

PERK: LONG-CONTEXT REASONING AS PARAMETER-EFFICIENT TEST-TIME LEARNING

Zeming Chen Angelika Romanou Gail Weiss Antoine Bosselut

EPFL

email {zeming.chen, antoine.bosselut}@epfl.ch

ABSTRACT

Long-context reasoning requires accurately identifying relevant information in extensive, noisy input contexts. In this work, we propose PERK (Parameter Efficient Reasoning over Knowledge), a scalable approach for learning to encode long contexts using gradient updates at test time. Specifically, PERK employs two nested optimization loops in a meta-training phase. The inner loop rapidly encodes contexts into a low-rank adapter (LoRA) that serves as a parameter-efficient memory module for the base model. Concurrently, the outer loop learns to use the updated adapter to accurately recall and reason over relevant information from the encoded long context. Our evaluations on several long-context reasoning tasks show that PERK significantly outperforms the standard long-context finetuning, achieving average absolute performance gains of up to 20% for Qwen-2.5 (0.5B & 7B) on synthetic and real-world long-context reasoning. PERK also maintains its advantages across model scales and families. Compared to specialized long-context LLMs, PERK matches or surpasses their performance. Finally, our analyses show PERK is more robust to reasoning complexity, length extrapolation, and the positions of relevant information in contexts. <https://perk-long-context.web.app>

1 INTRODUCTION

Large language models (LLMs) struggle with identifying and reasoning over relevant information in long, potentially noisy, input contexts (*i.e.*, long-context reasoning; Anil et al., 2022; Lee et al., 2024; Li et al., 2024; Liu et al., 2023; Xu et al., 2024), including modern models with substantial context windows (Kuratov et al., 2024; Levy et al., 2024; Hsieh et al., 2024). These challenges arise because longer contexts contain more irrelevant *distractor* information, and may also be attached to more complex reasoning problems (*i.e.*, multiple hops), both of which degrade LLM performance (Chen et al., 2023b). Furthermore, these challenges are exacerbated by positional biases in long-context models, which focus attention on the beginning or end of contexts, neglecting intermediate information (Liu et al., 2024).

In this work, we address these challenges of long-context reasoning by reframing it as test-time learning, enabling models to encode context at inference time using gradient-based parameter adaptation. Our approach, Parameter-Efficient Reasoning over Knowledge (PERK), processes and compresses lengthy sequences as a *batch* (or multiple batches) of shorter segments of the long context, rather than processing the full sequence token-by-token. The updated model can then use the newly stored parametric information (and its initial parametric knowledge) to answer relevant questions. The original context is then discarded in this answering step as it is now encoded parametrically by the model. To train PERK, we operate two nested optimization loops: the inner loop, which is also used at test time, encodes the long context segment batches into model parameters, while the outer loop learns to respond to queries using the updated parameters representing the context.

However, bi-level optimization algorithms for training test-time learning methods entail prohibitive memory overhead (Finn et al., 2017; Chen et al., 2023b; Hu et al., 2023), particularly when updating all model parameters, because they require backpropagating through a complete multi-step optimization trajectory. To reduce memory overhead, PERK (1) internalizes the context into a Low-Rank Adapter

(LoRA; Hu et al., 2022) rather than the initial model parameters, and (2) backpropagates the outer loop over only a truncated trajectory from the final few inner-loop adaptation steps (Shaban et al., 2019). Consequently, PERK scales efficiently at training, supporting larger models and longer contexts.

We compare PERK to models finetuned on long contexts for classical in-context reasoning (FT-ICR) on multiple *Needle-in-a-Haystack* benchmarks, demonstrating PERK’s advantages. For Qwen-2.5-0.5B with a native 32K context window, PERK surpasses FT-ICR baselines by a 20% absolute improvement. PERK also outperforms specialized long-context LLMs (Gao et al., 2025a; Yang et al., 2025) with over 7B parameters that have undergone extensive long-context training, achieving the highest accuracy when generalizing to contexts up to 128k tokens. As relevant facts in *Needle-in-a-Haystack* problems are often easily distinguishable from distractor text, we also propose more challenging evaluation settings where relevant facts are surrounded by distributionally-similar distractor facts (e.g., *Drops-in-the-Ocean*, HotPotQA, TriviaQA). In these settings, PERK models again outperform the FT-ICR baselines, and even surpass commercial LLMs (e.g., GPT, Gemini) on benchmarks for long-context symbolic reasoning. PERK also works effectively across many model scales (0.1B, 0.5B, 8B) and families (GPT-2, Qwen-2.5, LLaMA-3.1/3.2).

Finally, our analysis also shows PERK’s advantages yield strong test-time length extrapolation. PERK trained on 8K-token contexts maintains decent performance on 64K and 128K windows, consistently surpassing FT-ICR length generalization, and even outperforming models with native 512K and 1M context windows. Finally, because PERK encodes sequences as a permutation-invariant *batch* of context window segments (rather than maintaining the full contiguous context), it generalizes across contexts where relevant information is variably distributed along context positions, unlike FT-ICR’s dramatic drops in performance (up to 90%) when relevant information shifts positions at test time.

2 BACKGROUND AND PRELIMINARIES

Causal Language Models. *Causal language models* (CLM) f_θ (Radford & Narasimhan, 2018) with parameters θ are typically trained with an autoregressive learning objective such as the *token-level negative log-likelihood* (NLL). For a dataset $\mathcal{D} = \{x\}$ with $N_{\mathcal{D}}$ sequences, this loss is defined as:

$$\mathcal{L}_{\text{NLL}}(\theta, \mathcal{D}) = -\frac{1}{N_{\mathcal{D}}} \sum_{x \in \mathcal{D}} \sum_{t=1}^{\tau} \log p_\theta(x_t | \mathbf{x}_{<t}). \quad (1)$$

The initial parameters θ are optimized iteratively via an optimization algorithm $\text{Alg} : (\mathcal{L}, \mathcal{D}, \theta, \mathbf{h}) \mapsto \theta'$, where \mathbf{h} are hyperparameters, such as optimizer choice and training length.

Reasoning with CLMs. We denote *reasoning* problems as tuples $r = (\mathcal{K}, \mathbf{q}, y)$ of context \mathcal{K} , question \mathbf{q} , and target response y (e.g., a question-answering dataset like HotpotQA; Yang et al., 2018). With the pretrained CLMs’s ability to reason over facts in input contexts with reasonable success (Clark et al., 2020), prompt-based *in-context reasoning* (ICR) has become popular for solving reasoning tasks. In this setting, \mathcal{K} and \mathbf{q} are combined into a single prompt $\mathbf{x}(r)$. The model’s parameters are optimized to predict the tokens of the correct answer:

$$\mathcal{L}_{\text{reason}}(\theta, \mathcal{D}) = -\frac{1}{N_{\mathcal{D}_y}} \sum_{r=(\mathcal{K}, \mathbf{q}, y) \in \mathcal{D}} \sum_{t=1}^{\tau_y} \log p_\theta(y_t | \mathbf{x}(r), y_{<t}) \quad (2)$$

where \mathcal{D} is the training corpus, τ_y is the length of an example answer, and $N_{\mathcal{D}_y}$ is the total length of all answers in the dataset.

Reasoning as Test-time Learning. In this setting, models simulate reasoning by first encoding the given context \mathcal{K} into a CLM’s parameters θ , and then using the updated parameters $\phi(\theta, \mathcal{K})$ to generate a response to the given question \mathbf{q} (Chen et al., 2023b). The context is encoded by adapting the model’s parameters using a CLM objective over the context: $\phi(\theta, \mathcal{K}) = \text{Alg}(\mathcal{L}_{\text{NLL}}, \mathcal{K}, \theta, \mathbf{h})$. We then use f_ϕ to generate answers for any relevant questions \mathbf{q} .

Test-time learning methods for reasoning stem from a key observation: CLMs can store and retrieve diverse information in their parameters from pre-training (Roberts et al., 2020; Bayazit et al., 2024). Yet these same models struggle with long-context problems (Hsieh et al., 2024; Kuratov et al., 2024; Xu et al., 2024; Levy et al., 2024), despite the contexts containing far less information than what is seen during pre-training. CLMs may reason better with knowledge in their parameters than in context, motivating methods that encode new context parametrically for reasoning.

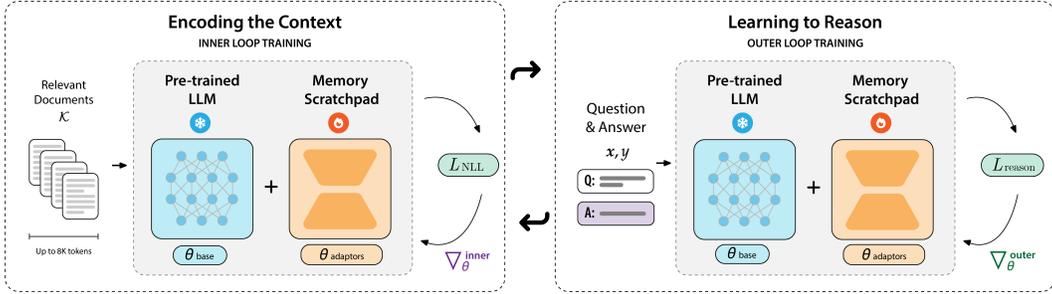


Figure 1: **Meta-learning PERK for long-context reasoning.** The training procedure involves a nested inner and outer loop. The inner loop optimizes the likelihood of a batch of long context segments with respect to the parameters of the LoRA-based memory scratchpad. In the outer loop, the model uses the encoded information in the memory scratchpad to answer questions. In both cases, only the memory scratchpad parameters are updated while the base LLM’s parameters are frozen.

Training a CLM for Test-time Learning. Given a distribution R of possible reasoning problems, the optimal parameters θ^* for a test-time learning CLM minimize the expected reasoning loss of the CLM’s adaptation to each reasoning problem $r \sim R$:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{r \sim R} \left[\mathcal{L}_{\text{reason}} \left(\phi(\theta, \mathcal{K}), \{(q, y)\} \right) \right]$$

where $\phi(\theta, \mathcal{K})$ is a CLM adaptation of θ to the context \mathcal{K} : $\phi(\theta, \mathcal{K}) = \text{Alg}(\mathcal{L}_{\text{NLL}}, \mathcal{K}, \theta, \mathbf{h})$.

Meta-learning algorithms, such as MAML (Finn et al., 2017), can be used for training effective test-time learning methods. In MAML, we obtain a meta-model f_{θ} via bi-level optimization: the *outer loop* samples batches of reasoning problems from a training dataset of reasoning problems, the *inner loop* obtains the adaptations of θ to each problem’s context \mathcal{K} , and the *outer loop* then optimizes the meta-parameters based on the performance of each $f_{\phi(\theta, \mathcal{K})}$ on the corresponding questions. The outer loop optimization backpropagates through the entire inner loop, an operation that becomes increasingly expensive with more parameters and inner loop update steps.

3 PERK: META-LEARNING PARAMETER-EFFICIENT REASONING OVER KNOWLEDGE

We propose PERK, a bi-level optimization algorithm for model reasoning via test-time adaptation for any given long context. At inference time, PERK adapts a base model to encode a context via CLM optimization on that context. Subsequently, the adapted model can be used to answer relevant questions about that context. To scale PERK’s adaptation to larger and more capable base LLMs for effective long-context reasoning, we use parameter-efficient methods to reduce both the size and amount of gradient unrolling required during optimization.

3.1 PARAMETER-EFFICIENT META-LEARNING WITH LOW-RANK ADAPTATION

During training, PERK optimizes a parameter-efficient adapter, LoRA (Hu et al., 2022), instead of the complete model. Specifically, PERK encodes a context into the adapter through gradient-based updates, leaving the base model parameters unchanged. We denote by θ_{base} and θ_{adapter} the parameters of the base model and the LoRA adapter, respectively. θ_{base} is never optimized, while θ_{adapter} is meta-learned to a good initial state for test-time adaptations.

Test-Time Learning and Inner Loop Training. We set the test time learning objective (thus also the inner loop objective) to causal language modeling of the context: $\mathcal{L}_{\text{NLL}}(\mathcal{K}, (\theta_{\text{base}}, \theta_{\text{adapter}}))$, and set the adaptation algorithm $\text{Alg}(\mathcal{L}_{\text{NLL}}, \mathcal{K}, (\theta_{\text{base}}, \theta_{\text{adapter}}), \mathbf{h})$ to only update θ_{adapter} . The adaptation algorithm computes the gradient $\nabla \mathcal{L}_{\text{NLL}}$ on a batch of sub-sequences from the full context sequence \mathcal{K} . This compression as a parallel *batch* (or multiple batches) allows us to process lengthy sequences beyond the model’s native context window. The exact manner by which \mathcal{K} is broken into batches is a hyperparameter of the algorithm, and for efficiency reasons, we set it to one batch during training.

Outer Loop Training. In the outer loop, the model learns to reason over the encoded information in the adapter (memory scratchpad) given a problem q without access to the long context \mathcal{K} explicitly in the text space. We optimize the meta parameters θ_{adapter} to learn from R , the corresponding distribution of reasoning problems. The optimal θ_{adapter} minimize the expected reasoning loss of the adapter’s adaptation to each reasoning problem $r = (\mathcal{K}, q, y) \sim R$:

$$\theta_{\text{adapter}}^* = \arg \min_{\theta_{\text{adapter}}} \mathbb{E}_{r \sim R} \left[\mathcal{L}_{\text{reason}} \left(\theta_{\text{base}}, \phi_{\text{adapter}}((\theta_{\text{base}}, \theta_{\text{adapter}}), \mathcal{K}), \{(q, y)\} \right) \right], \quad (3)$$

where $\phi_{\text{adapter}}((\theta_{\text{base}}, \theta_{\text{adapter}}), \mathcal{K}) = \text{Alg}(\mathcal{L}_{\text{NLL}}, \mathcal{K}, (\theta_{\text{base}}, \theta_{\text{adapter}}), \mathbf{h})$ is the parameter-efficient CLM adaptation of θ_{adapter} to the context \mathcal{K} obtained via the test-time adaptation. We may use any optimizer for θ_{adapter} , and in our experiments, we use AdamW (Loshchilov & Hutter, 2017). We describe further hyperparameters, training & inference details, and long-context data preprocessing in Appendix B.

3.2 SCALABLE META-LEARNING WITH TRUNCATED GRADIENT UNROLLING

To optimize Equation (3) with gradient-based methods, we need to differentiate through Alg , which involves higher-order derivatives and saving the complete trajectory to compute the meta-gradient (i.e., the gradient of $\mathcal{L}_{\text{reason}}$). However, this creates high memory costs, capping the method’s applicability. To reduce memory costs, we truncate the backpropagation for the meta-gradient computation. We briefly describe this approach here, presenting the case where Alg optimizes θ_{adapter} via direct gradient descent for simplicity (though we use more complicated optimizers for the inner loop in practice).

Gradient Unrolling (GU). We denote an N -step inner-loop optimization’s intermediate states by:

$$\phi_{\text{adapter}}^{(0)} = \theta_{\text{adapter}}, \quad \phi_{\text{adapter}}^{(n+1)} = \phi_{\text{adapter}}^{(n)} - \alpha g^{(n)}, \quad \phi_{\text{adapter}}^* = \phi_{\text{adapter}}^{(N)}$$

where $g^{(n)} = \nabla_{\phi_{\text{adapter}}^{(n)}} \mathcal{L}_{\text{NLL}}(\theta_{\text{base}}, \phi_{\text{adapter}}^{(n)}, \mathcal{K})$ is the gradient at step n , and α is the learning rate.

In vanilla MAML, to compute the gradient of the outer-loop loss $\mathcal{L}_{\text{reason}}$ with respect to the meta-parameters θ_{adapter} , one backpropagates through all N inner steps. By the chain rule, this equates to accumulating the Jacobian $J^{(n)} = \frac{\partial \phi_{\text{adapter}}^{(n+1)}}{\partial \phi_{\text{adapter}}^{(n)}}$ from each step:

$$\nabla_{\theta_{\text{adapter}}} \mathcal{L}_{\text{reason}} = \frac{\partial \mathcal{L}_{\text{reason}}}{\partial \phi_{\text{adapter}}^{(N)}} J^{(N-1)} J^{(N-2)} \dots J^{(0)} \frac{\partial \phi_{\text{adapter}}^{(0)}}{\partial \theta_{\text{adapter}}} = \frac{\partial \mathcal{L}_{\text{reason}}}{\partial \phi_{\text{adapter}}^{(N)}} \prod_{n=0}^{N-1} J^{(n)}, \quad (4)$$

where $J^{(n)}$ can be computed via: $J^{(n)} = I - \alpha H^{(n)}$, $H^{(n)} = \nabla_{\phi_{\text{adapter}}^{(n)}}^2 \mathcal{L}_{\text{NLL}}$. Each Jacobian $J^{(n)}$ requires storing the computation graph of step n in memory, so the cost of the backward pass scales linearly with N , creating a high memory requirement for gradient unrolling.

Truncated Gradient Unrolling (TGU). To reduce memory, we follow (Shaban et al., 2019) and run the inner loop for all N specified update steps, but only store the computational graph for the last $T \ll N$ steps. This substantially reduces the computational memory requirements of the method, at the cost of slightly less accurate meta-gradients for the outer-loop optimization of θ_{adapter} . To implement this approximation, we do not retain the inner loop optimization’s computation graph until the start of the *retention window*: the last T inner-loop optimization steps. Since all Jacobians before step n_{start} are treated as constants, the Jacobian product in Equation (4) is effectively cut:

$$\nabla_{\theta_{\text{adapter}}} \mathcal{L}_{\text{reason}} \approx \frac{\partial \mathcal{L}_{\text{reason}}}{\partial \phi_{\text{adapter}}^{(N)}} \underbrace{\prod_{n=N-T}^{N-1} (J^{(n)})}_{\text{last } T \text{ steps}} \cancel{\prod_{n=0}^{N-T-1} (J^{(n)})}. \quad (5)$$

4 EXPERIMENTS

We empirically evaluate PERK’s long-context reasoning performance in three scenarios: reasoning over (1) Needles-in-a-Haystack (NIAH, Kuratov et al., 2024), (2) Open-domain, multi-document QA (Multi-Doc), and (3) Drops-in-an-Ocean (DIO), a novel long context evaluation setting we propose.

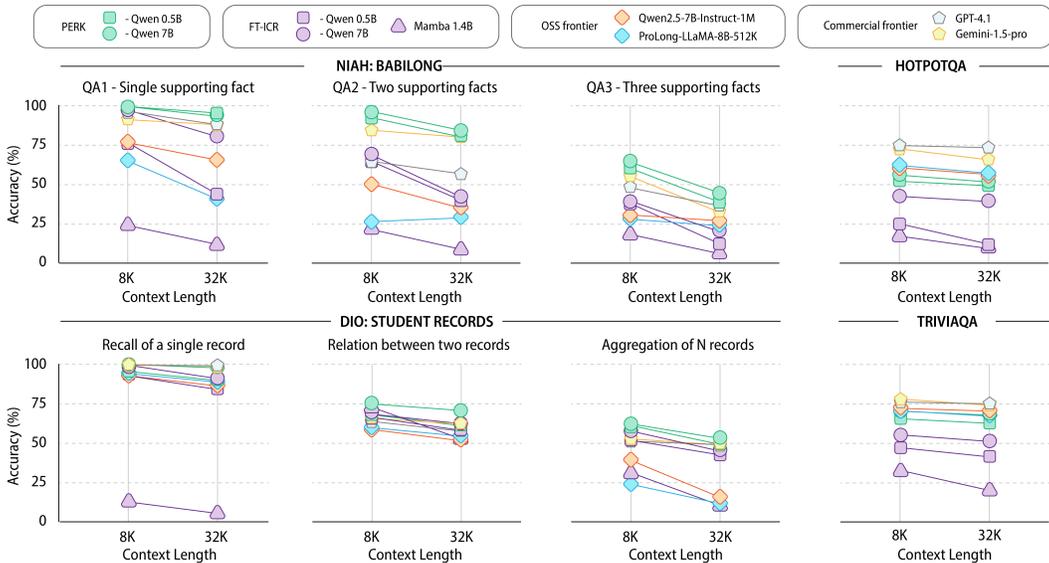


Figure 2: **Performance on Long-context Reasoning.** We show the evaluation results on NIAH with BabiLong, Multi-Doc with HotpotQA & TriviaQA, and DIO with Student Records. All PERK and FT-ICR models (including Mamba) are trained on contexts with 8K tokens. When evaluated on the out-of-distribution contexts with 32K tokens, they must extrapolate to a new context length. Note that we use a substring exact match as the Accuracy metric for HotpotQA & TriviaQA.

4.1 EXPERIMENTAL SETUP

Datasets. We use the BabiLong framework (Kuratov et al., 2024) for reasoning over Needles-in-a-Haystack (NIAH), with single (QA1), two-hop (QA2), and three-hop (QA3) reasoning over facts scattered across very long documents. We sample documents from FineWeb-Edu (Lozhkov et al., 2024) postdating our base model’s knowledge cutoff. To address NIAH’s fundamental limitation, where target information is stylistically distinct from distractors and thus artificially simpler to identify, we evaluate in two additional settings.

First, we evaluate on HotpotQA (Yang et al., 2018) and TriviaQA (Joshi et al., 2017) for real-world long-context QA. Following Yen et al. (2025), we inject retrieved distractors (Zhang et al., 2024a) at random positions until reaching the context window limit. We report substring exact match (SubEM) as accuracy following (Asai et al., 2023).

Second, we propose *Drops-in-the-Ocean* (DIO), a novel evaluation with distributionally similar relevant/irrelevant information. We create **Student Records**, a synthetic dataset where each context contains multiple student records (ID, name, school, major, grade). Context length scales with record count. We define three tasks: Recall (retrieve attributes by ID), Relation (compare attributes between IDs), and Aggregate (compute max/min/average grades). Test entities and attribute combinations are disjoint from training to ensure generalization. Examples are available in Appendix A.

Baselines. We compare PERK against multiple categories of baselines. First, following conventional practice (Fu et al., 2024; Gao et al., 2025b), we finetune open-weight models for prompting-based in-context reasoning. Our baselines include: (1) GPT-2-127M (Radford & Narasimhan, 2018) with 1K context window, extended via positional interpolation (Karypis et al., 2024) for longer contexts; (2) Qwen-2.5 (0.5B, 7B; Qwen et al., 2025) with 32K context; (3) LLaMA-3.2-1B and LLaMA-3.1-8B (Grattafiori et al., 2024) with 128K context; (4) Mamba-1.4B (Gu & Dao, 2024) with unlimited context length. We also compare to commercial models with long-context abilities, GPT-4.1 (OpenAI, 2025) and Gemini-1.5-Pro (Team et al., 2024), and specialized open-source models: Qwen2.5-7B-Instruct-1M (1M context via long-context pretraining, progressive length extension, and long-instruction post-training; Yang et al., 2025) and ProLong-Llama-3-8B-512k-Instruct (512K context; Gao et al., 2025a). We use 4-shot in-context learning as the inference method, following the prompt templates of Kuratov et al. (2024) and Yen et al. (2025)

4.2 LONG-CONTEXT REASONING RESULTS

NIAH. In the top row, left section of Figure 2, we show the evaluation results on BabiLong. PERK outperforms all baselines on all three of the **QA1**, **QA2**, and **QA3** tasks, including much larger models and specialized long-context models that support much larger context-windows than 32K tokens (512K - 1M). When comparing similar training settings (*i.e.*, training on 8k contexts), PERK (Qwen-7B) achieves a higher accuracy across all task complexities than FT-ICR (Qwen-7B) with an in-distribution generalization window of 8K at evaluation time. When extrapolating to a 32K context window, PERK is much more robust, outperforming FT-ICR by 23% points on average.

Multi-Doc. Next, we evaluate PERK’s performance on multi-document open-domain QA tasks HotpotQA and TriviaQA. When trained on long contexts with 8K tokens, PERK (Qwen) consistently outperforms FT-ICR (Qwen) for both HotpotQA and TriviaQA, achieving a 20% absolute gain at the 0.5B scale and 15% gain at the 7B scale. When extrapolating to contexts with 32K tokens, PERK shows stronger generalization over FT-ICR by 30% (0.5B) and 14% (7B). Compared to open-source long-context models (Qwen-1M and ProLong) with extensive training on long-context instruction data, PERK still matches the performance with a very narrow performance gap (3% for HotpotQA and 1.5% for TriviaQA). While not quite at the level of commercial LLMs, GPT-4.1 and Gemini-1.5-pro, PERK’s performance is closer to these leading performances than FT-ICR. Our results demonstrate PERK’s strong potential for retrieving from and reasoning over multiple long-context documents.

Drops-in-the-Ocean. Results on the Student Records dataset in Figure 2 mirror the trends seen in BabiLong. PERK (Qwen) outperforms FT-ICR baselines across task difficulties and context lengths and shows a smaller drop than FT-ICR baselines when increasing the evaluation context length to 32k tokens. The performance differential between PERK and FT-ICR baselines also widens monotonically with task difficulty (Aggregate > Relation > Recall), suggesting that PERK enhances complex reasoning capabilities over long contexts rather than merely improving simple retrieval operations. Finally, as before, both PERK (Qwen) 7B and 0.5B outperform specialized long-context models despite their large parameter size and optimized training strategies for long sequences. PERK (Qwen-7B) also achieves performance that matches or exceeds the commercial models, Gemini-1.5-pro and GPT-4.1, even on the Aggregate task.

Model Generalizability. We further analyze whether PERK’s performance can generalize across diverse model families and scales in Figure 3. While FT-ICR shows highly variable accuracy ranging from 18.2% on GPT-2 to 89.8% on LLaMA-8B, PERK maintains consistently high accuracy across all tested models. PERK’s performance scales positively with model size, achieving 4% increasing on Qwen-7B compared to Qwen-0.5B, and reaching 99.1% accuracy on LLaMA-8B. FT-ICR, despite improving with model size, still lags behind PERK by 9.3 percentage points even on the largest model tested. The generalization across model families is also clear. PERK achieves strong performance on models from the GPT-2, Qwen-2.5, and LLaMA-3.1/3.2 families. Our results demonstrate that PERK’s advantages are not limited to specific architectures or model scales, but rather represent a fundamentally effective approach for long-context modeling.

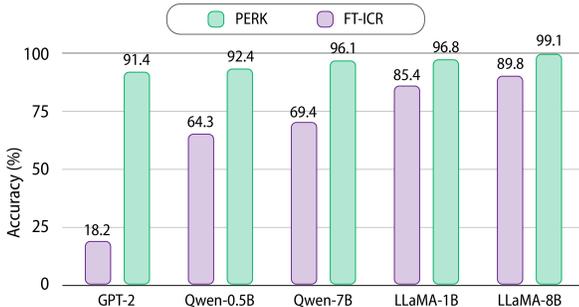


Figure 3: **PERK with diverse model scales and families** on BabiLong QA2 with an 8K train & test context length. PERK’s performance generalizes across models and scales, consistently outperforms FT-ICR baselines.

5 ANALYSIS

In this section, we analyze PERK’s robustness in two challenging settings: test-time length robustness, where PERK is evaluated on longer (or shorter) contexts than those seen during long-context training (Press et al., 2022), and positional variation, where relevant contextual information is distributed throughout the long input context window (Liu et al., 2024).

5.1 TEST-TIME LENGTH ROBUSTNESS

Setup. Focusing on BabiLong tasks, which allow us to generate contexts of arbitrarily long lengths, we train a Qwen-2.5-0.5B base model using PERK on fixed-length contexts ranging from 1K to 8K tokens (as well as an FT-ICR baseline). Subsequently, we test on contexts ranging from 1K to 32K tokens (Figure 4). For analyzing length generalization on extreme context windows, we also extrapolate on contexts with 64K and 128K tokens (Table 1). We also compare PERK to recent training-free length generalization methods, Yarn (Peng et al., 2023b) and Dual Chunk Attention (DCA; An et al., 2024a) by applying these techniques to the FT-ICR Qwen model at inference time.

PERK generalizes to new context lengths at test time. Figure 4 shows that PERK outperforms FT-ICR on the BabiLong QA1 (single-hop) and QA2 (two-hop) tasks for all settings. FT-ICR’s performance drops drastically as the inference context length exceeds the training context length, especially for shorter training contexts. While PERK’s accuracy also decreases (from a higher starting point), the degradation is dampened at longer test lengths. Meanwhile, for interpolation, FT-ICR performance drops on shorter contexts when finetuned on longer ones, matching observations from prior work (Gao et al., 2025b). In contrast, PERK maintains (and even increases) its performance.

To stress-test the limit of PERK models’ length extrapolation ability, we test on sequences with context lengths beyond the context window, 32k, of the Qwen-2.5 model – 64K, and 128K tokens in Table 1, and demonstrate that PERK consistently extrapolates to extreme contexts lengths more robustly than long-context finetuning (FT-ICR). With 128K token context windows, PERK’s accuracy experiences a noticeable drop to 61.4% and 44.4% accuracy for QA1 and QA2, respectively, but these scores remain substantially better than the FT-ICR baseline at this same context length. After applying the training-free length generalization methods, Yarn and DCA, to FT-ICR Qwen models, the baselines show improved generalization on the extreme contexts. However, Yarn + DCA’s extrapolation performance is still far below PERK for both QA1 and QA2 (around 35% lower on average). Compared to the specialized frontier models, PERK outperforms gemini-1.5-pro and the two open-source models, and achieves a narrow performance gap to GPT-4.1 on both 64K and 128K contexts. We note, as well, that PERK is only trained on contexts with 8K tokens, while the Qwen-1M and ProLong models have been trained on much longer contexts with 256K and 512K tokens. Overall, PERK demonstrates strong length generalization, extrapolating to longer contexts while preserving high performance on shorter ones.

| Model | (a) Single-hop reasoning | | | (b) Two-hop reasoning | | |
|---|--------------------------|-------------|-------------|-----------------------|-------------|-------------|
| | 32K | 64K | 128K | 32K | 64K | 128K |
| <i>Commercial Frontier Models</i> | | | | | | |
| GPT-4.1 | 87.8 | 78.7 | 69.4 | 80.6 | 66.4 | 48.2 |
| Gemini-1.5-pro | 87.7 | 82.3 | 73.1 | 56.4 | 48.8 | 40.2 |
| <i>Trained on contexts with >256K tokens</i> | | | | | | |
| Qwen2.5-7B-Instruct-1M | 65.7 | 30.3 | 21.4 | 35.3 | 21.0 | 12.2 |
| ProLong-8B-Instruct-512K | 41.0 | 33.2 | 24.3 | 29.0 | 18.5 | 17.7 |
| <i>Trained on contexts with 8K tokens</i> | | | | | | |
| FT-ICR (Qwen2.5-0.5B) | 43.8 | 34.5 | 0 | 39.4 | 11.7 | 0 |
| w/ Yarn + DCA | 59.2 | 35.5 | 25.4 | 42.5 | 26.3 | 18.5 |
| PERK (Qwen2.5-0.5B) | 95.3 | 88.9 | 61.4 | 80.9 | 62.5 | 44.4 |

Table 1: **Test-time length extrapolation beyond 32K** on BabiLong QA1 (a) and QA2 (b). PERK and FT-ICR are trained on 8K-token sequences. The context length for inference grows from 64K to 128K. PERK extrapolates substantially better than FT-ICR.

5.2 ROBUSTNESS TO POSITIONAL BIASES

Prior work has shown that the position of relevant information in long contexts affects a language model’s ability to use that information (Liu et al., 2024). In this experiment, we test PERK’s robustness to positional biases by evaluating its test-time performance on contexts where the position of relevant information is distributed differently from those seen in training.

Setup. We use API documents from Patil et al. (2024), where each document has an associated API call, and create long contexts containing multiple API-call-to-document pairs. In this task, the model must retrieve the correct API call based on user instructions and documents. Using fixed 4k and 8k token contexts, we train PERK and FT-ICR models on contexts with the target document at varying positions and test them on contexts distributed across all positions.

PERK generalizes regardless of information position in the context. Figure 5 compares PERK to FT-ICR under different position distributions (at train and test time) and context lengths. When trained on contexts with relevant documents randomly located (**Rnd**) throughout the context, PERK

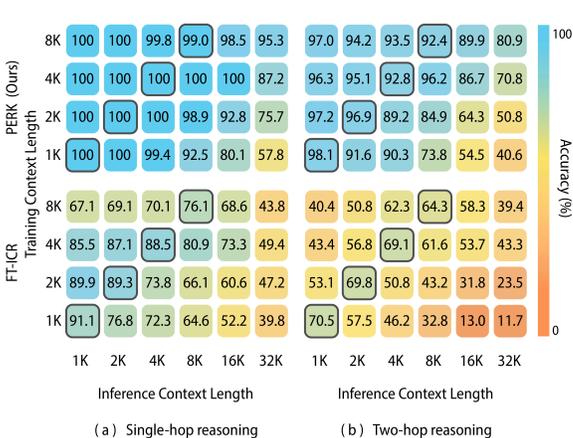


Figure 4: **Test-time context length robustness evaluation** on BabiLong QA1 (a) and QA2 (b) between PERK and FT-ICR on the Qwen-2.5-0.5B model. The y-axis represents the training context lengths, while the x-axis indicates various test-time context lengths. We test for both test lengths shorter than the training length, and test lengths longer than the training length. Bordered cells denote evaluation on context lengths equal to those in training. PERK shows stronger robustness across both settings.

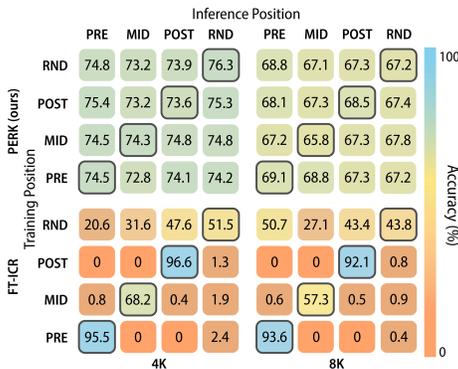


Figure 5: **Positional Bias.** Comparison of PERK and FT-ICR on 4K and 8K contexts, on Qwen-2.5-0.5B. We train on problems where the relevant information appears in the beginning (**Pre**), middle (**Mid**), or end (**Post**) of the context, and evaluate on all three positional settings. We also train models on contexts where the relevant information is randomly located (**Rnd**), testing these on all four positional distributions (**Pre**, **Post**, **Mid**, **Rnd**). Bordered cells show in-distribution performances. PERK demonstrates strong positional robustness.

outperforms FT-ICR by 24%. Furthermore, FT-ICR shows large performance drops when the relevant document *consistently* appears at different locations (**Pre**, **Mid**, and **Post**) at test time. Meanwhile, shifting relevant positions at test-time shows minimal effect (within 1-2%) on PERK.

To further stress-test positional generalization, we also force relevant documents into specific positions during training and testing, spanning the full context. We find that FT-ICR easily overfits to the position pattern, completely failing to generalize to test-time position changes (performance drops to close to 0% when the position shifts at test time). Although PERK underperforms FT-ICR when the train-test positions are the same and the beginning (**Pre**) or end (**Post**) of the context, performance remains consistent regardless of the positional distribution shift between train and test time. We attribute PERK’s robustness in this setting to the fact that documents in the long context are encoded as members of a permutation-invariant *batch* into the parameter space, limiting the overall impact of absolute position in the sequence. Meanwhile, FT-ICR directly aligns specific positions with answers during training.

5.3 INFERENCE LATENCY AND MEMORY USAGE

We evaluate PERK’s efficiency on long contexts by measuring its inference memory cost and run-time, compared to in-context reasoning with finetuned models (FT-ICR).

Setup We complete all runs on a single Nvidia H100 GPU with 93 GB of VRAM and with Huggingface’s Transformers library version 4.51.3. We perform four gradient update steps for the test-time inner loop adaptations for PERK. At inference time, all methods generate 64 tokens through greedy decoding using Huggingface’s text generation pipeline. We measure runtime complexity in seconds (wall-clock time) and memory complexity in gigabytes. Both methods undergo 10 hardware warm-up iterations.

As explained in Section 3, PERK splits long sequences into batches of shorter segments. In this experiment, each batch has segments with an effective context length of 128 tokens. Since PERK performs gradient-based encoding over these batches, we can apply gradient accumulation to trade

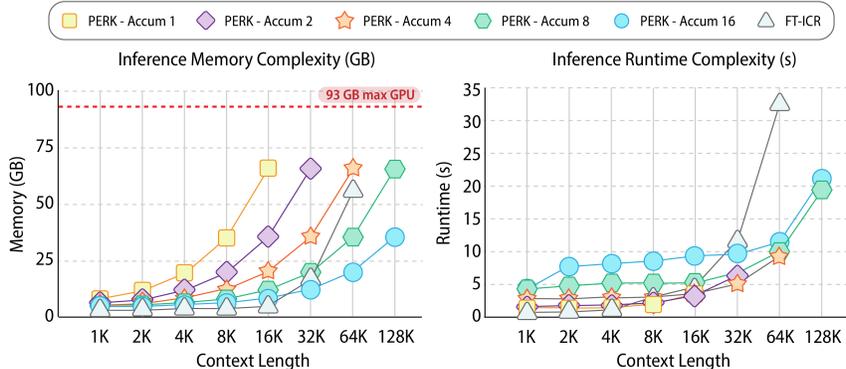


Figure 6: **Inference-time memory footprint and wall-clock runtime** as context length increases (up to 128K tokens), comparing PERK with FT-ICR. Curves that terminate before 128K indicate the method failed with an OOM error at longer context lengths, preventing further measurement. PERK demonstrates more efficient scaling in both memory and runtime, particularly for longer sequences.

runtime for reduced memory consumption. We thus analyze PERK 's efficiency across varying gradient accumulation steps, ranging from 2 to 16.

PERK scales more efficiently in memory and runtime on extremely long contexts. Figure 6 shows that PERK provides more efficient memory and runtime scaling for extremely long contexts compared to FT-ICR. While FT-ICR is initially more efficient, its memory and runtime grow rapidly, leading to OOM errors at a context length of 128K. In contrast, PERK can manage the long sequences through gradient accumulation, which, while increasing runtime, reduces the memory footprint. For instance, at 8K tokens, increasing the accumulation steps from 1 to 16 reduces memory usage from 35.2GB to 5.9GB, although the runtime increases from 1.9s to 8.5s. Crucially, even with this runtime trade-off, PERK runs faster than FT-ICR at very long contexts where both can operate. At 64K tokens, FT-ICR takes 32.6s and 55.7GB, whereas PERK with 16 steps of accumulation completes in 11.4s, using only 19.6GB. With 8 steps, it is even faster at 9.7s, using 35.2 GB. Ultimately, at 128K tokens, where FT-ICR fails, PERK with 16 steps successfully processes the context using 35.2GB in 20.9s, showing that PERK provides a practical path to handle extreme context lengths efficiently in both memory and runtime compared to standard approaches.

6 RELATED WORK

Test-Time Learning. MAML (Finn et al., 2017) introduces optimization-based model-agnostic meta-learning for few-shot adaptation. Chen et al. (2023b) applies MAML to GPT models, integrating knowledge into model parameters through test-time gradient updates for downstream reasoning. Other methods achieve test-time learning in the model architecture. Sun et al. (2025) proposes RNN layers with learnable hidden states updated via self-supervised gradients on test inputs. Titans (Behrouz et al., 2024) incorporates a differentiable memory with transformers for rapid inference-time memorization. Building on Titan, ATLAS (Behrouz et al., 2025) proposes a new transformer architecture that contains the differentiable memory module as part of the network to support large-scale distributed pretraining. Other approaches adapt through context alone. MetaICL (Min et al., 2022) meta-learns from in-context examples without gradient updates. Our method, PERK , introduces parameter-efficient meta-learning using lightweight adapters, improving scalability for long contexts.

Long-Context Language Models. Developing long-context models has become a major research topic of its own. Many works explore extending transformer-based language models' context windows with minimal training, either by position extrapolation (Chen et al., 2023a; Peng et al., 2023c; Xiao et al., 2024; Yen et al., 2024) or improving the transformer's attention mechanism (Bertsch et al., 2023; Mohtashami & Jaggi, 2023; Shen et al., 2023; Jin et al., 2024). Other works show that using the default attention, applying simple position extrapolation, and fine-tuning the model on long documents (Fu et al., 2024; Lu et al., 2024) or synthetically-generated long data (An et al., 2024b; Xiong et al., 2024) yields much stronger results. The FT-ICR baseline in our work follows this paradigm. To reduce the expansive computation cost of long-context training, Chen et al. (2024)'s

work combines LoRA with trainable embedding and normalization to achieve efficient finetuning on long contexts. Other LLM works (Lieber et al., 2024; Grattafiori et al., 2024; Gao et al., 2025b) propose achieving long-context capabilities through a long-context continued training stage between standard pre-training and supervised fine-tuning. In contrast to all the approaches above, PERK achieves long-context reasoning through the ability to encode long sequences using test-time gradient updates on a lightweight model adapter.

Long-Context Architectures. In parallel, there have been many efforts in designing more efficient architectures, for example, linear attention with RNNs (Chandrayan et al., 2023; Peng et al., 2023a; Gu & Dao, 2024; Dao & Gu, 2024; Yang et al., 2024; Beck et al., 2024) and alternative attention architectures (Rubin & Berant, 2024; Sun et al., 2024; Han et al., 2023). However, these new architectures often require training from scratch, and many have inherent limitations in terms of long-context recall (Jelassi et al., 2024; Arora et al., 2025). Recent works also explore hybrid models (Lieber et al., 2024; De et al., 2024) or distilling existing LLMs into hybrid models (Wang et al., 2024) and show promising results. In comparison, PERK directly augments off-the-shelf pre-trained LLMs without requiring additional architecture modifications or extensive pre-training, while naturally being parameter-efficient.

Long-Context Evaluation. Many benchmarks have been proposed to evaluate language models’ long-context abilities (Krishna et al., 2023; Shaham et al., 2023; Bai et al., 2024; Zhang et al., 2024b; Hsieh et al., 2024; Yen et al., 2025; Ye et al., 2025; Xia et al., 2025). Some works focus on the impact of extending input lengths with irrelevant contexts (Kuratov et al., 2024; Levy et al., 2024) or changing relevant information positions (Liu et al., 2024) on LLM reasoning. Many works also reframe natural-language-processing tasks, such as retrieval (Lee et al., 2024; Qiu et al., 2025), summarization (Kim et al., 2024), and in-context learning (Li et al., 2024; Xu et al., 2024; Bertsch et al., 2025), as long-context challenges. In this work, we define *Drops-in-the-Ocean* as a new long-context evaluation that makes distractors distributionally similar to relevant information.

7 CONCLUSION

In this work, we introduce PERK, a parameter-efficient meta-learning approach that learns to perform long-context reasoning by encoding information into a lightweight adapter through test-time gradient updates. Our experiments on multiple long-context benchmarks (Needle-in-the-Haystack, HotpotQA, TriviaQA, Drops-in-the-Ocean, and Code API retrieval) demonstrate substantial performance gains over models finetuned on long contexts for classical in-context reasoning across reasoning complexity and context lengths. We demonstrate that PERK’s performance robustly transfers across model scales (0.5B to 8B) and model families (GPT-2, Qwen2.5, LLaMA-3). Our generalization analysis highlights PERK’s strong test-time length extrapolation ability, where the model can generalize to unseen contexts $32\times$ longer than the training data. When relevant information shifts positions at test-time, PERK robustly generalizes across different positions. Overall, PERK excels in performance, generalization, and robustness for long-context reasoning.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of the Swiss National Science Foundation (No. 215390), Innosuisse (PFFS-21-29), the EPFL Center for Imaging, Sony Group Corporation, and a Meta LLM Evaluation Research Grant.

REFERENCES

- Chenxin An, Fei Huang, Jun Zhang, Shansan Gong, Xipeng Qiu, Chang Zhou, and Lingpeng Kong. Training-free long-context scaling of large language models, 2024a. URL <https://arxiv.org/abs/2402.17463>.
- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, Jian-Guang Lou, and Weizhu Chen. Make your llm fully utilize the context. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 62160–62188. Curran Associates, Inc., 2024b. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/71c3451f6cd6a4f82bb822db25cea4fd-Paper-Conference.pdf.

- Cem Anil, Yuhuai Wu, Anders Andreassen, Aitor Lewkowycz, Vedant Misra, Vinay Ramasesh, Ambrose Slone, Guy Gur-Ari, Ethan Dyer, and Behnam Neyshabur. Exploring length generalization in large language models. *Advances in Neural Information Processing Systems*, 35:38546–38556, 2022.
- Simran Arora, Sabri Eyuboglu, Michael Zhang, Aman Timalsina, Silas Alberti, Dylan Zinsley, James Zou, Atri Rudra, and Christopher Ré. Simple linear attention language models balance the recall-throughput tradeoff. *arXiv preprint arXiv:2402.18668*, 2025. URL <https://arxiv.org/abs/2402.18668>.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to retrieve, generate, and critique through self-reflection, 2023. URL <https://arxiv.org/abs/2310.11511>.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longbench: A bilingual, multitask benchmark for long context understanding, 2024. URL <https://arxiv.org/abs/2308.14508>.
- Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss, and Antoine Bosselut. Discovering knowledge-critical subnetworks in pretrained language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 6549–6583, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.376. URL <https://aclanthology.org/2024.emnlp-main.376/>.
- Maximilian Beck, Korbinian Pöppel, Markus Spanring, Andreas Auer, Oleksandra Prudnikova, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. xlstm: Extended long short-term memory. *arXiv preprint arXiv:2405.04517*, 2024. URL <https://arxiv.org/abs/2405.04517>.
- Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. Titans: Learning to memorize at test time. *arXiv preprint arXiv:2501.00663*, 2024. URL <https://arxiv.org/abs/2501.00663>.
- Ali Behrouz, Zeman Li, Praneeth Kacham, Majid Daliri, Yuan Deng, Peilin Zhong, Meisam Razaviyayn, and Vahab Mirrokni. Atlas: Learning to optimally memorize the context at test time, 2025. URL <https://arxiv.org/abs/2505.23735>.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew Gormley. Unlimiformer: Long-range transformers with unlimited length input. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 35522–35543. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/6f9806a5adc72b5b834b27e4c7c0df9b-Paper-Conference.pdf.
- Amanda Bertsch, Maor Ivgi, Emily Xiao, Uri Alon, Jonathan Berant, Matthew R. Gormley, and Graham Neubig. In-context learning with long-context models: An in-depth exploration. *arXiv preprint arXiv:2405.00200*, 2025. URL <https://arxiv.org/abs/2405.00200>.
- Kedar Chandrayan, Lance Martin, Greg Kamradt, Lazaro Hurtado, Arkady Arkhangorodsky, Lazaro Hurtado, and Pavel Kral. Needle in a haystack - pressure testing llms. 2023.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*, 2023a. URL <https://arxiv.org/abs/2306.15595>.
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. Longlora: Efficient fine-tuning of long-context large language models, 2024. URL <https://arxiv.org/abs/2309.12307>.
- Zeming Chen, Gail Weiss, Eric Mitchell, Asli Celikyilmaz, and Antoine Bosselut. Reckoning: reasoning through dynamic knowledge encoding. *Advances in Neural Information Processing Systems*, 36:62579–62600, 2023b.

- Peter Clark, Oyvind Tafjord, and Kyle Richardson. Transformers as soft reasoners over language. In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 3882–3890. ijcai.org, 2020. doi: 10.24963/IJCAI.2020/537. URL <https://doi.org/10.24963/ijcai.2020/537>.
- Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024. URL <https://arxiv.org/abs/2405.21060>.
- Soham De, Samuel L. Smith, Anushan Fernando, Aleksandar Botev, George Cristian-Muraru, Albert Gu, Ruba Haroun, Leonard Berrada, Yutian Chen, Srivatsan Srinivasan, Guillaume Desjardins, Arnaud Doucet, David Budden, Yee Whye Teh, Razvan Pascanu, Nando De Freitas, and Caglar Gulcehre. Griffin: Mixing gated linear recurrences with local attention for efficient language models. *arXiv preprint arXiv:2402.19427*, 2024. URL <https://arxiv.org/abs/2402.19427>.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hannaneh Hajishirzi, Yoon Kim, and Hao Peng. Data engineering for scaling language models to 128k context. *arXiv preprint arXiv:2402.10171*, 2024. URL <https://arxiv.org/abs/2402.10171>.
- Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively), 2025a. URL <https://arxiv.org/abs/2410.02660>.
- Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively). *arXiv preprint arXiv:2410.02660*, 2025b. URL <https://arxiv.org/abs/2410.02660>.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, JELMER van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale,

Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangrabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojuan Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao,

- Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2024. URL <https://arxiv.org/abs/2312.00752>.
- Insu Han, Rajesh Jayaram, Amin Karbasi, Vahab S. Mirrokni, David P. Woodruff, and Amir Zandieh. Hyperattention: Long-context attention in near-linear time. *ArXiv*, abs/2310.05869, 2023. URL <https://api.semanticscholar.org/CorpusID:263831240>.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekish, Fei Jia, Yang Zhang, and Boris Ginsburg. Ruler: What’s the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654*, 2024. URL <https://arxiv.org/abs/2404.06654>.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- Nathan Hu, Eric Mitchell, Christopher D. Manning, and Chelsea Finn. Meta-learning online adaptation of language models. *arXiv preprint arXiv:2403.04317*, 2023. URL <https://arxiv.org/abs/2305.15076>.
- Samy Jelassi, David Brandfonbrener, Sham M. Kakade, and Eran Malach. Repeat after me: Transformers are better than state space models at copying. *arXiv preprint arXiv:2402.01032*, 2024. URL <https://arxiv.org/abs/2402.01032>.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. Llm maybe longlm: Self-extend llm context window without tuning. *arXiv preprint arXiv:2401.01325*, 2024. URL <https://arxiv.org/abs/2401.01325>.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147/>.
- Damjan Kalajdzievski. A rank stabilization scaling factor for fine-tuning with lora. *arXiv preprint arXiv:2312.03732*, 2023. URL <https://arxiv.org/abs/2312.03732>.
- Petros Karypis, Julian McAuley, and George Karypis. Extending input contexts of language models through training on segmented sequences. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 3040–3052, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.191. URL <https://aclanthology.org/2024.findings-naacl.191/>.
- Yekyung Kim, Yapei Chang, Marzena Karpinska, Aparna Garimella, Varun Manjunatha, Kyle Lo, Tanya Goyal, and Mohit Iyyer. Fables: Evaluating faithfulness and content selection in book-length summarization, 2024. URL <https://arxiv.org/abs/2404.01261>.
- Kalpesh Krishna, Erin Bransom, Bailey Kuehl, Mohit Iyyer, Pradeep Dasigi, Arman Cohan, and Kyle Lo. Longeval: Guidelines for human evaluation of faithfulness in long-form summarization, 2023. URL <https://arxiv.org/abs/2301.13298>.
- Yury Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack. *Advances in Neural Information Processing Systems*, 37:106519–106554, 2024.

- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Taffjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*, 2025. URL <https://arxiv.org/abs/2411.15124>.
- Jinhyuk Lee, Anthony Chen, Zhuyun Dai, Dheeru Dua, Devendra Singh Sachan, Michael Boratko, Yi Luan, Sébastien M. R. Arnold, Vincent Perot, Siddharth Dalmia, Hexiang Hu, Xudong Lin, Panupong Pasupat, Aida Amini, Jeremy R. Cole, Sebastian Riedel, Iftekhar Naim, Ming-Wei Chang, and Kelvin Guu. Can long-context language models subsume retrieval, rag, sql, and more? *arXiv preprint arXiv:2406.13121*, 2024. URL <https://arxiv.org/abs/2406.13121>.
- Mosh Levy, Alon Jacoby, and Yoav Goldberg. Same task, more tokens: the impact of input length on the reasoning performance of large language models. *arXiv preprint arXiv:2402.14848*, 2024. URL <https://arxiv.org/abs/2402.14848>.
- Tianle Li, Ge Zhang, Quy Duc Do, Xiang Yue, and Wenhui Chen. Long-context llms struggle with long in-context learning. *arXiv preprint arXiv:2404.02060*, 2024. URL <https://arxiv.org/abs/2404.02060>.
- Opher Lieber, Barak Lenz, Hofit Bata, Gal Cohen, Jhonathan Osin, Itay Dalmedigos, Erez Safahi, Shaked Meir, Yonatan Belinkov, Shai Shalev-Shwartz, Omri Abend, Raz Alon, Tomer Asida, Amir Bergman, Roman Glozman, Michael Gokhman, Avashalom Manevich, Nir Ratner, Noam Rozen, Erez Shwartz, Mor Zusman, and Yoav Shoham. Jamba: A hybrid transformer-mamba language model. *arXiv preprint arXiv:2403.19887*, 2024. URL <https://arxiv.org/abs/2403.19887>.
- Bingbin Liu, Jordan Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. Exposing attention glitches with flip-flop language modeling. *Advances in Neural Information Processing Systems*, 36:25549–25583, 2023.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. Fineweb-edu: the finest collection of educational content. 2024. doi: 10.57967/hf/2497. URL <https://huggingface.co/datasets/HuggingFaceFW/fineweb-edu>.
- Yi Lu, Jing Nathan Yan, Songlin Yang, Justin T. Chiu, Siyu Ren, Fei Yuan, Wenting Zhao, Zhiyong Wu, and Alexander M. Rush. A controlled study on long context extension and generalization in llms. *arXiv preprint arXiv:2409.12181*, 2024. URL <https://arxiv.org/abs/2409.12181>.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. MetaICL: Learning to learn in context. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (eds.), *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2791–2809, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.201. URL <https://aclanthology.org/2022.naacl-main.201/>.
- Amirkeivan Mohtashami and Martin Jaggi. Random-access infinite context length for transformers. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=7eHn64wOVy>.
- OpenAI. Gpt-4.1. <https://openai.com>, 2025. Large language model.
- Shishir G Patil, Tianjun Zhang, Xin Wang, and Joseph E Gonzalez. Gorilla: Large language model connected with massive apis. *Advances in Neural Information Processing Systems*, 37: 126544–126565, 2024.

- Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Stella Biderman, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, Kranthi Kiran GV, Xuzheng He, Haowen Hou, Jiaju Lin, Przemyslaw Kazienko, Jan Kocon, Jiaming Kong, Bartlomiej Koptyra, Hayden Lau, Krishna Sri Ipsit Mantri, Ferdinand Mom, Atsushi Saito, Guangyu Song, Xiangru Tang, Bolun Wang, Johan S. Wind, Stanislaw Wozniak, Ruichong Zhang, Zhenyuan Zhang, Qihang Zhao, Peng Zhou, Qinghua Zhou, Jian Zhu, and Rui-Jie Zhu. Rwkv: Reinventing rns for the transformer era. *arXiv preprint arXiv:2305.13048*, 2023a. URL <https://arxiv.org/abs/2305.13048>.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models, 2023b. URL <https://arxiv.org/abs/2309.00071>.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*, 2023c. URL <https://arxiv.org/abs/2309.00071>.
- Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2022. URL <https://arxiv.org/abs/2108.12409>.
- Yifu Qiu, Varun Embar, Yizhe Zhang, Navdeep Jaitly, Shay B. Cohen, and Benjamin Han. Eliciting in-context retrieval and reasoning for long-context large language models. *arXiv preprint arXiv:2501.08248*, 2025. URL <https://arxiv.org/abs/2501.08248>.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2025. URL <https://arxiv.org/abs/2412.15115>.
- Alec Radford and Karthik Narasimhan. Improving language understanding by generative pre-training. 2018. URL <https://api.semanticscholar.org/CorpusID:49313245>.
- Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5418–5426, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.437. URL <https://aclanthology.org/2020.emnlp-main.437/>.
- Ohad Rubin and Jonathan Berant. Retrieval-pretrained transformer: Long-range language modeling with self-retrieval. *Transactions of the Association for Computational Linguistics*, 12:1197–1213, 2024. doi: 10.1162/tacl_a_00693. URL <https://aclanthology.org/2024.tacl-1.66/>.
- Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1723–1732. PMLR, 16–18 Apr 2019. URL <https://proceedings.mlr.press/v89/shaban19a.html>.
- Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. ZeroSCROLLS: A zero-shot benchmark for long text understanding. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 7977–7989, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.536. URL <https://aclanthology.org/2023.findings-emnlp.536/>.
- Kai Shen, Junliang Guo, Xu Tan, Siliang Tang, Rui Wang, and Jiang Bian. A study on relu and softmax in transformer. *arXiv preprint arXiv:2302.06461*, 2023. URL <https://arxiv.org/abs/2302.06461>.

- Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Sanmi Koyejo, Tatsunori Hashimoto, and Carlos Guestrin. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv preprint arXiv:2407.04620*, 2025. URL <https://arxiv.org/abs/2407.04620>.
- Yutao Sun, Li Dong, Yi Zhu, Shaohan Huang, Wenhui Wang, Shuming Ma, Quanlu Zhang, Jianyong Wang, and Furu Wei. You only cache once: Decoder-decoder architectures for language models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 7339–7361. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/0df38cd13520747e1e64e5b123a78ef8-Paper-Conference.pdf.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, Andrea Tacchetti, Colin Gaffney, Samira Daruki, Olcan Sercinoglu, Zach Gleicher, Juliette Love, Paul Voigtlaender, Rohan Jain, Gabriela Surita, Kareem Mohamed, Rory Blevins, Junwhan Ahn, Tao Zhu, Kornnaphop Kawintiranon, Orhan Firat, Yiming Gu, Yujing Zhang, Matthew Rahtz, Manaal Faruqui, Natalie Clay, Justin Gilmer, JD Co-Reyes, Ivo Penchev, Rui Zhu, Nobuyuki Morioka, Kevin Hui, Krishna Haridasan, Victor Campos, Mahdis Mahdieh, Mandy Guo, Samer Hassan, Kevin Kilgour, Arpi Vezer, Heng-Tze Cheng, Raoul de Liedekerke, Siddharth Goyal, Paul Barham, DJ Strouse, Seb Noury, Jonas Adler, Mukund Sundararajan, Sharad Vikram, Dmitry Lepikhin, Michela Paganini, Xavier Garcia, Fan Yang, Dasha Valter, Maja Trebacz, Kiran Vodrahalli, Chulayuth Asawaroengchai, Roman Ring, Norbert Kalb, Livio Baldini Soares, Siddhartha Brahma, David Steiner, Tianhe Yu, Fabian Mentzer, Antoine He, Lucas Gonzalez, Bibo Xu, Raphael Lopez Kaufman, Laurent El Shafey, Junhyuk Oh, Tom Hennigan, George van den Driessche, Seth Odoom, Mario Lucic, Becca Roelofs, Sid Lall, Amit Marathe, Betty Chan, Santiago Ontanon, Luheng He, Denis Teplyashin, Jonathan Lai, Phil Crone, Bogdan Damoc, Lewis Ho, Sebastian Riedel, Karel Lenc, Chih-Kuan Yeh, Aakanksha Chowdhery, Yang Xu, Mehran Kazemi, Ehsan Amid, Anastasia Petrushkina, Kevin Swersky, Ali Khodaei, Gowoon Chen, Chris Larkin, Mario Pinto, Geng Yan, Adria Puigdomenech Badia, Piyush Patil, Steven Hansen, Dave Orr, Sebastien M. R. Arnold, Jordan Grimstad, Andrew Dai, Sholto Douglas, Rishika Sinha, Vikas Yadav, Xi Chen, Elena Gribovskaya, Jacob Austin, Jeffrey Zhao, Kaushal Patel, Paul Komarek, Sophia Austin, Sebastian Borgeaud, Linda Friso, Abhimanyu Goyal, Ben Caine, Kris Cao, Da-Woon Chung, Matthew Lamm, Gabe Barth-Maron, Thais Kagohara, Kate Olszewska, Mia Chen, Kaushik Shivakumar, Rishabh Agarwal, Harshal Godhia, Ravi Rajwar, Javier Snaider, Xerxes Dotiwalla, Yuan Liu, Aditya Barua, Victor Ungureanu, Yuan Zhang, Bat-Orgil Batsaikhan, Mateo Wirth, James Qin, Ivo Danihelka, Tulsee Doshi, Martin Chadwick, Jilin Chen, Sanil Jain, Quoc Le, Arjun Kar, Madhu Gurusurthy, Cheng Li, Ruoxin Sang, Fangyu Liu, Lampros Lamprou, Rich Munoz, Nathan Lintz, Harsh Mehta, Heidi Howard, Malcolm Reynolds, Lora Aroyo, Quan Wang, Lorenzo Blanco, Albin Cassirer, Jordan Griffith, Dipanjan Das, Stephan Lee, Jakub Sygnowski, Zach Fisher, James Besley, Richard Powell, Zafarali Ahmed, Dominik Paulus, David Reitter, Zalan Borsos, Rishabh Joshi, Aedan Pope, Steven Hand, Vittorio Selo, Vihan Jain, Nikhil Sethi, Megha Goel, Takaki Makino, Rhys May, Zhen Yang, Johan Schalkwyk, Christina Butterfield, Anja Hauth, Alex Goldin, Will Hawkins, Evan Senter, Sergey Brin, Oliver Woodman, Marvin Ritter, Eric Noland, Minh Giang, Vijay Bolina, Lisa Lee, Tim Blyth, Ian Mackinnon, Machel Reid, Obaid Sarvana, David Silver, Alexander Chen, Lily Wang, Loren Maggiore, Oscar Chang, Nithya Attaluri, Gregory Thornton, Chung-Cheng Chiu, Oskar Bunyan, Nir Levine, Timothy Chung, Evgenii Eltyshev, Xiance Si, Timothy Lillicrap, Demetra Brady, Vaibhav Aggarwal, Boxi Wu, Yuanzhong Xu, Ross McIlroy, Kartikeya Badola, Paramjit Sandhu, Erica Moreira, Wojciech Stokowiec, Ross Hemsley, Dong Li, Alex Tudor, Pranav Shyam, Elahe Rahimtoroghi, Salem Haykal, Pablo Sprechmann, Xiang Zhou, Diana Mincu, Yujia Li, Ravi Addanki, Kalpesh Krishna, Xiao Wu, Alexandre Frechette, Matan Eyal, Allan Dafoe, Dave Lacey, Jay Whang, Thi Avrahami, Ye Zhang, Emanuel Taropa, Hanzhao Lin, Daniel Toyama, Eliza Rutherford, Motoki Sano, HyunJeong Choe, Alex Tomala, Chalence Safranek-Shrader, Nora Kassner, Mantas Pajarskas, Matt Harvey, Sean Sechrist, Meire Fortunato, Christina Lyu, Gamaleldin Elsayed, Chenkai Kuang, James Lottes, Eric Chu, Chao Jia, Chih-Wei Chen, Peter Humphreys, Kate Baumli, Connie Tao, Rajkumar Samuel, Cicero Nogueira dos Santos, Anders Andreassen, Nemanja Rakićević, Dominik Grewe, Aviral Kumar, Stephanie Winkler, Jonathan Caton, Andrew Brock, Sid Dalmia, Hannah Sheahan, Iain Barr, Yingjie Miao, Paul Natsev, Jacob Devlin, Feryal Behbahani, Flavien Prost, Yanhua Sun, Artiom Myaskovsky,

Thanumalayan Sankaranarayana Pillai, Dan Hurt, Angeliki Lazaridou, Xi Xiong, Ce Zheng, Fabio Pardo, Xiaowei Li, Dan Horgan, Joe Stanton, Moran Ambar, Fei Xia, Alejandro Lince, Mingqiu Wang, Basil Mustafa, Albert Webson, Hyo Lee, Rohan Anil, Martin Wicke, Timothy Dozat, Abhishek Sinha, Enrique Piqueras, Elahe Dabir, Shyam Upadhyay, Anudhyan Boral, Lisa Anne Hendricks, Corey Fry, Josip Djolonga, Yi Su, Jake Walker, Jane Labanowski, Ronny Huang, Vedant Misra, Jeremy Chen, RJ Skerry-Ryan, Avi Singh, Shruti Rijhwani, Dian Yu, Alex Castro-Ros, Beer Changpinyo, Romina Datta, Sumit Bagri, Arnar Mar Hrafnkelsson, Marcello Maggioni, Daniel Zheng, Yury Sulsky, Shaobo Hou, Tom Le Paine, Antoine Yang, Jason Riesa, Dominika Rogozinska, Dror Marcus, Dalia El Badawy, Qiao Zhang, Luyu Wang, Helen Miller, Jeremy Greer, Lars Lowe Sjos, Azade Nova, Heiga Zen, Rahma Chaabouni, Mihaela Rosca, Jiepu Jiang, Charlie Chen, Ruibo Liu, Tara Sainath, Maxim Krikun, Alex Polozov, Jean-Baptiste Lespiau, Josh Newlan, Zeyncep Cankara, Soo Kwak, Yunhan Xu, Phil Chen, Andy Coenen, Clemens Meyer, Katerina Tsihlias, Ada Ma, Juraj Gottweis, Jinwei Xing, Chenjie Gu, Jin Miao, Christian Frank, Zeynep Cankara, Sanjay Ganapathy, Ishita Dasgupta, Steph Hughes-Fitt, Heng Chen, David Reid, Keran Rong, Hongmin Fan, Joost van Amersfoort, Vincent Zhuang, Aaron Cohen, Shixiang Shane Gu, Anhad Mohananey, Anastasija Ilic, Taylor Tobin, John Wieting, Anna Bortsova, Phoebe Thacker, Emma Wang, Emily Caveness, Justin Chiu, Eren Sezener, Alex Kaskasoli, Steven Baker, Katie Millican, Mohamed Elhawy, Kostas Aisopos, Carl Lebsack, Nathan Byrd, Hanjun Dai, Wenhao Jia, Matthew Wiethoff, Elnaz Davoodi, Albert Weston, Lakshman Yagati, Arun Ahuja, Isabel Gao, Golan Pundak, Susan Zhang, Michael Azzam, Khe Chai Sim, Sergi Caelles, James Keeling, Abhanshu Sharma, Andy Swing, YaGuang Li, Chenxi Liu, Carrie Grimes Bostock, Yamini Bansal, Zachary Nado, Ankesh Anand, Josh Lipschultz, Abhijit Karmarkar, Lev Proleev, Abe Ittycheriah, Soheil Hassas Yeganeh, George Polovets, Aleksandra Faust, Jiao Sun, Alban Rrustemi, Pen Li, Rakesh Shivanna, Jeremiah Liu, Chris Welty, Federico Lebron, Anirudh Baddepudi, Sebastian Krause, Emilio Parisotto, Radu Soricut, Zheng Xu, Dawn Bloxwich, Melvin Johnson, Behnam Neyshabur, Justin Mao-Jones, Renshen Wang, Vinay Ramasesh, Zaheer Abbas, Arthur Guez, Constant Segal, Duc Dung Nguyen, James Svensson, Le Hou, Sarah York, Kieran Milan, Sophie Bridgers, Wiktor Gworek, Marco Tagliasacchi, James Lee-Thorp, Michael Chang, Alexey Guseynov, Ale Jakse Hartman, Michael Kwong, Ruizhe Zhao, Sheleem Kashem, Elizabeth Cole, Antoine Miech, Richard Tanburn, Mary Phuong, Filip Pavetic, Sebastien Cevy, Ramona Comanescu, Richard Ives, Sherry Yang, Cosmo Du, Bo Li, Zizhao Zhang, Mariko Iinuma, Clara Huiyi Hu, Aurko Roy, Shaan Bijwadia, Zhenkai Zhu, Danilo Martins, Rachel Saputro, Anita Gergely, Steven Zheng, Dawei Jia, Ioannis Antonoglou, Adam Sadovsky, Shane Gu, Yingying Bi, Alek Andreev, Sina Samangooei, Mina Khan, Tomas Kocisky, Angelos Filos, Chintu Kumar, Colton Bishop, Adams Yu, Sarah Hodkinson, Sid Mittal, Premal Shah, Alexandre Moufarek, Yong Cheng, Adam Bloniarz, Jaehoon Lee, Pedram Pejman, Paul Michel, Stephen Spencer, Vladimir Feinberg, Xuehan Xiong, Nikolay Savinov, Charlotte Smith, Siamak Shakeri, Dustin Tran, Mary Chesus, Bernd Bohnet, George Tucker, Tamara von Glehn, Carrie Muir, Yiran Mao, Hideto Kazawa, Ambrose Slone, Kedar Soparkar, Disha Shrivastava, James Cobon-Kerr, Michael Sharman, Jay Pavagadhi, Carlos Araya, Karolis Misiunas, Nimesh Ghelani, Michael Laskin, David Barker, Qiuqia Li, Anton Briukhov, Neil Houlsby, Mia Glaese, Balaji Lakshminarayanan, Nathan Schucher, Yunhao Tang, Eli Collins, Hyeontaek Lim, Fangxiaoyu Feng, Adria Recasens, Guangda Lai, Alberto Magni, Nicola De Cao, Aditya Siddhant, Zoe Ashwood, Jordi Orbay, Mostafa Dehghani, Jenny Brennan, Yifan He, Kelvin Xu, Yang Gao, Carl Saroufim, James Molloy, Xinyi Wu, Seb Arnold, Solomon Chang, Julian Schrittwieser, Elena Buchatskaya, Soroush Radpour, Martin Polacek, Skye Giordano, Ankur Bapna, Simon Tokumine, Vincent Hellendoorn, Thibault Sottiaux, Sarah Cogan, Aliaksei Severyn, Mohammad Saleh, Shantanu Thakoor, Laurent Shefey, Siyuan Qiao, Meenu Gaba, Shuo yiin Chang, Craig Swanson, Biao Zhang, Benjamin Lee, Paul Kishan Rubenstein, Gan Song, Tom Kwiatkowski, Anna Koop, Ajay Kannan, David Kao, Parker Schuh, Axel Stjerngren, Golnaz Ghiasi, Gena Gibson, Luke Vilnis, Ye Yuan, Felipe Tiengo Ferreira, Aishwarya Kamath, Ted Klimentko, Ken Franko, Kefan Xiao, Indro Bhattacharya, Miteyan Patel, Rui Wang, Alex Morris, Robin Strudel, Vivek Sharma, Peter Choy, Sayed Hadi Hashemi, Jessica Landon, Mara Finkelstein, Priya Jhakra, Justin Frye, Megan Barnes, Matthew Mauger, Dennis Daun, Khuslen Baatarsukh, Matthew Tung, Wael Farhan, Henryk Michalewski, Fabio Viola, Felix de Chaumont Quiry, Charline Le Lan, Tom Hudson, Qingze Wang, Felix Fischer, Ivy Zheng, Elspeth White, Anca Dragan, Jean baptiste Alayrac, Eric Ni, Alexander Pritzel, Adam Iwanicki, Michael Isard, Anna Bulanova, Lukas Zilka, Ethan Dyer, Devendra Sachan, Srivatsan Srinivasan, Hannah Muckenhirn, Honglong Cai, Amol Mandhane, Mukarram Tariq, Jack W. Rae, Gary Wang, Kareem Ayoub, Nicholas FitzGerald, Yao Zhao, Woohyun Han, Chris Alberti, Dan Garrette,

Kashyap Krishnakumar, Mai Gimenez, Anselm Levskaya, Daniel Sohn, Josip Matak, Inaki Iturrate, Michael B. Chang, Jackie Xiang, Yuan Cao, Nishant Ranka, Geoff Brown, Adrian Hutter, Vahab Mirrokni, Nanxin Chen, Kaisheng Yao, Zoltan Egyed, Francois Galilee, Tyler Liechty, Praveen Kallakuri, Evan Palmer, Sanjay Ghemawat, Jasmine Liu, David Tao, Chloe Thornton, Tim Green, Mimi Jasarevic, Sharon Lin, Victor Cotruta, Yi-Xuan Tan, Noah Fiedel, Hongkun Yu, Ed Chi, Alexander Neitz, Jens Heitkaemper, Anu Sinha, Denny Zhou, Yi Sun, Charbel Kaed, Brice Hulse, Swaroop Mishra, Maria Georgaki, Sneha Kudugunta, Clement Farabet, Izhak Shafran, Daniel Vlasic, Anton Tsitsulin, Rajagopal Ananthanarayanan, Alen Carin, Guolong Su, Pei Sun, Shashank V, Gabriel Carvajal, Josef Broder, Iulia Comsa, Alena Repina, William Wong, Warren Weilun Chen, Peter Hawkins, Egor Filonov, Lucia Loher, Christoph Hirsenschall, Weiyi Wang, Jingchen Ye, Andrea Burns, Hardie Cate, Diana Gage Wright, Federico Piccinini, Lei Zhang, Chu-Cheng Lin, Ionel Gog, Yana Kulizhskaya, Ashwin Sreevatsa, Shuang Song, Luis C. Cobo, Anand Iyer, Chetan Tekur, Guillermo Garrido, Zhu Yun Xiao, Rupert Kemp, Huaixiu Steven Zheng, Hui Li, Ananth Agarwal, Christel Ngani, Kati Goshvadi, Rebeca Santamaria-Fernandez, Wojciech Fica, Xinyun Chen, Chris Gorgolewski, Sean Sun, Roopal Garg, Xinyu Ye, S. M. Ali Eslami, Nan Hua, Jon Simon, Pratik Joshi, Yelin Kim, Ian Tenney, Sahitya Potluri, Lam Nguyen Thiet, Quan Yuan, Florian Luisier, Alexandra Chronopoulou, Salvatore Scellato, Praveen Srinivasan, Minmin Chen, Vinod Koverkathu, Valentin Dalibard, Yaming Xu, Brennan Saeta, Keith Anderson, Thibault Sellam, Nick Fernando, Fantine Huot, Junehyuk Jung, Mani Varadarajan, Michael Quinn, Amit Raul, Maigo Le, Ruslan Habalov, Jon Clark, Komal Jalan, Kalesha Bullard, Achintya Singhal, Thang Luong, Boyu Wang, Sujeevan Rajayogam, Julian Eisenschlos, Johnson Jia, Daniel Finchelstein, Alex Yakubovich, Daniel Balle, Michael Fink, Sameer Agarwal, Jing Li, Dj Dvijotham, Shalini Pal, Kai Kang, Jaclyn Konzelmann, Jennifer Beattie, Olivier Dousse, Diane Wu, Remi Crocker, Chen Elkind, Siddhartha Reddy Jonnalagadda, Jong Lee, Dan Holtmann-Rice, Krystal Kallarackal, Rosanne Liu, Denis Vnukov, Neera Vats, Luca Invernizzi, Mohsen Jafari, Huanjie Zhou, Lilly Taylor, Jennifer Prendki, Marcus Wu, Tom Eccles, Tianqi Liu, Kavya Koppurapu, Francoise Beaufays, Christof Angermueller, Andreea Marzoca, Shourya Sarcar, Hilal Dib, Jeff Stanway, Frank Perbet, Nejc Trdin, Rachel Sterneck, Andrey Khorlin, Dinghua Li, Xihui Wu, Sonam Goenka, David Madras, Sasha Goldshtein, Willi Gierke, Tong Zhou, Yaxin Liu, Yannie Liang, Anais White, Yunjie Li, Shreya Singh, Sanaz Bahargam, Mark Epstein, Sujoy Basu, Li Lao, Adnan Ozturel, Carl Crous, Alex Zhai, Han Lu, Zora Tung, Neeraj Gaur, Alanna Walton, Lucas Dixon, Ming Zhang, Amir Globerson, Grant Uy, Andrew Bolt, Olivia Wiles, Milad Nasr, Ilia Shumailov, Marco Selvi, Francesco Piccinno, Ricardo Aguilar, Sara McCarthy, Misha Khalman, Mrinal Shukla, Vlado Galic, John Carpenter, Kevin Vilella, Haibin Zhang, Harry Richardson, James Martens, Matko Bosnjak, Shreyas Rammohan Belle, Jeff Seibert, Mahmoud Alnahlawi, Brian McWilliams, Sankalp Singh, Annie Louis, Wen Ding, Dan Popovici, Lenin Simicich, Laura Knight, Pulkit Mehta, Nishesh Gupta, Chongyang Shi, Saaber Fatehi, Jovana Mitrovic, Alex Grills, Joseph Pagadora, Tsendsuren Munkhdalai, Dessie Petrova, Danielle Eisenbud, Zhishuai Zhang, Damion Yates, Bhavishya Mittal, Nilesh Tripuraneni, Yannis Assael, Thomas Brovelli, Prateek Jain, Mihajlo Velimirovic, Canfer Akbulut, Jiaqi Mu, Wolfgang Macherey, Ravin Kumar, Jun Xu, Haroon Qureshi, Gheorghe Comanici, Jeremy Wiesner, Zhitao Gong, Anton Ruddock, Matthias Bauer, Nick Felt, Anirudh GP, Anurag Arnab, Dustin Zelle, Jonas Rothfuss, Bill Rosgen, Ashish Shenoy, Bryan Seybold, Xinjian Li, Jayaram Mudigonda, Goker Erdogan, Jiawei Xia, Jiri Simsa, Andrea Michi, Yi Yao, Christopher Yew, Steven Kan, Isaac Caswell, Carey Radebaugh, Andre Elisseeff, Pedro Valenzuela, Kay McKinney, Kim Paterson, Albert Cui, Eri Latorre-Chimoto, Solomon Kim, William Zeng, Ken Durden, Priya Ponnappalli, Tiberiu Sosea, Christopher A. Choquette-Choo, James Manyika, Brona Robenek, Harsha Vashisht, Sebastien Pereira, Hoi Lam, Marko Velic, Denese Owusu-Afriyie, Katherine Lee, Tolga Bolukbasi, Alicia Parrish, Shawn Lu, Jane Park, Balaji Venkatraman, Alice Talbert, Lambert Rosique, Yuchung Cheng, Andrei Sozanschi, Adam Paszke, Praveen Kumar, Jessica Austin, Lu Li, Khalid Salama, Bartek Perz, Wooyeol Kim, Nandita Dukkupati, Anthony Baryshnikov, Christos Kaplanis, XiangHai Sheng, Yuri Chervonyi, Caglar Unlu, Diego de Las Casas, Harry Askham, Kathryn Tunyasuvunakool, Felix Gimeno, Siim Poder, Chester Kwak, Matt Miecznikowski, Vahab Mirrokni, Alek Dimitriev, Aaron Parisi, Danyi Liu, Tomy Tsai, Toby Shevlane, Christina Kouridi, Drew Garmon, Adrian Goedeckemeyer, Adam R. Brown, Anitha Vijayakumar, Ali Elqursh, Sadegh Jazayeri, Jin Huang, Sara Mc Carthy, Jay Hoover, Lucy Kim, Sandeep Kumar, Wei Chen, Courtney Biles, Garrett Bingham, Evan Rosen, Lisa Wang, Qijun Tan, David Engel, Francesco Pongetti, Dario de Cesare, Dongseong Hwang, Lily Yu, Jennifer Pullman, Srini Narayanan, Kyle Levin, Siddharth Gopal, Megan Li, Asaf Aharoni, Trieu Trinh, Jessica Lo, Norman Casagrande, Roopali Vij, Loic Matthey, Bramandia Ramadhana,

- Austin Matthews, CJ Carey, Matthew Johnson, Kremena Goranova, Rohin Shah, Shereen Ashraf, Kingshuk Dasgupta, Rasmus Larsen, Yicheng Wang, Manish Reddy Vuyyuru, Chong Jiang, Joana Ijazi, Kazuki Osawa, Celine Smith, Ramya Sree Boppana, Taylan Bilal, Yuma Koizumi, Ying Xu, Yasemin Altun, Nir Shabat, Ben Bariach, Alex Korchemniy, Kiam Choo, Olaf Ronneberger, Chimezie Iwuanyanwu, Shubin Zhao, David Soergel, Cho-Jui Hsieh, Irene Cai, Shariq Iqbal, Martin Sundermeyer, Zhe Chen, Elie Bursztein, Chaitanya Malaviya, Fadi Biadisy, Prakash Shroff, Inderjit Dhillon, Tejasi Latkar, Chris Dyer, Hannah Forbes, Massimo Nicosia, Vitaly Nikolaev, Somer Greene, Marin Georgiev, Pidong Wang, Nina Martin, Hanie Sedghi, John Zhang, Praseem Banzal, Doug Fritz, Vikram Rao, Xuezhi Wang, Jiageng Zhang, Viorica Patraucean, Dayou Du, Igor Mordatch, Ivan Jurin, Lewis Liu, Ayush Dubey, Abhi Mohan, Janek Nowakowski, Vlad-Doru Ion, Nan Wei, Reiko Tojo, Maria Abi Raad, Drew A. Hudson, Vaishakh Keshava, Shubham Agrawal, Kevin Ramirez, Zhichun Wu, Hoang Nguyen, Ji Liu, Madhavi Sewak, Bryce Petri, DongHyun Choi, Ivan Philips, Ziyue Wang, Ioana Bica, Ankush Garg, Jarek Wilkiewicz, Priyanka Agrawal, Xiaowei Li, Danhao Guo, Emily Xue, Naseer Shaik, Andrew Leach, Sadh MNM Khan, Julia Wiesinger, Sammy Jerome, Abhishek Chakladar, Alek Wenjiao Wang, Tina Ornduff, Folake Abu, Alireza Ghaffarkhah, Marcus Wainwright, Mario Cortes, Frederick Liu, Joshua Maynez, Andreas Terzis, Pouya Samangouei, Riham Mansour, Tomasz Kępa, François-Xavier Aubet, Anton Algyrn, Dan Banica, Agoston Weisz, Andras Orban, Alexandre Senges, Ewa Andrejczuk, Mark Geller, Niccolo Dal Santo, Valentin Anklin, Majid Al Merey, Martin Baeuml, Trevor Strohmman, Junwen Bai, Slav Petrov, Yonghui Wu, Demis Hassabis, Koray Kavukcuoglu, Jeff Dean, and Oriol Vinyals. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context, 2024. URL <https://arxiv.org/abs/2403.05530>.
- Junxiong Wang, Daniele Paliotta, Avner May, Alexander M. Rush, and Tri Dao. The mamba in the llama: Distilling and accelerating hybrid models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 62432–62457. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/723933067ad315269b620bc0d2c05cba-Paper-Conference.pdf.
- Menglin Xia, Victor Rühle, Saravan Rajmohan, and Reza Shokri. Minerva: A programmable memory test benchmark for language models. *arXiv preprint arXiv:2502.03358*, 2025. URL <https://arxiv.org/abs/2502.03358>.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2024. URL <https://arxiv.org/abs/2309.17453>.
- Zheyang Xiong, Vasilis Papageorgiou, Kangwook Lee, and Dimitris Papailiopoulos. From artificial needles to real haystacks: Improving retrieval capabilities in llms by finetuning on synthetic data. *arXiv preprint arXiv:2406.19292*, 2024. URL <https://arxiv.org/abs/2406.19292>.
- Xiaoyue Xu, Qinyuan Ye, and Xiang Ren. Stress-testing long-context language models with lifelong icl and task haystack. *arXiv preprint arXiv:2407.16695*, 2024. URL <https://arxiv.org/abs/2407.16695>.
- An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, Junyang Lin, Kai Dang, Kexin Yang, Le Yu, Mei Li, Minmin Sun, Qin Zhu, Rui Men, Tao He, Weijia Xu, Wenbiao Yin, Wenyuan Yu, Xiafei Qiu, Xingzhang Ren, Xinlong Yang, Yong Li, Zhiying Xu, and Zipeng Zhang. Qwen2.5-1m technical report, 2025. URL <https://arxiv.org/abs/2501.15383>.
- Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. *arXiv preprint arXiv:2312.06635*, 2024. URL <https://arxiv.org/abs/2312.06635>.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL <https://aclanthology.org/D18-1259/>.

- Han-Jia Ye and Wei-Lun Chao. How to train your MAML to excel in few-shot classification. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=49h_IkpJtaE.
- Xi Ye, Fangcong Yin, Yinghui He, Joie Zhang, Howard Yen, Tianyu Gao, Greg Durrett, and Danqi Chen. Longproc: Benchmarking long-context language models on long procedural generation. *arXiv preprint arXiv:2501.05414*, 2025. URL <https://arxiv.org/abs/2501.05414>.
- Howard Yen, Tianyu Gao, and Danqi Chen. Long-context language modeling with parallel context encoding. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2588–2610, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.142. URL <https://aclanthology.org/2024.acl-long.142/>.
- Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat, and Danqi Chen. Helmet: How to evaluate long-context language models effectively and thoroughly. *arXiv preprint arXiv:2410.02694*, 2025. URL <https://arxiv.org/abs/2410.02694>.
- Xin Zhang, Yanzhao Zhang, Dingkun Long, Wen Xie, Ziqi Dai, Jialong Tang, Huan Lin, Baosong Yang, Pengjun Xie, Fei Huang, Meishan Zhang, Wenjie Li, and Min Zhang. mgte: Generalized long-context text representation and reranking models for multilingual text retrieval, 2024a. URL <https://arxiv.org/abs/2407.19669>.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. ∞ Bench: Extending long context evaluation beyond 100K tokens. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15262–15277, Bangkok, Thailand, August 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.814. URL <https://aclanthology.org/2024.acl-long.814/>.

A DATASETS

In this section, we provide details on the datasets used for each task in our experiments and analysis. For each dataset, we visualize an example in Figure 7.

A.1 BABILONG

We construct the train and evaluation data using the data generation framework proposed by Kuratov et al. (2024). As we mentioned in the experiment setup, we mainly focus on tasks involving single-hop (QA1), two-hop (QA2), and three-hop (QA3) reasoning.

Tasks The main task for BabiLong is to reason over facts distributed in a long context. For each QA task, we have a set of facts, including distractors and supporting facts. Each fact describes an agent’s action or movement (which also describes the agent’s location). For example, *Mary journeyed to the garden* is a movement, where *Mary* is the agent, and *garden* is Mary’s current location. As an example for actions, *John picked up the apple there* describes John’s action of picking up an apple. Together with multiple agents performing multiple movements and actions, the set of facts became a temporal sequence of events.

QA1 focuses on one-hop reasoning, where the problems can be answered by using a single supporting fact. For example, in a sequence of actions and movements, the agent *John*’s last movement is *John moved to the bedroom*. Then the one-hop question is “*Where is John located now?*”. By the last movement alone, we can find the answer “*bedroom*”.

QA2 focuses on one-hop reasoning, where the problems can be answered by using two supporting facts, where one fact is a movement and the other one is an action. For example, in the event sequence, we have *John moved to the garden* and *John dropped the apple there*. Using these two facts, we can answer “*Where is the apple now?*”, which is “*garden*”.

QA3 focuses on one-hop reasoning, where the problems can be answered by using three supporting facts, as a combination of movements and actions. For example, we have three supporting facts: *Daniel went back to the office*, *Daniel went back to the bedroom*, and *Daniel left the milk* in order. We can infer the answer to the question “*Where was the milk before the bedroom?*” is “*office*”.

Build Long Context To build the long context, we use the framework to construct the tasks with long natural documents sampled from FineWeb-edu (Lozhkov et al., 2024), specifically from the *CC-MAIN-2024-42* subset, ensuring the content postdates our base model’s knowledge cutoff. Given a particular context window, such as 8192, we sample that many tokens from the corpus as the *Haystack* long context.

We need to convert a long context into a batch of subsequences, such that PERK can encode them in parallel through test-time gradient-based adaptation. In practice, we set the effective context length (i.e., the number of tokens of a subsequence on average) as 256. We divide the long context into chunks where each chunk is approximately 256 tokens long. Then, we randomly distribute the set of facts into these chunks. For each fact, we randomly select a chunk and then insert this fact into the chunk at a random position. To preserve the temporal order of the facts (since they come from a sequence of events with temporal order), we index each fact with a prefix. Finally, we have a batch of subsequences where each sequence might contain some facts at random positions.

A.2 STUDENT RECORDS

In the main paper, we propose *Drops-in-the-Ocean* (DIO), a new evaluation setting where long contexts are formed from structurally similar documents. We aim to create problems in which the relevant and irrelevant information will likely be distributionally similar, making it more difficult to identify critical facts. We construct a synthetic dataset, **Student Records**, where each context simulates a database containing multiple student records. Each record includes several attributes: ID, name, school, major, and grade.

Tasks The task operates on top of a database that consists of multiple records. Each record holds information about a student. The record is structured as follows:

Student Id: 3109998, Student Name: Alison Keith, Year: Sophomore

School: School of Computer Information and Data Sciences, Major: Computer Science, Grade: 72

We generate the student IDs and grades using a random number generator. The student names are generated from the Name package of Python. We define a pool of years, schools, and courses. Then, we randomly select each entity for each student. Crucially, to test generalization, specific entities (e.g., IDs, names) and attribute value combinations encountered in the test set are disjoint from those used during training.

Based on the database of student records, we construct three main tasks: **Recall** (retrieving attributes for a specific ID), **Relation** (comparing attributes between two IDs), and **Aggregate** (calculating the maximum, minimum, and average grade across all students).

Recall tests a model’s ability to retrieve an attribute of a student given a specific student ID number. For example, we can ask “*What major does student 1778924 study?*”. The model should answer “*What major does student 1778924 study?*”. The questions cover all attributes for testing coverage.

Relation tests a model’s ability to compare the attributes of two students. For example, for the attribute of grade, we can ask “*Does student 1778924 have a higher grade than student 7077015?*”. The model should be able to compare their grades and infer “*Yes*”. Or we can also ask “*Do student 1778924 and student 7077015 study the same major?*”, the model should recognize that their majors are different and answer “*No*”.

Aggregation is the most difficult task in this dataset. This task requires many-hop reasoning ability, where the model needs to aggregate attributes from all students in a long context to make an accurate conclusion. We focus on three types of numerical aggregation among the students’ grades: (1) Maximization, (2) Minimization, and (3) Averaging. For example, we ask the model “*What is the average grade of all students?*”, the model should be able to compute the sum of all student grades, count the total number of students in the context, and then compute the average, “*49.3*”.

Build Long Context Context length scales directly with the number of student records included. For a given context window, we continuously append student records to the context until the number of tokens exceeds the context window. To convert the long context into a batch of subsequences for PERK, we simply treat each student record as an individual subsequence in the batch, since they are already independent documents by design.

A.3 OPEN-DOMAIN MULTI-DOC QA

We select HotpotQA (Yang et al., 2018) and TriviaQA (Joshi et al., 2017) as the main evaluation for open-domain QA over long-context documents. HotpotQA tests the model’s ability to perform 2-hop reasoning over two gold documents randomly located in a long list of documents. TriviaQA tests the model’s ability to do reading comprehension over very long documents where the target information can be scattered around the long context. We refer to the original papers for QA construction details, data sources, and task examples of these two benchmarks.

Build Long Context To build long contexts for these two benchmarks, given a specific context window (for example, 32768 tokens), we follow Yen et al. (2025)’s process for their RAG evaluation. We extend the original context by injecting retrieved distractor documents (Zhang et al., 2024a) at random positions until reaching the context window limit.

A.4 LOST-IN-THE-MIDDLE API RETRIEVAL

To analyze a model’s robustness against relevant information positions, we follow the setting proposed by *Lost in the Middle* (Liu et al., 2024). We use API documents from Patil et al. (2024), where each document has an associated API call, and create long contexts containing multiple API-call-to-document pairs. We place the relevant API document in a particular user instruction at different locations in the context (beginning, middle, end, and random position).

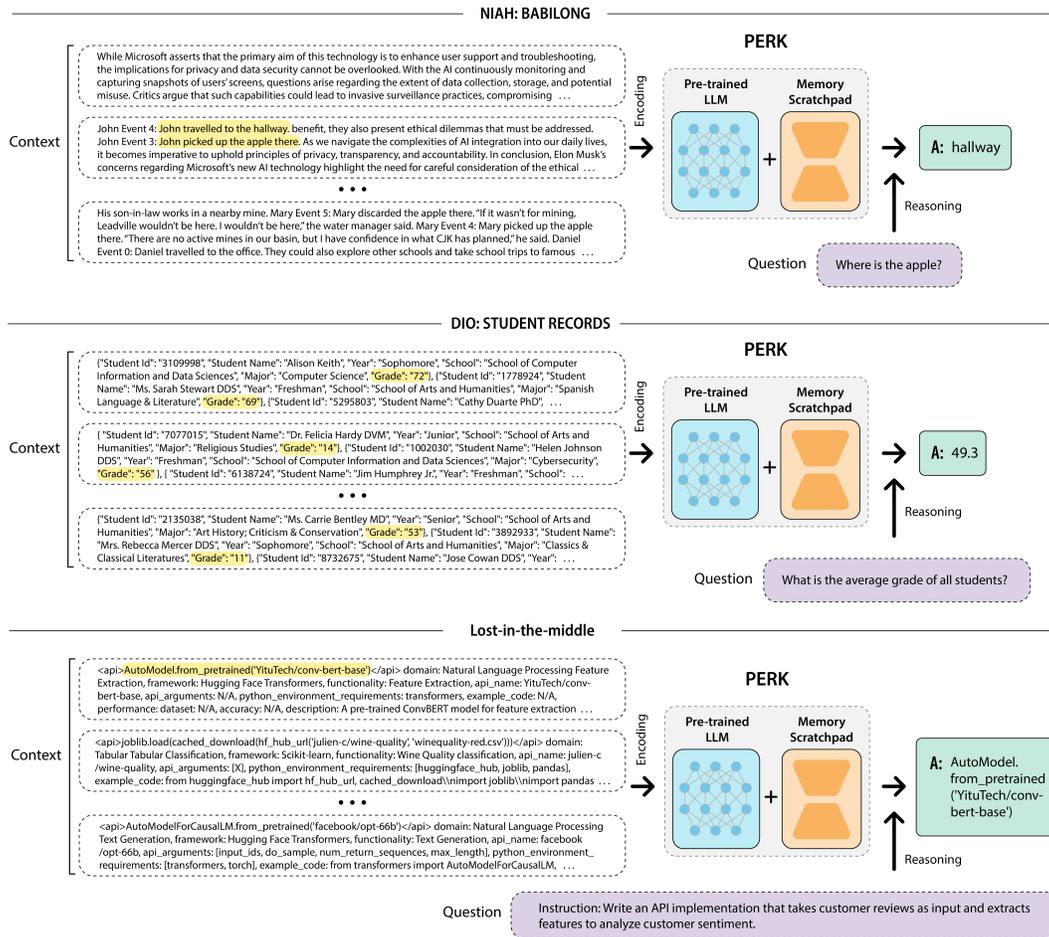


Figure 7: **Dataset Examples:** Here we show an example from the three datasets we used: (Top) BabiLong, (Middle) Student Records, and (Bottom) Lost-in-the-Middle API Retrieval. PERK first encode the contexts into the memory scratchpad through self-supervised test-time learning. Next, the user can ask any number of questions about that context, and the LLM can reason over the updated memory scratchpad to generate the answers.

Task We formulate the main task as gold API retrieval. In this task, the model must retrieve the correct API call based on user instructions and documents. Using fixed 4k and 8k token contexts, we train PERK and FT-ICR models on contexts with the target document at varying positions and test them on contexts distributed across all positions.

Build Long Context We follow Student Records and scale the context length with the number of API documents included. For a given context window, we continuously append API documents to the context until the number of tokens exceeds the context window. Similarly, since each API document is independent, we treat a document as a subsequence in the batch for PERK. Since the documents are processed in parallel as a batch, they do not preserve any ordering information. Thus, by design, PERK is robust to the position of the relevant information in a context.

B TRAINING AND INFERENCE DETAILS

Hardware Information All training runs are conducted on a single Nvidia H100 GPU with 93GB of VRAM. The training jobs run on an internal research cluster managed by Slurm. Each job is provided with 64 CPUs and 64GB of RAM. The GPU is completely utilized without sharing with other concurrent jobs.

Standard Hyperparameters Each training (for PERK and baselines) sets 2 epochs as the maximum runtime with early stopping based on validation loss. In practice, most training runs typically stop early, within one epoch of runtime. We set the baseline training and PERK’s outer loop learning rate to $1e-5$, and the weight decay to 0.01. We use a cosine learning rate scheduler with 0.03 of the total optimization steps as the warm-up phase. We use PyTorch Lightning¹ as the trainer for PERK. The Qwen and GPT-2 FT-ICR baselines are trained using the Open-Instruct (Lambert et al., 2025) trainer from AI2, and the Mamba FT-ICR baselines are trained using the Huggingface Trainer². For the stress test of length generalization beyond 64K tokens, we change the default maximum position of 32768 to 131072.

LoRA Hyperparameters We set LoRA rank to 256 for both PERK GPT-2 and PERK Qwen. We use the default values for the other hyperparameters, including setting the alpha to 16 and the dropout to 0.1. We apply LoRA to all modules in the model architecture. For 1B scale and below, we apply LoRA to all layers in the model. However, for 7B and 8B models, we apply it only to the top 4 layers in order to reduce the memory overhead of meta-learning. We set the scaling factor α to 16 initially and enable rank-stabilized LoRA fine-tuning Kalajdziewski (2023), which results in a target α of 256, corresponding to a rank of 256.

Meta-Learning Hyperparameters For the inner loop adaptation, we use 4 inner loop gradient steps. For the truncated gradient unrolling, we truncate steps 1 and 2 for input context length $\leq 4K$, and we truncate steps 1, 2, and 3 for context length $> 8K$. We follow Ye & Chao (2022)’s practice on how to train a good MAML model and meta-learn a set of per-layer-per-step dynamic learning rates for the inner loop as part of the outer loop learning, using $5e-5$ as the initial value. This approach learns a separate learning rate for each network layer and each adaptation step in the inner loop. These learning rates are optimized during the meta-training process alongside the model parameters. Note that these per-layer-per-step learning rates are learned parameters, not hyperparameters that require manual tuning. During meta-training, our algorithm automatically learns to adjust these learning rates for each layer and inner loop step dynamically based on the meta-objective. This approach reduces the hyperparameter tuning burden compared to manually setting the adaptive learning rates. For the language modeling loss in the inner loop, we use a weighted version similar to CaMeLS (Hu et al., 2023). The token-level weights are predicted via a weighting model implemented as a two-layer MLP network. The parameters of the weighting model are co-meta-learned along with the LoRA parameters and the adaptive per-layer-per-step learning rates. We use the differentiable version of the AdamW (Loshchilov & Hutter, 2017) optimizer for the inner loop optimization, implemented by the Higher³ library for higher-order optimization. We compute higher-order derivatives using the automatic differentiation mechanism implemented by PyTorch⁴. The outer loop optimizer is the standard PyTorch implementation of AdamW.

Long-context Data Collation for PERK At training time, to collate a reasoning problem with a long context and a set of associated question-answer pairs, we first convert a long context into a batch of subsequences, with relatively even lengths. In practice, we set the effective context length (i.e., the number of tokens of a subsequence on average) as 256, and divide the long context into N chunks where each chunk is approximately 256 tokens long. For a long context with 8192 tokens, the batch size is thus $N = 32$. PERK’s inner loop adaptation encodes the batch of subsequences in parallel through gradient updates. At inference time, if we want to test PERK on a long context with 32K

¹<https://lightning.ai/docs/pytorch/stable/>

²https://huggingface.co/docs/transformers/en/main_classes/trainer

³<https://github.com/facebookresearch/higher>

⁴https://docs.pytorch.org/tutorials/beginner/blitz/autograd_tutorial.html

tokens, the effective context length stays the same as approximately 256 tokens, while the batch size increases to $N = 128$.

Note that splitting a long context into a batch of shorter sequences can either be permutation-invariant or maintain the original ordering. PERK induces ordering through explicit indexing in the prompt formatting, which allows the model to learn implicit ordering during the encoding phase. Specifically, when chunking a long context, we can preserve ordering information by indexing each chunk. For example, if we split a long document into segments A, B, C, D, E , we format them as: ["Document 1: A", "Document 2: B", ..., "Document 5: E"]. During the inner loop adaptation with the NLL objective, the LoRA parameters learn to associate these indices with their content, effectively encoding the ordering information.

Evaluation Hyperparameters At inference time, we generate the answer for a reasoning problem via greedy decoding (temperature is set to 0). The maximum number of tokens that can be generated is 512. The EOS token is set to the default one defined in the tokenizer. The generation is handled by the text-generation pipeline implemented by Huggingface⁵. For the length generalization analysis, we apply recent training-free methods Yarn (Peng et al., 2023b) and Dual Chunk Attention (DCA, An et al. (2024a)) to the finetuned Qwen models. We use vLLM⁶ (0.10.2) as the inference engine to apply Yarn + DCA. For Yarn, we apply a factor of 4.0 to the original max position embeddings, 32768, and scale the context window to 131072. For DCA, we use vLLM’s *DUAL_CHUNK_FLASH_ATTN* backend for generation.

C ADDITIONAL ANALYSIS

C.1 IMPACT OF LORA RANK IN PERK

In this section, we analyze the impact of LoRA’s rank on PERK’s long-context reasoning performance. We focus on the BabiLong tasks with a 1024 context window for an efficient ablation study. We apply PERK to Qwen-2.5-0.5B with LoRA rank of 16 and 256 and evaluate on the BabiLong tasks.

We show the results in Table 2. We also include the performance from the FT-ICR Qwen baseline as a reference. The results show that PERK’s performance only decreases slightly when the rank is reduced to 16. However, PERK still maintains a much stronger performance than the FT-ICR Qwen baseline. The analysis shows that PERK can maintain its strong performance with a more parameter-efficient setting, highlighting its potential for more efficient long-context reasoning.

| Method | QA1 | QA2 | QA3 |
|---------------|------------|-------------|-------------|
| FT-ICR | 91.1 | 70.5 | 48.6 |
| PERK LoRA-16 | 99.0 | 96.9 | 84.7 |
| PERK LoRA-256 | 100 | 98.1 | 89.1 |

Table 2: Analyzing LoRA rank’s impact on PERK’s performance.

C.2 INNER LOOP ADAPTATION STEPS

In this section, we analyze the impact of inner loop adaptation steps on PERK’s long-context reasoning performance. Similarly, we focus on the BabiLong tasks with a 1024 context window. We again apply PERK to Qwen-2.5-0.5B for 2 and 4 inner loop steps. With 4 steps, we truncate the first 2 steps in backpropagation. With 2 steps, we truncate the first step in backpropagation. We show the results in Table 3. The results show that reducing the number of inner loop adaptation steps decreases the performance more noticeably than reducing LoRA rank. However, PERK still outperforms FT-ICR Qwen on all three tasks. The results here suggest that we should select a higher number of inner loop adaptation steps to maximize the performance. If in settings where GPU memory might be limited, reducing the number of inner loop steps can reduce the performance, but it

| Method | QA1 | QA2 | QA3 |
|--------------|------------|-------------|-------------|
| FT-ICR | 91.1 | 70.5 | 48.6 |
| PERK 2 Steps | 99.2 | 94.5 | 82.4 |
| PERK 4 Steps | 100 | 98.1 | 89.1 |

Table 3: Analyzing the impact of the number of inner loop adaptation steps on PERK’s performance.

⁵https://huggingface.co/docs/transformers/v4.43.4/en/main_classes/pipelines

⁶<https://docs.vllm.ai/en/latest>

is still able to outperform the baseline. However, the analysis is done with limited hardware resources, which limits our analysis to 4 steps. To further analyze the impact of the inner loop step counts, more systematic experiments in a scalable hardware setting are needed for future work.

C.3 GRADIENT UNROLLING TRUNCATION STEPS

In this section, we analyze the impact of the truncation steps we use when performing gradient unrolling (GU) for computing the meta-gradients in the outer loop of the meta-learning process. Recall that to reduce the memory cost of full gradient unrolling, which requires saving the computation graph for each inner loop adaptation step, we truncate the unrolling by detaching the early steps’ graph to save memory. Bounded by memory constraints, we can only train on 1K-token contexts with full GU without any truncation. Thus, we conduct our analysis under the 2K context window. Given 4 inner loop adaptation steps, we truncate 0-4 steps with Qwen2.5-0.5B on the BabiLong-QA1 task. The analysis shows that we can truncate up to 3 steps and only unroll the last inner loop adaptation while maintaining a high accuracy without significant degradation. However, truncating all steps, which results in a first-order approximation of the meta-gradients, shows a large performance drop, even on this very simple setting with QA1-1K.

| Truncation Length | Accuracy |
|-------------------|----------|
| TGU (0) | 100 |
| TGU (1) | 100 |
| TGU (2) | 100 |
| TGU (3) | 98.8 |
| TGU (4) | 92.4 |

Table 4: Truncation length ablations on the BabiLong QA1 task with 1K tokens. We conduct 4 steps of inner loop adaptation and truncate 0–4 steps. PERK maintains the accuracy up to 3 truncation steps, but drops when all steps are truncated.

D TRAINING MEMORY COST ANALYSIS

We show PERK’s training scalability by measuring its peak memory use across different context lengths (compared to a prior test-time learning method, RECKONING; Chen et al., 2023b).

Setup We complete all runs on a single Nvidia H100 GPU with 93 GB of VRAM using Huggingface’s Transformers library version 4.51.3. We perform four gradient update steps for the training-time inner loop adaptations for PERK. We test four truncation levels (§3.2): truncate 0 (full gradient unrolling without truncation), 1, 2, and 3 steps. Both methods undergo 10 hardware warm-up iterations.

PERK is more memory efficient at training.

Figure 8a shows that, as the training context length grows from 1K to 8K tokens, RECKONING quickly runs into Out-Of-Memory (OOM) errors at context length 2K. At the same time, PERK, with increasing truncation steps, continues to fit the maximum GPU memory, successfully validating PERK’s training scalability on long contexts.

E DETAILED RESULTS

We show detailed quantitative performance of PERK, FT-ICR baselines, and frontier long-context LLMs on the long-context reasoning benchmarks we evaluate on in Table 5.

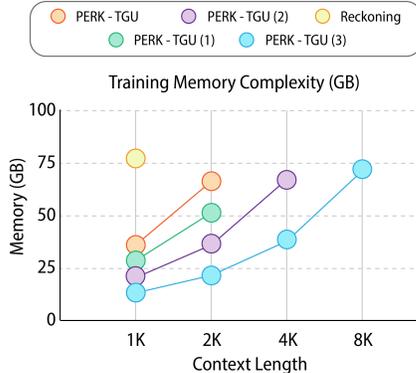


Figure 8: **Peak GPU memory usage during training** with context lengths ranging from 1K to 8K tokens for RECKONING and PERK. While PERK successfully scales to 8K tokens, RECKONING encounters out-of-memory (OOM) errors at shorter context lengths.

| Model | BBL-QA1 | | BBL-QA2 | | BBL-QA3 | | HotpotQA | | TriviaQA | | SR-Recall | | SR-Relation | | SR-Aggregate | | AVG | |
|---|------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|
| | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K | 8K | 32K |
| <i>Commercial Frontier Models</i> | | | | | | | | | | | | | | | | | | |
| Gemini-1.5-pro | 95.6 | 87.7 | 64.0 | 56.4 | 49.4 | 36.7 | 75.3 | 73.4 | 76.4 | 75.1 | 100 | 99.4 | 63.3 | 58.8 | 50.9 | 48.7 | 71.7 | 67.0 |
| GPT-4.1 | 91.4 | 87.8 | 84.5 | 80.6 | 54.9 | 33.1 | 72.4 | 65.8 | 78.3 | 74.3 | 100 | 98.8 | 66.5 | 62.5 | 52.4 | 49.4 | 75.1 | 69.0 |
| <i>Trained on contexts with >256K tokens</i> | | | | | | | | | | | | | | | | | | |
| Qwen2.5-7B-1M | 77.4 | 65.7 | 50.2 | 35.3 | 30.5 | 27.2 | 60.3 | 55.8 | 72.8 | 70.9 | 92.7 | 86.7 | 58.2 | 52.1 | 39.4 | 15.8 | 60.2 | 51.2 |
| ProLong-8B-512K | 65.3 | 41.0 | 26.2 | 29.0 | 27.8 | 24.4 | 62.4 | 57.3 | 70.4 | 68.2 | 94.5 | 89.2 | 60.1 | 54.4 | 24.3 | 12.1 | 53.9 | 47.0 |
| <i>Trained on contexts with 8K tokens</i> | | | | | | | | | | | | | | | | | | |
| FT-ICR (Mamba-1.4B) | 24.2 | 12.4 | 22.2 | 8.5 | 18.5 | 6.8 | 17.2 | 9.8 | 32.1 | 20.5 | 12.1 | 5.4 | 73.0 | 52.3 | 30.5 | 10.9 | 28.7 | 15.8 |
| FT-ICR (Qwen2.5-0.5B) | 76.1 | 43.8 | 64.3 | 39.4 | 37.5 | 12.1 | 24.4 | 11.9 | 47.4 | 41.7 | 93.6 | 84.5 | 65.3 | 57.8 | 51.8 | 42.4 | 57.6 | 41.7 |
| FT-ICR (Qwen2.5-7B) | 97.5 | 80.9 | 69.4 | 42.7 | 39.4 | 20.8 | 42.4 | 39.8 | 55.4 | 51.7 | 98.8 | 91.5 | 69.7 | 62.4 | 57.6 | 45.8 | 66.3 | 54.5 |
| <i>PERK (Ours), trained to encode contexts with 8K tokens</i> | | | | | | | | | | | | | | | | | | |
| PERK (Qwen2.5-0.5B) | 99.0 | 95.3 | 92.4 | 80.9 | 60.4 | 38.7 | 52.2 | 49.4 | 65.5 | 62.8 | 95.8 | 90.4 | 67.9 | 60.9 | 60.1 | 49.5 | 74.2 | 66.0 |
| PERK (Qwen2.5-7B) | 100 | 94.2 | 96.1 | 84.4 | 65.2 | 44.5 | 55.6 | 52.2 | 70.8 | 69.4 | 100 | 98.5 | 75.5 | 70.3 | 62.4 | 53.3 | 78.2 | 70.9 |

Table 5: **Model Performance Across Long-context Reasoning Benchmarks and Context Lengths (8K, 32K)**: We list detailed performance of our method PERK, FT-ICR baselines, and the frontier specialized long-context LLMs: Qwen2.5-7B-Instruct-1M (Yang et al., 2025), ProLong-LLaMA3-8B-Instruct-512K (Gao et al., 2025a), GPT-4.1 (OpenAI, 2025), Gemini-1.5-pro (Team et al., 2024). **BBL** here stands for BabiLong (Kuratov et al., 2024). SR stands for Student Records, our proposed dataset for Drops-in-the-Ocean reasoning. PERK and FT-ICR models have only seen training data contexts with 8K tokens, requiring them to generalize to the unseen 32K context window. PERK achieves the highest performance on average for both 8K-token and 32K-token contexts.