

Packing Analysis: Packing Is More Appropriate for Large Models or Datasets in Supervised Fine-tuning

Shuhe Wang[♠], Guoyin Wang[♦], Yizhong Wang[▲]
Jiwei Li[▼], Eduard Hovy^{♣♣}, Chen Guo[►]

Abstract

Packing, initially utilized in the pre-training phase, is an optimization technique designed to maximize hardware resource efficiency by combining different training sequences to fit the model’s maximum input length. Although it has demonstrated effectiveness during pre-training, there remains a lack of comprehensive analysis for the supervised fine-tuning (SFT) stage on the following points: (1) whether packing can effectively enhance training efficiency while maintaining performance, (2) the suitable size of the model and dataset for fine-tuning with the packing method, and (3) whether packing unrelated or related training samples might cause the model to either excessively disregard or over-rely on the context.

In this paper, we perform extensive comparisons between SFT methods using padding and packing, covering SFT datasets ranging from 69K to 1.2M and models from 8B to 70B. This provides the first comprehensive analysis of the advantages and limitations of packing versus padding, as well as practical considerations for implementing packing in various training scenarios. Our analysis covers various benchmarks, including knowledge, reasoning, and coding, as well as GPT-based evaluations, time efficiency, and other fine-tuning parameters. We also open-source our code for fine-tuning and evaluation and provide checkpoints fine-tuned on datasets of different sizes, aiming to advance future research on packing methods.

1 Introduction

Supervised fine-tuning (SFT) refers to adapting a pre-trained model to perform specific tasks by train-

ing it on a labeled conversation dataset consisting of (instruction, answer) pairs (Wang et al., 2022a; Taori et al., 2023; Chiang et al., 2023; Almazrouei et al., 2023; Wang et al., 2023b; Bai et al., 2023; Wang et al., 2023a; Cai et al., 2024; Young et al., 2024). As models and datasets grow, the cost of fine-tuning rises. Identifying cost-effective methods, optimizing resource utilization, and alleviating the financial burden of large-scale training present new challenges in the SFT process (Hu et al., 2021; Zhang et al., 2023a; Dettmers et al., 2024).

Packing addresses these challenges by combining multiple training samples into a single sample. Originally used during the pre-training phase, packing extends each training sequence to the model’s maximum input length, optimizing hardware resource usage, such as GPUs, and improving training efficiency (Brown et al., 2020; Rae et al., 2021; Chowdhery et al., 2022; OpenAI, 2023; Touvron et al., 2023; Dubey et al., 2024). Despite its proven effectiveness during the pre-training phase, for SFT, a thorough analysis is still lacking on: (1) whether packing can effectively enhance training efficiency while maintaining performance, (2) the suitable size of the model and dataset for fine-tuning with packing, and (3) whether packing unrelated or related training samples might cause model to either excessively disregard or over-rely on the context.

To address these concerns, this paper provides a thorough analysis of packing during the supervised fine-tuning (SFT) stage. Specifically, we perform extensive comparisons between SFT methods using padding and packing, covering SFT datasets ranging from 69K to 1.2M and models from 8B to 70B. Our comparisons include various benchmarks, such as knowledge, reasoning, and coding, GPT-based evaluations, time efficiency, and other fine-tuning parameters, concluding that:

- Models using packing generally perform better on average compared to those using

Email: shuhewang@student.unimelb.edu.au

♠The University of Melbourne, ♦Alibaba, ▲University of Washington, ▼Zhejiang University, ♣Language Technologies Institute, Carnegie Mellon University, ►China Mobile Group Device Co., Ltd

Our codes, models, and datasets are available at: <https://github.com/ShuheWang1998/Packing-Analysis?tab=readme-ov-file>

padding across various benchmarks.

- As the model size grows, the performance gap between padding and packing-based models on the benchmark increases.
- Tailoring the packing of specific training samples may result in desired performance.
- Compared to padding, the packing method greatly reduces training time, making it possible to fine-tune large models on large datasets.
- Using longer training samples increases the time required for the packing method to process each sample, making it less suitable for training on particularly small datasets.
- In packing mode, the batch size is no longer directly proportional to the learning rate.
- Applying packing to datasets with only single-turn conversations may lead to substantial decline in performance on few-shot benchmarks.

Building on these findings, we provide the first comprehensive analysis of the advantages and limitations of packing compared to padding, as well as practical considerations for implementing packing in various training scenarios. Additionally, we have open-sourced our code for fine-tuning and evaluation and released checkpoints fine-tuned on datasets of varying sizes, contributing to future research on packing methods.

2 Related Work

Supervised fine-tuning (SFT) in large language models (LLMs) involves additional training on a dataset of (instruction, answer) pairs. This approach helps align the LLMs' training goal of predicting the next word with users' expectations for the models to follow human instructions more accurately (Mishra et al., 2021; Wei et al., 2021; Rosenbaum et al., 2022; Ouyang et al., 2022; Wang et al., 2022b; Dwivedi-Yu et al., 2022; Longpre et al., 2023; Zhang et al., 2023b; Qi et al., 2023; Chung et al., 2024; Liu et al., 2024).

The initial step in SFT is to create annotated data, but current SFT datasets are often constrained by their limited quantity, diversity, and creativity (Mukherjee et al., 2023; Xu et al., 2023; Lu et al., 2023; Song et al., 2024; Wang et al., 2023b; Zhou et al., 2024). To address this issue and rich resources for research, in one line, researchers distilled data from powerful large models (e.g., GPT-4 (OpenAI, 2023)) (Chiang et al., 2023; Ding et al., 2023; Zhao et al., 2024). On the other line, some

researchers are working on methods to enable pre-trained models to self-generate useful SFT data (Wang et al., 2022a; Sun et al., 2024). Once high-quality SFT datasets are created, the next step is to use them for fine-tuning pre-trained models. Many studies are dedicated to minimizing the costs of fine-tuning, including GPU usage and time, while maintaining performance, such as light-weight fine-tuning (Hu et al., 2021; Dettmers et al., 2024), speeding up attention algorithm for transformer-based LLMs (Dao et al., 2022), and efficient distributed fine-tuning (Rajbhandari et al., 2020)².

In this paper, we are analyzing one of the efficient fine-tuning techniques: packing, which packs multiple training samples into a single sample to maximize the utilization of hardware resources and enhance fine-tuning efficiency, providing a comprehensive understanding of its effectiveness and potential risks. Below, we will thoroughly analyze its effectiveness and potential risks.

3 Methods: Padding and Packing

Padding and packing are two distinct methods to organize training samples, shown in Figure 1. In this section, we first define a set of mathematical symbols, followed by a detailed explanation of one padding method and two packing methods: random packing and greedy packing.

We assume that $\{C\} = \{c^1, \dots, c^N\}$ denotes the training conversations, where N is the size of the training conversation, and $c^i = \{(x_1^i, y_1^i), \dots, (x_m^i, y_m^i)\}$, $m \geq 1$ denotes an instruction x and answer y pair of length m . In this paper, we also use the terms "*single-turn conversation*" and "*multi-turn conversation*" to distinguish between training conversations that consist of a single instruction and answer pair ($m = 1$) and those that include multiple contextual instruction and answer pairs ($m \geq 2$).

3.1 Padding

Padding refers to extending the length of training conversations to a consistent size to match the input requirements of the pre-trained LLM, following are its detailed process, strengths, and weaknesses:

3.1.1 Process of Padding

Specifically, for a training sample comprising m -turn ($m \geq 1$) conversations

²<https://github.com/microsoft/DeepSpeed>

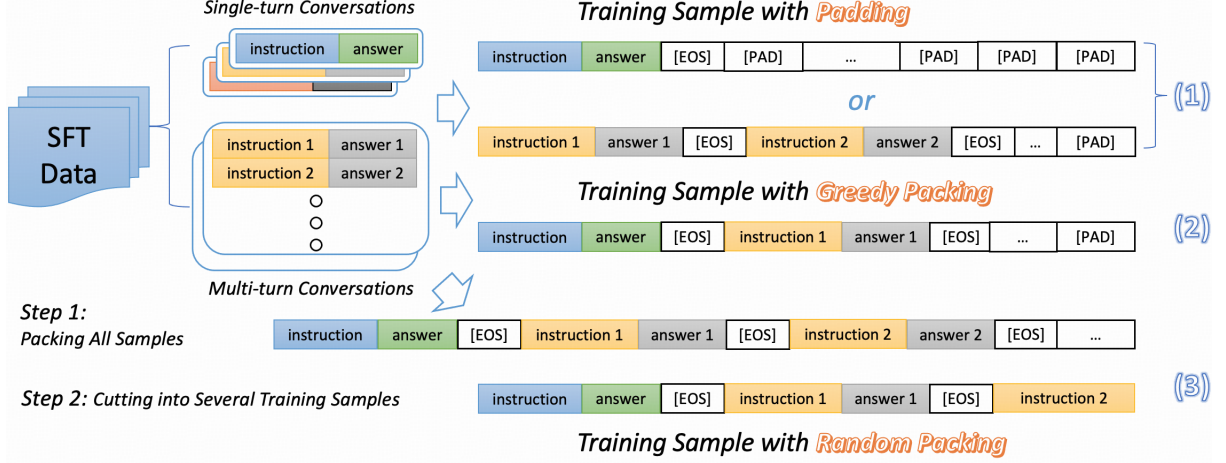


Figure 1: An example for the process of padding and packing methods: **(1) Padding**: Each training sample is appended with the special token “[PAD]” to meet the requirement of the model’s input length; **(2) Random Packing**: Each training sample is packed together as much as possible according to its length; and **(3) Greedy Packing**: firstly, all training samples are packed into one single sample, and then the single sample is cut into several short training samples according to the maximum input length of the model.

Model		WildChat (GPT-4), 69K	TULU, 326K	WildChat, 652K	Open-source 1M, 1.2M
LLaMA-3-8B	<i>padding</i>	28.86	19.11	21.06	18.38
	<i>random packing</i>	27.89(-0.97)	20.84(+1.73)	20.73(-0.33)	20.42(+2.04)
	<i>greedy packing</i>	29.81(+0.95)	20.73(+1.62)	21.34(+0.28)	21.9(+3.52)
LLaMA-3-70B	<i>padding</i>	37.0	22.84	30.69	34.95
	<i>random packing</i>	39.92(+2.92)	23.93(+1.09)	30.76(+0.07)	35.21(+0.26)
	<i>greedy packing</i>	41.09(+4.09)	24.46(+1.62)	31.26(+0.57)	35.81(+0.86)

Table 1: Results of different size models and datasets on the WildBench benchmark. We highlight the highest score in bold and use blue and red to indicate whether the score has increased or decreased compared to the padding method. Note that the WildChat (GPT-4) dataset is composed entirely of real user interactions with GPT-4, and its internal data is the foundation of WilBench research. Therefore, it is reasonable for a model trained on the WildChat dataset to achieve a high score.

$\{(x_1, y_1), \dots, (x_m, y_m)\}$, we start by concatenating each turn of the conversation into a single sequence, using the special token $[EOS]$ to separate each instruction x and response y :

$$\begin{aligned}
 & x_1 y_1 [EOS] x_2 y_2 [EOS] \dots x_m y_m [EOS] \\
 & = \\
 & w_1 w_2 \dots [EOS] w_t w_{t+1} \dots w_T [EOS]
 \end{aligned}$$

where $w_i, 1 \leq i \leq T$ denotes the word of the concatenated sequence, and T denotes the number of words in this sequence.

During the training stage, LLMs often process data in batches, and each batch must consist of input sequences of equal length. This uniformity is essential because the fundamental mathematical operations required, such as matrix multiplications, necessitate tensors of consistent sizes. In this condition, for one batch with l concatenated sequences,

we must ensure that all sequences in this batch are limited to the minimum of the maximum input length of the pre-trained model (for example, 8192 for LLaMA-3-8B-base (Dubey et al., 2024)) and the length of the longest sequence within the batch:

$$T = \min(\text{Maximum Input Length}, \arg \max_{1 \leq i \leq l} T_i)$$

For sequences that exceeding length T , we truncate them, and for those shorter than T , we fill them up with non-informative special token $[PAD]$:

$$\begin{aligned}
 & w_1^1 w_2^1 \dots [EOS] w_{t_1}^1 \dots [EOS] \dots w_T^1 [EOS] \\
 & w_1^2 w_2^2 \dots [EOS] w_{T-2}^2 [PAD] [PAD] [EOS] \\
 & \dots \\
 & w_1^l w_2^l \dots [EOS] w_{t_l}^l \dots w_{T-1}^l [PAD] [EOS]
 \end{aligned}$$

3.1.2 Weaknesses of Padding

While straightforward to implement, excessive padding can lead to inefficiencies. If there are too

Model	MMLU Exact Match (5-shot)	GSM8K Exact Match (4-shot)	MATH Exact Match (4-shot)	BBH Exact Match (3-shot)	IFEval (prompt-loose-accuracy) (0-shot)	HumanEval pass@1 (0-shot)	Avg
WildChat (GPT-4) Dataset, Size: 69K							
<i>LLaMA-3-8B</i>							
Padding	63.99	58.76	14.72	60.71	56.01	43.29	49.58
Random Packing	63.5(-0.44)	61.18(+2.42)	15.58(+0.86)	61.04(+0.33)	51.57(-4.44)	43.9 (+0.61)	49.46(-0.12)
Greedy Packing	64.71(+0.72)	60.88(+2.12)	15.6(+0.88)	62.59(+1.88)	57.12(+1.11)	42.68(-0.61)	50.6(+1.02)
<i>LLaMA-3-70B</i>							
Padding	73.47	79.3	28.8	78.33	51.76	57.32	61.50
Random Packing	75.16(+1.69)	82.38(+3.08)	31.46(+2.66)	79.94(+1.61)	61.00(+9.24)	65.85(+8.53)	65.97(+4.47)
Greedy Packing	74.77(+1.3)	81.61(+2.31)	32.84(+4.04)	80.98(+2.65)	64.33(+12.57)	60.98(+3.66)	65.92(+4.42)
TULU Dataset, Size: 326K							
<i>LLaMA-3-8B</i>							
Padding	62.26	57.32	14.6	60.14	41.77	44.24	46.72
Random Packing	63.94(+1.68)	58.83(+1.51)	13.94(-0.66)	61.11(+0.97)	42.51(+0.74)	45.61(+1.37)	47.66(+0.94)
Greedy Packing	62.14(-0.12)	60.8(+3.48)	14.74(+0.14)	61.26(+1.12)	46.40(+4.63)	44.51(+0.27)	48.31(+1.59)
<i>LLaMA-3-70B</i>							
Padding	73.2	81.18	29.02	78.06	47.32	62.95	61.96
Random Packing	73.48(+0.28)	81.73(+0.55)	29.42(+0.4)	78.35(+0.29)	47.29(-0.03)	60.37(-2.58)	61.77(-0.19)
Greedy Packing	73.43(+0.23)	81.2(+0.02)	30(+0.18)	77.54(-0.52)	53.05(+5.73)	68.9(+5.95)	64.02(+2.06)
WildChat Dataset, Size: 652K							
<i>LLaMA-3-8B</i>							
Padding	64.52	61.83	14.21	61.88	51.36	40.12	48.99
Random Packing	64.46(-0.06)	62.77(+0.94)	14.44(+0.23)	62(+0.12)	50.28(-1.08)	40.24(+0.12)	49.03(+0.04)
Greedy Packing	65.07(+0.55)	61.41(-0.42)	15.08(+0.87)	62.83(+0.95)	52.68(+1.32)	48.17(+8.05)	50.87(+1.88)
<i>LLaMA-3-70B</i>							
Padding	74.82	79.26	29.44	76.31	52.19	63.7	62.62
Random Packing	75.67(+0.85)	80.1(+0.84)	30.37(+0.93)	76.74(+0.43)	52.43(+0.24)	65.26(+1.56)	63.43(+0.81)
Greedy Packing	75.36(+0.46)	79.45(+0.19)	31.28(+1.84)	77.47(+1.16)	53.60(+1.41)	64.02(+0.32)	63.53(+0.91)
Open-source 1M Dataset, Size: 1.2M							
<i>LLaMA-3-8B</i>							
Padding	63.7	77.08	27.96	63.45	48.39	45.22	54.3
Random Packing	63.96(0.26)	77.26(+0.16)	28.4(+0.44)	64.83(+1.38)	49.54(+1.15)	45.73(+0.51)	54.95(+0.65)
Greedy Packing	63.63(-0.07)	77.48(+0.4)	28.26(+0.3)	63.01(-0.44)	51.57(+3.28)	46.34(+1.12)	55.05(+0.75)
<i>LLaMA-3-70B</i>							
Padding	74.97	85.23	41.82	78.65	54.33	61.74	66.12
Random Packing	76.38(+1.41)	86.14(+0.91)	42.73(+0.91)	79.42(+0.77)	55.9(+1.57)	62.98(+1.24)	67.26(+1.14)
Greedy Packing	75.69(+0.72)	86.88(+1.65)	42.92(+1.1)	79.94(+1.29)	56.82(+2.49)	62.98(+1.24)	67.54(+1.42)

Table 2: Results of different size models and datasets on various benchmarks. We highlight the highest score in bold and use blue and red to indicate whether the score has increased or decreased compared to the padding method.

many padding tokens relative to actual data, it can lead to increased computation without corresponding benefits in learning. In order to alleviate this weakness, below, we introduce two packing-based methods: Random Packing and Greedy Packing.

3.2 Random Packing

Unlike the padding method, which extends shorter sequences with the special token $[PAD]$, random packing combines multiple training conversations into a single sequence randomly, to maximize the model’s learning efficiency and effectiveness:

3.2.1 Process of Random Packing

Firstly, we concatenate all training conversations $\{C\} = \{c^1, \dots, c^N\}$ into one single sequence with the special token $[EOS]$ to separate each instruction x and response y :

$$c_1[EOS]c_2[EOS] \cdots [EOS]c_N[EOS]$$

$$=$$

$$w_1w_2 \cdots [EOS]w_tw_{t+1} \cdots w_m[EOS]$$

where n represents the total number of training instructions x and responses y , w_i denotes each word in the concatenated sequence, and m signifies the total number of words in that sequence.

Secondly, assuming that the maximum input length for the pre-trained model is T , we adjust to this limit by dividing the concatenated sequence into M subsequences, each with a length of T . In the end, we randomly pack these sequences into batches with each size of l , forming like:

$$w_1^1w_2^1 \cdots [EOS]w_t^1w_{t+1}^1 \cdots w_T^1[EOS]$$

$$\cdots$$

$$w_1^lw_2^l \cdots [EOS]w_t^lw_{t+1}^l \cdots w_T^l[EOS]$$

Model	Epoch	Total Steps	Total Training Time (s)↓	Steps per Second↑	Samples per Second↑
WildChat (GPT-4) Dataset, Size: 69K					
<i>LLaMA-3-8B</i>					
Padding	4	1964	1188.8449	0.165	21.13
Random Packing	4	728	445.28773(-743.55717)	0.163(-0.002)	20.934(-0.196)
Greedy Packing	4	492	308.33346(-880.51144)	0.16(-0.005)	20.48(-0.65)
<i>LLaMA-3-70B</i>					
Padding	3	2943	9533.42936	0.031	1.976
Random Packing	3	1092	3749.3016(-5784.12776)	0.029(-0.002)	1.865(-0.111)
Greedy Packing	3	741	2573.34781(-6960.08155)	0.029(-0.002)	1.84(-0.136)
TULU Dataset, Size: 326K					
<i>LLaMA-3-8B</i>					
Padding	4	9183	4906.59014	0.165	21.084
Random Packing	4	1928	1175.43583(-3731.15431)	0.164(-0.001)	20.977(-0.107)
Greedy Packing	4	1956	1328.12592(-3578.46422)	0.147(-0.018)	18.841(-2.243)
<i>LLaMA-3-70B</i>					
Padding	3	13761	40735.40051	0.034	2.162
Random Packing	3	2889	9758.68127(-30976.71924)	0.03(-0.004)	1.895(-0.267)
Greedy Packing	3	2931	10313.89593(-30421.50458)	0.028(-0.006)	1.82(-0.342)
WildChat Dataset, Size: 652K					
<i>LLaMA-3-8B</i>					
Padding	4	18340	11738.48881	0.156	20.183
Random Packing	4	5348	3422.97918(-8315.50963)	0.156	20.006(-0.177)
Greedy Packing	4	4780	3124.28736(-8614.20145)	0.153(-0.003)	19.58(-0.603)
<i>LLaMA-3-70B</i>					
Padding	3	27510	97893.95669	0.034	2.261
Random Packing	3	8025	28904.78592(-68989.17077)	0.030(-0.004)	2.083(-0.178)
Greedy Packing	3	7170	25124.6234(-72769.33329)	0.029(-0.005)	1.826(-0.435)
Open-source 1M Dataset, Size: 1.2M					
<i>LLaMA-3-8B</i>					
Padding	4	33064	19918.48664	0.168	21.413
Random Packing	4	5400	3253.07972(-16665.40692)	0.166(-0.002)	21.255(-0.158)
Greedy Packing	4	5104	3175.09395(-16743.39269)	0.161(-0.007)	20.571(-0.842)
<i>LLaMA-3-70B</i>					
Padding	3	49596	184709.04470	0.031	2.306
Random Packing	3	8103	29893.65963(-154815.38507)	0.03(-0.001)	2.193(-0.113)
Greedy Packing	3	7653	27426.66515(-157282.37955)	0.028(-0.003)	1.786(-0.52)

Table 3: The training time of models across various datasets, with blue indicating an improvement over the padding method, while red represents a decrease in performance compared to the padding method.

3.2.2 Strengths of Random Packing

Compared to padding, random packing enhances computational efficiency by densely packing each training batch, minimizing unused space and optimizing the use of the model’s capacity. Furthermore, this approach potentially boosts the model’s ability to generalize by exposing it to a broader range of contextual combinations in each training sample, thereby providing more diverse scenarios.

3.2.3 Weaknesses of Random Packing

There are two potential weaknesses of random packing, one is that it can lead to the concatenation of two distinct or similar samples, which may cause the model to either excessively ignore or rely

on the context. We have put more analyses about this issue in Section 5. The other potential issue is that combining all training conversations and then dividing them into sequences could result in a single conversation being split across two different sequences. For example, the instruction might end up at the tail end of one sequence, while the corresponding answer starts at the beginning of the next sequence, which is shown at the bottom of Figure 1. To alleviate this issue, below, we turn to the other packing-based method: Greedy Packing.

3.3 Greedy Packing

Instead of random packing that might result in a single conversation being split across two different

sequences, greedy packing starts by sorting and selecting training conversations based on their length:

3.3.1 Process of Greedy Packing

Formally, for a m -turn ($m \geq 1$) training conversation $\{(x_1, y_1), \dots, (x_m, y_m)\}$, we first use the special token $[EOS]$ to concatenate all instructions x and answers y into one single sequence s :

$$\begin{aligned} & x_1 y_1 [EOS] \cdots [EOS] x_m y_m [EOS] \\ & = \\ & w_1 w_2 \cdots [EOS] w_t w_{t+1} \cdots w_T [EOS] \end{aligned}$$

where $w_i, 1 \leq i \leq T$ denotes the word of the concatenated sequence, and T denotes the number of words in this sequence.

Then, we sort all of the concatenated sequences s , iterating from the longest sequence, and, in a greedy way, we pack as many sequences as possible without exceeding the maximum input length allowed by the pre-trained model. The full process is present in Appendix 1, which results in M packed sequences S :

$$\begin{aligned} & w_1^1 w_2^1 \cdots [EOS] w_{t_1}^1 w_{t_1+1}^1 \cdots w_{T_1}^1 [EOS] \\ & \cdots \\ & w_1^M w_2^M \cdots [EOS] w_{t_M}^M w_{t_M+1}^M \cdots w_{T_M}^M [EOS] \end{aligned}$$

Finally, similar to the padding method, we truncate packed sequences that exceed the maximum length allowed by the pre-trained model and pad shorter sequences with the special token $[PAD]$ to randomly batch the packed sequences.

3.3.2 Strengths of Greedy Packing

Greedy packing mainly serves as a modification of random packing, designed to reduce the risk of dividing relevant contexts across different batches. Simultaneously, it preserves the benefits of packing: enhancing the model’s generalization capabilities by exposing it to a wider variety of contextual combinations within each training sample, thus encompassing a more diverse set of scenarios.

3.3.3 Weaknesses of Greedy Packing

In addition to the issue associated with the packing method: it may cause the model to either excessively ignore or rely on the context by packing two distinct or similar training samples into one sequence. Another potential concern is the break in the random sampling of training data. Since greedy

packing entails sorting and organizing data prior to batching, it naturally diminishes the randomness in the distribution of sequences across batches. This can affect the diversity of data within each batch, as it is not entirely random but instead guided by the specific criteria (sequence length), for packing. However, despite these concerns in subjective analysis, our analysis and a series of experimental results in Section 5 have shown that using models trained with the greedy packing method does not result in any performance loss across various downstream benchmarks and GPT-based evaluations.

4 Experimental Setups

In this section, we sequentially describe our “Training Setups” in 4.1 and “Evaluation Setups” in 4.2.

4.1 Training Setups

4.1.1 Training Datasets

The development of packing was primarily aimed at maximizing hardware resource utilization and minimizing training duration. To demonstrate these training differences between packing and padding, we analyze four SFT datasets with different sizes: (1) WildChat-GPT-4 ($\sim 69k$) (Zhao et al., 2024), (2) TULU ($\sim 326k$) (Wang et al., 2023b), (3) WildChat ($\sim 652k$) (Zhao et al., 2024), and (4) Open-source 1M ($\sim 1.2M$). Due to the page limit, we put more details about these four datasets on Appendix B.1.

4.1.2 Model Training Details

Pre-trained Models. In this paper, our experiments utilize the LLaMA-3-8B and LLaMA-3-70B (Dubey et al., 2024) models, which are among the largest and most advanced pre-trained models currently accessible to the research community.

Chat Template. Following Dubey et al. (2024), we format all datasets to follow a chat template to unify the varied styles and formats of instruction-answer, which details are put in Appendix C.

Training Details. During the training stage, we follow the default settings in Wang et al. (2023b) masking loss belonging to the input (instruction), and only computing the loss after the special token “<|start_header_id>assistant<|end_header_id>”. All experiments were deployed the cluster with 4 nodes, each node containing 8 NVIDIA A800 80GB GPUs. For parameters, shown in Appendix 5, we keep the same between padding and packing.

4.2 Evaluation Setups

Following Wang et al. (2023b) that general-purpose models should be able to perform some core tasks before they can generalize to satisfy various practical needs, we first assess the core capabilities of our fine-tuned models using specific benchmarks: MMLU (Hendrycks et al., 2020), MATH (Hendrycks et al., 2021), GSM8K (Cobbe et al., 2021), BBH (Suzgun et al., 2022), HumanEval (Chen et al., 2021) and IFEval (Zhou et al., 2023). Subsequently, we employ evaluations based on GPT-4 (WildBench (Lin et al., 2024)) to gauge their proficiency in following instructions and aligning with human preferences. Due to the page limit, we put details of these benchmarks on Appendix C.2.

5 Results and Analysis

In this section, we provide our experimental results as well as analysis based on them.

5.1 Analysis on Various Benchmarks

Table 2 and Table 1 show results of different size models and datasets on various benchmarks, from that we can observe that:

(1) Models using packing generally perform better on average compared to those using padding across various benchmarks. Compared to the padding method, the packing method exposes models to a wider variety of contextual combinations within each training sample, offering a more diverse set of scenarios and enhancing the models’ ability to generalize. For example, 61.50 (Padding) v.s. 65.97 (Random Packing) on the LLaMA-3-70B for the WildChat (GPT-4) dataset.

(2) As the model size grows, the performance gap between padding and packing-based models on the benchmark increases. This is due to enhanced contextual efficiency. As the model size increases, its ability to effectively utilize extended contexts improves, thereby magnifying the advantages of the diverse contextual combinations brought by packing. For example, on the WildChat (GPT-4) dataset, the average score is 49.58 (Padding) v.s. 50.6 (Greedy Packing) on the 8B model, while on the 70B model, it is amplified to 61.50 (Padding) v.s. 65.92 (Greedy Packing).

(3) Models that use greedy packing generally perform better than those employing random packing across most benchmarks, with particularly strong results on the IFEval and WildBench, which assess instruction-following capabilities.

This advantage of the greedy packing lies in its ability to maintain the coherence of multi-turn conversations, unlike random packing which may split such conversations across different training samples. By preserving the integrity of multi-turn conversations, greedy packing helps models better learn when to use prior context effectively. This improves the model’s ability to selectively access relevant information in instructions, resulting in enhanced performance on benchmarks such as IFEval and WildBench that evaluate instruction-following capabilities, for example, the IFEval score improved from 49.54 to 51.57 for the Open-source 1M dataset based on the model LLaMA-3-8B. This interesting observation also suggests a new direction: **tailoring the packing of specific training samples to achieve desired performance**, where we will put more efforts in the future.

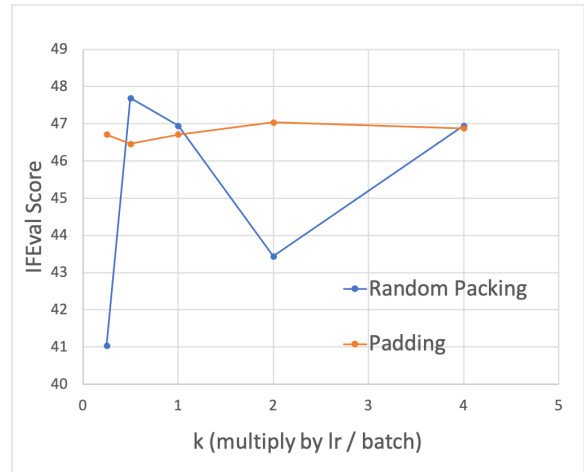


Figure 2: The results of fine-tuning the LLaMA-3-8B model on the TULU dataset using different linear combinations of batch size and learning rate.

5.2 Analysis on Training Time

Table 3 shows the training time of different size models on various training datasets, from that:

(1) Compared to padding, packing greatly reduces training time, making it possible to fine-tune large models on large datasets. The packing method significantly decreases training time by efficiently utilizing the available computational resources, for example, 40735s (Padding) v.s. 9758s (Random Packing) on the model LLaMA-3-70B for the TULU dataset. This reduction is particularly beneficial for scaling up model training and enables the effective handling of larger models and more extensive datasets, for example, from 184709s (Padding) significantly dropping to

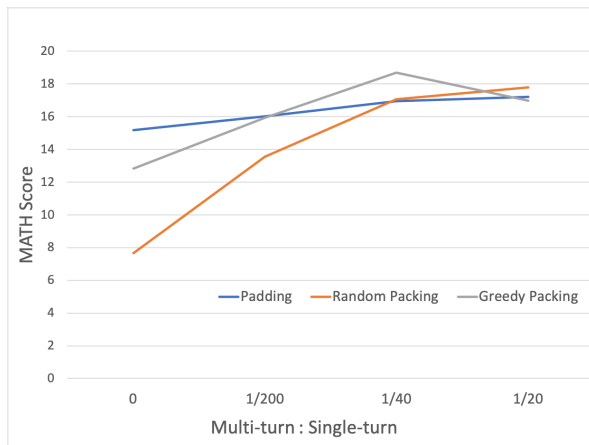


Figure 3: The results of fine-tuning the LLaMA-3-8B model by varying the ratio of multi-turn conversations and single-turn conversations.

27426s (Greedy Packing) on the LLaMA-3-70B for the Open-source 1M dataset, thus enhancing the overall training throughput and allowing for more complex and comprehensive model fine-tuning.

(2) Using longer training samples increases the time required for the packing method to process each sample, making it less suitable for training on particularly small datasets. Compared to padding, the packing method results in a lower number of samples processed per second, for example, 2.162 (Padding) v.s. 1.895 (Random Packing) for the TULU dataset based on the model LLaMA-3-70B. Therefore, if your goal is to fine-tune a small model (e.g., 6B, 8B, or 9B) on a small dataset (e.g., 20K or 30K), using the padding method might be more time-efficient.

5.3 Other Analysis

In addition to the analysis provided, we conducted additional experiments and concluded that:

(1) In packing mode, the batch size is no longer directly proportional to the learning rate. Previous research indicates that when increasing the batch size by a factor of k , the learning rate should also be multiplied by k to maintain a constant variance in the gradient expectation (Goyal, 2017). This raises the question of whether the linear relationship between batch size and learning rate still holds when using the packing method. To determine this, we compare the padding method and the random packing method by separately fine-tuning the LLaMA-3-8B model on the TULU dataset using different linear combinations of batch size and learning rate. Results are shown in Figure

2, where the IFEval (Zhou et al., 2023) score is the primary evaluation metric. The results reveal that while the batch size and learning rate adhere to a linear relationship in the padding method, this is not the case with the packing method. This discrepancy is due to the nature of packing: it does not ensure that each training sample consistently contains the same number of training conversations. Consequently, when the batch size is increased by a factor of k , the actual number of training conversations is not necessarily scaled up by the same factor, disrupting the linear relationship between batch size and learning rate.

(2) Applying packing to datasets with only single-turn conversations may lead to a significant decrease in performance on few-shot benchmarks. In packing methods, training samples that lack contextual connections may be combined to create what could be considered "fake" multi-turn conversations. When the training dataset includes multi-turn conversations, this approach allows the model to learn when to consider the context and when not to. However, if the training dataset only consists of single-turn conversations, there's a risk that the model might become less effective at utilizing context, potentially leading to a decline in performance on few-shot benchmarks. To investigate this, we fine-tuned the LLaMA-3-8B model separately using the packing and padding methods on the filtered 200K OpenHermes 2.5 dataset³, which only consists of single-turn conversations. The results, shown in Figure 3, reveal a significant drop in performance on the MATH (Hendrycks et al., 2021) benchmark. Then, when we added multi-turn conversations into the fine-tuning dataset, scaling from 1/40 to 1/20 of multi-turn conversations was sufficient to restore performance to normal levels.

Interestingly, our recent experiments with our internal 200K high-quality single-turn dataset did not show any decline in performance on few-shot benchmarks. We attribute this to differences in data quality and plan to conduct further analysis.

6 Conclusion

In this paper, we conduct a thorough comparison of SFT methods using padding and packing, analyzing datasets from 69K to 1.2M and models ranging from 8B to 70B. This provides the first detailed examination of the advantages and limitations of

³<https://huggingface.co/datasets/teknium/OpenHermes-2.5>

packing versus padding, as well as practical considerations for implementing packing in various training scenarios. Our evaluation spans a range of benchmarks, including knowledge, reasoning, and coding, and includes GPT-based assessments, time efficiency, and other fine-tuning factors. We also open-source our code, along with checkpoints fine-tuned on datasets of varying sizes, to support future research into packing techniques.

7 Limitations

The main limitations of this paper can be summarized in two aspects:

(1) Exploration of Additional Packing Methods.

Due to page constraints, this paper does not cover a wider range of packing strategies that could potentially enhance model performance. For instance, we did not explore methods such as grouping similar samples together into a single packed instance or deliberately combining highly dissimilar samples to examine their effects on training dynamics. These alternative packing approaches may yield different trade-offs in terms of generalization, stability, and computational efficiency. In future work, we intend to systematically investigate these and other advanced packing techniques to identify the most effective strategies for different learning scenarios.

(2) In-Depth Analysis of Packing and Data Quality. While we demonstrate the effectiveness of certain packing techniques, we do not conduct a detailed investigation into key aspects that could further refine our approach. Specifically, we did not explore how tailoring the packing strategy to specific types of training samples can optimize performance for different tasks. Additionally, the impact of data quality in padding and packing remains an open question—understanding how variations in data quality influence training outcomes could lead to more robust and adaptive packing strategies. Future research will focus on these aspects, aiming to develop a more nuanced understanding of how different packing methods interact with data quality and task-specific requirements.

References

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine

Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.

Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. *GitHub repository*.

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.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov,

- and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- 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*.
- OpenCompass Contributors. 2023. Opencompass: A universal evaluation platform for foundation models.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi Zheng, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. *arXiv preprint arXiv:2305.14233*.
- 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*.
- Jane Dwivedi-Yu, Timo Schick, Zhengbao Jiang, Maria Lomeli, Patrick Lewis, Gautier Izacard, Edouard Grave, Sebastian Riedel, and Fabio Petroni. 2022. Editeval: An instruction-based benchmark for text improvements. *arXiv preprint arXiv:2209.13331*.
- P Goyal. 2017. Accurate, large minibatch sg d: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- 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. *arXiv preprint arXiv:2103.03874*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Wing Lian, Bleys Goodson, Eugene Pentland, Austin Cook, Chanvichet Vong, and "Teknium". 2023. Openorca: An open dataset of gpt augmented flan reasoning traces. *HuggingFace repository*.
- Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. Wildbench: Benchmarking llms with challenging tasks from real users in the wild. *arXiv preprint arXiv:2406.04770*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. pages 26296–26306.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The flan collection: Designing data and methods for effective instruction tuning. pages 22631–22648.
- 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.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2021. Cross-task generalization via natural language crowdsourcing instructions. *arXiv preprint arXiv:2104.08773*.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint arXiv:2306.02707*.
- OpenAI. 2023. Gpt-4 technical report.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susanah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.
- Nazneen Rajani, Lewis Tunstall, Edward Beeching, Nathan Lambert, Alexander M. Rush, and Thomas Wolf. 2023. No robots. *Hugging Face repository*.

- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE.
- Andy Rosenbaum, Saleh Soltan, Wael Hamza, Yannick Versley, and Markus Boese. 2022. Linguist: Language model instruction tuning to generate annotated utterances for intent classification and slot tagging. *arXiv preprint arXiv:2209.09900*.
- Shivalika Singh, Freddie Vargus, Daniel Dsouza, Börje F Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura OMahony, et al. 2024. Aya dataset: An open-access collection for multilingual instruction tuning. *arXiv preprint arXiv:2402.06619*.
- Feifan Song, Bowen Yu, Hao Lang, Haiyang Yu, Fei Huang, Houfeng Wang, and Yongbin Li. 2024. Scaling data diversity for fine-tuning language models in human alignment. *arXiv preprint arXiv:2403.11124*.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Zhiqing Sun, Yikang Shen, Qinhong Zhou, Hongxin Zhang, Zhenfang Chen, David Cox, Yiming Yang, and Chuang Gan. 2024. Principle-driven self-alignment of language models from scratch with minimal human supervision. *Advances in Neural Information Processing Systems*, 36.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models.
- Shuhe Wang, Beiming Cao, Shengyu Zhang, Xiaoya Li, Jiwei Li, Fei Wu, Guoyin Wang, and Eduard Hovy. 2023a. Sim-gpt: Text similarity via gpt annotated data. *arXiv preprint arXiv:2312.05603*.
- Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Chandu, David Wadden, Kelsey MacMillan, Noah A Smith, Iz Beltagy, et al. 2023b. How far can camels go? exploring the state of instruction tuning on open resources. *Advances in Neural Information Processing Systems*, 36:74764–74786.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022a. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al. 2022b. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. *arXiv preprint arXiv:2204.07705*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- 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*.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhengguo Li, Adrian Weller, and Weiyang Liu. 2023. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*.
- Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng,

Weizhu Chen, and Tuo Zhao. 2023a. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*.

Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023b. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.

Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. Wildchat: 1m chatgpt interaction logs in the wild. *arXiv preprint arXiv:2405.01470*.

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 Algorithm For Greedy Packing

B Experimental Setups

B.1 Training Datasets

The development of packing methods was primarily aimed at maximizing hardware resource utilization and minimizing training duration. To demonstrate these training differences between packing and padding, we analyze four SFT datasets with different sizes:

WildChat (GPT-4). WildChat (Zhao et al., 2024) is a corpus comprising roughly 652k real-world interactions between users and ChatGPT, noted for its variety of languages and diverse user prompts. This dataset was created by providing users free access to ChatGPT and GPT-4 in return for their consent to collect chat histories. WildChat (GPT-4) is the smallest dataset in our experiments consisting of approximately 69k real-world interactions, selected specifically to include interactions between users and the GPT-4 model.

Algorithm 1: Greedy Packing Training Sequences.

```

Data:  $s_1, \dots, s_N$ 
% Concatenated training sequences
Result:  $S_1, \dots, S_M$ 
% Packed sequences
% Sort training sequences based on their lengths
1 SORT ( $s_1, \dots, s_N$ )
% Initialize the max input length of the pre-trained model
2  $MaxLength \leftarrow$  The max input length of the pre-trained model
% Initialize the index of the packed sequence  $S_j$ 
3  $j \leftarrow 1$ 
4 for  $i = N, \dots, 1$  do
% Skip the training sequence that has been visited
5 if  $i$  has not been visited then
6   if length of  $(S_j + s_i) \leq MaxLength$  then
% Pack the training sequence  $s_i$  into  $S_j$ 
7      $S_j \leftarrow S_j + s_i$ 
% Mark the training sequence  $i$  as the visited state
8     Visited  $i$ 
9   else
10    if length of  $S_j \neq 0$  then
% Skip the packed sequence that is longer than the max input length of the pre-trained model
11       $j \leftarrow j + 1$ 
12    end
% Initialize the packed sequence  $S_j$  with the training sequence  $s_i$ 
13     $S_j \leftarrow s_i$ 
14    Visited  $i$ 
15  end
16 end
17 end

```

TULU. TULU (Wang et al., 2023b) is a dataset consisting of around 326k conversations, sourced both from real-world interactions between users and open large LLMs as well as from manually annotated dialogues. As a synthetic dataset, TULU (Wang et al., 2023b) aims to combine the benefits of various open resources, enhancing the performance of models fine-tuned on it to deliver the highest general performance.

WildChat (Full). WildChat (Full) includes the entire 652k training conversations from the WildChat (Zhao et al., 2024) corpus. Utilizing such a large dataset allows us to confirm that the performance differences between padding and packing methods are statistically significant and not merely random fluctuations. Additionally, it provides an opportunity to assess the scalability and consistency of the padding and packing methods as the dataset is processed over time.

Open-source 1M. The larger the dataset, the more reliable the conclusions that can be drawn from the experiments, particularly in terms of how each method handles memory and computational resources at varying scales. To facilitate this analysis, we create a large data mixture named "open-source 1M", which consists of approximately 1.2M conversations collected from several high-quality open resources such as ShareGPT (Chiang et al., 2023), FLAN V2 (Longpre et al., 2023), Alpaca (Taori et al., 2023), among others. A complete list of these resources is detailed in Table 4.

C Chat Template.

Following Dubey et al. (2024), we format all datasets to follow a chat template to unify the varied styles and formats of the instruction datasets:

As highlighted in red, we add special tokens “<|start_header_id|>user<|end_header_id|>” and “<|start_header_id|>assistant<|end_header_id|>” before instructions and answers separately, and an end flag “<|eot_id|>” at the end of each instruction and answer, which at inference time, will guide the model to stop generating responses.

The diagram illustrates the chat template structure. It shows a sequence of messages: a user message (blue bubble) and an assistant message (green bubble). Each message is preceded by a header token (highlighted in red) and followed by an end-of-turn token (highlighted in red). The user message is: <|start_header_id|>user<|end_header_id|> Resistance is due to collision with protons, and pretty much everything contains protons. So technically is everything a resistor? (Or at least, can anything be a resistor?) <|eot_id|>. The assistant message is: <|start_header_id|>assistant<|end_header_id|> Actually, there are lots of materials which don't contain protons outside of atomic nuclei — e.g. steel, glass, oxygen — but all these do have resistance. Dominant factors determining resistance vary from material to material: these can be scattering of electrons/holes by motion of atomic nuclei (i.e. by phonons), by point-like defects in the crystals, etc. Scattering by protons you mention is a very specific case of scattering by impurities (crystal defects). <|eot_id|>. This pattern repeats for a second user and assistant message pair.

C.1 Training Details

	LLaMA-3-8B		LLaMA-3-70B	
	padding	packing	padding	packing
Batch (per GPU)	2	2	1	1
Gradient Acc	2	2	2	2
Learning Rate	1e-5	1e-5	1e-5	1e-5
Deepspeed	Stage-3	Stage-3	Stage-3	Stage-3
Max Seq Length	4096	4096	4096	4096
Warmup Ratio	0.2	0.2	0.2	0.2
Epochs	4	4	3	3
Offload Optimization	No	No	Yes	Yes
Flash-Attention (Dao et al., 2022)	Yes	Yes	Yes	Yes

Table 5: Training parameters for our experiments.

C.2 Evaluation Setups

Following the argument in (Wang et al., 2023b) that general-purpose models should be able to perform some core tasks before they can generalize to satisfy various practical needs, we first assess the core capabilities of our fine-tuned models using a set of specific benchmarks. Subsequently, we employ evaluations based on GPT-4 to gauge their proficiency in following instructions and aligning with human preferences.

C.2.1 Specific Benchmarks

We evaluate our models on the following benchmarks:

Dataset	# Instance	# Generator
Aya (Singh et al., 2024)	202K	Human
ChatArena (Zheng et al., 2024)	33K	Open LLMs
LIMA (Zhou et al., 2024)	1K	Human
MetaMathQA (Yu et al., 2023)	395K	GPT
No Robots (Rajani et al., 2023)	9.5K	Human
ShareGPT (Chiang et al., 2023)	114K	GPT
UltraChat 200K (Ding et al., 2023)	200K	GPT and Human
WildChat (GPT-4) (Zhao et al., 2024)	69k	GPT
Evol-Instruct (Xu et al., 2023)	143K	GPT
FLAN (Longpre et al., 2023)	100K	Human-LLMs Mixtures
Alpaca GPT-4 (Peng et al., 2023)	20K	GPT
Code Alpaca (Chaudhary, 2023)	20K	Model Self-generation
OpenOrca (Lian et al., 2023)	30K	GPT-4
Open-source 1M	1.3M	Mixture

Table 4: Details of the collected Open-source 1M dataset.

MMLU. Massive Multitask Language Understanding (MMLU) (Hendrycks et al., 2020) consists of 14079 questions covering 57 tasks including elementary mathematics, US history, computer science, law, and more. The wide range of subjects and complex questions make MMLU suitable for testing the model’s language comprehension and decision-making capabilities.

MATH and GSM8K. MATH (Hendrycks et al., 2021) and GSM8K (Cobbe et al., 2021) are two distinct mathematical datasets utilized for evaluating different aspects of model capabilities. The MATH (Hendrycks et al., 2021) dataset comprises 12,500 complex competition-level mathematics problems, primarily designed to assess the ability of models to tackle challenging and advanced mathematical questions typically encountered at the college level. Conversely, the GSM8K (Cobbe et al., 2021) dataset contains 8,500 high-quality elementary school math problems, aimed at testing the basic mathematical reasoning abilities of models.

BBH. BBH, short for BIG-Bench Hard (Suzgun et al., 2022), is a subset of the BIG-Bench (Srivastava et al., 2022) dataset comprising 23 challenging tasks. These tasks were selected because they consistently proved too difficult for current large language models to handle effectively. Requiring complex, multi-step reasoning, the BBH dataset is primarily utilized to assess the general reasoning capabilities of models, testing their ability to

navigate and solve intricate problems.

HumanEval. HumanEval (Chen et al., 2021) consists of 164 programming problems, including language comprehension, algorithms, and simple mathematics, with some comparable to simple software interview questions. The primary purpose of this dataset is to assess the ability of models to generate correct programs based on provided docstrings.

IFEval. IFEval (Zhou et al., 2023) consists of 500 prompts, each containing specific instructions like "write an article with more than 800 words" or "enclose your response in double quotation marks." This dataset is used to test the ability of large language models to accurately follow given instructions.

C.2.2 Evaluations Based on GPT-4

While human-based evaluation provides important insights into user preferences, it suffers from significant drawbacks like high labor costs and lack of real-time feedback. To address these limitations, we employ WildBench (Lin et al., 2024), an automated evaluation framework based on GPT-4. WildBench consists of 1,024 tasks manually selected from over one million human-chatbot conversation logs. It employs advanced LLMs (e.g., GPT-4-turbo) alongside specifically tailored checklists to systematically evaluate models’ outputs and provide structured explanations supporting scoring and comparisons.

For settings, we use WildBench-v2⁴ as the test set and gpt-4o-2024-05-13⁵ as the annotator. We use the OpenCompass toolkit (Contributors, 2023), which is a one-stop platform for large model evaluations, and official prompts to make sure our results can be comparable to those on the open leaderboard⁶.

⁴To safeguard against the leakage of test data, WildBench periodically releases new versions of its test set. WildBench-v2 is the version in: <https://huggingface.co/datasets/WildEval/WildBench-V2>

⁵Most advanced GPT-4 model, which can be found in: <https://platform.openai.com/docs/models/gpt-4o>

⁶<https://huggingface.co/spaces/allenai/WildBench>