

---

# Adam-mini: Use Fewer Learning Rates To Gain More

---

Yushun Zhang<sup>\*12</sup> Congliang Chen<sup>\*12</sup> Ziniu Li<sup>12</sup> Tian Ding<sup>2</sup>  
Chenwei Wu<sup>3</sup> Yinyu Ye<sup>4</sup> Zhi-Quan Luo<sup>12</sup> Ruoyu Sun<sup>12</sup>

## Abstract

We propose Adam-mini, an optimizer that achieves on-par or better performance than AdamW with 45% to 50% less memory footprint. Adam-mini reduces memory by cutting down the learning rate resources in Adam (i.e.,  $1/\sqrt{v}$ ). We find that  $\geq 90\%$  of these learning rates in  $v$  could be harmlessly removed if we (1) carefully partition the parameters into blocks following our proposed principle on Hessian structure; (2) assign a single but good learning rate to each parameter block. We further find that, for each of these parameter blocks, there exists a single high-quality learning rate that can outperform Adam, provided that sufficient resources are available to search it out. We then provide one cost-effective way to find good learning rates and propose Adam-mini. Empirically, we verify that Adam-mini performs on par or better than AdamW on various language models sized from 125M to 7B for pre-training, supervised fine-tuning, and RLHF. The reduced memory footprint of Adam-mini also alleviates communication overheads among GPUs, thereby increasing throughput. For instance, Adam-mini achieves 49.6% higher throughput than AdamW when pre-training Llama2-7B on  $2\times$  A800-80GB GPUs, which saves 33% wall-clock time for pre-training. Our implementation of Adam-mini is available at <https://github.com/zyushun/Adam-mini>.

## 1. Introduction

Adam(W) (Kingma & Ba, 2014; Loshchilov & Hutter, 2017) has become the de-facto optimizer for training large language models (LLMs) (e.g., (Vaswani et al., 2017; Achiam et al., 2023; Touvron et al., 2023; Team et al., 2023)). De-

---

<sup>\*</sup>Equal contribution <sup>1</sup>The Chinese University of Hong Kong, Shenzhen, China <sup>2</sup>Shenzhen Research Institute of Big Data <sup>3</sup>Duke University <sup>4</sup>Stanford University. Correspondence to: Ruoyu Sun <sunruoyu@cuhk.edu.cn>.

spite its superior performance, Adam is expensive to use. Specifically, Adam requires the memory for its optimizer states: the first-order momentum  $m$ , and the second-order momentum  $v$ . These in total take at least  $2\times$  the memory of the model size<sup>1</sup>. This memory consumption has become a major burden in LLM training. For instance, to train a 7B model, Adam alone requires about 56 GB per card for  $m$  and  $v$ , and with the gradients included, a total of 86 GB is needed. This is expensive even for cutting-edge graphics cards (e.g., A100-80GB). To support the high-memory algorithm, CPU-offload and sharding (Rajbhandari et al., 2020) must be used in practice, which unfortunately increases the latency and slows down the training (Rajbhandari et al., 2021). The situation further deteriorates when training larger models like PaLM with 540B parameters (Chowdhery et al., 2023). In this case, Adam alone occupies more than 50 GPUs and becomes a major overhead for pre-training.

It is intriguing to design effective optimizers that require less memory. **First**, the reduction in memory can ease the burden of CPU offloading and alleviate the need to shard model parameters. Both of these reduce the communication among GPUs and CPUs, which in turn, can enhance the throughput and accelerate the training process. **Second**, it allows practitioners to use fewer GPUs to train a model with a desired size, leading to substantial savings in both cost and energy. **Third**, it lowers the threshold of training LLMs and encourages participation from more researchers with limited GPU resources.

It is challenging to modify Adam without sacrificing its performance. One primary reason is that we still lack understanding of the role of Adam’s  $m$  and  $v$  (Zhang et al., 2020; Kunstner et al., 2023). It remains uncertain which components in Adam are indispensable for superior performance, and conversely, which components could be re-designed or improved. One notable attempt is Adafactor (Shazeer & Stern, 2018), which cuts down memory by low-rank factorization on  $v$ . However, Adafactor is widely reported to suffer degenerated performance in LLM training (e.g. (Luo et al., 2023)). One possible reason is that the current  $v$  in Adam is crucial for effectiveness and cannot be reduced

---

<sup>1</sup>We restate the update rules of Adam and AdamW in Appendix D.

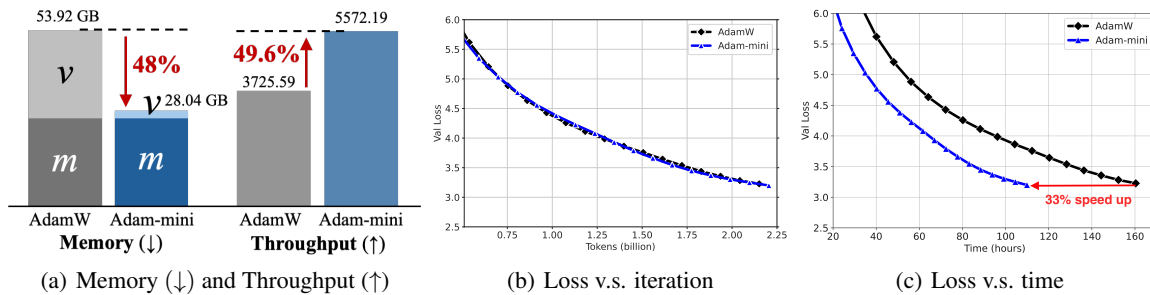


Figure 1: Results for Llama2-7B pre-training. (a) Adam-mini takes less memory and can reach higher throughput (# tokens per second). The throughput is tested on  $2 \times$  A800-80GB GPUs. (b, c) Adam-mini performs on-par with AdamW, but takes 33% less time for processing the same # tokens.

without compromising performance. This is possible as most existing Adam variants that attempt to modify  $v$  to varying extents have been reported to perform worse than Adam (Orabona, 2020). Another possible reason is that there is potential to cut down  $v$ , but Adafactor did not use the most suitable way: matrix factorization is a generic approach that could be applied broadly, but it does not leverage much problem-specific structure, thus it does not work well on specific neural-net tasks.

In this work, we find it is possible to significantly reduce the usage of  $v$ . Currently, Adam’s  $v$  assigns an individual learning rate for each parameter, i.e.,  $i$ -th parameter receives learning rate  $\frac{\eta}{\sqrt{v_i}}$ , where  $v_i$  is the  $i$ -th component of  $v$ . For a billion-sized model, Adam designs billions of learning rates. We argue that it is possible to achieve on-par or better performance with much fewer learning rate resources. We notice that Transformer’s Hessian has a near-block-diagonal structure consisting of different sizes of dense sub-blocks. We find that, for each of these dense sub-blocks, there exists a single high-quality learning rate that outperforms Adam, provided that we have enough resources to search it out. Since the number of dense sub-blocks is much fewer than the number of parameters, our findings imply that it is possible to achieve good performance with much fewer learning rates. The remaining question is how to find them efficiently.

We then propose a cheap and simple way to find good learning rates that are sufficient to perform on-par or better than Adam. We introduce the proposed design principle here: we first partition the gradient vector into  $B$  sub-vectors according to the dense Hessian sub-blocks, and call it  $g_b$  for  $b = [B]$ . For each  $g_b$ , we calculate the quantity below.

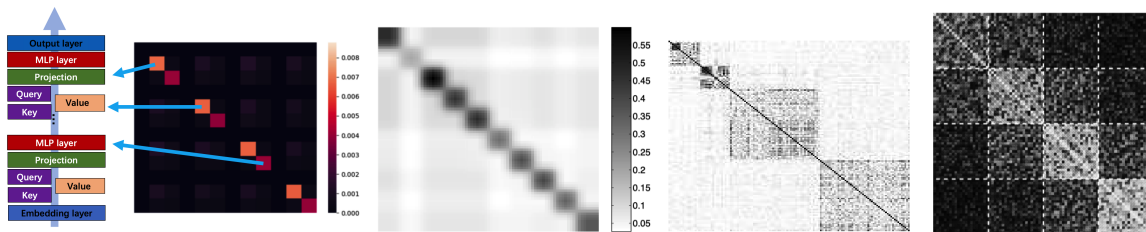
$$v_b = (1 - \beta_2) * \text{mean}(g_b \circ g_b) + \beta_2 * v_b, \quad b = 1, \dots, B$$

We then use  $\eta/\sqrt{v_b}$  as the learning rate for the parameters associated with  $g_b$ . We call the corresponding method Adam-mini. As a result, Adam-mini changes almost all Adam’s  $v$  to a handful of scalars and thus significantly reduce the memory cost. We summarize our main contribution as follows.

- **New optimizer.** We propose a new optimizer called

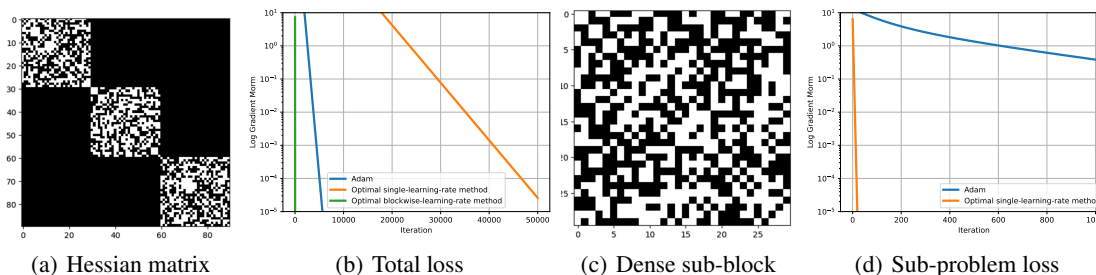
Adam-mini. First, Adam-mini partitions the model parameters based on the principle we established upon the Hessian structure. Then, it chooses a single learning rate for each block using the average of Adam’s  $v$  in that block. Adam-mini has the following advantages.

- **Lightweightness:** By design, Adam-mini largely reduces the number of learning rates used in Adam. For mainstream LLMs, Adam-mini could cut down  $\geq 90\%$  proportion of Adam’s  $v$ , which saves 45% to 50% of the memory cost of Adam.
- **Effectiveness:** Despite the memory cut down, we empirically verify that Adam-mini performs on par or even better than AdamW on various language models sized from 125M to 7B, including pre-training, supervised fine-tuning (SFT), and reinforcement learning from human feedback (RLHF). Adam-mini also achieves better performance on non-LLM tasks such as training diffusion models, vision models, and graph neural networks.
- **Efficiency:** Adam-mini can reach higher throughput than AdamW. We observe that Adam-mini reaches 49.6% higher throughput of AdamW when pre-training Llama2-7B on  $2 \times$  A800-80GB, which saves 33.1% wall-clock time for pre-training. The efficiency comes from two factors. First, Adam-mini does not introduce extra computation in per-step updates. Second, the memory cut-down allows larger batch sizes per GPU, and at the same time, it eases the burden of communication among GPUs, which is usually a major overhead.
- **Partition principle.** A key component in Adam-mini is the strategy for parameter partition. We propose to partition parameters based on the smallest dense sub-block in Hessian. This principle can apply to generic problems: for problems with block diagonal structure, we find that more learning rates do not necessarily bring extra gain. In particular, for the problem associated with each dense sub-block, a single (but good) learning rate suffices to bring better performance.
- **Hessian structure of Transformers.** We empirically apply the above principle to Transformers. We find that



(a) Transformer from (Zhang et al., 2024), (b) MLP from (Collobert, 2004), (c) MLP from (Roux et al., 2007), (d) MLP from (Martens & Grosse, 2015)

Figure 2: The Hessian of neural nets have near-block-diagonal structure. This is widely reported in the literature on Transformers (a) and various multi-layer perceptrons (MLPs) (b)(c)(d).



(a) Hessian matrix

(b) Total loss

(c) Dense sub-block

(d) Sub-problem loss

Figure 3: **(a)**: The Hessian of a three-block random quadratic problem. **(b)**: Training curves for the problem associated with the full Hessian in (a). The optimal single (blockwise) learning rate is chosen based on the full (blockwise) Hessian in (a). **(c)**: The 1st dense Hessian sub-blocks in (a). **(d)**: Training curves for the new problem associated with the Hessian in (c).

Transformer Hessian’s smallest dense blocks are: Query and Key (by heads); Value and Projection (as a whole); and MLP (by layer). We emphasize that our Hessian-based partition principle is crucial for good performance. In contrast, naive or default partitions would oversimplify the problem and cause training instability on LLMs.

## 2. Method

### 2.1. Motivations and Observations

Now we discuss our observations that motivate the design of Adam-mini. We start by investigating the role of Adam’s  $v$  and explore possibilities for improvement. In Adam,  $v$  provides an individual learning rate for each parameter, i.e.,  $i$ -th parameter receives the learning rate  $\frac{\eta}{\sqrt{v_i}}$ , where  $v_i$  is the  $i$ -th component of  $v$ . Recently, Zhang et al. (2024) pointed out that such design is crucial because Transformers need different learning rates for different blocks. They provide two pieces of evidence: First, the Hessian of Transformers and various neural nets are near-block-diagonal (restated in Figure 2). Second, for Transformers, each block has a dramatically different eigenvalue distribution (restated in Appendix D.2). As such, Transformers need different learning rates for different blocks to handle the eigenvalue heterogeneity. This could be provided by Adam’s  $v$ .

The findings in (Zhang et al., 2024) suggest that it is necessary to use a different learning rate for *each block*. Nonetheless, Adam does much more than that: it assigns an indi-

vidual learning rate not just for each block, but for *each parameter*. Note that the number of parameters (could be billions) is much larger than the number of blocks (usually hundreds). This begs the question:

Is it necessary to use an individual learning rate for each parameter? If not, how much can we save?

We first explore this question on generic optimization problems and then discuss Transformers. Following the findings in Figure 2, we conduct a case study on random quadratic minimization problem with block-diagonal Hessian. We visualize the Hessian in Figure 3 (a) and relegate the detailed setup in Appendix E.2. We conduct the coordinate-wise learning-rate method, i.e., Adam, and the single-learning-rate method, i.e., gradient descent (GD). We choose quadratic minimization because the optimal learning rate has a close form. We have the following findings.

- **(1)**: as shown in Figure 3 (a) and (b), Adam outperforms the optimal single-learning-rate method. This is expected since Adam deploys different learning rates to different parameters.
- **(2)**: as shown in Figure 3 (c) and (d), we consider a new problem whose Hessian is a dense sub-block of (a). We consider the optimal single learning-rate method for this new problem and find it outperforms Adam, even though Adam assigns much more learning rates. Similar phenomena apply to all the three sub-blocks of (a).

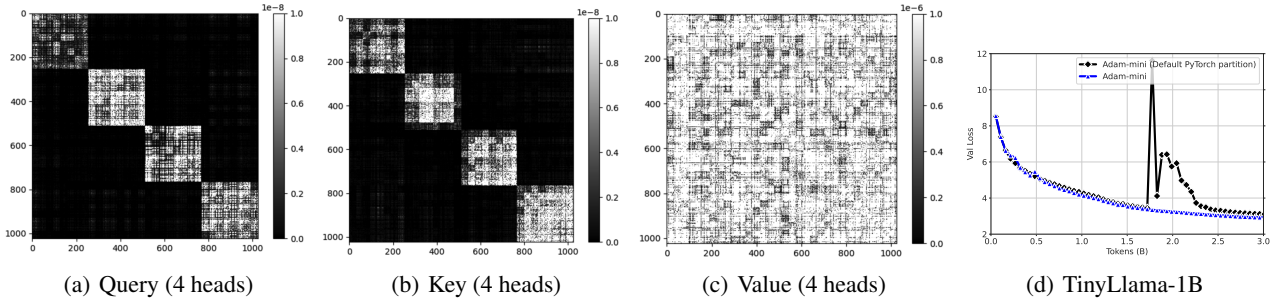


Figure 4: (a,b,c): The Hessian of different parameter blocks in a 4-layer Transformer at initialization. More figures in other blocks and detailed setup are shown in Appendix D.3 and E.2. (d) Training curves on TinyLlama-1B. When using the PyTorch default partition, Adam-mini could suffer loss spikes. The spike disappears when we change the partition strategy to **Algorithm 2**.

- **(3):** If we collect these optimal learning rates in (2) and apply them to a “blockwise” version of GD, it would be faster than Adam on the original problem (the green line in Figure 3 (b)).

In summary, for generic problems with block-diagonal Hessian, more learning rates do not necessarily bring extra gain. In particular, **for each dense sub-block, a single (but good) learning rate suffices to bring better performance** than using tens or hundreds more. Why would this happen? One possible explanation is the following. Adam can be viewed as a diagonal preconditioned method, but it may not be a good preconditioner and thus cannot effectively reduce the condition number of the dense sub-matrix (Forsythe & Straus, 1955; Young, 1954; Sun & Ye, 2021).

We also observe the similar phenomena on Transformers. For a small Transformer, we find that: block-wise learning rate methods can outperform Adam, given sufficient grid search to find the good learning rates. Due to the space constraint, we relegate these results to Appendix A.

To summarize all the phenomena above, we find that: for Transformers, it is possible to reach similar or better performance with much fewer learning rates than Adam. The remaining issue is how to find them without grid-search. Following this finding, in the next part, we propose a simple and cost-effective method called Adam-mini, which could bring comparable or even better performance than Adam, but with 90% to 99% fewer learning rates.

## 2.2. Proposed Method: Adam-mini

Based on the discussion above, we propose Adam-mini in **Algorithm 1**. Adam-mini aims to cut down the learning rate resources in Adam without laboriously grid-searching the learning rates as in Section 2.1. Adam-mini consists of two steps. **Step 1** is only required at the initialization.

**Step 1-1.** We partition the model parameters into blocks. For Transformers, we use **Algorithm 2**: “Partition for Trans-

formers”, which partitions all Queries and Keys by heads and uses the default PyTorch partition for the rest. For other networks, we use the default PyTorch partition, and we call it **Algorithm 3**: “Partition for non-Transformers”. We discuss its design principle later in Section 2.3.

**Step 1-2.** We choose `embd_blocks` using **Algorithm 4**. For Transformers, it includes the embedding layer and output layer. For other networks, no parameters will be chosen.

**Step 2.** For each parameter block that is outside the `embd_blocks`, we use a single learning rate. To efficiently choose a suitable learning rate in each block, Adam-mini simply replaces  $g \circ g$  in vanilla Adam by its mean value. We adopt the moving average on these mean values as in Adam.

**A simple example.** We use a simple example to illustrate the key design of Adam-mini. For a problem with 5 parameters  $w \in \mathbb{R}^5$ , Adam and Adam-mini both perform  $w = w - u \circ m$ , where  $m$  is the 1st-order momentum and  $u$  has different forms as follows:

- For Adam:  $u_{\text{Adam}} = \left( \frac{\eta}{\sqrt{v_1}}, \frac{\eta}{\sqrt{v_2}}, \frac{\eta}{\sqrt{v_3}}, \frac{\eta}{\sqrt{v_4}}, \frac{\eta}{\sqrt{v_5}} \right)$ .
- For Adam-mini: suppose the partition is **(1, 2, 3)** and **(4, 5)** then

$$u_{\text{mini}} = \left( \frac{\eta}{\sqrt{v_{123}}}, \frac{\eta}{\sqrt{v_{123}}}, \frac{\eta}{\sqrt{v_{123}}}, \frac{\eta}{\sqrt{v_{45}}}, \frac{\eta}{\sqrt{v_{45}}} \right)$$

where  $v_{123} = (v_1 + v_2 + v_3)/3$  and  $v_{45} = (v_4 + v_5)/2$ .

Note that the number of effective elements  $u_{\text{mini}}$  equals # blocks, which could be significantly smaller than that of  $u_{\text{Adam}}$ , which equals # parameters. For LLMs, we will show that this would free  $\geq 90\%$  elements in  $v$ .

**Remark on the “embd\_blocks”.** For Transformers, “embd\_blocks” refer to the embedding layer and output layer. Here, we explain why **Step 2** does not apply to these blocks. This is because these blocks do not reconcile with

the average operation in **Step 2**: By design of the embedding layer, many rows will be zero if the tokens corresponding to these rows do not appear in the current data minibatch. In other words, the only non-zero row is the row that represents the current input word (Press & Wolf, 2016). As such, taking an average over the entire embedding layer will absorb a significant amount of zeros, which leads to highly biased learning rates. A similar argument also applies to the output layer. In Figure 6 (a), we show that removing the `embd_blocks` causes training instability (see the curve labeled as “Adam-mini (`embd_blocks_removed`)”).

---

**Algorithm 1** Adam-mini in Pytorch style
 

---

```

1: Input weight-decay coefficient  $\lambda$  and
   current step  $t$ 
2: Choose param_blocks from
   Algorithm 2 or 3
3: Choose embd_blocks from
   Algorithm 4
4: for name, param in param_blocks do
5:   g = param.grad
6:   param = param -  $\eta_t * \lambda * \text{param}$ 
7:   m =  $(1 - \beta_1) * \text{g} + \beta_1 * \text{m}$ 
8:    $\hat{m} = \frac{\text{m}}{1 - \beta_1^t}$ 
9:   if param in embd_blocks then
10:    v =  $(1 - \beta_2) * \text{g} \circ \text{g} + \beta_2 * \text{v}$ 
11:   else
12:    v =  $(1 - \beta_2) * \text{mean}(\text{g} \circ \text{g}) + \beta_2 * \text{v}$ 
13:   end if
14:    $\hat{v} = \frac{\text{v}}{1 - \beta_2^t}$ 
15:   param = param -  $\eta_t * \frac{\hat{m}}{\sqrt{\hat{v} + \epsilon}}$ 
16: end for

```

---

**Algorithm 2** Partition for Transformers
 

---

```

1: param_blocks = {}
2: for name, param in parameters do
3:   if name is Query or Key then
4:     Partition param by heads as
       param[0..heads-1]
5:     for i = 0..heads-1 do
6:       param_blocks[name+i] = param[i]
7:     end for
8:   else
9:     param_blocks[name] = param
10:  end if
11: end for
12: Return param_blocks

```

---

**Algorithm 3** Partition for non-Transformers
 

---

```

1: param_blocks = {}
2: for name, param in parameters do
3:   param_blocks[name] = param
4: end for
5: Return param_blocks

```

---

**Algorithm 4** `Get_embd_blocks`


---

```

1: if Transformer then
2:   Return [Embedding layer,
           Output layer]
3: else
4:   Return []
5: end if

```

---

### 2.3. Principle for the Partition Strategy

We now discuss how to choose the parameter partition for Adam-mini. Based on the analysis in Section 2.1, a general principle is shown in **Principle 1** below. Building upon this principle, we now present the corresponding partitioning strategy in the context of neural networks.

**Principle 1:** We should partition parameters into blocks, s.t., each parameter block is associated with the smallest dense sub-blocks in Hessian.

Based on the block-diagonal structure reported in the literature (Figure 2), the default partition in PyTorch would be a reasonable candidate. In Section 3, we will show that this partition indeed works well widely on non-Transformer tasks such as ResNet, diffusion model, and graph models. We show this strategy in **Algorithm 3**: “Partition for non-Transformers”.

Unfortunately, we find that the default PyTorch partition does not always work well on Transformers. In particular, we find that Adam-mini encounters training instability on 1B models (see Figure 4 (d)). We suspect this is because the default PyTorch partition did not fully capture the Hessian structure. To find out, we explore the Hessian of a 4-layer Transformer as in Figure 8. Under the default PyTorch partition, we compute the exact Hessian at initialization for each parameter block (or equivalently, the principal sub-blocks in the whole Hessian). We report some representatives in Figure 4 and relegate the others in Appendix D. We find there are two classes of Hessian sub-blocks.

- **Class 1:** Similar to the whole Hessian, the Hessian sub-block itself further has a block-diagonal structure and consists of smaller dense matrices. This class includes `Query` and `Key`. We empirically find that the number of small dense sub-blocks equals the number of heads in multi-head attention.
- **Class 2:** The Hessian sub-blocks have dense structures and cannot be further partitioned into smaller ones. This class includes `Value`, attention projection, and MLP layers. Note that the Hessian structure of `Value` is different from that of `Query` and `Key`, although they all consist of four heads. This is probably because `Value` is positioned outside the softmax operator in the self-attention design, while `Query` and `Key` are not.

Based on above findings, we find the default Pytorch par-

tion is indeed not the best fit for Transformers: Querys and Keys can be further partitioned into different blocks by head. By **Principle 1**, different heads should belong to different blocks. Intuitively, different heads play a different role in understanding tokens, thus it seems reasonable that they require distinct learning rates.

We then introduce the resulting **Algorithm 2**: “Partition for Transformers”, which partitions Querys and Keys by heads. This procedure follows the principle based on the smallest Hessian sub-blocks. As shown in Figure 4 (d). This strategy indeed stabilizes the training and boosts the performance.

#### 2.4. Some Characteristics of Adam-mini

**Memory cut down.** By design, for Transformers, Adam-mini reduces the number of learning rates from the number of total paramters to the sum of the size of embedding layer, the size of output layer, and the number of blocks in other layers according to our procedure to partition Transformers. Thus, the proportion of reduced memory depends on the fraction of non-embedding parameters in the model. On Llama2-7B, this fraction is 96.2%. For mainstreamed LLMs, this fraction is usually  $\geq 90\%$  (see a summary in Table 5 in Appendix D). Note that this fraction monotonically approaches 100% as more and more layers are stacked into the architecture, e.g., for Llama3-70B, this proportion is 99.25%. Thus, Adam-mini could cut down more than 90% of  $v$ , which saves up to 45% to 50% of Adam’s memory.

Table 1: Memory cost of AdamW v.s. Adam-mini on Llama2-7B. Calculation is based on `float32`.

Model	Optimizer	Memory (GB)
GPT2-1.5B	AdamW	12.48
GPT2-1.5B	Adam-mini	6.56 (47.4% ↓)
TinyLlama-1B	AdamW	8.80
TinyLlama-1B	Adam-mini	4.92 (44.10% ↓)
Llama2-7B	AdamW	53.92
Llama2-7B	Adam-mini	28.04 (48.04% ↓)

Table 2: Throughput (↑) test on  $2 \times$  A800-80GB GPUs for Llama2-7B pre-training.  $\times$  means out of memory.

Optimizer	Micro_bs	Total_bs	Throughput (↑)
Adam-mini	4	256	<b>5572.19</b> (↑ 49.6%)
AdamW	2	256	$\times$
AdamW	1	256	3725.59

**Higher throughput.** Adam-mini can reach a higher throughput than AdamW, especially under limited hardware resources. There are two reasons. First, Adam-mini does not introduce extra computation in its update rules. The averaging operation in **Algorithm 1** incurs negligible cost. Additionally, Adam-mini significantly reduces the number of tensor-square-root and tensor-division operations in AdamW. Second, thanks to the memory cut-down, Adam-mini can support larger batch sizes per GPU, and at the same

time, it reduces the communication among GPUs, which is known to be a major overhead (Rajbhandari et al., 2021).

Owing to these properties, Adam-mini could reduce the overall time for pre-training. We provide evidence in Table 2. Here, Micro\_bs and Total\_bs denote per-device and total batch size. When pre-training Llama2-7B on  $2 \times$  A800-80GB GPUs, we find Adam-mini could reach 49.6% higher throughput than AdamW. Recall that throughput refers to # processed tokens per second, so Adam-mini could save 33.1% time on processing the same amount of tokens for pre-training.

**Has room to improve.** Adam-mini designs the learning rate for each dense Hessian sub-block using the average of Adam’s  $v$  in that block. Such a design achieves cheap computation, but it might not be optimal. Fortunately, the current design is sufficient to achieve the same or marginally better performance than Adam with much less memory. We believe there is great room to improve the learning rate design, yet it requires more fine-grained analysis of each dense Hessian sub-block. We leave it as an important future direction.

### 3. Experiments

We now verify the efficacy of Adam-mini. We primarily focus on LLM tasks, including pre-training, supervised fine-tuning (SFT) and reinforcement learning from human feedback (RLHF). We will also evaluate Adam-mini on non-LLM tasks. All LLM experiments are conducted on four NVIDIA A800-80GB GPUs and the rest are conducted on four V100 GPUs. All the experimental details are explained in Appendix E.1.

#### 3.1. Pre-training

**Setups.** We pre-train open-sourced LLMs including GPT2 series and Llama series. We train these models on mainstream English Corpus from scratch. In particular, We train GPT2 (Radford et al., 2019) series on Openwebtext (Gokaslan et al., 2019). We train TinyLlama-1B, Llama2-7B (Touvron et al., 2023) on CommonCrawl<sup>2</sup>. We train models with 2B, 3B, and 25B tokens (as recommended by Chinchilla’s law (Hoffmann et al., 2022)). We choose model configurations (e.g. context length) by their standard protocols. We compare Adam-mini with AdamW as well as popular memory-efficient methods including Adafactor (Shazeer & Stern, 2018), CAME (Luo et al., 2023), and SM3 (Anil et al., 2019). For Adafactor and SM3, we incorporate momentum with  $\beta_1 = 0.9$  to ensure a fair comparison with other methods. We tune the learning rate for all methods within the same budget and report the best performance.

**Llama series.** Figure 5 (a) shows the validation loss curve for pre-training TinyLlama-1B. The training curve

<sup>2</sup><https://commoncrawl.org/>.

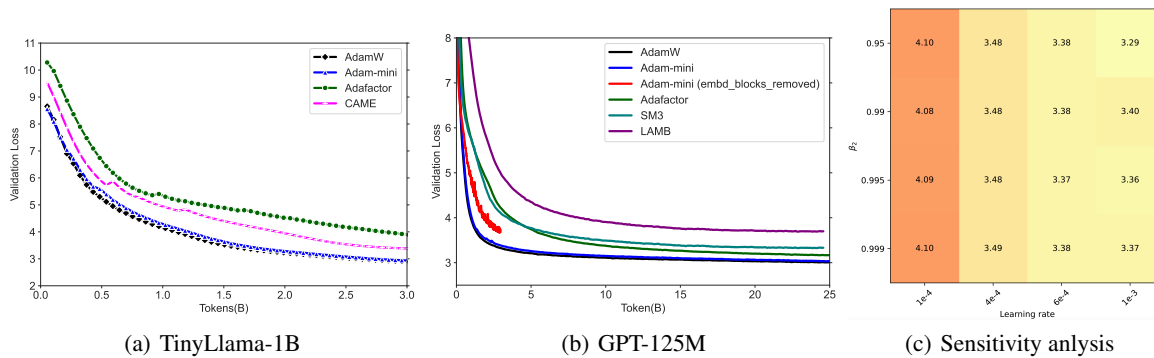


Figure 5: Training curves of (a) TinyLlama-1B. (b) GPT2-125M. Adam-mini performs on par as AdamW with less memory, while other methods perform worse on these tasks. (c): Adam-mini seems not sensitive to hyperparameters.

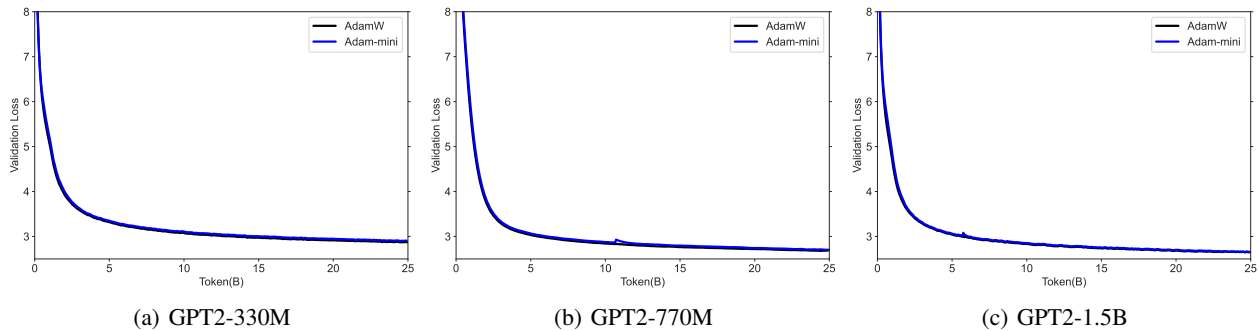


Figure 6: Training curves of GPT2 pre-training. Adam-mini reaches the same performance as AdamW with less memory.

for Llama2-7B is presented in Figure 1 (c) in Section 1. For both TinyLlama-1B and Llama2-7B, we find Adam-mini performs on par with AdamW with less memory.

**GPT2 series.** Figure 5 (b) shows the validation loss curve for GPT2-125M and Figure 6 shows the validation loss curve for GPT2 sized from 330M to 1.5B. We find that Adam-mini matches the good performance of AdamW with less memory, while other methods perform worse. In Figure 5 (b), we further consider Adam-mini (embd\_blocks\_removed), which sets the `embd_blocks` =  $\emptyset$ . That is, we use one single learning rate for the whole embedding (output) layer. We find that Adam-mini (embd\_blocks\_removed) performs poorly, as expected from the analysis in Section 2.1. We stop the trial since it shows clear unstable behavior.

**Sensitivity analysis.** On GPT2-125M pre-training task, we test the sensitivity of Adam-mini to hyperparameters. We report the validation loss after training with 2.5B tokens. As shown in Figure 5, Adam-mini seems not sensitive to hyperparameters.

### 3.2. Supervised Fine-tuning and RLHF

In this section, we evaluate the effectiveness of Adam-mini for downstream fine-tuning tasks. Specifically, we consider two representative tasks: supervised fine-tuning (SFT) and reinforcement learning from human feedback

(RLHF). We use the Llama-2-7b pretrained model (released by Meta (Touvron et al., 2023)) for our study. We use the `ultrafeedback` dataset<sup>3</sup> and implement the RLHF workflow from (Ouyang et al., 2022). Specifically, we train an SFT model with 40% of the chosen data and train a reward model using the remaining 60%. Then, we apply the reinforcement learning algorithm ReMax (Li et al., 2023), a memory-efficient alternative to PPO (Schulman et al., 2017), to optimize the preference reward.

We evaluate the alignment performance in terms of chat ability using the MT-Bench (Zheng et al., 2024), where GPT-4 assesses multi-turn chatting capabilities and assigns a score from 0 to 10 (higher is better). Our results, presented in Table 3, demonstrate that Adam-mini can outperform AdamW, even though Adam-mini uses a single learning rate and is more memory-efficient.

Table 3: GPT-4 evaluation score ( $\uparrow$ ) of SFT and RLHF on the MT-Bench (Zheng et al., 2024).

	SFT (LoRA)		SFT		RLHF	
	AdamW	Adam-mini	AdamW	Adam-mini	AdamW	Adam-mini
MT-Bench	4.23	<b>4.41</b>	5.37	<b>5.40</b>	5.54	<b>5.68</b>

<sup>3</sup><https://huggingface.co/datasets/argilla/ultrafeedback-binarized-preferences-cleaned>

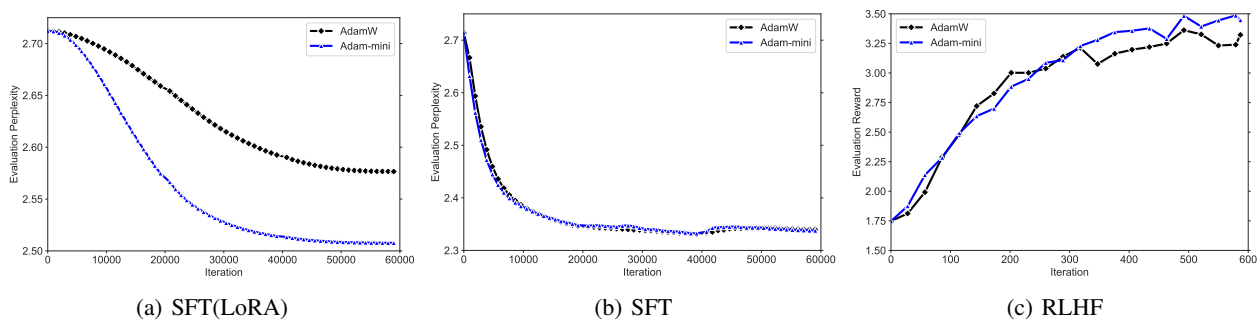


Figure 7: Training curves of SFT (LoRA), SFT and RLHF when aligning LLMs. Adam-mini reaches better performance (smaller perplexity, higher reward) than AdamW with less memory.

### 3.3. Non-LLM Tasks

We also evaluate Adam-mini on popular non-LLM tasks including vision and graph tasks. We find that Adam-mini performs on-par or better than Adam, but with less memory. We relegate these results to Appendix B.

## 4. Concluding Remarks

We proposed Adam-mini, an optimizer that saves 45% to 50% memory of Adam. We remark that there is great room to improve the design of Adam-mini: currently Adam-mini uses a simple and cost-effective way to design a learning rate for each dense Hessian sub-block, but it might not be an optimal way. We leave the development of stronger designs as a future direction.

## References

- Abdin, M., Jacobs, S. A., Awan, A. A., Aneja, J., Awadallah, A., Awadalla, H., Bach, N., Bahree, A., Bakhtiari, A., Behl, H., et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Ahn, K., Zhang, Z., Kook, Y., and Dai, Y. Understanding adam optimizer via online learning of updates: Adam is ftrl in disguise. In *Forty-first International Conference on Machine Learning*.
- Anil, R., Gupta, V., Koren, T., and Singer, Y. Memory efficient adaptive optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- Collobert, R. Large scale machine learning. Technical report, Université de Paris VI, 2004.
- Cui, G., Yuan, L., Ding, N., Yao, G., Zhu, W., Ni, Y., Xie, G., Liu, Z., and Sun, M. Ultrafeedback: Boosting language models with high-quality feedback, 2023.
- Da Silva, A. B. and Gazeau, M. A general system of differential equations to model first-order adaptive algorithms. *The Journal of Machine Learning Research*, 21(1):5072–5113, 2020.
- Dauphin, Y. N., Agarwala, A., and Mobahi, H. Neglected hessian component explains mysteries in sharpness regularization. *arXiv preprint arXiv:2401.10809*, 2024.
- Defazio, A., Mehta, H., Mishchenko, K., Khaled, A., Cutkosky, A., et al. The road less scheduled. *arXiv preprint arXiv:2405.15682*, 2024.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Dettmers, T., Lewis, M., Shleifer, S., and Zettlemoyer, L. 8-bit optimizers via block-wise quantization. In *International Conference on Learning Representations*, 2021.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- Forsythe, G. E. and Straus, E. G. On best conditioned matrices. *Proceedings of the American Mathematical Society*, 6(3):340–345, 1955.



- Ghorbani, B., Krishnan, S., and Xiao, Y. An investigation into neural net optimization via hessian eigenvalue density. In *International Conference on Machine Learning*, pp. 2232–2241. PMLR, 2019.
- Ginsburg, B., Castonguay, P., Hrinchuk, O., Kuchaiev, O., Lavrukhin, V., Leary, R., Li, J., Nguyen, H., Zhang, Y., and Cohen, J. M. Training deep networks with stochastic gradient normalized by layerwise adaptive second moments. 2019.
- Gokaslan, A., Cohen, V., Pavlick, E., and Tellex, S. Openwebtext corpus, 2019.
- Gur-Ari, G., Roberts, D. A., and Dyer, E. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Casas, D. d. L., Hendricks, L. A., Welbl, J., Clark, A., et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Jiang, K., Malik, D., and Li, Y. How does adaptive optimization impact local neural network geometry? *Advances in Neural Information Processing Systems*, 36, 2023.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016.
- Kunstner, F., Chen, J., Lavington, J. W., and Schmidt, M. Noise is not the main factor behind the gap between sgd and adam on transformers, but sign descent might be. *arXiv preprint arXiv:2304.13960*, 2023.
- Li, B., Chen, J., and Zhu, J. Memory efficient optimizers with 4-bit states. *Advances in Neural Information Processing Systems*, 36, 2024.
- Li, Z., Xu, T., Zhang, Y., Yu, Y., Sun, R., and Luo, Z.-Q. Remax: A simple, effective, and efficient method for aligning large language models. *arXiv preprint arXiv:2310.10505*, 2023.
- Liao, Z. and Mahoney, M. W. Hessian eigenspectra of more realistic nonlinear models. *Advances in Neural Information Processing Systems*, 34:20104–20117, 2021.
- Liu, H., Li, Z., Hall, D., Liang, P., and Ma, T. Sophia: A scalable stochastic second-order optimizer for language model pre-training. *arXiv preprint arXiv:2305.14342*, 2023.
- Liu, Y., Bernstein, J., Meister, M., and Yue, Y. Learning by turning: Neural architecture aware optimisation. In *International Conference on Machine Learning*, pp. 6748–6758. PMLR, 2021.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Luo, Q., Yu, H., and Li, X. Badam: A memory efficient full parameter training method for large language models. *arXiv preprint arXiv:2404.02827*, 2024.
- Luo, Y., Ren, X., Zheng, Z., Jiang, Z., Jiang, X., and You, Y. Came: Confidence-guided adaptive memory efficient optimization. *arXiv preprint arXiv:2307.02047*, 2023.
- Lv, K., Yan, H., Guo, Q., Lv, H., and Qiu, X. Adalomo: Low-memory optimization with adaptive learning rate. *arXiv preprint arXiv:2310.10195*, 2023a.
- Lv, K., Yang, Y., Liu, T., Gao, Q., Guo, Q., and Qiu, X. Full parameter fine-tuning for large language models with limited resources. *arXiv preprint arXiv:2306.09782*, 2023b.
- Malladi, S., Gao, T., Nichani, E., Damian, A., Lee, J. D., Chen, D., and Arora, S. Fine-tuning language models with just forward passes. *Advances in Neural Information Processing Systems*, 36:53038–53075, 2023.
- Martens, J. and Grosse, R. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pp. 2408–2417. PMLR, 2015.
- Orabona, F. Neural networks (maybe) evolved to make adam the best optimizer. 2020. URL <https://parameterfree.com/2020/12/06/neural-network-maybe-evolved-to-make-adam-the-best/>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

- Pan, Y. and Li, Y. Toward understanding why adam converges faster than sgd for transformers. *arXiv preprint arXiv:2306.00204*, 2023.
- Papayan, V. The full spectrum of deepnet Hessians at scale: Dynamics with sgd training and sample size. *arXiv preprint arXiv:1811.07062*, 2018.
- Papayan, V. Measurements of three-level hierarchical structure in the outliers in the spectrum of deepnet Hessians. *arXiv preprint arXiv:1901.08244*, 2019.
- Papayan, V. Traces of class/cross-class structure pervade deep learning spectra. *The Journal of Machine Learning Research*, 21(1):10197–10260, 2020.
- Pearlmutter, B. A. Fast exact multiplication by the Hessian. *Neural computation*, 6(1):147–160, 1994.
- Press, O. and Wolf, L. Using the output embedding to improve language models. *arXiv preprint arXiv:1608.05859*, 2016.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rajbhandari, S., Rasley, J., Ruwase, O., and He, Y. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.
- Rajbhandari, S., Ruwase, O., Rasley, J., Smith, S., and He, Y. Zero-infinity: Breaking the GPU memory wall for extreme scale deep learning. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*, pp. 1–14, 2021.
- Roux, N., Manzagol, P.-A., and Bengio, Y. Topmoumoute online natural gradient algorithm. *Advances in neural information processing systems*, 20, 2007.
- Sagun, L., Bottou, L., and LeCun, Y. Eigenvalues of the Hessian in deep learning: Singularity and beyond. *arXiv preprint arXiv:1611.07476*, 2016.
- Sagun, L., Evci, U., Guney, V. U., Dauphin, Y., and Bottou, L. Empirical analysis of the Hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.
- Sankar, A. R., Khasbage, Y., Vigneswaran, R., and Balasubramanian, V. N. A deeper look at the Hessian eigen-spectrum of deep neural networks and its applications to regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 9481–9488, 2021.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shazeer, N. and Stern, M. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pp. 4596–4604. PMLR, 2018.
- Sun, R. and Ye, Y. Worst-case complexity of cyclic coordinate descent:  $O(n^2)$  gap with randomized version. *Mathematical Programming*, 185:487–520, 2021.
- Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Team, G., Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju, S., Pathak, S., Sifre, L., Rivière, M., Kale, M. S., Love, J., et al. Gemma: Open models based on Gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y., et al. Graph attention networks. *stat*, 1050(20):10–48550, 2017.
- Wang, B., Zhang, Y., Zhang, H., Meng, Q., Ma, Z.-M., Liu, T.-Y., and Chen, W. Provable adaptivity in adam. *arXiv preprint arXiv:2208.09900*, 2022.
- Wu, Y., Zhu, X., Wu, C., Wang, A., and Ge, R. Dissecting Hessian: Understanding common structure of Hessian in neural networks. *arXiv preprint arXiv:2010.04261*, 2020.
- Xie, X., Zhou, P., Li, H., Lin, Z., and Yan, S. Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. *arXiv preprint arXiv:2208.06677*, 2022.
- Yao, Z., Gholami, A., Lei, Q., Keutzer, K., and Mahoney, M. W. Hessian-based analysis of large batch training and robustness to adversaries. *Advances in Neural Information Processing Systems*, 31, 2018.
- Yao, Z., Gholami, A., Keutzer, K., and Mahoney, M. W. PyHessian: Neural networks through the lens of the Hessian. In *2020 IEEE international conference on big data (Big data)*, pp. 581–590. IEEE, 2020.

- You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C.-J. Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962*, 2019.
- Young, D. Iterative methods for solving partial difference equations of elliptic type. *Transactions of the American Mathematical Society*, 76(1):92–111, 1954.
- Zhang, G., Li, L., Nado, Z., Martens, J., Sachdeva, S., Dahl, G., Shallue, C., and Grosse, R. B. Which algorithmic choices matter at which batch sizes? insights from a noisy quadratic model. *Advances in neural information processing systems*, 32, 2019a.
- Zhang, J., He, T., Sra, S., and Jadbabaie, A. Why gradient clipping accelerates training: A theoretical justification for adaptivity. *arXiv preprint arXiv:1905.11881*, 2019b.
- Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S., Kumar, S., and Sra, S. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33:15383–15393, 2020.
- Zhang, Y., Chen, C., Shi, N., Sun, R., and Luo, Z.-Q. Adam can converge without any modification on update rules. *Advances in Neural Information Processing Systems*, 35: 28386–28399, 2022.
- Zhang, Y., Chen, C., Ding, T., Li, Z., Sun, R., and Luo, Z.-Q. Why transformers need adam: A hessian perspective. *arXiv preprint arXiv:2402.16788*, 2024.
- Zhao, J., Zhang, Z., Chen, B., Wang, Z., Anandkumar, A., and Tian, Y. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint arXiv:2403.03507*, 2024.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zheng, S. and Kwok, J. T. Blockwise adaptivity: Faster training and better generalization in deep learning. *arXiv preprint arXiv:1905.09899*, 2019.

## A. More Observations for Motivating Adam-mini

Following the findings in Section 2.1, we now conduct similar experiments for Transformers. We consider a 4-layer Transformer in the codebase<sup>4</sup>. Under the default partition by PyTorch, We randomly choose one parameter block as the “left-out” block and change the coordinate-wise learning rate to a single-learning rate counter-part. We use Adam for the rest of the blocks. We grid-search the learning rate for the left-out block and apply the cosine decay schedule. We report the best result and call this method “Adam (leave-one-out)”. Figure 8 shows that Adam (leave-one-out) can reach similar or better performance than Adam for all randomly picked left-out blocks. A similar phenomenon is also observed when we randomly leave out up to three blocks and search three learning rates. We cannot afford to leave more blocks out since the cost of grid search grows exponentially.

To summarize all the phenomena above, we find that: for Transformers, it is possible to reach similar or better performance with much fewer learning rates than Adam. The remaining issue is how to find them without grid-search. Following this finding, in the next part, we propose a simple and cost-effective method called Adam-mini, which could bring comparable or even better performance than Adam, but with 90% to 99% fewer learning rates.

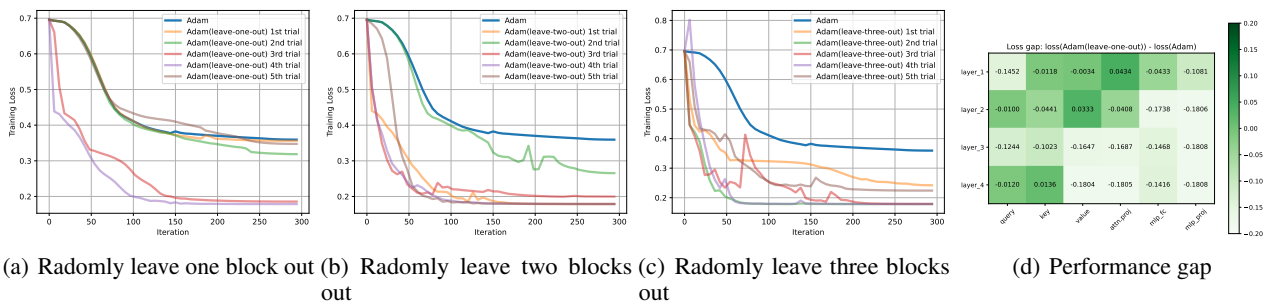


Figure 8: (a) (b) (c) Adam (leave- $x$ -out) can reach a similar or better performance than Adam for all randomly picked left-out blocks.  $x = 1, 2, 3$ . (d) The performance gap between Adam and Adam (leave-one-out) for all possible blocks. We find Adam (leave-one-out) always performs on par with Adam, and for most blocks, Adam (leave-one-out) performs better.

## B. Adam-mini for Non-LLM Tasks

Table 4: On popular Non-LLM tasks, Adam-mini performs on par or better than AdamW.

Domain	Model	Optimizer	Metric	25% steps	50% steps	75% steps	100% steps
Vision	ResNet18	AdamW	Test acc ( $\uparrow$ )	0.6149	0.6478	0.6613	0.6669
Vision	ResNet18	Adam-mini	Test acc ( $\uparrow$ )	0.6140	<b>0.6501</b>	<b>0.6629</b>	0.6667
Vision	Diffusion model	AdamW	Train loss ( $\downarrow$ )	0.0529	0.0497	0.0420	0.0394
Vision	Diffusion model	Adam-mini	Train loss ( $\downarrow$ )	<b>0.0525</b>	<b>0.0495</b>	<b>0.0416</b>	<b>0.0388</b>
Graph	GAT	AdamW	Val acc ( $\uparrow$ )	0.7277	0.7367	0.7399	0.7421
Graph	GAT	Adam-mini	Val acc ( $\uparrow$ )	<b>0.7378</b>	<b>0.7394</b>	<b>0.7403</b>	<b>0.7429</b>
Graph	GCN	AdamW	Val acc ( $\uparrow$ )	0.7347	0.7428	0.7379	0.7374
Graph	GCN	Adam-mini	Val acc ( $\uparrow$ )	<b>0.7406</b>	0.7427	<b>0.7380</b>	<b>0.7423</b>

We now evaluate Adam-mini on non-LLM tasks. Table 4 shows the results for training ResNet18 on ImageNet; Diffusion model on CelebA; Graph Convolution Net (GCN), Graph Attention Net (GAT) on OGB-arxiv. We find that Adam-mini reaches comparable or better performance than AdamW, but with less memory. Detailed setups are shown in Appendix E.1.

## C. Related works

**Understanding of Adam.** There is an active line of works trying to understand why Adam works well (Zhang et al., 2019b; Wu et al., 2020; Zhang et al., 2022; Wang et al., 2022; Pan & Li, 2023; Jiang et al., 2023; Kunstner et al., 2023; Zhang et al.,

<sup>4</sup><https://colab.research.google.com/drive/1SiF0KZJp75rUeetKOWqpsA8clMHP6jMg?usp=sharing>

2024; Ahn et al.). In contrast to these works, we point out that Adam’s  $v$  might not function at its full potential as effectively as we expected: sometimes fewer learning rates can reach the same or better results (due to the dense Hessian sub-blocks). Our findings might motivate stronger optimizers that better fit the neural-net Hessian structure.

**On the Hessian of Neural Nets.** Hessian matrix is crucial for the behaviors of gradient methods. There are several important attempts to study the Hessian of MLPs and CNNs (Collobert, 2004; Roux et al., 2007; Martens & Grosse, 2015; Sagun et al., 2016; 2017; Chaudhari et al., 2019; Pappan, 2020; Wu et al., 2020; Liao & Mahoney, 2021; Pappan, 2018; 2019; Sankar et al., 2021; Gur-Ari et al., 2018; Yao et al., 2018; Zhang et al., 2019a; Ghorbani et al., 2019; Yao et al., 2020; Dauphin et al., 2024). Inspired by these works, we explore the Hessian structure of Transformers and connect it to the behaviors of Adam. We then find room to improve and propose to slim down Adam into Adam-mini.

**Lightweight optimizers for general tasks.** There are several attempts to reduce the memory cost of Adam. Adafactor (Shazeer & Stern, 2018) and its variant CAME (Luo et al., 2023) conduct nonnegative low-rank factorization over Adam’s  $v$ . SM3 (Anil et al., 2019) is a lightweight version of AdaGrad (Duchi et al., 2011). SM3 chooses the learning rate of the  $i$ -th parameter by taking the minimal value in a certain candidate set, and each element in the candidate set is related to the maximal squared gradient under a predetermined cover. All these aforementioned methods could release almost all memory for  $v$  and save about 48% of Adam’s memory. However, we find that their performance degenerate in various experiments, while Adam-mini maintains as effective as AdamW (Figure 5).

After completing this work, we noticed two methods that share some of the ideas of Adam-mini: BAGM (Zheng & Kwok, 2019) and NovoGrad (Ginsburg et al., 2019). Both of them use block-wise or layer-wise adaptive learning rates to achieve robust performance and better generalization. We summarize their key differences with Adam-mini. BAGM partitions parameters to reach minimal-norm solutions and achieve provable robustness. Based on their theory, the resulting strategy is that: (I) they used the default PyTorch partition; (II) they did not introduce `embd.blocks`. Although these partition designs may have benefits on robustness, they overlook the Hessian structure and oversimplify the training problem (as we discussed in Section 2.2 and 2.3). As a result, these designs will lead to training instability in large-scale experiments. This is evident in our failed preliminary versions of Adam-mini in Figure 4 (d) and 5 (b). As for NovoGrad, it uses layer-wise learning rate design, which is coarser than the blockwise design in BAGM and thus would share similar drawbacks to BAGM. Further, NovoGrad introduces a different design to 1st-order momentum: instead of performing weighted-sum on the past gradients, it performs weighted-sum on “the current gradient divided by the 2nd-order momentum”. Such design is largely different from AdamW and Adam-mini. It seems unclear whether such design can work on large-scale tasks like LLMs.

In summary, these two methods have different designs and their partition strategies oversimplify the training problems. Consequently, they would cause training instability on large-scale experiments as evident in Figure 4 (d) and 5 (b). In contrast, Adam-mini carefully assigns learning rates following our proposed principle on Hessian structures. Such design principle is crucial for training stability and it works well on various LLMs including 7B models.

**Other orthogonal methods.** The idea of Adam-mini can be orthogonally combined with various existing approaches to gain more merits. For instance, GaLore (Zhao et al., 2024) is a new memory-efficient optimizer for LLMs. Given a gradient matrix  $g$ , GaLore calculates a low-rank gradient estimator  $\hat{g}$  and then calculates  $m$  and  $v$  based on this  $\hat{g}$ . Adam-mini can potentially be combined with GaLore to reach further memory reduction on  $v$ . The combined method, e.g., “GaLore-mini”, can further reduce about 40% memory on GaLore and about 81% on AdamW in total.<sup>5</sup> Additionally, GaLore-mini can ease the offload burden and enhance the throughput of GaLore, especially when training on customer-level GPUs with limited memory. Sophia (Liu et al., 2023) is another recent diagonal preconditioned optimizer. Just as Adam, Sophia requires memory for  $m$  and  $v$ . It is possible to combine Adam-mini and Sophia to get “Sophia-mini”, which saves up to 50% of memory in Sophia. Sophia-mini can also enhance throughput and further speed up Sophia on wall-clock time as in Table 2.

We list more potential combinations here. LoRA (Hu et al., 2021) is a memory-efficient method for SFT tasks. This method fine-tunes the model via additive low-rank adaptors and uses Adam to update these adaptors. Note that the Adam steps in LoRA can be replaced by Adam-mini. As a result, Adam-mini brings better performance (Figure 7). In parallel to our work, BAdam (Luo et al., 2024) conducts SFT in a block-coordinate-descent (BCD) fashion. This method requires repeated Adam steps to solve the sub-problem in BCD. Similarly as in LoRA, the Adam steps in BAdam can be replaced by Adam-mini to further reduce memory. Nero optimizer (Liu et al., 2021) also cuts down the memory of Adam. It removes the 1st-order momentum and uses a neuron-specific projected gradient-style update. According to (Liu et al., 2021), their design imposes constraints on weight matrices and has the advantage of “balanced excitation and inhibition”. Such design

<sup>5</sup>These results are calculated based on (Zhao et al., 2024, Table 1). We consider Llama2-7B and  $r = 1024$  in GaLore.

can potentially be combined with Adam-mini to further boost performance. To save the memory cost for fine-tuning LLMs, MeZO (Malladi et al., 2023) uses zeroth-order methods to approximate the gradient information. It is possible to combine this idea with Adam-mini to further save memory for SFT. Adam-mini can also potentially be combined with other diagonal preconditioned methods (such as AdaGrad (Duchi et al., 2011) and Adan (Xie et al., 2022)) as well as recent schedule-free optimizers such as SchedulefreeAdamW (Defazio et al., 2024).

There are several other tricks that ease GPU memory burden but are orthogonal to optimizer design. These tricks include gradient checkpointing (Chen et al., 2016), model offloading and sharding (Rajbhandari et al., 2020; 2021), quantization (Dettmers et al., 2021; Li et al., 2024), and fused update (Lv et al., 2023a;b). Adam-mini can be implemented upon these tricks.

Finally, we discuss another popular adaptive optimizer called LAMB (You et al., 2019) (see **Algorithm 7** in Appendix D.1). LAMB might be misunderstood as a similar optimizer to Adam-mini, but actually, it is not. We emphasize that Adam-mini is *significantly different* from LAMB. First, LAMB still keeps the same coordinate-wise learning-rate design  $1/\sqrt{v}$  as in Adam. Second, in addition to this  $1/\sqrt{v}$ , LAMB further re-scales the parameters in a layer-by-layer fashion. This re-scaling design is often known as the “layer-wise learning rates”, but to be precise, it is actually an additional “layer-wise scaling” besides the “coordinate-wise learning rates  $1/\sqrt{v}$ ”. As a result, LAMB does not save memory over Adam and its overall design is quite different from Adam-mini. This is understandable because LAMB was designed for large-batch training, not for memory saving. Numerically, we find that LAMB performs worse than Adam-mini on GPT2 pre-training (Figure 5 (b)).

## D. More Results

### D.1. Preliminaries on Adam, AdamW and LAMB

We introduce Adam, AdamW, and LAMB in **Algorithm 5**, **6**, and **7**. These methods need to track  $m$  and  $v$  along the training. Both  $m$  and  $v$  are vectors of the same size as # model parameter.

---

#### Algorithm 5 Adam in Pytorch style

---

```

1: Let  $\lambda$  be the weight decay coefficient
2: for param in parameter_blocks do
3:    $g = \text{param.grad}$ 
4:   if  $\lambda > 0$  then
5:      $g = g + \lambda * \text{param}$ 
6:   end if
7:    $\text{param} = \text{param} - \eta_t * \lambda * g$ 
8:    $m = (1 - \beta_1) * g + \beta_1 * m$ 
9:    $\hat{m} = \frac{m}{1 - \beta_1^t}$ 
10:   $v = (1 - \beta_2) * g \odot g + \beta_2 * v$ 
11:   $\hat{v} = \frac{v}{1 - \beta_2^t}$ 
12:   $\text{param} = \text{param} - \eta_t * \frac{\hat{m}}{\sqrt{\hat{v} + \epsilon}}$ 
13: end for

```

---

### D.2. Preliminary results in (Zhang et al., 2024)

We here restate (Zhang et al., 2024, Figure 3). This figure shows that: for Transformers, different parameter blocks have different Hessian eigenvalue distributions, while for CNNs, the eigenvalue distributions are similar among blocks. This suggests that Transformers need different learning rates for different blocks to handle the heterogeneity in eigenvalue distributions.

### D.3. More Hessian sub-blocks in Figure 4

In Figure 4, we present the Hessian of Query, Key, Value and MLP\_proj in the 2nd layer. We now comprehensively provide Hessian in all attention and MLP blocks in all layers in the 4-layer Transformer used in Figure 4. We plot the Hessian at initialization. Experimental details can be seen in Appendix E.2. As we can see below, almost all these Hessian sub-blocks belong to either **Class 1** or **Class 2**. The only exception is the MLP\_fc block in the 4th layer, which shows a diagonal pattern.

**Algorithm 6** AdamW in Pytorch style

```

1: Let  $\lambda$  be the weight decay coefficient
2: for param in parameter_blocks do
3:    $g = \text{param.grad}$ 
4:    $\text{param} = \text{param} - \eta_t * \lambda * g$ 
5:    $m = (1 - \beta_1) * g + \beta_1 * m$ 
6:    $\hat{m} = \frac{m}{1 - \beta_1^t}$ 
7:    $v = (1 - \beta_2) * g \odot g + \beta_2 * v$ 
8:    $\hat{v} = \frac{v}{1 - \beta_2^t}$ 
9:    $\text{param} = \text{param} - \eta_t * \frac{\hat{m}}{\sqrt{\hat{v} + \epsilon}}$ 
10: end for

```

**Algorithm 7** LAMB in Pytorch style

```

1: Let  $\lambda$  be the weight decay coefficient, let  $\phi$  be a scaling function.
2: for param in all_layers do
3:    $g = \text{param.grad}$ 
4:    $\text{param} = \text{param} - \eta_t * \lambda * g$ 
5:    $m = (1 - \beta_1) * g + \beta_1 * m$ 
6:    $\hat{m} = \frac{m}{1 - \beta_1^t}$ 
7:    $v = (1 - \beta_2) * g \odot g + \beta_2 * v$ 
8:    $\hat{v} = \frac{v}{1 - \beta_2^t}$ 
9:    $r = \frac{\hat{m}}{\sqrt{\hat{v} + \epsilon}}$ 
10:   $\text{param} = \text{param} - \eta_t * \frac{\phi(\|\text{param}\|)}{\|\text{r} + \lambda * \text{param}\|} * r$ 
11: end for

```

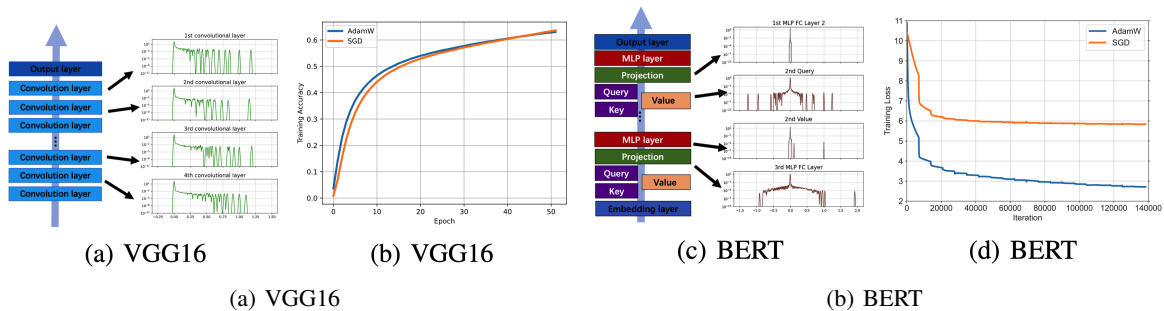


Figure 9: Figure 3 in (Zhang et al., 2024). The eigenvalues distribution are similar among blocks for CNNs, while they differ significantly across blocks for Transformers. This indicates Transformers need different learning rates for different blocks to handle the heterogeneity in eigenvalue distributions.

This begs the question: does Adam significantly outperforms single-learning-rate methods on this diagonal block? We find the answer is still no. This is because the Hessian eigenvalues in this block is concentrated. We provide evidence in Figure 14. This suggests that the optimization problem in this block is well-conditioned and single-learning-rate methods can work well. As shown in Figure 14, optimal single-learning-rate still significantly outperforms Adam, similarly as what happens in other dense blocks. Another evidence lies in Figure 8 (d): when the MLP\_fc block in the 4th layer is chosen as the left-out block, Adam (leave-one-out) achieves 0.1416 lower loss than Adam.

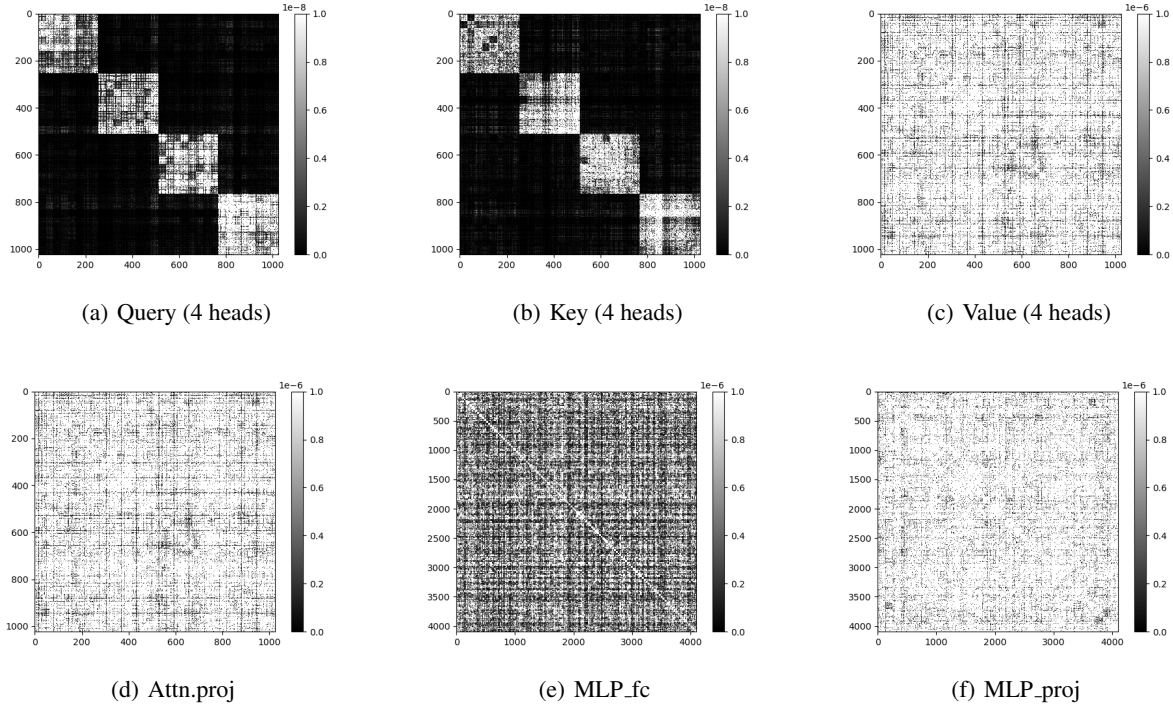


Figure 10: The initial Hessian structure of different parameter blocks in the 1st layer.



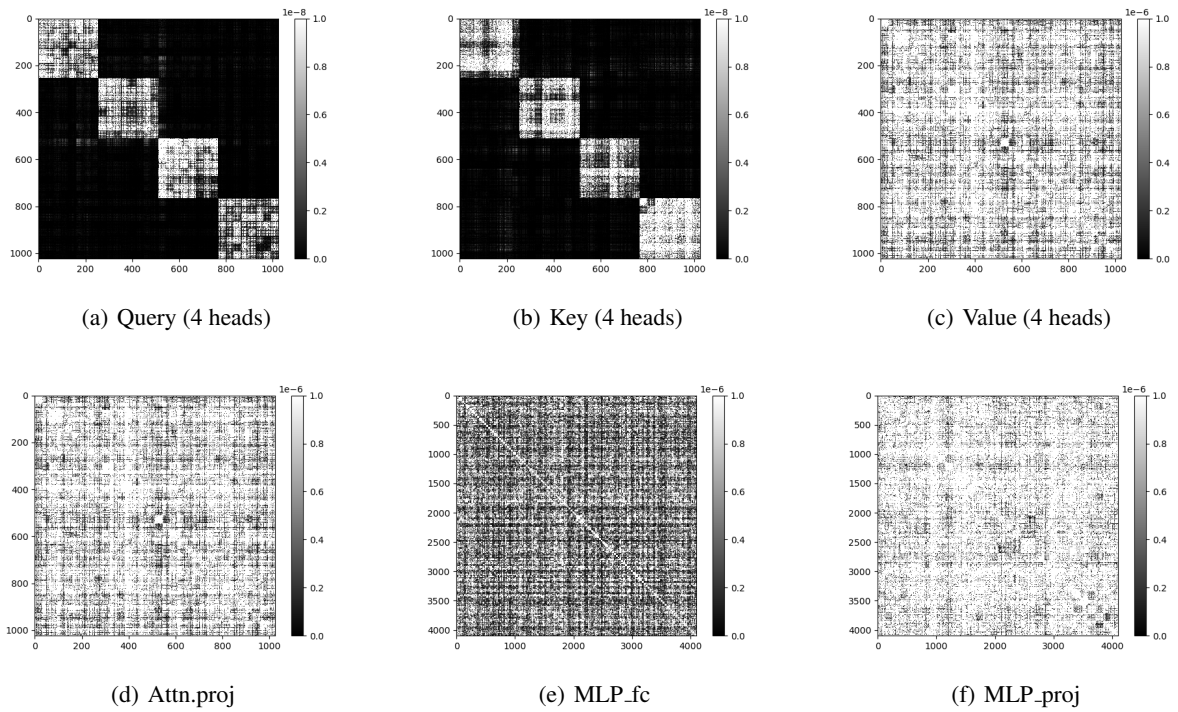


Figure 11: The initial Hessian structure of different parameter blocks in the 2nd layer.

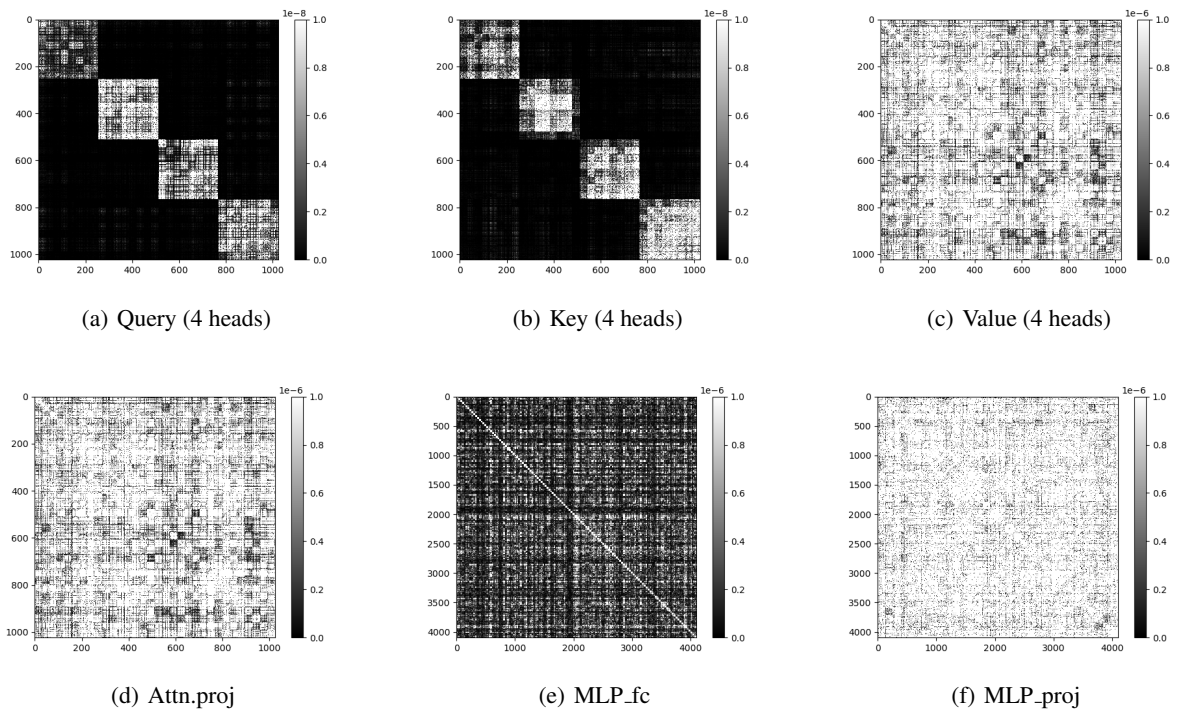


Figure 12: The initial Hessian structure of different parameter blocks in the 3rd layer.

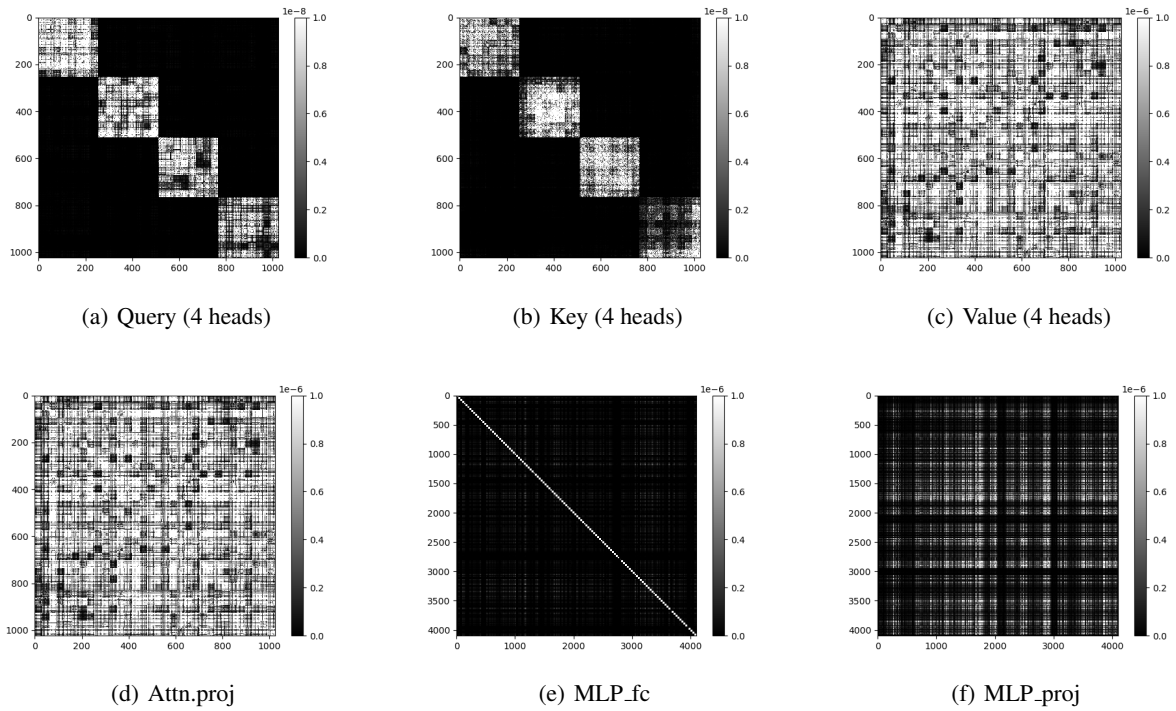


Figure 13: The initial Hessian structure of different parameter blocks in the 4th layer.

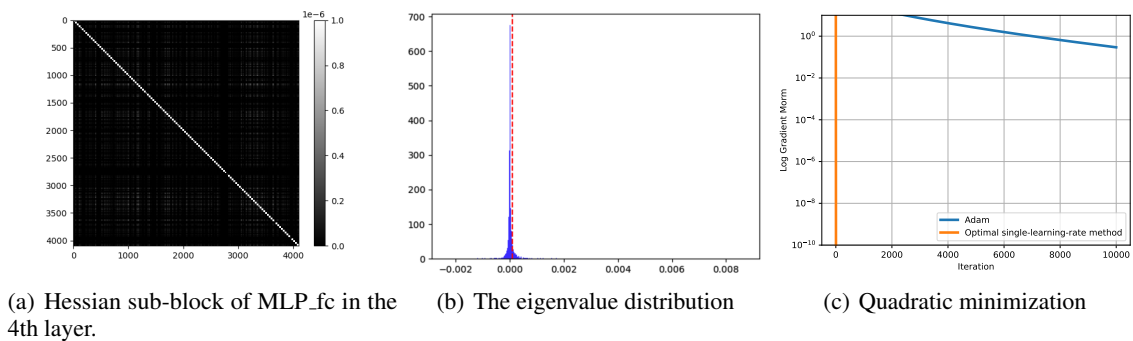


Figure 14: (a) The Hessian sub-block of MLP.fc in the 4th layer. There is a near-diagonal pattern. (b) The eigenvalues are concentrated in this Hessian sub-block. (c) When taking this matrix as the Hessian of a quadratic function, the optimal single-learning-rate method outperforms Adam.

#### D.4. Embedding Proportion in LLMs

Table 5 shows the proportion of the embedding & output parameters in mainstream LLMs including GPT2 (Radford et al., 2019), Llama series (Touvron et al., 2023), Phi-3 (Abdin et al., 2024), and Gemma (Team et al., 2024). We find that the embedding & output layer takes a fairly small proportion of the total parameters, usually  $\leq 10\%$ .

Table 5: Proportion of the embedding & output parameters in mainstream LLMs. Note that GPT2 and Gemma use weight-tying, so the parameters for the embedding & output layers are shared. We find that the embedding & output layer usually takes  $\leq 10\%$  of total parameters.

Model	Vocabulary size	Embedding dimension	# Parameters for the embedding & output layer	# Total parameters	Proportion
GPT2-1.5B	50304	1600	50304 * 1600	1.5B	5.17%
TinyLlama-1B	32000	2048	32000*2048*2	1.1B	11.80%
Llama2-7B	32000	4096	32000*4096*2	6.7B	3.80%
Llama2-13B	32000	4096	32000*4096*2	13B	2.02%
Llama2-33B	32000	4096	32000*4096*2	33B	0.79%
Llama2-65B	32000	4096	32000*4096*2	65B	0.40%
Llama3-8B	128256	4096	128256*4096*2	8B	13%
Llama3-70B	128256	4096	128256*4096*2	70B	0.15%
Phi-3	32064	3072	32064*3072*2	3.8B	5.18%
Gemma-7B	256128	3072	256128*3072	8.5B	9.83%

#### D.5. Further Comparison with Popular Memory-Efficient Optimizers

**Similar trajectory to AdamW.** As a bonus finding, Adam-mini could generate similar trajectories to that of AdamW, while other popular memory-efficient methods cannot. We train a 6-layer Transformer sized 11M on Openwebtext and launch AdamW, Adam-mini, and other memory-efficient optimizers under the same random seed and same learning rate  $1e-5$ . We save the model weights for every 250 iterations and compare their Euclidean distance to the weights along AdamW’s trajectory. As shown in Figure 15, the model weights obtained by Adam-mini is close to those by AdamW throughout training. In contrast, other optimizers deviate from AdamW. This is because Adam-mini makes fewer modifications over AdamW than other methods.

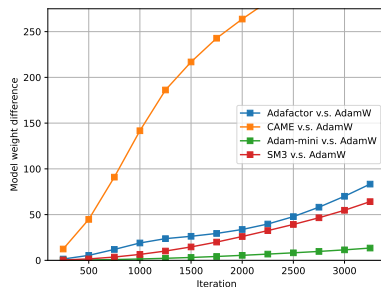


Figure 15: On a 6-layer Transformer, we compare the model checkpoints along the trajectory of different methods. We find the model weights obtained by Adam-mini is close to those by AdamW throughout training.

#### D.6. Sample Responses from LLMs trained by Adam-mini

The results are shown in Figure 16.

---



---

**Prompt: Compose an engaging travel blog post about a recent trip to Hawaii, highlighting cultural experiences and must-see attractions.**

---



---

**AdamW:** Sure! Here’s a travel blog post about a recent trip to Hawaii: ..... Whether you’re interested in history, nature, or simply soaking up the sun, Hawaii has something for everyone. Until next time, Aloha!

**Adam-mini:** Sure! Here’s a travel blog post about a recent trip to Hawaii: ..... I highly recommend visiting this beautiful archipelago for a truly unforgettable adventure.

---



---

**Prompt: Rewrite your previous response. Start every sentence with the letter A.**

---



---

**AdamW:** A trip to Hawaii is a must-visit destination for any traveler, offering a variety of cultural experiences and must-see attractions. From the Polynesian Cultural Center to the USS Arizona Memorial, and from Lanikai Beach to the lively night markets, every corner of the islands reveals a new surprise and a new way to connect with the spirit of Hawaii.

**Adam-mini:** A trip to Hawaii is a must-do for any travel enthusiast. A trip to the beautiful Hawaiian Islands is a paradise on earth. A visit to Oahu is a great way to start your adventure. A helicopter tour of the island is both educational and breathtaking. A visit to Maui is a must-see. A visit to Haleakala National Park is a highlight of your trip. A visit to the Polynesian Cultural Center is a great way to learn about the history and traditions of the islands. A traditional luau is a must-attend event. Overall, a trip to Hawaii is an unforgettable experience.

---

Figure 16: Response Sample of Llama2-7B supervised fine-tuned by AdamW and Adam-mini to the #81 test question from the MT-Bench (Zheng et al., 2024) dataset. Since there is little information in the first round response, we omit the answers from the models.

## E. More Experimental Details

### E.1. Training configurations for Section 3

For all experiments, we choose the model configurations (e.g. context length) by their standard protocols. We choose the learning rates by the recommendation from open-source platforms if applicable. For instance, for GPT2 series, we use the recommended learning rates by (Liu et al., 2023), which are reported to be optimal by grid search. Unless mentioned otherwise, Adam-mini, Adafactor, SM3, and LAMB use the same learning rate as the recommended ones of AdamW. We choose the learning rate of CAME following the suggestions from the authors<sup>6</sup>. If there is no public recommended learning rate for AdamW, we tune the learning rate for all optimizers within the same computational budget and report the best performance. For other hyperparameters, we follow the recommendation from open-source platforms or by their default setting. For SM3 and Adafactor, we incorporate momentum with  $\beta_1 = 0.9$  to offer a fair comparison with other optimizers and the rest of the hyperparameters are set as default. The detailed configurations are explained as follows.

**GPT2 pre-training.** We use the nanoGPT codebase<sup>7</sup> to train GPT2 sized 125M (small), 330M (medium), 770M (large), and 1.5B (XL) on Openwebtext. For all models, we use context length = 1024, batch size = 480, weight decay coefficient  $\lambda = 0.1$ ,  $\epsilon = 1e-8$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ . We use cosine-decay learning rate schedule with 2000 iterations of warm-up. For GPT2-small, medium, and large, we use the recommended peak learning rate by (Liu et al., 2023), which are reported to be the optimal ones found by grid search. For GPT2-XL, we use the recommended peak learning rate by the Levanter<sup>8</sup>. The chosen peak learning rates are  $6e-4$ ,  $3e-4$ ,  $2e-4$ ,  $1e-4$  for GPT2-small, medium, large, XL, respectively. The minimal learning rate is chosen as  $3e-5$ ,  $6e-5$ ,  $1e-5$ ,  $1e-5$  for these models.

**Llama pre-training.** We use TinyLlama codebase<sup>9</sup> to train TinyLlama-1B and Llama2-7B on CommonCrawl<sup>10</sup>. We use batch size = 40. We preprocess the directionary ’2019-30’ of CommonCrawl into 10850 bins, which in total gives 85GB of tokens. We leave the last 20 bins out as the validation set. For both models, we use weight decay coefficient  $\lambda = 0.1$ ,  $\epsilon = 1e-8$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ . We use cosine-decay learning rate schedule with 2000 iterations of warm-up. For TinyLlam-1B, we use context length = 2048, batch size = 512, peak learning rate =  $2e-4$ , and minimal learning rate =  $2e-5$ . For Llama2-7B,

<sup>6</sup><https://github.com/yanlguo7/CAME>

<sup>7</sup><https://github.com/karpathy/nanoGPT/tree/master>

<sup>8</sup>[https://github.com/stanford-crfm/levanter/blob/e183ec80ec5971b12d4a3fb08a160268de342670/config/gpt2\\_xl.yaml](https://github.com/stanford-crfm/levanter/blob/e183ec80ec5971b12d4a3fb08a160268de342670/config/gpt2_xl.yaml)

<sup>9</sup><https://github.com/jzhang38/TinyLlama>

<sup>10</sup><https://commoncrawl.org/>

we use context length = 4096, batch size = 256, peak learning rate =  $8e-5$ , and minimal learning rate =  $8e-6$ .

**SFT and RLHF.** The implementation of SFT and RLHF code is based on the ReMax codebase<sup>11</sup>. We use DeepSpeed ZeRO-2 in our training. GPT-4 evaluation template in Table 3 is from the codebase<sup>12</sup>. In the reward optimization stage, We use ReMax, a memory-efficient alternative to PPO. We use UltraFeedback dataset (Cui et al., 2023) and use 40% data for SFT and 60% data for ReMax.

**SFT.** We use 80 samples in a batch and train the model for 3 epochs. By searching the learning rate from  $\{1e-5, 2e-5, 5e-6\}$  based on validation loss, we use  $5e-6$  with cosine annealing for both AdamW and Adam-mini. We apply LoRA for all layers except the embedding layer. The rank of LoRA is set to 128. After selecting the learning rate from the same set as the full parameter tuning, we use  $5e-6$  for both AdamW and Adam-mini when LoRA is applied. The weight decay coefficient is set to 0 as recommended by LlamaFactory<sup>13</sup>. The rest of the hyperparameters of AdamW and Adam-mini are  $\epsilon = 1e-8$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ .

**ReMax.** We use 48 samples in a batch and train the model for 1 epoch. By searching the peak learning rate from  $\{1e-6, 2e-6, 5e-7\}$  based on validation reward, AdamW uses  $1e-6$  while Adam-mini selects  $5e-7$  as the peak learning rate. The weight decay coefficient is set to 0. The rest of the hyperparameters of AdamW and Adam-mini are  $\epsilon = 1e-8$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ .

**ResNet.** We use the PyTorch official implementation codebase<sup>14</sup> to train ResNet18 (He et al., 2016) on ImageNet (Deng et al., 2009). We use cosine-decay learning rate, epoch = 90,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1e-8$ . For ResNet18, we use batch size = 256, peak learning rate = 0.005. For ViT-base, we use batch size = 128, peak learning rate = 0.0001. These configurations are used for both Adam-mini and AdamW.

**Diffusion models.** We use the codebase<sup>15</sup> to train diffusion models. The image size is 64 and the training objective is to predict the noise as in (Ho et al., 2020). We use the default U-Net architecture hyper-parameters and the dimension multiply in U-Net is (1, 2, 4, 8). We use the CelebA dataset<sup>16</sup> and train the diffusion model with a learning rate  $5 \times 10^{-5}$  with cosine decay. The batch size is 128 and the training epoch is 50.

**Graph Neural Networks.** We use the DGL implementation<sup>17</sup> of Graph Convolution Networks (GCN) (Kipf & Welling, 2016) and Graph Attention Networks (GAT) (Velickovic et al., 2017) for OGBN-arxiv<sup>18</sup> dataset. All configurations as default. For both Adam-mini and AdamW, we use the default learning rate = 0.005 for GCN and the default learning rate = 0.002 for GAT.

## E.2. Detailed Setup for Other Experiments

**Configuration for Figure 3.** For each dense sub-block  $H_l, l = 1, 2, 3$ , we use random positive definite matrices. We fix the choose the eigenvalues of each  $H_l$  as follows: for  $l = 1$ , we independently sample from  $\{1, 2, 3\}$  for 30 times; for  $l = 2$ , we repeat this procedure for  $\{99, 100, 101\}$ ; for  $l = 3$ , we repeat this procedure for  $\{4998, 4999, 5000\}$ . For the single (blockwise) learning rate method, we use GD with optimal constant learning rate  $2/(L + \mu)$ , where  $L, \mu$  are the largest and smallest eigenvalue of the (blockwise) Hessian. We use Adam with  $\beta_1 = 0$ . This helps us focus on the effect of coordinatewise learning rate in Adam. We also set  $\beta_2 = 1$  to the time-varying learning rate. This is necessary because, for any  $\beta_2 < 1$ , Adam with constant learning rate will oscillate on quadratic functions. This is theoretically proved in (Da Silva & Gazeau, 2020, Proposition 12, Figure 1) and empirically observed in (Zhang et al., 2024, Section 3.3).

<sup>11</sup><https://github.com/liziniu/ReMax>

<sup>12</sup>[https://github.com/lm-sys/FastChat/tree/main/fastchat/llm\\_judge](https://github.com/lm-sys/FastChat/tree/main/fastchat/llm_judge)

<sup>13</sup><https://github.com/hiyouga/LLaMA-Factory>

<sup>14</sup><https://github.com/pytorch/examples/blob/main/imagenet/main.py>

<sup>15</sup><https://github.com/lucidrains/denoising-diffusion-pytorch>

<sup>16</sup><https://cseweb.ucsd.edu/~weijian/static/datasets/celeba/>

<sup>17</sup><https://github.com/dmlc/dgl/tree/master/examples/pytorch/ogb/ogbn-arxiv>

<sup>18</sup><https://ogb.stanford.edu/docs/nodeprop/>

**Configuration for Figure 4.** We use the codebase of <sup>19</sup>. We consider a 4-layer Transformer with  $n\_emb = 32$ ,  $n\_head = 4$ . The dataset is a randomly generated binary sequence of length 42. With the help of auto-differentiation framework, we calculate the Hessian on the whole dataset with two passes of backpropagation (Pearlmutter, 1994) and the calculation is exact. In Figure 4, we first set the largest value of color bar  $v_{max}$  as  $1e-6$  for all blocks, which is roughly the maximal value in most blocks. But then we find the whole query and key blocks are black, so we decrease  $v_{max}$  to  $1e-8$  for these blocks and then the block diagonal pattern becomes clear.

**Throughput Comparison in Table 2.** The results are tested on  $2 \times$  A800-80GB GPUs using Tinyllama codebase. We did not turn on CPU offload. We report the throughput from the summary file of Wandb log.

---

<sup>19</sup><https://colab.research.google.com/drive/1SiF0KZJp75rUeetKOWqpsA8clmHP6jMg?usp=sharing>