, and explain the algorithmic steps and time complexity involved in the process.</p> <p>Round2:Reverse the alphanumeric string provided in the input, detailing the recursive algorithmic steps and analyzing the time complexity in Big O notation.</p> <p>Round3:Reverse the hexadecimal string provided in the input, detailing the recursive algorithmic steps and analyzing the time complexity in Big O notation, considering the impact of string length on performance.</p>
<p>Auto Evol-Instruct</p> <p>Round1:Given an input string, reverse it only if its length exceeds 5 characters; otherwise, return the original string.</p> <p>Round2:Given an input string, reverse it and convert to uppercase only if its length exceeds 5 characters and it contains only alphabetic characters; otherwise, return the original string or a message indicating non-alphabetic characters.</p> <p>Round3:Given an input string, reverse it and convert to uppercase only if its length exceeds 5 characters, it contains only alphabetic characters, more than 2 vowels, and the substring "abc"; otherwise, return the original string or a message indicating the specific reason for non-conversion.</p>
<p>Tag-Instruct</p> <p>Tag 1:Reverse the string given in the input and multiply its length by 2.</p> <p>Tag subset:['basic math calculations']</p> <p>Tag 3:Reverse the string given in the input, ensuring it contains both letters and numbers. Process the form to validate the input, and if the string length is even, append the sum of the digits in the string to the reversed output.</p> <p>Tag subset:['form processing', 'letters and numbers', 'basic math calculations']</p> <p>Tag 5:Reverse the string given in the input, but first, process the input form to ensure it contains both letters and numbers. For each letter, convert it to its corresponding ASCII value, and for each number, double it. Then, calculate the position of the third character in the modified string and append this position to the end of the reversed string. Provide a coding example to demonstrate the process.</p> <p>Tag subset:['form processing', 'letters and numbers', 'coding example', 'basic math calculations', 'third character']</p>

Table 9: Multi-round case on code domain for Evol-Instruct, Auto Evol-Instruct and Tag-Instruct

Original: Expand $(x - 2)(x + 2)(x^2 + 4)$.

Evol-Instruct

Round1: Expand $(x - 2)(x + 2)(x^2 + 4)$ and explain the significance of the resulting polynomial in the context of algebraic identities.

Round2: Expand $(x - 2)(x + 2)(x^2 + 4)$, explain the significance of the resulting polynomial in the context of algebraic identities, and discuss its applications in solving higher-degree polynomial equations.

Round3: Expand $(x - 2)(x + 2)(x^2 + 4)$, explain how the resulting polynomial exemplifies the sum and difference of squares identity, and discuss its applications in factoring and solving quartic equations.

Auto Evol-Instruct

Round1: Expand $(x - a)(x + 2 + ax)(x^2 + 4a)$ where a is a non-zero integer.

Round2: Expand $(x - a)(x + 2 + b + ax)(x^2 + 4a)$ where $a = x + 1$, b is a positive integer, and a is a non-zero integer.

Round3: Expand $(x - a)(x + 2 + b + ax)(x^2 + 4a)$ where $a = x + 1$, $b = 2x + 1$, $c = x^2 + b$, a is a non-zero integer, b is a positive integer, and $c > x$.

Tag-Instruct

Tag 1: Expand $(x - 2)(x + 2)(x^2 + 4)$, given that x is an integer such that $x > 0$.

Tag subset:['conditions on variables']

Tag 3: Expand $(x - 2)(x + 2)(x^2 + 4)$ for $x > 0$, then use the result to find the perimeter of a square whose side length is the square root of the expanded expression evaluated at $x = 3$.

Tag subset:['conditions on variables', 'sequential operations', 'polygon perimeter']

Tag 5: Expand $(x - 2)(x + 2)(x^2 + 4)$, where x is an integer such that $-3 \leq x \leq 3$. Factor the expanded polynomial, if possible, and then evaluate $\sin(\pi \cdot \text{expanded polynomial})$. Find the sum of all integer values of x for which the result is an integer.

Tag subset:['conditions on variables', 'factor', 'evaluating trigonometric functions', 'sequential operations', 'integer sum problem']

Table 10: Multi-round case on math domain for Evol-Instruct, Auto Evol-Instruct and Tag-Instruct