

# Beyond Attention: Breaking the Limits of Transformer Context Length with Recurrent Memory

Aydar Bulatov<sup>1</sup>, Yuri Kuratov<sup>2,1</sup>, Yermek Kapushev<sup>2</sup>, Mikhail Burtsev<sup>3</sup>

<sup>1</sup>Neural Networks and Deep Learning Lab, MIPT, Dolgoprudny, Russia

<sup>2</sup>AIRI, Moscow, Russia

<sup>3</sup>London Institute for Mathematical Sciences, London, UK  
{bulatov.as, yurii.kuratov}@phystech.edu, mb@lims.ac.uk

## Abstract

A major limitation for the broader scope of problems solvable by transformers is the quadratic scaling of computational complexity with input size. In this study, we investigate the recurrent memory augmentation of pre-trained transformer models to extend input context length while linearly scaling compute. Our approach demonstrates the capability to store information in memory for sequences of up to an unprecedented two million tokens while maintaining high retrieval accuracy. Experiments with language modeling tasks show perplexity improvement as the number of processed input segments increases. These results underscore the effectiveness of our method, which has significant potential to enhance long-term dependency handling in natural language understanding and generation tasks, as well as enable large-scale context processing for memory-intensive applications.

## Introduction

Transformer-based models show their effectiveness across multiple domains and tasks. The self-attention allows to combine information from all sequence elements into context-aware representations. However, global and local information has to be stored mostly in the same element-wise representations. Moreover, the length of an input sequence is limited by quadratic computational complexity of self-attention.

In this work, we propose and study a memory-augmented segment-level recurrent Transformer (Recurrent Memory Transformer or RMT). Memory allows to store and process local and global information as well as to pass information between segments of the long sequence with the help of recurrence. We implement a memory mechanism with no changes to Transformer model by adding special memory tokens to the input or output sequence. Then Transformer is trained to control both memory operations and sequence representations processing.

In this study we show that by using simple token-based memory mechanism introduced in (Bulatov, Kuratov, and Burtsev 2022) can be combined with pretrained transformer models like BERT (Devlin et al. 2019) and GPT-2 (Radford et al. 2019) with full attention and full precision operations.

**Contributions** 1. We expand application of RMT to encoder-only and decoder-only pre-trained language models. Proposed segment wise curriculum learning allows to fine-tune majority of pre-trained transformer based models for processing potentially unlimited sequences.

2. To benchmark generalization capabilities of RMT we propose a set of novel memory acquisition and retention tasks scalable to extremely long sequences of million tokens.

3. We demonstrate the unparalleled ability of RMT to generalize memory operations, successfully detecting and storing information about facts for up to two million tokens. To the best of our knowledge, this establishes a record for the longest sequence task processed by any existing deep neural network. Furthermore, we identify no technical limitations that prevent further scaling.

4. We compare computational complexity of RMT vs. other transformer models and demonstrate the significant advantage of RMT due to its linear scaling of inference operations and constant memory.

The code is available on GitHub<sup>1</sup>. The paper version with supplementary materials is available on arXiv<sup>2</sup>.

## Related Work

Our work revolves around the concept of memory in neural architectures. Memory has been a recurrent theme in neural network research, dating back to early works (McCulloch and Pitts 1943; Stephen 1956) and significantly advancing in the 1990s with the introduction of the *Back-propagation Through Time* learning algorithm (Werbos 1990) and *Long-Short Term Memory* (LSTM) neural architecture (Hochreiter and Schmidhuber 1997). Contemporary memory-augmented neural networks (MANNs) typically utilize some form of recurrent external memory separate from the model’s parameters. *Neural Turing Machines* (NTMs) (Graves, Wayne, and Danihelka 2014) and *Memory Networks* (Weston, Chopra, and Bordes 2015) are equipped with storage for vector representations accessible through an attention mechanism. Memory Networks (Weston, Chopra, and Bordes 2015; Sukhbaatar et al. 2015) were designed to enable reasoning through sequential attention over memory content.

<sup>1</sup><https://github.com/booydar/t5-experiments/tree/aaai24>

<sup>2</sup><https://arxiv.org/abs/2304.11062>

NTMs, followed by *Differentiable Neural Computer* (DNC) (Graves et al. 2016) and *Sparse DNC* (Rae et al. 2016), are implemented as recurrent neural networks capable of writing to memory storage over time. All these models are differentiable and trainable via backpropagation through time (BPTT). Parallel research lines extend recurrent neural networks, such as LSTM, with data structures like stacks, lists, or queues (Joulin and Mikolov 2015; Grefenstette et al. 2015). MANN architectures with more advanced addressing mechanisms, such as address-content separation and multi-step addressing, have been proposed in (Gulcehre et al. 2016; Gulcehre, Chandar, and Bengio 2017; Meng and Rumshisky 2018). The Global Context Layer model (Meng and Rumshisky 2018) employs address-content separation to address the challenge of training content-based addressing in canonical NTMs.

Memory is often combined with Transformers in a recurrent approach. Long inputs are divided into smaller segments, processed sequentially with memory to access information from past segments. Transformer-XL (Dai et al. 2019) preserves previous hidden states for reuse in subsequent segments, while Compressive Transformer (Rae et al. 2020) adds new compressed memory. Ernie-Doc (Ding et al. 2021) enhances contextual information flow by employing same-layer recurrence instead of attending to previous layer outputs of preceding segments. Memformer (Wu et al. 2022a) introduces a dedicated memory module to store previous hidden states in summarized representations. Using a similar approach to Memformer, MART (Lei et al. 2020) and Block-Recurrent Transformer (Hutchins et al. 2022) adopt memory update rules analogous to LSTM (Hochreiter and Schmidhuber 1997) and GRU (Cho et al. 2014). FeedBack Transformer (Fan et al. 2020) implements full recurrence beyond the segment level and merges low and high layers representations into a memory state.

A drawback of most existing recurrent methods is the need for architectural modifications that complicate their application to various pre-trained models. In contrast, the Recurrent Memory Transformer can be built upon any model that uses a common supported interface.

Some approaches redesign the self-attention mechanism to reduce computational complexity while minimizing input coverage loss. *Star-Transformer* (Guo et al. 2019), *Longformer* (Beltagy, Peters, and Cohan 2020), *GMAT* (Gupta and Berant 2020), *Extended Transformer Construction* (ETC) (Ainslie et al. 2020), and *Big Bird* (Zaheer et al. 2020) limit attention distance and employ techniques such as global representations to preserve long-range dependencies. *Memory Transformer* (Burtsev et al. 2020) introduces memory by extending the unchanged model input with special memory tokens.

A common constraint of these methods is that memory requirements grow with input size during both training and inference, inevitably limiting input scaling due to hardware constraints. In contrast, recurrent approaches have constant memory complexity during inference. The longest Longformer, Big Bird, and Long T5 (Guo et al. 2022), LongNet (Ding et al. 2023) models reported in their respective papers have a maximum length of less than 33,000 to-

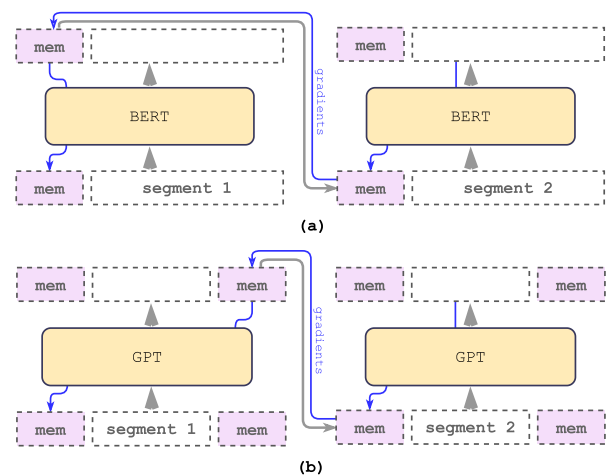


Figure 1: Recurrent memory mechanism. Memory is passed to Transformer along input sequence embeddings, and memory output is passed to the next segment. (a) For encoder-only models there is a single memory and (b) for decoder models with causal attention mask we add an additional write memory at the end. During training gradients flow from the current segment through memory to the previous segment.

kens. CoLT5 (Ainslie et al. 2023) can handle up to 64,000 tokens. LongNet with sparse delayed attention can potentially handle up to 1B tokens per batch in distributed setting. However, CoLT5 and LongNet hit the GPU limit and run out of memory with longer inputs. Furthermore, LongNet is unable to scale to millions of tokens on single GPU due to memory constraints (Ding et al. 2023). Memorizing Transformers (Wu et al. 2022b) and Unlimiformer (Bertsch et al. 2023) further extend memory through k-NN.

Another line of recent related work focuses on models with an alternative to traditional attention mechanism: S4 (Gu, Goel, and Re 2021), Hyena (Poli et al. 2023), RWKV (Peng et al. 2023), RetNet (Sun et al. 2023). They aim to combine the best of convolutions and recurrence – high parallelism during training and linear scaling with sequence length. As architectures change, these methods require training models from scratch and are unable to reuse plethora of already pre-trained transformer-based models.

## Recurrent Memory Transformer

Starting from the initial Recurrent Memory Transformer (Bulatov, Kuratov, and Burtsev 2022) (RMT), we adapted it for a plug-and-play approach as a wrapper for a range of popular Transformers.

This adaptation augments its backbone with memory, composed of  $m$  real-valued trainable vectors (Figure 1). The lengthy input is divided into segments, and memory vectors are prepended to the first segment embeddings and processed alongside the segment tokens. For encoder-only models like BERT, memory is added only once at the beginning of the segment, unlike (Bulatov, Kuratov, and Burtsev 2022), where decoder-only models separate memory into read and




Figure 2: RMT inference scales linearly with respect to the input sequence length. We estimate the required FLOP increase for the forward pass compared to running models on sequences with 512 tokens. (a): lengths from 512 to 32,000 tokens, (b): lengths from 32,000 to 2,048,000 tokens. The RMT segment length is fixed at 512 tokens. While larger models (OPT-30B, OPT-175B) tend to exhibit near-linear scaling on relatively short sequences up to 32,000, they reach quadratic scaling on longer sequences. The angles in the log-log scale indicate a power of a polynomial function. Full attention transformer-based OPT models are much closer to quadratic scaling (dash-dot line). Smaller models (OPT-125M, OPT-1.3B) demonstrate quadratic scaling even on shorter sequences. On sequences with 2,048,000 tokens, RMT can run OPT-175B with  $\times 29$  fewer FLOPs and with  $\times 295$  fewer FLOPs than OPT-135M.

write sections. For the time step  $\tau$  and segment  $H_\tau^0$ , the recurrent step is performed as follows:

$$\begin{aligned} \tilde{H}_\tau^0 &= [H_\tau^{mem} \circ H_\tau^0], \bar{H}_\tau^N = \text{Transformer}(\tilde{H}_\tau^0), \\ [\bar{H}_\tau^{mem} \circ H_\tau^N] &:= \bar{H}_\tau^N, \end{aligned}$$

here  $N$  is a number of Transformer layers.

After the forward pass,  $\bar{H}_\tau^{mem}$  contains updated memory tokens for the segment  $\tau$ .

We can estimate the required FLOPs for RMT and Transformer models of different sizes and sequence lengths. We took configurations (vocabulary size, number of layers, hidden size, intermediate hidden size, and number of attention heads) for the OPT model family (Zhang et al. 2022) and computed the number of FLOPs for the forward pass following (Hoffmann et al. 2022). We also modified FLOP estimates to account for the effect of RMT recurrence.

Segments of the input sequence are processed sequentially. To enable the recurrent connection, we pass the outputs of the memory tokens from the current segment to the input of the next one:

$$H_{\tau+1}^{mem} := \bar{H}_\tau^{mem}, \tilde{H}_{\tau+1}^0 = [H_{\tau+1}^{mem} \circ H_{\tau+1}^0].$$

Both memory and recurrence in the RMT are based only on global memory tokens. This allows the backbone Transformer to remain unchanged, making the RMT memory augmentation compatible with any Transformer-based model.

Figure 2 shows that RMT scales linearly for any model size if the segment length is fixed. We achieve linear scaling by dividing an input sequence into segments and computing the full attention matrix only within segment boundaries. Larger Transformer models tend to exhibit slower quadratic scaling with respect to sequence length because of compute-heavy FFN layers (which scale quadratically with respect to hidden size). However, on extremely long sequences  $> 32,000$ , they fall back to quadratic scaling. RMT requires

fewer FLOPs than non-recurrent models for sequences with more than one segment ( $> 512$  in this study) and can reduce the number of FLOPs by up to  $\times 295$  times. RMT provides a larger relative reduction in FLOPs for smaller models, but in absolute numbers, a  $\times 29$  times reduction for OPT-175B models is highly significant. Additional experimental comparison of computational efficiency can be found in the Appendix.

## Memorization Tasks

To test memorization abilities, we constructed synthetic datasets that require memorization of simple facts and basic reasoning. The task input consists of one or several facts and a question that can be answered only by using all of these facts. To increase the task difficulty, we added natural language text unrelated to the questions or answers. This text acts as noise, so the model’s task is to separate facts from irrelevant text and use them to answer the questions. The task is formulated as a multi-class classification, with each class representing a separate answer option.

Facts are generated using the bAbI dataset (Weston et al. 2016), while the background text is sourced from questions in the QuALITY (Pang et al. 2022) long QA dataset.

Background text example: ” ... *He was a big man, broad-shouldered and still thin-waisted. Eddie found it easy to believe the stories he had heard about his father ...* ”

The first task tests the ability of RMT to write and store information in memory for an extended time (Figure 3, top). In the simplest case, the fact is always located at the beginning of the input, and the question is always at the end. The amount of irrelevant text between the question and answer is gradually increased, so that the entire input does not fit into a single model input. Example: ”*Fact: Daniel went back to the hallway. Question: Where is Daniel? Answer: hallway*”

Fact detection increases the task difficulty by moving the fact to a random position in the input (Figure 3, middle).

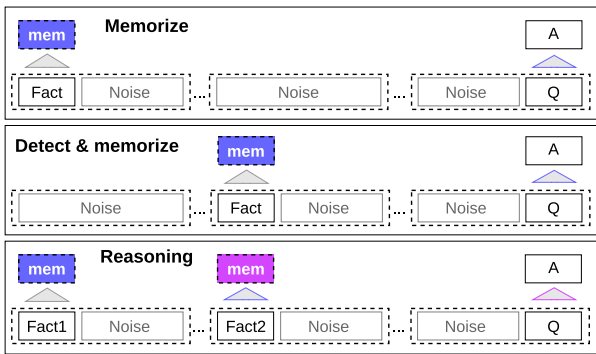


Figure 3: Memory-intensive synthetic tasks. Synthetic tasks and the required RMT operations to solve them are presented. In the Memorize task, a fact statement is placed at the start of the sequence. In the Detect and Memorize task, a fact is randomly placed within a text sequence, making its detection more challenging. In the Reasoning task, two facts required to provide an answer are randomly placed within the text. For all tasks, the question is at the end of the sequence. 'mem' denotes memory tokens, 'Q' represents the question, and 'A' signifies the answer.

This requires the model to first distinguish the fact from irrelevant text, write it to memory, and later use it to answer the question located at the end.

Another important operation with memory is being able to operate with several facts and current context. To evaluate this function, we use a more complicated task called "reasoning", where two facts are generated and positioned randomly within the input sequence (Figure 3, bottom). The question posed at the end of the sequence is formulated in a way that any of the facts must be used to answer the question correctly (i.e., the *Two Argument Relation* bAbI task). Example: "Fact1: The hallway is east of the bathroom. Fact2: The bedroom is west of the bathroom. Question: What is the bathroom east of? Answer: bedroom"

### Learning Memory Operations

We use the pretrained models from Hugging Face Transformers (Wolf et al. 2020) as backbones for RMT in our experiments. All models are augmented with memory and trained using the AdamW optimizer (Loshchilov and Hutter 2019) with linear learning rate scheduling and warmup. The technical details of training and the full set of hyperparameters are available in the Appendix and training scripts in the GitHub repository.

To improve training stability of the original RMT we introduce curriculum learning. At the beginning of training, RMT is fine-tuned on shortest one segment version of the task, and upon convergence, the task length is increased by adding one more segment. The curriculum learning continues until the desired input length is reached.

In our experiments, we begin with sequences that fit in a single segment. The practical segment size is 499, as 3 special tokens of BERT and 10 placeholders for memory are reserved from the model input, sized 512. We notice that




Figure 4: Generalization of memory retrieval. Evaluation of checkpoints trained on 1-7 segment tasks with memory size 10 on varying input lengths. (a): Memorization task, (b): Detection & memorization, (c): Reasoning. Models trained on more than 5 segments generalize well on longer tasks.

after training on shorter tasks, it is easier for RMT to solve longer versions as it converges faster to the perfect solution.

How well does RMT generalize to different sequence lengths? To answer this question, we evaluate models trained on a varying number of segments to solve tasks of larger lengths (Figure 4). We observe that most models tend to perform well on shorter tasks. The only exception is the single-segment reasoning task, which becomes hard to solve once the model is trained on longer sequences. One possible ex-




Figure 5: Recurrent Memory Transformer retains information across up to  $2 \times 10^6$  tokens. By augmenting a pre-trained BERT model with recurrent memory (Bulatov, Kuratov, and Burtsev 2022), we enabled it to store task-specific information across 7 segments of 512 tokens each. During inference, the model effectively utilized memory for up to 4,096 segments with a total length of 2,048,000 tokens—significantly exceeding the largest input size reported for transformer models (64K tokens for CoLT5 (Ainslie et al. 2023), and 32K tokens for GPT-4 (OpenAI 2023), and 100K tokens for Claude). This augmentation maintains the base model’s memory size at 3.6 GB in our experiments.

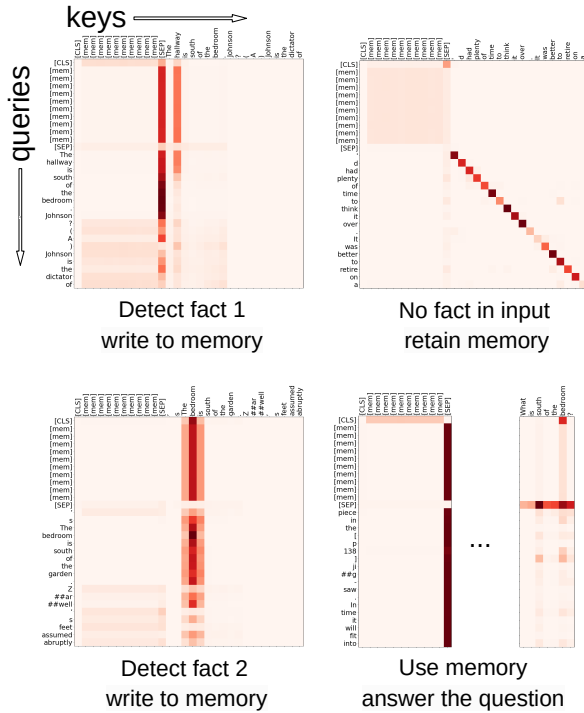


Figure 6: Attention maps for operations with memory. These heatmaps show operations performed during specific moments of a 4-segment reasoning task. The darkness of each pixel depends on the attention value between the corresponding key and value. From left to right: RMT detects the first fact and writes its content to memory ([mem] tokens); the second segment contains no information, so the memory keeps the content unchanged; RMT detects the second fact in reasoning tasks and appends it to memory; CLS reads information from the memory to answer the question.

planation is that since the task size exceeds one segment, the model stops expecting the question in the first segment, leading to quality degradation.

Interestingly, the ability of RMT to generalize to longer sequences also emerges with a growing number of training segments. After being trained on 5 or more segments, RMT can generalize nearly perfectly for tasks twice as long. To test the limits of generalization, we increase the validation task size up to 4096 segments or 2,043,904 tokens (Figure 5). RMT holds up surprisingly well on such long sequences, with Detect & Memorize being the easiest and Reasoning task the most complex.

By examining the RMT attention on specific segments, as shown in Figure 6, we observe that memory operations correspond to particular patterns in attention. Furthermore, the high extrapolation performance on extremely long sequences, as presented on the Fig. 5, demonstrates the effectiveness of learned memory operations, even when used thousands of times. The RMT does not have any specific memory read/write modules and Transformer learns how to operate with memory recurrently. This is particularly impressive, considering that these operations were not explicitly motivated by the task loss.

### Natural and Formal Language Modeling

To study the contribution of recurrent memory for long text understanding, we focus on the long range language modeling task. To capture long-term dependencies in text, memory is required to find and store various type of information between segments. We train the GPT-2 Hugging Face checkpoint with 2 memory tokens using the recurrent memory approach on the ArXiv documents from The Pile (Gao et al. 2020). The dataset is preprocessed by splitting each document into non-overlapping segments of fixed length, which are prepended with their respective histories that consist of several segments. During both training and evaluation we process history and target segments one by one and calculate

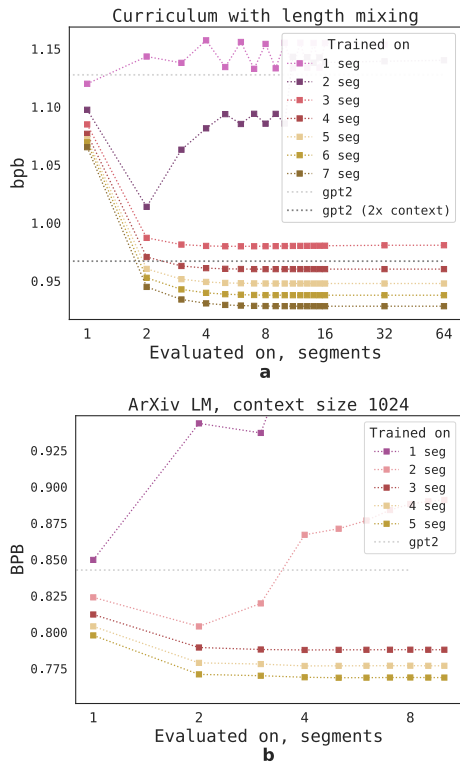


Figure 7: Generalization of memory on language modeling task. Models with input sizes (a): 128 and (b): 1024 trained with RMT show better performance and generalization across longer sizes of context. Perplexity improvement from training RMT with memory size 2 compared to training the baseline GPT-2 for the same number of steps.

loss and perplexity only for the last target segment. Similarly to memorization tasks, we employ curriculum learning for training, starting without history and then gradually increasing context size. We also find that mixing the number of segments on each curriculum step leads to much better generalization on other sequence lengths. We discuss curriculum procedures and usage of parameter-efficient methods in the Appendix.

As expected, increasing the effective context size leads to an improvement in perplexity (Figure 7). RMT trained for an equal number of steps as the baseline GPT-2 displays substantially lower perplexity values. With increasing number of segments in training RMT starts exhibiting better tolerance to longer history sizes. Performance of memory models trained without history suffers when applied to long contexts, but improves after multi-segment training.

To understand how memory is utilized during generation of the sequence we measured perplexity for every position in it (see Figure 8). Baseline shows low prediction quality at the beginning of the sequence due to short context available to condition generation. On the other hand, RMT ensures equally good prediction for all tokens due to carryover of information from the previous segment.

To test our approach in a different domain we fine-tune




Figure 8: Memory improves prediction at a beginning of a segment. As we can see, there is an increase in the loss for tokens at the beginning for GPT-2 (context size 0), showing that it struggles to predict the first tokens since they have no context. The RMT keeps information about previous segments in memory tokens, which helps it to improve tokens predictions. However, showing the model the exact previous context (context size 128 and 768) allows for larger loss gains, but at a higher inference cost.

RMT on a complex mathematical task: generating a proof for a given mathematical theorem in formal language. For our experiments, we utilized Lean 3 (de Moura et al. 2015) and its library, Mathlib (mathlib Community 2020), which contains a range of formalized theories.

Each proof relies on known results, referred to as lemmas. To ensure an effective model, it must accurately assess the relevance of a lemma to the given proof. Subsequently, it should memorize the lemma’s name and incorporate it within the proof. To construct our dataset, we organized each sample into a sequence format. The sequence comprises the theorem statement at the beginning, followed by a randomly ordered list of relevant and irrelevant lemmas, and concludes with the human-written proof. By adjusting the presence of irrelevant lemmas, we control the sequence length. We further divide the sequence into non-overlapping segments of fixed size.

For training and evaluation, we calculate the loss and perplexity of the entire sequence. Similar to memorization tasks, we train the RMT model and gradually increase size of the sequences. As our backbone, we employ GPT-Neo (Black et al. 2021) with 1.3B parameters. We incorporate 10 memory tokens and set the segment size to 2028.

To assess the performance of the RMT model, we compare it with GPTNeo without memory trained on a sequences of 2 segments (first segment always contains the theorem statement and the second contains the proof). GPT-Neo undergoes fine-tuning using the same number of tokens as RMT with 2 segments. Figure 9 shows the results of the RMT model. The RMT model improves perplexity compared to the memory-less model.

However, training with 4 or more segments does not enhance predictions for longer sequences. According to how the sequence is constructed and split into segments, we hypothesize that the model is more concentrated on learning to remember the beginning of the last lemma in the previous

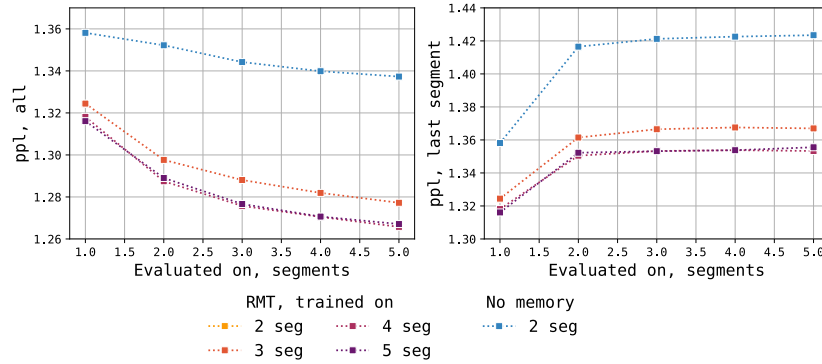


Figure 9: Lemmas memorization for theorem proving. Evaluation of the RMT model and backbone model without memory. Two metrics are calculated: perplexity on all sequence tokens (left) and on the last segment of the sequence (right). RMT model shows better quality.

segment to predict its end in the subsequent segment. The effect of detecting and memorizing relevant lemmas and utilizing them in proof generation is less notable. We believe that the results can be improved by more careful loss construction and data preparation.

## Conclusion

The problem of long input scaling in Transformers has been extensively studied since the introduction of this architecture. Our research has presented a series of significant advancements in augmenting and training of Transformer language models. The work expands the conventional capabilities of pre-trained encoder-only and decoder-only transformers to an unprecedented level of scalability through the integration of token-based memory storage and segment-level recurrence using recurrent memory (RMT).

We have shown that by employing the RMT combined with curriculum learning, even models pre-trained on shorter sequences can be effectively adapted to manage tasks involving significantly longer sequences. This demonstrates that the input length originally designed for the model does not necessarily restrict its potential capabilities, thus offering a new perspective on the adaptability of Transformers.

Our work further uncovered the remarkable adaptability of the trained RMT models in extrapolating to tasks of varying lengths. The results obtained showcased the RMT’s ability to handle sequences exceeding 1 million tokens. Importantly, the computational requirements scaled linearly, thereby maintaining computational efficiency even as task length drastically increased. This is a substantial contribution that could lead to broader applications and improved performance in handling large-scale data. Through an analysis of attention patterns, we provided insight into the operations RMT engages to manipulate memory.

Overall, our research contributes significantly to the understanding and enhancement of pre-trained Transformer language models. It offers a promising direction for future work, particularly in terms of handling longer sequences and improving the adaptability of these models.

## Limitations and Discussion

The curriculum procedure has a substantial impact on the generalization abilities of RMT. Consequently, careful consideration and implementation of curriculum is needed, in contrast to straightforward training of regular Transformers.

We demonstrate scaling to extremely long sequences such as 2M tokens only on specialized tasks. Unfortunately, there are currently no established benchmarks for NLP tasks with such lengths. However, there are no technical limitations to use the proposed methods on tasks with 2M+ tokens lengths.

Training with BPTT is less computationally expensive than full attention, but still requires a significant amount of computation. In our experiments, BPTT with a maximum unroll of 7 segments was sufficient to show generalization on much longer sequences. However, larger models would be more expensive to train with BPTT and some tasks may require more segments to generalize. Techniques such as gradient checkpointing, truncated BPTT or parameter efficient training can reduce the amount of required resources.

Another point is that with unlimited resources and general-purpose information to remember, full attention models might still have an edge in performance. We can think of full-attention models as an upper bound for RMT, since RMT has to operate only on memory states that represent compressed information, not on actual exact hidden states of the past. Recurrent-based approaches, on the other hand may be useful in complex step-by-step reasoning tasks, with specialized memory-intensive tasks or in cases where current models are limited (Liu et al. 2023).

## Acknowledgements

We are thankful to SberDevices for granting us access to additional computational resources. A.B. and Y.K.’s work was supported by a grant for research centers in the field of artificial intelligence, provided by the Analytical Center for the Government of the Russian Federation in accordance with the subsidy agreement (agreement identifier 000000D730321P5Q0002) and the agreement with the Moscow Institute of Physics and Technology dated November 1, 2021 No. 70-2021-00138.

## References

- Ainslie, J.; Lei, T.; de Jong, M.; Ontañón, S.; Brahma, S.; Zemlyanskiy, Y.; Uthus, D.; Guo, M.; Lee-Thorp, J.; Tay, Y.; Sung, Y.-H.; and Sanghai, S. 2023. CoLT5: Faster Long-Range Transformers with Conditional Computation. *arXiv:2303.09752*.
- Ainslie, J.; Ontanon, S.; Alberti, C.; Pham, P.; Ravula, A.; and Sanghai, S. 2020. ETC: Encoding Long and Structured Data in Transformers. *arXiv:2004.08483*.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Bertsch, A.; Alon, U.; Neubig, G.; and Gormley, M. R. 2023. Unlimiformer: Long-Range Transformers with Unlimited Length Input. *arXiv preprint arXiv:2305.01625*.
- Black, S.; Gao, L.; Wang, P.; Leahy, C.; and Biderman, S. 2021. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow.
- Bulatov, A.; Kuratov, Y.; and Burtsev, M. 2022. Recurrent Memory Transformer. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 11079–11091. Curran Associates, Inc.
- Burtsev, M. S.; Kuratov, Y.; Peganov, A.; and Sapunov, G. V. 2020. Memory transformer. *arXiv preprint arXiv:2006.11527*.
- Cho, K.; van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, 103–111. Doha, Qatar: Association for Computational Linguistics.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2978–2988. Florence, Italy: Association for Computational Linguistics.
- de Moura, L.; Kong, S.; Avigad, J.; Van Doorn, F.; and von Raumer, J. 2015. The Lean theorem prover (system description). In *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings 25*, 378–388. Springer.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Ding, J.; Ma, S.; Dong, L.; Zhang, X.; Huang, S.; Wang, W.; and Wei, F. 2023. Longnet: Scaling transformers to 1,000,000,000 tokens. *arXiv preprint arXiv:2307.02486*.
- Ding, S.; Shang, J.; Wang, S.; Sun, Y.; Tian, H.; Wu, H.; and Wang, H. 2021. ERNIE-Doc: A Retrospective Long-Document Modeling Transformer. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2914–2927. Online: Association for Computational Linguistics.
- Fan, A.; Lavril, T.; Grave, E.; Joulin, A.; and Sukhbaatar, S. 2020. Addressing some limitations of transformers with feedback memory. *arXiv preprint arXiv:2002.09402*.
- Gao, L.; Biderman, S.; Black, S.; Golding, L.; Hoppe, T.; Foster, C.; Phang, J.; He, H.; Thite, A.; Nabeshima, N.; Presser, S.; and Leahy, C. 2020. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. *arXiv preprint arXiv:2101.00027*.
- Graves, A.; Wayne, G.; and Danihelka, I. 2014. Neural Turing machines. *arXiv preprint arXiv:1410.5401*.
- Graves, A.; Wayne, G.; Reynolds, M.; Harley, T.; Danihelka, I.; Grabska-Barwińska, A.; Colmenarejo, S. G.; Grefenstette, E.; Ramalho, T.; Agapiou, J.; Badia, A. P.; Hermann, K. M.; Zwols, Y.; Ostrovski, G.; Cain, A.; King, H.; Summerfield, C.; Blunsom, P.; Kavukcuoglu, K.; and Hassabis, D. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626): 471–476.
- Grefenstette, E.; Hermann, K. M.; Suleyman, M.; and Blunsom, P. 2015. Learning to Transduce with Unbounded Memory. *arXiv:1506.02516*.
- Gu, A.; Goel, K.; and Re, C. 2021. Efficiently Modeling Long Sequences with Structured State Spaces. In *International Conference on Learning Representations*.
- Gulcehre, C.; Chandar, S.; and Bengio, Y. 2017. Memory augmented neural networks with wormhole connections. *arXiv preprint arXiv:1701.08718*.
- Gulcehre, C.; Chandar, S.; Cho, K.; and Bengio, Y. 2016. Dynamic neural Turing machine with soft and hard addressing schemes. *arXiv preprint arXiv:1607.00036*.
- Guo, M.; Ainslie, J.; Uthus, D.; Ontanon, S.; Ni, J.; Sung, Y.-H.; and Yang, Y. 2022. LongT5: Efficient Text-To-Text Transformer for Long Sequences. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 724–736. Seattle, United States: Association for Computational Linguistics.
- Guo, Q.; Qiu, X.; Liu, P.; Shao, Y.; Xue, X.; and Zhang, Z. 2019. Star-Transformer. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 1315–1325. Minneapolis, Minnesota: Association for Computational Linguistics.
- Gupta, A.; and Berant, J. 2020. GMAT: Global memory augmentation for transformers. *arXiv preprint arXiv:2006.03274*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 9(8): 1735–1780.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L. A.; Welbl, J.; Clark, A.; Hennigan, T.; Noland, E.; Millican, K.; van den Driessche, G.; Damoc, B.; Guy, A.; Osindero, S.; Simonyan, K.; Elsen, E.; Vinyals, O.; Rae, J.; and Sifre, L.



2022. An empirical analysis of compute-optimal large language model training. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 30016–30030. Curran Associates, Inc.
- Hutchins, D.; Schlag, I.; Wu, Y.; Dyer, E.; and Neyshabur, B. 2022. Block-Recurrent Transformers. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Joulin, A.; and Mikolov, T. 2015. Inferring Algorithmic Patterns with Stack-Augmented Recurrent Nets. arXiv:1503.01007.
- Lei, J.; Wang, L.; Shen, Y.; Yu, D.; Berg, T. L.; and Bansal, M. 2020. MART: Memory-Augmented Recurrent Transformer for Coherent Video Paragraph Captioning. arXiv:2005.05402.
- Liu, N. F.; Lin, K.; Hewitt, J.; Paranjape, A.; Bevilacqua, M.; Petroni, F.; and Liang, P. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- mathlib Community, T. 2020. The lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*. ACM.
- McCulloch, W. S.; and Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4): 115–133.
- Meng, Y.; and Rumshisky, A. 2018. Context-aware neural model for temporal information extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 527–536.
- OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774.
- Pang, R. Y.; Parrish, A.; Joshi, N.; Nangia, N.; Phang, J.; Chen, A.; Padmakumar, V.; Ma, J.; Thompson, J.; He, H.; and Bowman, S. 2022. QuALITY: Question Answering with Long Input Texts, Yes! In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5336–5358. Seattle, United States: Association for Computational Linguistics.
- Peng, B.; Alcaide, E.; Anthony, Q.; Albalak, A.; Arcadinho, S.; Cao, H.; Cheng, X.; Chung, M.; Grella, M.; GV, K. K.; et al. 2023. RWKV: Reinventing RNNs for the Transformer Era. *arXiv preprint arXiv:2305.13048*.
- Poli, M.; Massaroli, S.; Nguyen, E.; Fu, D. Y.; Dao, T.; Bacchus, S.; Bengio, Y.; Ermon, S.; and Re, C. 2023. Hyena Hierarchy: Towards Larger Convolutional Language Models. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 28043–28078. PMLR.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners.
- Rae, J. W.; Hunt, J. J.; Harley, T.; Danihelka, I.; Senior, A.; Wayne, G.; Graves, A.; and Lillicrap, T. P. 2016. Scaling Memory-Augmented Neural Networks with Sparse Reads and Writes. arXiv:1610.09027.
- Rae, J. W.; Potapenko, A.; Jayakumar, S. M.; Hillier, C.; and Lillicrap, T. P. 2020. Compressive Transformers for Long-Range Sequence Modelling. In *International Conference on Learning Representations*.
- Stephen, C. 1956. Kleene. Representation of events in nerve nets and finite automata. *Automata studies*.
- Sukhbaatar, S.; Szlam, A.; Weston, J.; and Fergus, R. 2015. End-To-End Memory Networks. arXiv:1503.08895.
- Sun, Y.; Dong, L.; Huang, S.; Ma, S.; Xia, Y.; Xue, J.; Wang, J.; and Wei, F. 2023. Retentive network: A successor to transformer for large language models. *arXiv preprint arXiv:2307.08621*.
- Werbos, P. J. 1990. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10): 1550–1560.
- Weston, J.; Bordes, A.; Chopra, S.; and Mikolov, T. 2016. Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Weston, J.; Chopra, S.; and Bordes, A. 2015. Memory Networks. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, 38–45.
- Wu, Q.; Lan, Z.; Qian, K.; Gu, J.; Geramifard, A.; and Yu, Z. 2022a. Memformer: A Memory-Augmented Transformer for Sequence Modeling. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, 308–318. Online only: Association for Computational Linguistics.
- Wu, Y.; Rabe, M. N.; Hutchins, D.; and Szegedy, C. 2022b. Memorizing Transformers. In *International Conference on Learning Representations*.
- Zaheer, M.; Guruganesh, G.; Dubey, K. A.; Ainslie, J.; Alberti, C.; Ontanon, S.; Pham, P.; Ravula, A.; Wang, Q.; Yang, L.; and Ahmed, A. 2020. Big Bird: Transformers for Longer Sequences. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 17283–17297. Curran Associates, Inc.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; Mihaylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.; Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models. *ArXiv*, abs/2205.01068.