

# Efficient Many-Shot In-Context Learning with Dynamic Block-Sparse Attention

Emily Xiao, Chin-Jou Li, Yilin Zhang, Graham Neubig, Amanda Bertsch

Language Technologies Institute, Carnegie Mellon University

{emilyx, chinjou1, jasonzh3, gneubig, abertsch}@andrew.cmu.edu

## Abstract

Many-shot in-context learning has recently shown promise as an alternative to finetuning, with the major advantage that the same model can be served for multiple tasks. However, this shifts the computational burden from training-time to inference-time, making deployment of many-shot ICL challenging to justify in-practice. This cost is further increased if a custom demonstration set is retrieved for each inference example. We present Dynamic Block-Sparse Attention, a training-free framework for retrieval-based many-shot in-context learning. By combining carefully designed block-sparse attention and retrieval of cached *groups* of demonstrations, we achieve comparable per-example latency to finetuning while maintaining on average  $>95\%$  of the best method’s accuracy across strong ICL and finetuning baselines. We hope that this will further enable the deployment of many-shot ICL at scale.<sup>1</sup>

## 1 Introduction

When adapting large language models (LLMs) to a specific task, practitioners typically use finetuning or in-context learning (ICL). ICL has the advantage of convenience—it requires no parameter updates, adapts easily to new tasks and datasets, and does not require serving task-specific models in a production setting. In the standard few-shot setting, ICL typically underperforms finetuning (Mosbach et al., 2023) and requires a careful demonstration curation process (Dong et al., 2024b).

Recently, many-shot ICL—which uses thousands of demonstrations—shows performance comparable to or even surpassing finetuning for various tasks (Bertsch et al., 2024; Agarwal et al., 2024). However, scaling ICL to long contexts introduces a new computational tradeoff: many-shot ICL shifts the computational burden from fine-tuning to inference, creating significant efficiency challenges. If

<sup>1</sup>Data and code are available at <https://github.com/millix19/dbsa>

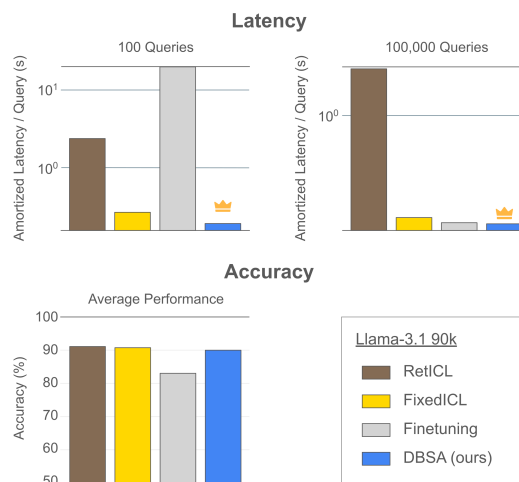


Figure 1: DBSA maintains high accuracy while achieving the best overall efficiency when compared to many-shot ICL baselines and finetuning, even under high request volumes.

the same fixed set of demonstrations is used for all requests, some computation can be cached; however, performance is higher when a custom set of demonstrations is retrieved. Processing thousands of demonstrations per request is orders of magnitude more expensive than zero-shot inference with a fine-tuned model, making many-shot ICL impractical for high-throughput applications.

In this paper, we propose Dynamic Block-Sparse Attention (DBSA), a training-free inference framework that minimizes many-shot ICL latency while maintaining  $>95\%$  of the best accuracy on average across all baselines, including many-shot ICL, retrieval ICL, and finetuning. We introduce key optimizations to both demonstration pre-encoding and inference.

During pre-encoding, we apply a structured block-sparse streaming attention pattern (Xiao et al., 2024b), where each demonstration attends only to a fixed number of others and a global attention sink. Then, during inference, we integrate retrieval ICL with KV cache reuse, dynamically

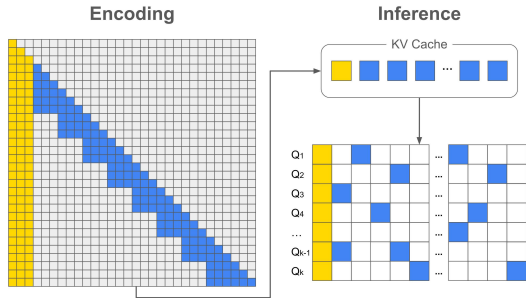


Figure 2: DBSA pre-encodes the many-shot demonstration pool with block-sparse attention, and dynamically selects relevant KV chunks during inference for each test query.

selecting *groups* of relevant demonstrations that were pre-encoded together for each test query. The test query attends only to the KV cache of selected groups to save computation.

Unlike standard key-value (KV) cache reuse, with DBSA the encoding time for new demonstrations remains constant as the demonstration pool grows, a benefit for practical scenarios where the demonstration pool may change over time. Importantly, this mechanism works without any additional training or optimization. And unlike standard retrieval ICL, DBSA reuses a pre-encoded cache, combining the efficiency benefits of pre-encoding a fixed demonstration set with the performance benefits of retrieving custom demonstrations.

We evaluate DBSA on five datasets using Llama-2 and Llama-3.1 models with up to 90k tokens of in-context demonstrations. Our second-stage inference speed up is proportional to the retrieval ratio – at a 30% retrieval ratio, our method achieves more than 2× speedup over cached many-shot ICL; when the selected KV cache is within 30k tokens, DBSA reaches inference latencies comparable to a fine-tuned Llama-3.1, with >1000x less setup time and better accuracy. Our results demonstrate many-shot ICL can be the most competitive solution in practice, even in scenarios with large request volumes.

## 2 Methodology

Dynamic Block-Sparse Attention treats many-shot ICL as a two-stage process. First, there is a one-time setup cost to encode the many-shot demonstrations; and second, a dynamic KV-cache is selected for each answer generation.

### 2.1 Stage 1: Many-Shot Encoding

In the first stage, we aim to encode the full demonstration set with minimal cost. Note that the full

demonstration set may exceed the maximum context size of our model, even when using a long-context model (e.g. if there are tens of thousands of training examples).

Prior research (Bertsch et al., 2024; Acharya et al., 2024; Lu et al., 2024) suggests that letting each example attend to some number of other examples is beneficial, but attending to the entire demonstration pool is often unnecessary to achieve high performance. We use a structured sparse attention pattern for efficiency.

Given a set of demonstrations  $D = \{d_1, d_2, \dots, d_n\}$ , where each demonstration  $d_i$  consists of a query-answer pair  $(q_i, a_i)$ , we first partition  $D$  into  $\frac{n}{k}$  blocks:  $D = [b_1, b_2, \dots, b_{n/k}]$ , where each block  $b_i$  contains  $k$  demonstrations. We encode  $D$  using a block-wise streaming attention mechanism, where each block  $b_i$  attends to three key components: (1) an anchor block  $b_1$ , (2) a fixed number ( $j$ ) of previous blocks  $\{b_{i-j}, \dots, b_{i-1}\}$ , and (3) itself using standard causal attention. Because not all examples attend to each other, the demonstrations do not all need to be encoded in the same forward pass or even at the same time; thus, new demonstrations can be added to the demonstration pool at any point by encoding an additional block with this attention pattern.

Following the same approach as StreamingLLM (Xiao et al., 2024b), we encode  $D$  using sequential position ids  $[0, \dots, n-1]$  and cache the KV states for each block before applying the rotary position transformation. However, unlike StreamingLLM, we retain all KV cache entries, as each test query benefits from a distinct set of demonstrations.

The many-shot encoding stage is implemented using Flex Attention (Dong et al., 2024a), which skips the computation for masked blocks and achieves a proportional performance speedup compared to FlashAttention-v2 (Dao, 2024).

### 2.2 Stage 2: Dynamic Demonstration Selection and Answer Generation

Prior research has demonstrated that retrieval ICL can achieve strong performance using only a subset of the available demonstration pool (Luo et al., 2024). However, achieving peak performance still requires hundreds of examples, which remains computationally expensive with a naive implementation that re-encodes the context for every test query. Our method reduces this inference overhead by reusing the precomputed KV caches from Stage 1.

Given a test query  $q^*$ , a retrieval method is ap-

plied to select a subset of demonstration blocks  $D' \subset D$ , where  $D' = \{b'_1, b'_2, \dots, b'_m\}$ ,  $m < n$ . The anchor block  $b_1$  is always included as  $b'_1$  to serve as the attention sink. We concatenate the KV caches of the selected demonstrations, and re-apply relative positional encoding based on new in-order position IDs ranging from  $[0, |D'| - 1]$ . Finally, we encode  $q^*$  with full attention to the selected KV cache, and generate the answer  $a^*$  autoregressively.

DBSA provides the flexibility to plugin any retrieval method of choice, including text-based similarity, cosine similarity, or diversity-focused retrievers (Dong et al., 2024b; Luo et al., 2024). The choice of retrieval method also presents an efficiency-performance tradeoff; for this work, we use the inexpensive and relatively strong BM25 retriever (Robertson and Zaragoza, 2009).

### 3 Experimental Setup

**Datasets.** We evaluate DBSA on 5 language classification datasets—TREC (Hovy et al., 2001), TREC-fine (Hovy et al., 2001), NLU (Xingkun Liu and Rieser, 2019), Banking-77 (Casanueva et al., 2020), and Clinc-150 (Larson et al., 2019). These datasets span a diverse range of domains and label spaces, are commonly used in prior work in ICL (Han et al., 2022; Ratner et al., 2023), and have been explored in the many-shot ICL setting (Bertsch et al., 2024; Yen et al., 2024). For further evaluation, we also include results on MMLU (a multitask QA benchmark) (Hendrycks et al., 2021) and SAMSum (a generative summarization task) (Gliwa et al., 2019) in Appendix B.

We consider two context lengths, 30k and 90k, and construct the demonstration pool  $D$  by randomly selecting the maximum number of demonstrations that fit in-context; for these classification datasets, which have short input/output pairs, thousands of demonstrations fit in 90k context. Table 1 provides summary statistics for each dataset.

Dataset	Domain	# Labels	# Demo (30k Context)	# Demo (90k Context)
TREC	Questions	6	1050	3150
TREC Fine	Questions	50	1000	3000
NLU	Conversational	68	1150	3450
Banking-77	Financial	77	800	2400
Clinc-150	Multiple	151	1000	3000

Table 1: Dataset statistics.

**Models.** We consider 2 models in the Llama model series for evaluation. Llama-2-7B (32k) (To-

getherAI, 2023) is the base Llama-2-7B model finetuned with an extended 32k context window. We also consider the more recent Llama-3.1-8B (Meta AI, 2024), which supports a 128k context window and incorporates Grouped-Query Attention (GQA) for efficient long-context inference.

**Baselines.** We compare our method against three baselines: Fixed ICL, Retrieval ICL, and finetuning. *Fixed ICL* uses the entire 30k or 90k demonstration set as context for every test query. We consider the efficient implementation where the context is encoded once and its KV cache reused across test queries. *Retrieval ICL* (RetICL) dynamically selects a subset of relevant demonstrations from the 30k or 90k demonstration pool for each test query. In this scenario, a distinct context has to be encoded for each test query at inference time. As retrieval strategy is a flexible choice, we use BM25 retriever (Bassani, 2023; Robertson and Zaragoza, 2009) for both this baseline and our method to ensure a fair comparison. *Finetuning* uses LoRA (Hu et al., 2022) with lora rank 8. We choose the best setup and hyperparameters based on ablation results in Bertsch et al. (2024). For additional details, see Appendix C. For all ICL methods, we use constrained decoding to output only valid labels. For finetuning, we finetune with a classification head.

**Evaluation.** Following prior work, we randomly sample 250 test examples from each dataset. We evaluate both accuracy and efficiency metrics, averaged over 10 runs. We compare our method to the baseline using FlashAttention-2 (Dao, 2024). All methods with 30k context are run on a single L40S GPU (48GB) and all methods with 90k context are run on a single A100 GPU (80GB).

**Configuration.** We use 50 examples per block, with demonstration blocks formed through random grouping. Across all context lengths, we designate the first block as the attention sink and allow each block to attend to the two preceding blocks as locally attended context. For both the Retrieval ICL baseline and our method, we retrieve 30% of the total demonstration pool for each test query. This retrieval ratio is used for all datasets.

## 4 Main Results

We compare DBSA to the baselines along two dimensions: efficiency and accuracy.

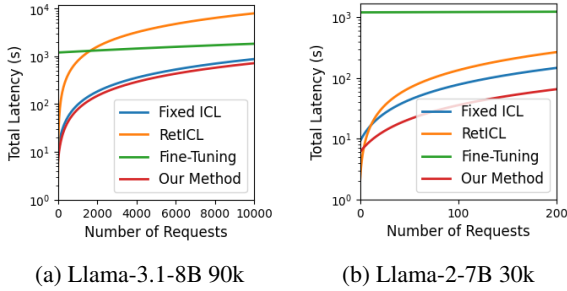


Figure 3: The total inference latency for each method, including setup and per-request inference latency. We use inference batch size 1.

#### 4.1 Efficiency Comparison

Efficiency is the primary motivation behind our method, which aims to close the inference latency gap between many-shot ICL and zero-shot inference with a finetuned model. To evaluate efficiency, we consider setup overhead, inference latency, and storage cost of each method.

Figure 3 and 1 shows that DBSA is the most efficient overall, even with more than 100,000 requests. The latency comparison in Table 2 breaks down the efficiency gains, showing that DBSA can bring inference costs close to those of zero-shot inference while maintaining a significantly lower setup time than fine-tuning.

**Setup overhead.** Retrieval ICL incurs the lowest setup time (in our case, only the time to construct the BM25 retriever index). Fixed ICL is more expensive at setup time, as setup requires encoding the full 30k or 90k demonstration set with full attention. Fine-tuning has the highest setup cost, as it requires training the model with LoRA for approximately 30 epochs to reach high performance. In practice, finetuning took about 30 minutes with a 30k demo pool, though here we exclude the time for evaluation steps in our reported time.

DBSA involves both building the retriever index and encoding the full demonstration pool. However, it still achieves a speedup over Fixed ICL thanks to the structured sparse attention. While setup cost is a one-time fixed cost with a static demonstration pool, real-world demonstrations are often non-static. In these cases, the speedup from sparse encoding becomes particularly valuable.

**Inference latency,** or the per-example cost, is key to the practicality of deployment. Compared to Fixed ICL, DBSA lowers latency by attending to a reduced KV cache context length. Theoretically, inference latency grows linearly with the size of the

Method	Setup Time (Relative)	Inference Latency (Relative)
<b>30k Context Length w/ Llama-2-7B w/ single L40S gpu</b>		
RetiCL (no cache)	<b>1x</b>	1x
Fixed ICL (cached)	4.5x	0.51x
Finetuning (LoRA)	> 600x	<b>0.12x</b>
DBSA (ours)	<u>3x</u>	<u>0.22x</u>
<b>30k Context Length w/ Llama-3.1-8B w/ single L40S gpu</b>		
RetiCL (no cache)	<b>1x</b>	1x
Fixed ICL (cached)	5x	0.11x
Finetuning (LoRA)	> 600x	<b>0.08x</b>
DBSA (ours)	<u>3x</u>	<u>0.10x</u>
<b>90k Context Length w/ Llama-3.1-8B w/ single A100 gpu</b>		
RetiCL (no cache)	<b>1x</b>	1x
Fixed ICL (cached)	6.5x	0.06x
Finetuning (LoRA)	> 1500x	<b>0.046x</b>
DBSA (ours)	<u>4x</u>	<u>0.053x</u>

Table 2: Relative latency of different methods compared to RetiCL baseline. Best method is bolded and second best is underlined. Since the absolute numbers depend on the computational environment, we present the relative speedups here; the raw numbers for our hardware setup are available in Appendix A

KV cache, a trend empirically shown in Agarwal et al. (2024). Therefore, smaller retrieval ratios can lead to further improvements.

While Retrieval ICL without cached encoding is always the most expensive in terms of latency per request, the latency gap between cached ICL and finetuning narrows significantly when using Llama-3.1 compared to Llama-2, making cached many-shot ICL increasingly competitive. One major factor contributing to this speedup is Grouped-Query Attention (GQA), which reduces the computational cost of handling long sequences. All ICL methods benefit from this speedup; however, as fixed ICL uses the most demonstrations in-context, it is disproportionately impacted. Despite this, DBSA is still slightly faster per-example.

In Figure 1 and 3, we consider an alternate metric: the total compute time necessary to accomplish the task, amortized over the total number of inference requests. This includes both setup overhead and the per-example latency. DBSA remains more efficient than finetuning even when more than 100,000 requests are made.

**Storage Cost Considerations.** Beyond computational efficiency, we also evaluate the storage cost.

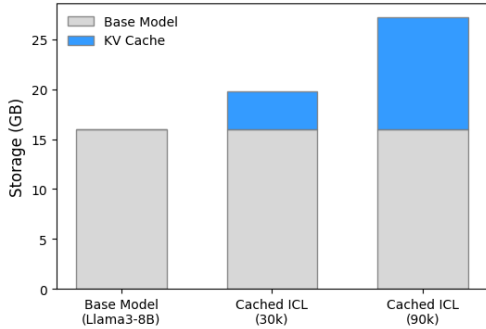


Figure 4: Storage cost for Cached many-shot ICL, which encompass Fixed ICL and DBSA

For Llama-3.1-8B, the KV cache requires 0.125 MiB per token, which is 3.7GB for a 30k demonstration pool and 11.1GB for 90k. While manageable for moderate demonstration sizes, this storage demand may require offloading or re-encoding for larger demonstration pools or multi-task settings. Note that this is a shared cost across all many-shot ICL methods with long context.

For a single task, the storage cost for LoRA fine-tuning is significantly smaller. With Llama-3.1-8B model and LoRA rank 8, the storage is 0.01GB per task. However, in settings where hundreds or thousands of tasks must be supported—such as user-specific tasks—the cumulative storage cost of fine-tuning also becomes substantial. Furthermore, task-switching and task addition introduce additional overhead (Xia et al., 2024; Sheng et al., 2023; Gabrielsson et al., 2024). In these cases, the fast setup and flexibility of ICL provide a strong advantage over fine-tuning— if storage cost is a limited factor, demonstration caches can be recomputed when needed instead of stored indefinitely.

## 4.2 Accuracy Comparison

Table 3 show that DBSA consistently achieves high accuracy across datasets, closely matching the best baseline method. This is particularly evident at 90k demonstration context, which is also where accuracy of these datasets saturate, making it the most relevant comparison point.

The success of our approach is largely due to the effectiveness of Retrieval ICL, which is often the best performing baseline. We set a fixed retrieval ratio of 30% across all datasets and context lengths. Even without optimizing the retrieval ratio, Retrieval ICL achieves accuracy comparable to Fixed ICL using the entire demonstration pool. Prior work suggests that Retrieval ICL can sur-

Dataset	Fixed ICL	Ret ICL	Fine-Tuning	DBSA (ours)
<b>30k Context Length w/ Llama-2-7B</b>				
TREC	0.93	0.92	<b>0.95</b>	0.91 (96%)
TREC Fine	0.76	<b>0.79</b>	0.77	0.77 (97%)
Banking77	0.81	<b>0.84</b>	0.80	0.80 (95%)
Clinic	<b>0.86</b>	0.83	0.60	0.80 (93%)
NLU	0.85	<b>0.86</b>	0.74	0.84 (98%)
Average	0.84	0.85	0.77	0.84 (97%)
<b>90k Context Length w/ Llama-3.1-8B</b>				
TREC	0.96	0.95	<b>0.96</b>	0.95 (99%)
TREC Fine	0.88	<b>0.89</b>	0.83	0.88 (99%)
Banking77	<b>0.91</b>	0.90	0.81	0.89 (98%)
Clinic	<b>0.92</b>	0.91	0.74	0.90 (98%)
NLU	0.89	<b>0.90</b>	0.82	0.88 (98%)
Average	0.91	0.91	0.83	0.90 (99%)

Table 3: Accuracy comparison. Bolded is the max accuracy for each dataset. Blue percentage is how DBSA compares to the best accuracy.

pass Fixed ICL when the retrieval ratio is tuned optimally (Luo et al., 2024; Bertsch et al., 2024). However, determining the ideal retrieval ratio is non-trivial— performance initially improves with more demonstrations, but eventually degrades due to noise from less relevant demonstrations. While we use a fixed retrieval ratio, our framework supports adaptive selection strategies to optimize retrieval for each query. Given that our method’s accuracy ceiling is constrained by the upper bound of Retrieval ICL, further improvements in retrieval quality would benefit our method as well.

At 30k context length, we observe a more noticeable accuracy gap between DBSA and Retrieval ICL. This comes from two key design choices. First, we always include the first 50-example block as the attention sink, which may not always be relevant to the test query. Second, we group blocks randomly and treat each block as a unified text chunk when computing relevant scores. This raises the risk of missing highly relevant examples. In Section 5.3 and Section 5.2, we analyze different block grouping strategies and discuss potential refinements that can balance retrieval relevance and diversity more effectively, and why block-level selection is still necessary.

We observe that finetuning underperforms significantly compared to many-shot ICL, even with 90k demonstration pool (3k demonstrations). This aligns with previous findings (Bertsch et al., 2024), which show that finetuning requires a larger number of demonstrations to achieve competitive or

Dataset	Fine-Tuning (90k)	Fine-Tuning (All)	DBSA (90k)
TREC	0.96	0.96	0.95
TREC Fine	0.83	0.89	0.88
Banking77	0.81	0.91	0.89
Clinc	0.74	0.89	0.90
NLU	0.82	0.87	0.88
<b>Avg</b>	<b>0.832</b>	<b>0.904</b>	<b>0.900</b>

Table 4: Accuracy comparison of finetuning with 90k demonstration pool, or using the full training dataset.

higher performance. However, even when trained on the maximum available data (5k to 20k demonstrations for our datasets), finetuning fails to substantially surpass the best accuracy achieved by many-shot ICL, as shown in Table 4.

## 5 Key Insights and Ablations

We conduct ablations to analyze the impact of various design decisions in our framework.

### 5.1 Sparse Attention for Many-shot ICL

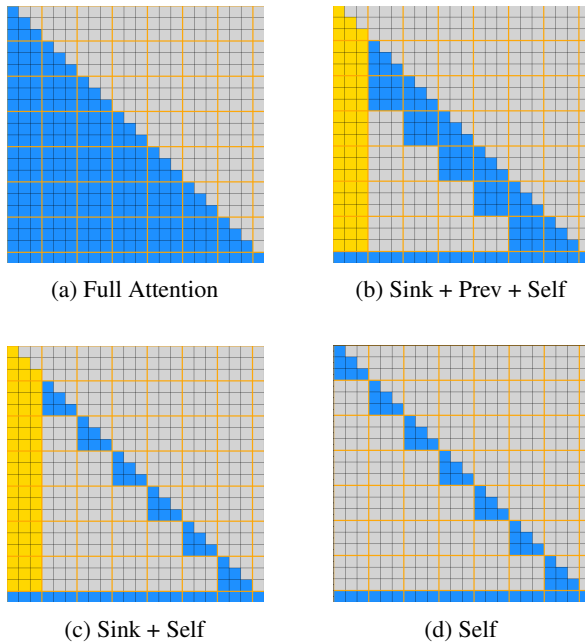


Figure 5: Visualization of block-sparse attention mechanisms for many-shot ICL in section 5.1.

The success of our sparse encoding approach shows that not all interactions within the demonstration set are necessary by default. However, sparse attention paradigms have been extensively explored since the introduction of self-attention, with varying applicability—some operate at inference time, while others require retraining for adaptation. Here,

we ablate over design decisions in the sparse attention to determine what makes this pattern effective without additional training.

While sparsity alone is not always optimal, block-wise sparsity considers I/O and memory efficiency, making it practical for real-world usage (Dong et al., 2024a; Dao, 2024; Guo et al., 2024). In this section, we compare variations of structured blockwise sparse attentions for many-shot encoding, illustrated in Figure 5:

- **Full Attention:** Each token attends to all other tokens.
- **Sink + Prev + Self:** Each block attends to an *anchor block*, a few *preceding blocks*, and *itself*. This is the pattern applied in DBSA.
- **Sink + Self:** Each block attends to a *sink block* and *itself* only.
- **Self-Only:** Each block attends *only to itself*.

These mechanisms have been applied to some extent in prior work, and are largely inspired by StreamingLLM (Xiao et al., 2024b). To evaluate these patterns, we allow test queries to attend to the full encoded context and compare the accuracies.

Dataset	Full	sink+prev+self	sink+self	self
Banking77	0.81	0.80	0.34	0.02
Clinc150	0.85	0.80	0.33	0.02
NLU	0.86	0.82	0.27	0.06
TREC	0.93	0.91	0.24	0.21
TREC Fine	0.76	0.76	0.17	0.151
<b>Avg</b>	<b>0.84</b>	<b>0.82</b>	<b>0.27</b>	<b>0.09</b>

Table 5: Comparison of accuracies using different sparse attention patterns, using Llama-2, 30k demo size, and block size of 50 demos.

Table 5 shows that both an attention sink and local context connections are necessary for sparse attention to achieve performance comparable to full attention without additional training.

These findings align with previous research that applied similar patterns. The "Sink + Self" pattern has been used in StarAttention (Acharya et al., 2024) for training-free inference speedup, but the results indicate that it performs well only when the number of blocks remains small (fewer than four). Other works also adopted this pattern (Lu et al., 2024) (Sun et al., 2024), but they required finetuning for adaptation. Meanwhile, the "Sink + Prev + Self" streaming attention (Xiao et al., 2024b), commonly used in KV eviction strategies (Xiao et al.,

2024a; Li et al., 2024), performs surprisingly well in our many-shot ICL setting. It remains scalable as demo size increase, as the same block size cause minimum degraation for both 30k and 90k setting. At 90k context, the attention mask is over 90% sparse while maintaining strong performance.

The scalability means that new demonstrations simply needs to attend to a fixed length of context as demo pool grows. Additionally, when KV cache storage is limited, applying this sparsity improves inference efficiency even in no-cache settings.

## 5.2 Block vs. Individual Demonstration Selection

In our framework, we group demonstrations into blocks for encoding and KV selection, rather than selecting individual examples like the Retrieval ICL baseline. In this section, we compare the two and separate the effects of using different selected examples and re-using KV cache segments.

Consider a set of demonstrations [1, 2, 3, 4, 5, 6, 7, 8] with sparse encoding, where [1] serves as the attention sink and each demonstration attends to itself and one preceding example (e.g., 5 attends to [1, 4, 5]). If we use example-level retrieval and retrieve [3, 6], 6 would not have attended to 3 during encoding. In contrast, if we select [3, 4] as a block, then it would be a contiguous KV cache segment. We hypothesize that block-level selection could better preserve intra-context relationships during encoding, and thus have better performance.

To test this hypothesis, we run block-level and example-level selection with DBSA, and see how each compare with the standard inference without sparse encoding and KV cache reuse. We make sure that the DBSA and standard inference use the same examples in context to isolate the effect of segmented KV cache reuse. Results in Table 6 show that example-level selection with DBSA is also effective, but block-level selection has slightly higher accuracy; it also has faster retrieval time and memory efficiency when handling KV caches.

Setting	example	example (DBSA)	diff	block	block (DBSA)	diff
30k - Llama-2	0.82	0.79	<b>0.03</b>	0.79	0.79	<b>0.00</b>
30k - Llama-3.1	0.86	0.82	<b>0.04</b>	0.84	0.82	<b>0.02</b>
90k - Llama-3.1	0.90	0.86	<b>0.04</b>	0.89	0.88	<b>0.01</b>

Table 6: Accuracy comparison of block-level and example-level retrieval, using standard inference vs. DBSA with sparse encoding and KV reuse. All use a 10% retrieval ratio.

## 5.3 Block Grouping

In our main experiments, we randomly group the demonstration pool into blocks before encoding and retrieval. However, the way demonstrations are grouped may impact accuracy. In this section, we explore whether alternative grouping strategies lead to improved performance.

We evaluate three grouping strategies: (1) random grouping, (2) k-means clustering based on lexical similarity (see Appendix F), and (3) starting with clusters, we swap 10% examples between clusters to introduce diversity within each block.

Table 7 shows that introducing some diversity within blocks improves performance compared to strictly clustered examples. Interestingly, random grouping remains competitive, sometimes outperforming clustering-based methods. A possibility is that random grouping naturally leads to diversity within a block, thus achieving the same goals.

Dataset	Random	Clustered	Clustered (w/ diversity)
Banking77	0.820	0.696*	0.740*
Clinic	0.796	0.792	0.796
NLU	0.840	0.792*	0.832
TREC	0.912	0.910	0.916
TREC Fine	0.766	0.692*	0.764
Avg	0.827	0.764	0.810

Table 7: Comparison of block grouping strategies with DBSA. \* denotes results significantly different than random grouping (using a t-test,  $p < 0.05$ ).

## 5.4 Dynamic Block Ordering

The "Lost in the Middle" phenomenon suggests that long-context models are sensitive to the position of relevant information, with content placed at the beginning or end of the context having a stronger influence (Liu et al., 2024b; Jin et al., 2024).

We evaluate whether, within our framework, the KV cache of demonstrations is robust to reordering and whether reordering strategies can be leveraged to enhance performance. Using BM25 relevance scores, we rank blocks for each test query and compare three reordering approaches: (1) maintaining the original encoded order, (2) re-arranging blocks from low to high relevance, and (3) reversing the encoded order. We keep the position of the sink block consistent in all experiments.

As shown in Table 8, reversing the order of cached embeddings significantly decreases perfor-

Dataset	DBSA			Non-cached		
	In.	L-to-H	Rev.	In.	L-to-H	Rev.
Banking	0.807	0.798	0.737*	0.821	0.828	0.811
Clinic	0.846	0.832	0.772*	0.868	0.863	0.854
NLU	0.832	0.838	0.804*	0.848	0.855	0.850
TREC	0.923	0.927	0.914	0.929	0.931	0.928
TREC Fine	0.759	0.782	0.738	0.845	0.850	0.840
Avg	0.834	0.835	0.794*	0.862	0.865	0.857

Table 8: Comparison of block ordering strategies with Llama-2 30k context. \* indicates statistical significance ( $p < 0.05$ ) in a T-test against the original encoded order (*In.*). *L-to-H* represents ordering from low to high relevance, and *Rev.* refers to reverse ordering.

mance. We believe this ordering dependence is because each block is permitted a small amount of local attention (i.e. to the previous two blocks) when encoding, imposing a loose ordering of blocks.

## 6 Related Work

**Many-shot ICL.** While our experiments focus on classification datasets, many-shot ICL has demonstrated potential across a wide range of tasks, including sequential parity, sentiment analysis, summarization, and translation (Agarwal et al., 2024; Li et al., 2023; Bertsch et al., 2024). Agarwal et al. (2024) showed it to be comparable to finetuning in translation, while Yin et al. (2024) found it more effective than finetuning in datasets with implicit structures. These work motivate many-shot ICL beyond our considered datasets and models.

**ICL demonstration selection.** ICL demonstration selection is an active area of research. Key approaches include demonstration augmentation, demonstration pool curation, retrieval ICL (Dong et al., 2024b). The effectiveness of retrieval ICL is influenced by multiple factors. However, Luo et al. (2024) shows that a simple and efficient BM25 retriever is effective across datasets, performing within 0.5% of a dual encoder-based retriever (GTR) across five tasks. To improve retrieval efficiency, CPU-based retrieval-augmented generation (RAG) techniques and approximate nearest neighbor (ANN) indexing (Malkov and Yashunin, 2020) can also be incorporated into DBSA.

**Sparse Attention Mechanisms** have been explored extensively in Transformers and LLMs to increase efficiency, with various patterns such as block-wise sparse attention (Child et al., 2019; Zaher et al., 2020; Wang et al., 2024; Acharya et al., 2024), hierarchical attention (Yang et al., 2016),

and layer or head-dependent sparsity. These methods primarily aim to improve training and inference efficiency together. However, our approach focus on efficiency without training.

A related training-free approach is Parallel Context Windows (Ratner et al., 2023), which also uses block-sparse attention. However, its goal is different: to extend the context length of limited-context models by reusing the same positional embeddings across blocks. This results in a weaker baseline compared to our approach, which operates on models that already support long contexts.

Among existing approaches, Lu et al. (2024) and Sun et al. (2024) are most similar to our method, but in the RAG setting. Like our approach, they incorporate a sparse attention encoding phase and optimize key-value (KV) cache re-utilization. However, a key difference is that these methods require fine-tuning to adapt to the sparse attention pattern. As discussed in Section 5.1, we hypothesize that this arises due to the absence of a local context component in their design.

**Key-value (KV) cache compression.** A common approach is token-level KV cache eviction to reduce memory overhead and speedup inference (Xiao et al., 2024b,a; Zhang et al., 2023). StreamingLLM (Xiao et al., 2024b) discovered the "attention sink" phenomenon and popularized token eviction by retaining only the initial and recent KV pairs. More recent methods (Li et al., 2024; Zhang et al., 2023; Liu et al., 2025) refine this approach by selectively retaining some KV pairs. However, eviction-based strategies are not well suited for many-shot ICL, as different test queries benefit from attending to distinct sets of demonstrations.

Beyond eviction, other KV cache compression methods focus on quantization and low-rank approximation to reduce storage requirements while retaining all tokens (Liu et al., 2024a; Zhang et al., 2024). These strategies can potentially be integrated with DBSA to achieve further memory reductions while maintaining inference quality.

## 7 Conclusion

We expect that many-shot ICL will be most useful in settings where the available data may change over time, either because of liability and privacy concerns (Liu et al., 2024c) or because of the continual addition of data in an online learning setting (Hoi et al., 2018). Unlike traditional cached ICL



or finetuning, the cost of setup when adding additional demonstrations to the demonstration pool using Dynamic Block-Sparse Attention is linear, and the increased cost at inference time is marginal, depending on the retrieval method used.

By improving inference efficiency while preserving accuracy, we hope that DBSA opens the door for adoption of many-shot ICL in real-world applications.

## 8 Limitations

The exact tradeoff between finetuning, Fixed ICL, and Dynamic Block-Sparse Attention depends on the context length and details of the model chosen (e.g. the use of GQA (Ainslie et al., 2023)); it is possible that there are some combinations of model, context length, and dataset where the marginal performance and efficiency gains of DBSA may be small or even nonexistent. For instance, because our method uses retrieval-based ICL, tasks where retrieval-based ICL is ineffective also will not benefit from our framework. In particular, tasks that require synthesizing information from all demonstrations (e.g. estimating the percentage of Amazon reviews that are positive (Shaham et al., 2023)) will not benefit from these efficiency speedups.

## Acknowledgments

We would like to thank Xiang Yue, Emmy Liu, and Vashisth Tiwari for useful discussions about this work.

AB was supported by a grant from the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE2140739. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors.

## References

Shantanu Acharya, Fei Jia, and Boris Ginsburg. 2024. Star attention: Efficient llm inference over long sequences. *arXiv:2411.17116*.

Rishabh Agarwal, Avi Singh, Lei Zhang, Bernd Bohnet, Luis Rosias, Stephanie Chan, Biao Zhang, Ankesh Anand, Zaheer Abbas, Azade Nova, John D. Co-Reyes, Eric Chu, Feryal Behbahani, Aleksandra Faust, and Hugo Larochelle. 2024. **Many-shot in-context learning**. In *Advances in Neural Information Processing Systems*, volume 37, pages 76930–76966. Curran Associates, Inc.

Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrón, and Sumit Sanghai. 2023. **Gqa: Training generalized multi-query transformer models from multi-head checkpoints**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901. Association for Computational Linguistics.

Elias Bassani. 2023. **retriv: A python search engine for the common man**. Released on 2023-05-28.

Amanda Bertsch, Maor Ivgi, Uri Alon, Jonathan Berant, Matthew R Gormley, and Graham Neubig. 2024. In-context learning with long-context models: An in-depth exploration. *arXiv preprint arXiv:2405.00200*.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. **Efficient intent detection with dual sentence encoders**. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, Online. Association for Computational Linguistics.

Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. **Generating long sequences with sparse transformers**. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11862–11872.

Tri Dao. 2024. **Flashattention-2: Faster attention with better parallelism and work partitioning**. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Juechu Dong, Boyuan Feng, Driss Guessous, Yanbo Liang, and Horace He. 2024a. **Flex attention: A programming model for generating optimized attention kernels**. *arXiv preprint arXiv:2412.05496*.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024b. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128. Association for Computational Linguistics.

Rickard Brüel Gabriëlsson, Jiacheng Zhu, Onkar Bhardwaj, Leshem Choshen, Kristjan Greenewald, Mikhail Yurochkin, and Justin Solomon. 2024. **Compress then serve: Serving thousands of loRA adapters with little overhead**.

Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. **Samsun corpus: A human-annotated dialogue dataset for abstractive summarization**. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. Association for Computational Linguistics.

Junxian Guo, Haotian Tang, Shang Yang, Zhekai Zhang, Zhijian Liu, and Song Han. 2024. **Block Sparse Attention**. <https://github.com/mit-han-lab/Block-Sparse-Attention>.

- Zhixiong Han, Yaru Hao, Li Dong, Yutao Sun, and Furu Wei. 2022. [Prototypical calibration for few-shot learning of language models](#). *arXiv preprint arXiv:2205.10183*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Steven C. H. Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. 2018. [Online learning: A comprehensive survey](#). *Preprint*, arXiv:1802.02871.
- Eduard Hovy, Laurie Gerber, Ulf Hermjakob, Chin-Yew Lin, and Deepak Ravichandran. 2001. Toward semantics-based answer pinpointing. In *Proceedings of the First International Conference on Human Language Technology Research*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Bowen Jin, Jinsung Yoon, Jiawei Han, and Sercan O. Arik. 2024. [Long-context llms meet rag: Overcoming challenges for long inputs in rag](#). *Preprint*, arXiv:2410.05983.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An evaluation dataset for intent classification and out-of-scope prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316, Hong Kong, China. Association for Computational Linguistics.
- Mukai Li, Shansan Gong, Jiangtao Feng, Yiheng Xu, Jun Zhang, Zhiyong Wu, and Lingpeng Kong. 2023. [In-context learning with many demonstration examples](#). *Preprint*, arXiv:2302.04931.
- Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. 2024. [Snapkv: Llm knows what you are looking for before generation](#). *arXiv preprint arXiv:2404.14469*.
- Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Gholamreza Haffari, and Bohan Zhuang. 2024a. [Minicache: Kv cache compression in depth dimension for large language models](#). *Preprint*, arXiv:2405.14366.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, and Yang Liu. 2024c. [Rethinking machine unlearning for large language models](#). *Preprint*, arXiv:2402.08787.
- Xiang Liu, Zhenheng Tang, Peijie Dong, Zeyu Li, Bo Li, Xuming Hu, and Xiaowen Chu. 2025. [Chunkkv: Semantic-preserving kv cache compression for efficient long-context llm inference](#). *arXiv preprint arXiv:2502.00299*.
- Songshuo Lu, Hua Wang, Yutian Rong, Zhi Chen, and Yaohua Tang. 2024. [Turborag: Accelerating retrieval-augmented generation with precomputed kv caches for chunked text](#). *ArXiv*, abs/2410.07590.
- Man Luo, Xin Xu, Yue Liu, Panupong Pasupat, and Mehran Kazemi. 2024. [In-context learning with retrieved demonstrations for language models: A survey](#). *arXiv preprint arXiv:2401.11624*.
- Yury A. Malkov and Dmitry A. Yashunin. 2020. [Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4):824–836.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. [Peft: State-of-the-art parameter-efficient fine-tuning methods](#). <https://github.com/huggingface/peft>.
- Meta AI. 2024. [Introducing llama 3.1: Our most capable models to date](#).
- Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. 2023. [Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12284–12314, Toronto, Canada. Association for Computational Linguistics.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Dov Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [Parallel context windows for large language models](#). In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. 2023. [ZeroSCROLLS: A zero-shot benchmark for long text understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7977–7989, Singapore. Association for Computational Linguistics.

- Ying Sheng, Shiyi Cao, Dacheng Li, Coleman Hooper, Nicholas Lee, Shuo Yang, Christopher Chou, Banghua Zhu, Lianmin Zheng, Kurt Keutzer, Joseph E. Gonzalez, and Ion Stoica. 2023. S-lora: Serving thousands of concurrent lora adapters. *arXiv preprint arXiv:2311.03285*.
- East Sun, Yan Wang, and Lan Tian. 2024. [Block-attention for efficient rag](#). *Preprint*, arXiv:2409.15355.
- TogetherAI. 2023. Llama-2-7b-32k-instruct - and fine-tuning for llama-2 models with together api.
- Haonan Wang, Qian Liu, Chao Du, Tongyao Zhu, Cunxiao Du, Kenji Kawaguchi, and Tianyu Pang. 2024. When precision meets position: Bfloat16 breaks down rope in long-context training. *arXiv preprint arXiv:2411.13476*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yifei Xia, Fangcheng Fu, Wentao Zhang, Jiawei Jiang, and Bin Cui. 2024. [Efficient multi-task llm quantization and serving for multiple lora adapters](#). In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Santa Clara, CA, USA. Poster Presentation.
- Guangxuan Xiao, Jiaming Tang, Jingwei Zuo, Junxian Guo, Shang Yang, Haotian Tang, Yao Fu, and Song Han. 2024a. Duoattention: Efficient long-context llm inference with retrieval and streaming heads. *arXiv*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024b. Efficient streaming language models with attention sinks. In *The Twelfth International Conference on Learning Representations*.
- Pawel Swietojanski Xingkun Liu, Arash Eshghi and Verena Rieser. 2019. Benchmarking natural language understanding services for building conversational agents. In *Proceedings of the Tenth International Workshop on Spoken Dialogue Systems Technology (IWSDS)*, Ortigia, Siracusa (SR), Italy. Springer.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1480–1489.
- Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat, and Danqi Chen. 2024. [Helmet: How to evaluate long-context language models effectively and thoroughly](#). *Preprint*, arXiv:2410.02694.
- Qingyu Yin, Xuzheng He, Luoao Deng, Chak Tou Leong, Fan Wang, Yanzhao Yan, Xiaoyu Shen, and Qiang Zhang. 2024. Deeper insights without updates: The power of in-context learning over fine-tuning. *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4138–4151.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 17283–17297.
- Yanqi Zhang, Yuwei Hu, Runyuan Zhao, John C. S. Lui, and Haibo Chen. 2024. [Unifying kv cache compression for large language models with leankv](#). *Preprint*, arXiv:2412.03131.
- Zhenyu (Allen) Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark W. Barrett, Zhangyang Wang, and Beidi Chen. 2023. [H2o: Heavy-hitter oracle for efficient generative inference of large language models](#). *ArXiv*, abs/2306.14048.

## A Efficiency metrics

In the main results § 4, we show the relative speedup. Here we report the raw metrics given our computational environment. The inference throughput is an average of 5 runs and rounded; for all settings, the standard deviation for this value is within 0.1. We evaluate inference with Hugging Face transformers (Wolf et al., 2020) using Flash Attention 2 (Dao, 2024), and applying Flex Attention (Dong et al., 2024a) for DBSA pre-encoding.

Method	Setup Time	Inference Throughput (queries/s)
<b>30k Context Length w/ Llama-2-7B w/ single L40S GPU</b>		
RetICL (no cache)	2 sec	0.8
Fixed ICL (cached)	9 sec	1.5
Finetuning	> 20 min	6.6
DBSA	6 sec	3.4
<b>30k Context Length w/ Llama-3.1-8B w/ single L40S GPU</b>		
RetICL (no cache)	2 sec	1.3
Fixed ICL (cached)	10 sec	11.6
Finetuning	> 20 min	15.8
DBSA	6 sec	12.8
<b>90k Context Length w/ Llama-3.1-8B w/ single A100 GPU</b>		
RetICL (no cache)	2 sec	0.4
Fixed ICL (cached)	13 sec	6.8
Finetuning	> 40 min	8.9
DBSA	7 sec	7.7

Table 9: Setup time and inference throughput (queries/s) across different methods and context lengths on Llama models with specified GPUs.

## B Additional Datasets

To further evaluate the effectiveness of DBSA, we include results on two additional datasets: MMLU (Hendrycks et al., 2021) and SAMSum (Gliwa et al., 2019). MMLU is a widely used benchmark for multitask multiple-choice question answering that covers a broad range of subjects. In addition to classification datasets, we also evaluate on SAMSum, a generative task on dialogue summarization.

Results in Tables 10 and 11 show strong performance using DBSA. Notably, DBSA outperforms fixedICL in terms of both task performance and efficiency (inference throughput).

<sup>2</sup>Experiment setup is as described in Section 3, which may differ from the official Llama 3 report, including prompt formatting, decoding strategy, and data preprocessing details.

Method	Accuracy	Throughput (q/s)
FixedICL (1 shot)	0.621	29.0
FixedICL (20 shot)	0.640	27.8
FixedICL (200 shot)	0.652	18.9
RetICL (20 from 200)	0.683	2.5
DBSA (20 from 200)	0.671	25.6

Table 10: MMLU results comparing DBSA with fixed and retrieval ICL baselines using a single A100 gpu. retICL and DBSA uses BM25 retrieval. DBSA uses 1 sink block, 2 preceding blocks, and 1 example per block.<sup>2</sup>

Method	ROUGE1/2/L	Throughput (q/s)
FixedICL (1 shot)	0.453 / 0.218 / 0.370	0.65
FixedICL (20 shot)	0.480 / 0.237 / 0.397	0.65
FixedICL (200 shot)	0.491 / 0.252 / 0.410	0.55
RetICL (20 from 200)	0.503 / 0.264 / 0.417	0.45
DBSA (20 from 200)	0.499 / 0.259 / 0.411	0.65

Table 11: SAMSum results comparing DBSA with fixed and retrieval ICL baselines using a single A100 gpu. retICL and DBSA uses BM25 retrieval. DBSA uses 1 sink block, 2 preceding blocks, and 1 example per block.

## C Finetuning Setup

We follow the same setup as Bertsch et al. (2024) for LoRA finetuning, using Hugging Face transformers (Wolf et al., 2020) along with the peft package (Mangrulkar et al., 2022). All finetuning is run on a single L40S GPU (48G). We finetune for 30 epochs, with LoRa rank  $r = 8$ , scaling factor  $\alpha = 32$ , and dropout 0.1. We use a learning rate of 0.001, a batch size of 32, and a weight decay of 0.01.

## D Accuracy Trend

In the main results (§ 4), we consider only the 30k (about 1k demos) or 90k (about 3k demos) demonstration pool. In Figures 6 and 7, we report two examples of the accuracy trend for each method as the size of demonstration pool increases. The trends we describe at 30k and 90k are consistent at other context lengths, although DBSA seems slightly less effective at very short context lengths; we do not believe this is a concern, as the relative efficiency gain from sparsity is minimal in short-context ICL.

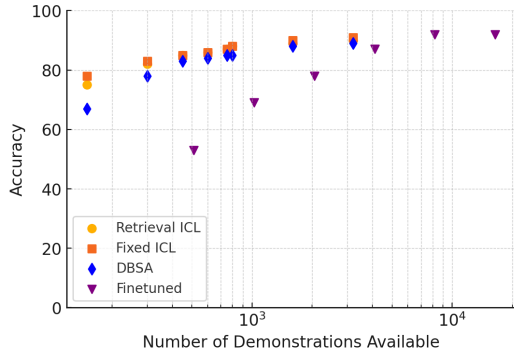


Figure 6: Performance trend for Banking 77.

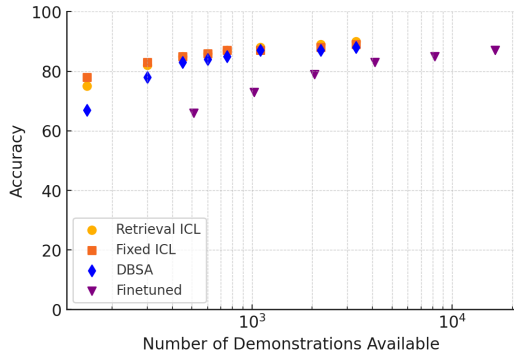


Figure 7: Performance trend for NLU.

as a query and obtain BM25 scores for each demonstration, resulting in a vector representation for each demonstration. This method is similar in concept to TF-IDF, which is commonly used in conjunction with k-means, especially before the popularity of semantic embeddings. We then run k-means clustering with the vector representations. To compare against this method, we randomly swap 10% examples between clusters to introduce diversity within each block.

## E Number of Preceding Blocks

In the main experiments the number of preceding blocks is set to 2. However, this is a hyperparameter choice, with more preceding blocks leading to higher accuracy. We show its impact on performance in Table 12.

Dataset	0	1	2	4	8	all (48-66)
banking77	0.828	0.875	0.883	0.881	0.895	0.913
clinc150	0.859	0.893	0.892	0.911	0.912	0.919
nlu	0.821	0.853	0.861	0.879	0.887	0.895
trec	0.875	0.929	0.941	0.955	0.964	0.967
trecfine	0.737	0.792	0.828	0.861	0.869	0.889
<b>Avg Acc</b>	0.824	0.869	0.881	0.897	0.905	0.917
<b>Recover (%)</b>	89.9	94.76	96.13	97.91	98.78	100

Table 12: Effect of changing number of preceding blocks in main experiment setting with Llama 3.1 8B 90k context length. Averaged over 3 runs.

## F Demonstration clustering implementation

In the ablations in Section 5.3, we use BM25 to construct a feature vector for each demonstration by first obtaining the vocabulary space of the demonstration pool. Then, we treat each vocabulary term