

COMPRESSING LLMs: THE TRUTH IS RARELY PURE AND NEVER SIMPLE

🍏 Ajay Jaiswal¹, Zhe Gan², Xianzhi Du², Bowen Zhang², Zhangyang Wang¹, Yinfei Yang²

¹University of Texas at Austin, ²Apple

ABSTRACT

Despite their remarkable achievements, modern Large Language Models (LLMs) face exorbitant computational and memory footprints. Recently, several works have shown significant success in *training-free* and *data-free* compression (pruning and quantization) of LLMs that achieve 50 - 60% sparsity and reduce the bit width to 3 or 4 bits per weight, with negligible degradation of perplexity over the uncompressed baseline. As recent research efforts are focused on developing increasingly sophisticated compression methods, our work takes a step back and re-evaluates the effectiveness of existing SoTA compression methods, which rely on a fairly simple and widely questioned metric, perplexity (even for dense LLMs). We introduce **Knowledge-Intensive Compressed LLM Benchmark (LLM-KICK)**, a collection of carefully curated tasks to redefine the evaluation protocol for compressed LLMs, which have significant alignment with their dense counterparts and perplexity fail to capture subtle change in their true capabilities. LLM-KICK unveils many favorable merits and unfortunate plights of current SoTA compression methods: all pruning methods suffer significant performance degradation, sometimes at trivial sparsity ratios (*e.g.*, 25-30%), and fail for N:M sparsity in knowledge-intensive tasks; current quantization methods are more successful than pruning; yet, pruned LLMs even at $\geq 50\%$ sparsity are robust in-context retrieval and summarization systems; among others. LLM-KICK is designed to holistically access compressed LLMs’ ability for language understanding, reasoning, generation, in-context retrieval, in-context summarization, *etc.* We hope our study can foster the development of better LLM compression methods. The reproduced codes are available at <https://github.com/VITA-Group/llm-kick>.

1 INTRODUCTION

Large Language Models (LLMs) are *omnipresent*, profoundly influencing not only the landscape of NLP (Ram et al., 2023; Liu et al., 2023a; Sawada et al., 2023; Qin et al., 2023; Zhuo, 2023; Lee et al., 2023), but also recently buttressing numerous computer vision (Lian et al., 2023; Wang et al., 2023; Lai et al., 2023; Lu et al., 2023) and graph neural networks (Ye et al., 2023; Chen et al., 2023; Qian et al., 2023; Duan et al., 2023) algorithms; achieving steller performance across various task leaderboards. Despite their numerous unprecedented capabilities, their democratization is primarily restricted by the presence of billions of parameters, which depends on astonishingly high computational and memory requirements. For example, GPT-175B requires 325 GB of GPU memory simply to load its model weights, and at least five A100 (80GB) GPUs with sophisticated parallelism techniques (Sheng et al., 2023).

To democratize LLMs, considerable efforts have been taking to mitigate their high computational cost, mainly divided into two research directions: *network pruning*, and *weight quantization*. The former shrinks network sizes by removing specific weights from the model – essentially setting them to zero, while the latter aims to quantize parameters into lower bit-level representations. Several recent success in network pruning (Sun et al., 2023; Frantar & Alistarh, 2023; Jaiswal et al., 2023a; Ma et al., 2023; Ji et al., 2023) and quantization (Liu et al., 2023c; Kim et al., 2023; Dettmers et al., 2023a; Frantar et al., 2022; Lin et al., 2023a; Dettmers et al., 2023c) (detailed related work

🍏 Work done during an internship at Apple.



Figure 1: **True Merits of SoTA Compression.** Top row indicates marginal increase in perplexity via using SoTA compression methods, when compared with simple magnitude-based pruning. Bottom row indicates the failure of compressed Vicuna-7B (Chiang et al., 2023) (via Magnitude, Wanda, SparseGPT, GPTQ) to respond correctly to knowledge-intensive factoid-based questions.

discussion in Appendix A.1) claim to retain the uncompressed LLM’s performance while achieving 50-60% sparsity or up to extreme 2-3 bit quantization. Although these advancements look fascinating, in most (if not all) cases, they heavily rely on **perplexity** as their primary metric to evaluate the performance claims. Such relatively restricted evaluations limit the scope for developing new compression methods, and are potentially ill-suited to identifying new and unexpected capabilities/limitations of compressed LLMs.

Perplexity, even in the case of dense LLMs, has been questioned as an unsatisfactory measure for comparing the true potential of LLMs, despite significant variations in model scales, training strategies, and architecture choices (Muhlgay et al., 2023). It is important to note that *all compressed models are derived from the same dense counterpart with high similarity*, and aforementioned differences don’t exist, making their evaluation more challenging. In this work, we revisit a widely known yet under-explored question: *How well does perplexity capture the change in capabilities of compressed LLMs that have significant alignment with their dense counterpart?* We focus on the case of compressed LLMs, because we observe comparatively more serious failure of perplexity to capture the delicate performance variations incurred across varying compression stages of LLMs, demanding a more fine-grained investigation.

In this work, we attempt to investigate the true promises and limitations of state-of-the-art compression algorithms for LLMs. We assemble the **first** comprehensive and diverse collection of tasks with varying difficulty levels to thoroughly study compressed LLMs under quantization and network pruning (structured and unstructured sparsity patterns). More specifically, we consider a broad range of tasks to evaluate subtle changes in pruned and quantized LLMs’ ability for *language understanding, reasoning, generation, in-context retrieval, long-context summarization, etc.* Note that none of the datasets in our multi-dimensional study of compressed LLMs was created from scratch, but we rely on existing datasets as they have been widely accepted by researchers, but unfortunately yet not been adopted to study the effect of compression. We rigorously measure the performance of SoTA quantization and pruning approaches (in their most common, default settings), to understand their potential for our challenging and interesting tasks with high practical value.

Our key observations and contributions can be unfolded as:

- We present *Knowledge-Intensive Compressed LLM Benchmark (LLM-KICK)*, to **re-define** the evaluation protocols for compressed LLMs and facilitate a comprehensive assessment of SoTA compression algorithms. The premise of our work is to develop a suite of challenging, realistic, and diverse tasks of high practical importance and datasets that can empower a systematic understanding of how *existing LLM compression strategies* truly perform in preserving performance

despite their similar perplexities, how they differ from each other, and how they compare against smaller LLMs of comparable parameter counts.

- LLM-KICK unveils many interesting and critical observations, that perplexity-based evaluations overlook. ① Most SoTA pruning methods suffer significant performance degradation, sometimes *at trivial sparsity ratios* (e.g., 25-30%), despite negligible changes in perplexity. ② All SoTA pruning methods do not work satisfactorily for structured N:M sparsity patterns on LLM-KICK. ③ Current SoTA LLM quantization methods are *more successful* in perpetuating performance in comparison to SoTA LLM pruning methods. ④ Compressed LLMs fail to generate knowledge-enriched and factually correct answers, despite the generated text is fluent, consistent, and coherent. ⑤ Compressed LLMs with larger architectures but same parameter counts perform poorer, which favors smaller dense models.
- We further investigate compressed LLMs’ ability for *in-context settings*, via adopting *in-context retrieval augmented question answering* (ICRA-QA) (Ram et al., 2023), and *text summarization with in-context learning* (IC-Sum) (Jain et al., 2023). To our surprise, pruned LLMs, even at non-trivial sparsity ratios (e.g., $\geq 50\%$), are **robust** retrieval systems, and can perform text summarization while maintaining similar performance as their dense counterpart. However, with increasing compression degrees, their ability to digest longer context is **affected more** than smaller context.

2 SoTA LLM COMPRESSION: PERPLEXITY, OR WHAT’S MORE?

Scaling neural networks, now LLMs, have achieved astonishing performance benefits on a wide array of tasks, but at the cost of gigantic computational and memory footprints. Network pruning and weight quantization are two popular remedies to mitigate these overheads due to billions of parameter counts in current LLMs. Despite numerous existing algorithms for pruning (Singh & Alistarh, 2020; Zhu & Gupta, 2017; Gale et al., 2019; Jaiswal et al., 2022; Lin et al., 2020; Liu et al., 2021a; Mostafa & Wang, 2019; Dettmers & Zettlemoyer, 2019; Evci et al., 2020) and quantization (Dong et al., 2022; Cardinaux et al., 2020; Kim et al., 2021; Liu et al., 2021b; Martinez et al., 2020), their ad-hoc adaptation for LLMs is restricted, due to the lack of luxury to perform iterative re-training to regain any performance drop during compression. Recently, several works have shown significant success in training-free and data-free compression of LLMs achieving 50-60% sparsity and reducing the bit-width down to 3 or 4 bits per weight, with negligible perplexity degradation relative to the uncompressed baseline.

Perplexity is a statistical measure of how confident a language model predicts a text sample and quantifies the “surprise” encoded within language models (the lower the perplexity, the better the model). Despite its popularity, perplexity has been widely questioned as an unsatisfactory measure to compare the true merits of two different LLMs (Muhlgay et al., 2023), even for dense models although they significantly vary in model scale, training strategies, and design choices (encoder only, decoder only, *etc.*). To address this issue, several works (Li et al., 2023; Kaddour et al., 2023; Muhlgay et al., 2023; Zhang et al., 2023; Valmeekam et al., 2022; Liu et al., 2023a; Sawada et al., 2023; Qin et al., 2023; Zhuo, 2023; Lee et al., 2023) attempt to go beyond perplexity, and evaluate the capabilities of **dense LLMs** across commonsense reasoning, language understanding, reading comprehension, programming, *etc.* However, it is **critically important** to note that *all compressed models are derived from the same dense counterpart with high similarity* sharing exactly the same scale, training strategies, design choices, *etc.* Surprisingly, unlike dense LLMs, no such effort has been carried out to understand subtle changes in the capabilities of compressed LLMs with varying compression strength. Orthogonal to the recent trend to develop new compression algorithms, our work provides the **first** attempt to assess the true merits and limitations of existing SoTA LLM compression algorithms, to provide a fair and detailed playground to develop better compression algorithms. We focus on the *case of compressed LLMs* because we observe the profound failure of perplexity in capturing the delicate performance variations across varying LLM compressions.

Figure 1(Top) illustrates the change in perplexity of SoTA compression methods (pruning and quantization), such as SparseGPT, Wanda, GPTQ and baseline one-shot magnitude-based pruning on Vicuna-7B, 13B, and 33B (Chiang et al., 2023). Clearly, the perplexity (\downarrow) of all models does not show any significant variation up to 45-60%, with a complete failure to capture subtle changes in the abilities of LLMs when compressed. It is also interesting to observe that to a certain degree of sparsity ($\sim 30\%$), all SoTA pruning methods have almost similar performance as the simple baseline of

one-shot magnitude-based pruning, which raises questions about their true merits within this sparsity range. Figure 1(Bottom) show the response of Vicuna-7B model when compressed with Magnitude, SparseGPT, and Wanda by 50% and quantized up to 4-bit. The uncompressed Vicuna-7B was successfully able to generate the correct answer, but all compressed versions failed to respond correctly, hallucinating with either wrong facts or irrelevant responses.

3 LLM-KICK: UNVEILING TRUE MERITS OF LLM COMPRESSION

LLM-KICK, short for Knowledge-intensive Compressed LLM Benchmark, is crafted to bring the attention of LLM compression community towards incompetence of perplexity to correctly reflect subtle changes in the ability of LLMs derived from dense counterparts with varying compression strength. LLM-KICK consists of a suite of challenging, realistic, and diverse task settings of high practical importance and datasets that can empower a systematic understanding of how existing LLM compression strategies truly perform in preserving performance despite having similar perplexity. Our work thoroughly investigates proclaimed merits/limitations of pruned and quantized LLMs for language understanding, reasoning, generation, in-context retrieval, in-context summarization,

Specifically, LLM-KICK consists of 3 broad task settings to study how compression impacts knowledge encoded during pre-training, how compressed LLMs perform tasks when required knowledge is augmented in-context, and how well compressed LLMs perform instruction following. To compartmentalize task difficulty and diversity, we include factoid-based QA, multiple-choice reasoning-based QA, in-context retrieval augmented QA, in-context text summarization, and instruction-based free-form text generation. Instead of creating new datasets, we carefully curate LLM-KICK from prior works and open-source GitHub repositories which have been widely accepted by researchers, but yet not explored by the LLM compression researchers. Our detailed prompt design strategies for different task settings can be found in Appendix A.2.

To reduce the expense of redundant experiments and clutter in results, our work primarily focuses on the top-2 existing training-free and data-free LLM pruning techniques SparseGPT (Frantar & Alistarh, 2023) and Wanda (Sun et al., 2023), along with the baseline of One-shot Magnitude-based Pruning (Han et al., 2016), plus a popular quantization technique (GPTQ) among recently available choices (Lin et al., 2023a; Frantar et al., 2022; Dettmers et al., 2023c). We consider two types of sparsities: i) Unstructured Sparsity: individual model weights are zeroed out independently, leading to irregular zero patterns (LeCun et al., 1990; Han et al., 2016); ii) Structured N:M Sparsity: a fine-grained sparsity pattern in which only weights are non-zero for every continuous M weights (Nvidia, 2020; Zhou et al., 2021). We use Vicuna models for experiments, which are open-source chatbot models trained by fine-tuning LLaMA (Chiang et al., 2023) on user-shared conversations collected from ShareGPT, and have demonstrated impressive 90% quality of OpenAI ChatGPT and Google Bard. Note that the aim of this work is not limited to identifying the failure cases of SoTA pruning methods, but instead provides an in-depth lookup of LLM's ability under compression, and bring new insights which include highlighting observations that work in favor of current SoTA compression methods.

Formally, we study the performance drop of LLMs after compression (without fine-tuning) with respect to their dense counterparts using a compression algorithm \mathcal{C} . For a pre-trained LLM $(x; \theta)$, a compressed LLM is a network $f_{\text{comp}}(x; \theta_c)$, which is a copy of $(x; \theta)$ with some weights fixed to 0 indicated by the pruning mask θ_c in the case of pruning, or quantized to k -bit using a quantization algorithm. Next, we define a matching compressed LLM.

Matching Compressed LLM: A compressed LLM $f_{\text{comp}}(x; \theta_c)$ is matching for a compression algorithm \mathcal{C} on task T , if it results in performance no less than (compression tolerance regime) in comparison with $(x; \theta; T)$. In this work, we consider θ_0 to be 5% of the performance $\text{df}(x; \theta; T)$.

Note that θ_0 is a simple indicator of the tolerance level of performance drop when we start compressing any LLM. Many prior works (Chen et al., 2020b; Jaiswal et al., 2023a) consider matching thresholds to be the same as the dense subnetwork performance or within the margins of 1%. However, in our work, we carefully relaxed it to 5% performance drop as an acceptable tolerance (before

Figure 2: Compressed LLMs for Factoid-based QA. Performance comparison of compressed LLMs on Factoid-QA task using FreebaseQA (Jiang et al., 2019). Results (average across 3 independent runs) presented are for structured (N:M sparsity), unstructured sparsity, and quantization.

calling the compressed model useless) keeping in mind that the performance of compressed LLM on any of our task categories/disciplines remains above the random guess.

3.1 SETTING 1: HOW WELL COMPRESSED LLM S ACCESS REMAINING KNOWLEDGE?

① Factoid-based Question Answering

Task Definition and Rationale. Factoid-based Question Answering (Factoid-QA) (Iyer et al., 2014), which asks precise facts about entities, is a long-standing problem in NLP. A typical Factoid-QA task aims to search for entities or entity attributes from a knowledge graph, and it is widely used as a tool in academia, commercial search engines, and conversational assistants. Modern LLMs are trained on gigantic text corpora ingesting a large amount of world knowledge about entities and their relationships during pre-training, and have unique abilities to generate factually correct responses to user queries. In this task setting, we aim to investigate how compression impacts LLMs' ability to answer natural language questions using facts, i.e., entities or attributes knowledge ingested within them during pre-training.

Dataset Details. We use FreebaseQA (Jiang et al., 2019) which is a dataset for open-domain QA over the Freebase knowledge graph. The QA pairs are collected from various sources, including the TriviaQA dataset (Joshi et al., 2017) and other trivia websites (QuizBalls, QuizZone, KnowQuiz), and are matched against Freebase to generate relevant subject-predicate-object triples that were further verified by human annotators. TriviaQA dataset shows rich linguistic variation and complexity, making it a good testbed for evaluating knowledge ingested within LLMs.

Results and Analysis. The results of various LLM compression methods are demonstrated in Figure 2. Our primary observations include: ① All SoTA LLM pruning methods seemingly fail to find matching sparse LLMs even at trivial sparsities such as 30-35%. While several methods maintain the matching performance at 20-25% sparsity, their performance starts to drop significantly after that undergoing catastrophic failures as sparsity ratio increases. This is in contrast with the claim made by SoTA pruning methods that pruning up to 50-60% of LLMs doesn't have any significant degradation on performance. ② All pruning methods doesn't work for fine-grained structured N:M sparsity patterns with performance drop as severe as 50%. ③ 8-10% drop in performance for non-aggressive 8-bit quantization indicates that along with chasing for aggressive quantization levels (1-2 bits), it is also important to focus on yet unsolved 8-bit quantization.

② Multiple-Choice Reasoning based Question Answering

Task Formulation and Rationale. Multiple-Choice Reasoning based QA (MCR-QA) uses a natural prompting approach to present the question and answer options to the LLMs jointly, and have it output the symbol (e.g., "A") associated with its chosen answer option. It allows the model to explicitly compare answer options. In this setting, we aim to investigate compressed LLMs' ability to understand natural language questions, effectively reason using knowledge remaining within them, and successfully associate the correct answer among the given answer options with the symbols that represent them; potentially minimizing the effect of tokenization and exact answer generation.

Dataset Details. We use the popular MMLU (Massive Multitask Language Understanding) benchmark which covers 50+ subjects across STEM, Humanities, Social Sciences, and more (Hendrycks et al., 2020). It ranges in difficulty from an elementary level to an advanced professional level, and it tests both world knowledge and problem-solving ability of LLMs. The granularity and breadth of subjects make it ideal for fine-grained evaluation of compressed LLMs' blind spots.

Figure 3: Compressed LLMs for Multiple-Choice Reasoning based QA Performance comparison of compressed LLMs on MCR-QA tasks using the MMLU benchmark (Hendrycks et al., 2020). Results (average across 3 independent runs) presented are for structured (N:M sparsity), unstructured sparsity, and quantization.

Results and Analysis. The results of various LLM compression methods are demonstrated in Figure 3. Our primary observations include: ① Despite a similar matching compression regime (20-40%) to Factoid-QA, the abrupt performance drop of all SoTA pruning methods for MMLU is comparatively subtle due to relaxing the task setting from exact answer generation to correct answer selection. ② No matching compressed LLMs are found for N:M structured sparsity. ③ SoTA LLM quantization is seemingly more successful than SoTA pruning methods: we found 8-bit and 4-bit compressed LLM to be matching for Vicuna-7B and Vicuna-13B, respectively. ④ Interestingly, both quantization and pruning have comparatively higher performance drop for Humanities and Social Science wrt. STEM, which indicates compression impacts some disciplines more than others. ⑤ Surprisingly, within the compression tolerance regime, simple one-shot magnitude pruning seems to perform quite well in comparison with SoTA pruning method, illustrating its high effectiveness.

3.2 SETTING 2: HOW WELL COMPRESSED LLM S SYNTHESIZE AUGMENTED KNOWLEDGE?

① In-context Retrieval Augmented Question Answering

Task Formulation and Rationale. In-context Retrieval-Augmented Question Answering (ICRA-QA) (Ram et al., 2023) grounds the LLM answer generation by conditioning on relevant documents retrieved from an external knowledge source using retrieval algorithms like BM25. Our ICRA-QA evaluation system includes two high-level components: ① document selection, selecting the set of documents upon which to condition; ② document reading, determining how to incorporate the selected documents into the LLM answer process, which requires extracting correct answer phrases from conditioned documents. To discount the impact of the lost encoded knowledge during compression, ICRA-QA augments the required relevant knowledge for QA task directly within the

Figure 4: Compressed LLMs for In-context Retrieval Augmented QA. Performance comparison of compressed LLMs on ICRA-QA task. We present head-to-head comparison of closed-book evaluation (no external knowledge is augmented in-context) with open-book evaluation (external knowledge is augmented in-context). Results (average across 3 independent runs) presented are for structured N:M sparsity, unstructured sparsity, and quantization.

prompt context. In this task setting, we aim to evaluate compressed LLMs' ability to synthesize long in-context knowledge provided within input prompts, and locate and retrieve correct answers within it. We also present a head-to-head comparison of how augmented knowledge can be used to supplement the lost knowledge under compression.

Dataset Details. We use TriviaQA (Joshi et al., 2017) for evaluation, a popular reading comprehension dataset which includes 95K question-answer pairs authored by trivia enthusiasts and independently gathered evidence documents, six per question on average, that provide high-quality distant supervision for answering the questions.

Results and Analysis. The results of various LLM compression methods are demonstrated in Figure 17. The closed-book setting differs from ICRA-QA (using the open-book setting) only in terms of whether conditioning on relevant documents retrieved from an external knowledge source. Our key findings are: ① When compressed LLMs are conditioned on external knowledge (open book) and assigned the task of in-context retrieval, extracting correct answer phrases from in-context knowledge, they perform significantly well even in extremely high compression regime. Vicuna-7B can remain matching till 40% sparsity and 8-bit quantization, while Vicuna-13B can remain matching up to 50% sparsity and 4-bit quantization. Our experimental results send a positive signal that even if high compression leads to significant knowledge loss, it doesn't leave LLMs completely useless and they still work as robust in-context retriever. ② Despite we observe a significant benefit while conditioning external knowledge, no matching compressed LLM can be identified for N:M sparsity. ③ Again, we observe surprisingly good performance of simple one-shot unstructured magnitude pruning wrt. SparseGPT (second-order pruning) and Wanda (activation-based pruning) that rely on calibration data.

② In-Context Text Summarization

Task Formulation and Details. Modern LLMs have shown astonishing success in summarizing long-context documents in both abstractive and extractive settings. However, we first explored how compression impacts LLMs' capability for summarization. In this task setting, we aim to investigate compressed LLMs' ability to hold onto consistency, coherence, fluency, and relevance when prompted to summarize textual information of varying length (small, medium, and large) in abstractive setting (Jain et al., 2023). For evaluation, similar to Zheng et al. (2023), we propose to use GPT-4 as a judge, which compares the compressed LLM generated summaries wrt. GPT-3.5 (text-davinci-003) generated summaries. Detailed evaluation settings can be found in Appendix A.3.

Dataset Details. We use a popular summarization dataset CNN/DailyMail (Chen et al., 2016) for evaluation, which is an English-language dataset containing just over 300k unique news articles written by journalists at CNN and DailyMail. We created 3 subset categories: small (470 words), medium (470 and 790 words), and large (790 words) of stories, each with 100 articles reflecting word distribution of CNN/DailyMail to minimize OpenAI API costs.

Figure 5: Compressed LLMs for In-Context Summarization. Performance comparison of compressed Vicuna-7B for in-context summarization of small, medium, and large stories while preserving coherence, consistency, fluency, and relevance. Results (average across 3 independent runs) presented are for structured (2:4 sparsity - Row 3), unstructured sparsity, and quantization.

Figure 6: Compressed LLMs for Instruction Following. LLM-as-a-Judge: GPT-4 based evaluation of compressed Vicuna-7B response wrt. ChatGPT (Vic-003). (Left) unstructured sparsity; (middle) structured N:M sparsity; (c) comparison of average unique token counts generated by compressed Vicuna-7B for 80 prompts across 10 different categories.

Results and Analysis. Results are summarized in Figure 5. We summarize our main observations as: ① All pruning and quantization methods tend to perform surprisingly well for in-context summarization, preserving high consistency, coherence, fluency, and relevance in generated summaries, which is an encouraging observation in favor of compression. ② With increasing context length (i.e., long stories), we observe a sharper performance drop for compressed LLMs, which highlights that compression impacts LLMs' ability to synthesize and summarize longer context length. ③ Quantization again seems to perform better than SoTA pruning methods, and surprisingly benefits positively over the dense model performance. ④ No matching compressed LLM can be identified for 2:4 structured sparsity.

3.3 SETTING 3: HOW WELL COMPRESSED LLM S PERFORM INSTRUCTION FOLLOWING?

Task Formulation and Rationale. In this task setting, we investigate compressed LLMs' ability to answer open-ended questions and evaluate their multi-turn conversational and instruction-following ability – two critical elements for human preference evaluation of AI chatbots is a challenging task, as it requires examining language understanding, reasoning, and context awareness. To compare the performance of compressed LLMs' responses, we closely follow the prompt design setting in MT-Bench (Zheng et al., 2023) using GPT-4 as a judge. We prompt GPT-4 to rate the answers gen-

erated by compressed LLMs wrt. GPT-3.5 (text-davinci-003) model based on varying metrics (correctness, helpfulness, logic, accuracy) on a scale of [0-10] with detailed explanations.

Dataset Details. We rely on the 80 high quality multi-turn questions identified in MT-Bench (Zheng et al., 2023). This setting covers common-use human-centric interaction with LLMs, and focuses on challenging questions to differentiate models. We used 8 common categories of user prompts to guide the prompt construction to interact with compressed LLMs: writing, roleplay, extraction, reasoning, math, coding, etc. For each category, we adopted manually designed 10 multi-turn questions from MT-Bench to evaluate our compressed models. Details can be found in Appendix A.4.

Results and Analysis. Results are summarized in Figure 6. Our primary observations are: ① Unlike in-context text summarization, in this task setting, compressed LLMs have to access the knowledge to respond to conversations maintaining high helpfulness, relevance, accuracy, and detail. We again observe that compressed LLMs with various pruning methods matching only up to sparsity ratio of 25%. ② Surprisingly, in the matching regime, the simple baseline of one-shot magnitude pruning performs comparable or slightly better than SoTA pruning methods. ③ No matching subnetwork can be identified for N:M sparsity. ④ Interestingly, our average generated unique token analysis in Figure 6(c) illustrates that compressed LLMs lose the ability to generate distinct unique content, instead, they can only produce more repetitive texts.

4 ADDITIONAL RESULTS AND DISCUSSIONS

Small-Dense vs. Large-Sparse: which is favorable? We attempt to understand an interesting question: if pruned LLMs with larger architecture (Large-Sparse) is better than smaller dense models with similar parameter count (Small-Dense)? Pruning large LLMs doesn't come for free, and it is important to investigate if the cost of pruning can be reflected in the performance benefit of Large-Sparse models. To our surprise, in comparison with dense Vicuna-7B (MMLU accuracy 46.7%), we found compressed Vicuna-13B with exactly similar parameter count (46.16% sparsity) of 7 billion using one-shot magnitude, Wanda, SparseGPT can only achieve MMLU accuracy of 31.7%, 45.3%, and 46.3%, respectively. This is a clear indication that current sparsity algorithms are up to a stage where the cost of pruning can be justified by performance benefits obtained from large-sparse compressed models.

How many calibration data samples are needed? We attempt to analyze how calibration dependent pruning methods (Wanda and SparseGPT) perform with varying amount of calibration samples. Figure 7 illustrates the zero-shot performance of 50% & 70% pruned Vicuna-7B using Wanda and SparseGPT on knowledge-intensive MMLU benchmark. It is interesting to observe that calibration sample count plays a vital role in preserving the performance of SparseGPT unlike Wanda. Note that at high sparsity ratio (70%), Wanda cannot recover any performance; SparseGPT surprisingly benefits noticeably from calibration. This suggests that carefully selected calibration samples can play a vital role in designing better pruning algorithms to compress LLMs even up to significantly high sparsity.

Figure 7: Zero-shot performance of 50% & 70% pruned Vicuna-7B wrt. calibration sample counts.

5 CONCLUSION AND LIMITATIONS

In this paper, we propose to explore the effectiveness of SoTA compression methods beyond perplexity to address the inability of perplexity to capture the subtle variations incurred during the derivation of compressed LLMs from their dense counterparts. Our work introduces Knowledge-Intensive Compressed LLM Benchmark (LLM-KICK) to facilitate a fair and holistic evaluation by unveiling many merits and pitfalls of SoTA compression methods. Our study reveals that compression significantly impacts the knowledge encoded in LLMs during pre-training, compressed LLMs perform quite well with knowledge augmented in-context settings. We primarily restrict our evaluation to Vicuna (decoder-only architecture) due to its open-source license, high performance, and instruction-following ability. For future work, we aim to investigate how the lost knowledge due to compression can be recovered using parameter-efficient fine-tuning methods, LoRA (Hu et al., 2021) and QLoRA (Dettmers et al., 2023b).

REFERENCES

- Fabien Cardinaux, Stefan Uhlich, Kazuki Yoshiyama, Javier Alonso Garcia Mauch, Stephen Tiedemann, Thomas Kemp, and Akira Nakamura. Iteratively training look-up tables for network quantization. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):860–870, 2020.
- Danqi Chen, Jason Bolton, and Christopher D Manning. A thorough examination of the cnn/daily mail reading comprehension task. *arXiv preprint arXiv:1606.02858*, 2016.
- Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. The lottery ticket hypothesis for pre-trained bert networks. *Advances in neural information processing systems*, 33:15834–15846, 2020a.
- Xiaohan Chen, Yu Cheng, Shuohang Wang, Zhe Gan, Zhangyang Wang, and Jingjing Liu. Earlybert: Efficient bert training via early-bird lottery tickets. *arXiv preprint arXiv:2101.00063*, 2020b.
- Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Hongzhi Wen, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, Wenqi Fan, Hui Liu, et al. Exploring the potential of large language models (llms) in learning on graphs. *arXiv preprint arXiv:2307.03393*, 2023.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332, 2022.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient netuning of quantized llms. *ArXiv*, abs/2305.14314, 2023a. URL <https://api.semanticscholar.org/CorpusID:258841328>.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient netuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023b.
- Tim Dettmers, Ruslan Svirschevski, Vage Egiazarian, Denis Kuznedelev, Elias Frantar, Saleh Ashkboos, Alexander Borzunov, Torsten Hoefer, and Dan Alistarh. Spqr: A sparse-quantized representation for near-lossless llm weight compression. *ArXiv*, abs/2306.03078, 2023c. URL <https://api.semanticscholar.org/CorpusID:259076379>.
- Runpei Dong, Zhanhong Tan, Mengdi Wu, Linfeng Zhang, and Kaisheng Ma. Finding the task-optimal low-bit sub-distribution in deep neural networks. *International Conference on Machine Learning*, pp. 5343–5359. PMLR, 2022.
- Keyu Duan, Qian Liu, Tat-Seng Chua, Shuicheng Yan, Wei Tsang Ooi, Qizhe Xie, and Junxian He. Simteg: A frustratingly simple approach improves textual graph learning. *arXiv preprint arXiv:2308.02565*, 2023.
- Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. *International Conference on Machine Learning*, pp. 2943–2952. PMLR, 2020.
- Elias Frantar and Dan Alistarh. Sparsegpt: Massive language models can be accurately pruned in one-shot, 2023.
- Elias Frantar, Eldar Kurtic, and Dan Alistarh. M-fac: Efficient matrix-free approximations of second-order information. *Advances in Neural Information Processing Systems*, 34:14873–14886, 2021.

- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *ArXiv, abs/2210.17323*, 2022. URL <https://api.semanticscholar.org/CorpusID:253237200>.
- Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. preprint *arXiv:1902.09574*, 2019.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representation*, 2016.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Mohit Iyyer, Jordan L. Boyd-Graber, Leonardo Max Batista Claudino, Richard Socher, and Hal Daumé. A neural network for factoid question answering over paragraphs. *Conference on Empirical Methods in Natural Language Processing*, 2014. URL <https://api.semanticscholar.org/CorpusID:216034672>.
- Sameer Jain, Vaishakh Keshava, Swarnashree Mysore Sathyendra, Patrick Fernandes, Pengfei Liu, Graham Neubig, and Chunting Zhou. Multi-dimensional evaluation of text summarization with in-context learning. *arXiv preprint arXiv:2306.01200*, 2023.
- Ajay Jaiswal, Shiwei Liu, Tianlong Chen, and Zhangyang Wang. The emergence of essential sparsity in large pre-trained models: The weights that matter. *arXiv preprint arXiv:2306.03805*, 2023a.
- Ajay Kumar Jaiswal, Haoyu Ma, Tianlong Chen, Ying Ding, and Zhangyang Wang. Training your sparse neural network better with any mask. *International Conference on Machine Learning*, pp. 9833–9844. PMLR, 2022.
- Ajay Kumar Jaiswal, Shiwei Liu, Tianlong Chen, Ying Ding, and Zhangyang Wang. Instant soup: Cheap pruning ensembles in a single pass can draw lottery tickets from large models. *International Conference on Machine Learning*, pp. 14691–14701. PMLR, 2023b.
- Yupeng Ji, Yibo Cao, and Jiucui Liu. Pruning large language models via accuracy prediction. preprint *arXiv:2309.09507*, 2023.
- Kelvin Jiang, Dekun Wu, and Hui Jiang. Freebaseqa: A new factoid qa data set matching trivia-style question-answer pairs with freebase. *North American Chapter of the Association for Computational Linguistics*, 2019. URL <https://api.semanticscholar.org/CorpusID:174800890>.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*, 2023.
- Jeonghoon Kim, Jung Hyun Lee, Sungdong Kim, Joonsuk Park, Kang Min Yoo, Se Jung Kwon, and Dongsoo Lee. Memory-efficient re-tuning of compressed large language models via sub-4-bit integer quantization. *ArXiv, abs/2305.14152*, 2023. URL <https://api.semanticscholar.org/CorpusID:258841104>.
- Sehoon Kim, Amir Gholami, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. I-bert: Integer-only bert quantization. *International conference on machine learning*, pp. 5506–5518. PMLR, 2021.

- Eldar Kurtic, Daniel Campos, Tuan Nguyen, Elias Frantar, Mark Kurtz, Benjamin Fineran, Michael Goin, and Dan Alistarh. The optimal bert surgeon: Scalable and accurate second-order pruning for large language models. *arXiv preprint arXiv:2203.07259*, 2022.
- François Lagunas, Ella Charlaix, Victor Sanh, and Alexander M Rush. Block pruning for faster transformers. *arXiv preprint arXiv:2109.04838*, 2021.
- Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language models. *arXiv preprint arXiv:2308.00692*, 2023.
- Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. *Advances in neural information processing systems*, pp. 598–605, 1990.
- Noah Lee, Na Min An, and James Thorne. Can large language models infer and disagree like humans? *ArXiv*, abs/2305.13788, 2023. URL <https://api.semanticscholar.org/CorpusID:258841424>.
- Xiang Li, Yiqun Yao, Xin Jiang, Xuezhi Fang, Xuying Meng, Siqi Fan, Peng Han, Jing Li, Li Du, Bowen Qin, et al. Flm-101b: An open llm and how to train it with 100 k budget. *arXiv preprint arXiv:2309.03852*, 2023.
- Zonglin Li, Chong You, Srinadh Bhojanapalli, Daliang Li, Ankit Singh Rawat, Sashank J Reddi, Ke Ye, Felix Chern, Felix Yu, Ruiqi Guo, et al. Large models are parsimonious learners: Activation sparsity in trained transformers. *arXiv preprint arXiv:2210.06313*, 2022.
- Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. *arXiv preprint arXiv:2305.13652*, 2023.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. Awq: Activation-aware weight quantization for llm compression and acceleration. *ArXiv*, abs/2306.00978, 2023a. URL <https://api.semanticscholar.org/CorpusID:258999941>.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. Awq: Activation-aware weight quantization for llm compression and acceleration. *arXiv preprint arXiv:2306.00978*, 2023b.
- Tao Lin, Sebastian U. Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. Dynamic model pruning with feedback. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJem8ISFwB>.
- Junling Liu, Chao Liu, Peilin Zhou, Qichen Ye, Dading Chong, Kang Zhou, Yueqi Xie, Yuwei Cao, Shoujin Wang, Chenyu You, et al. Llmrec: Benchmarking large language models on recommendation task. *arXiv preprint arXiv:2308.12241*, 2023a.
- Shiwei Liu, Tianlong Chen, Xiaohan Chen, Zahra Atashgahi, Lu Yin, Huanyu Kou, Li Shen, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Sparse training via boosting pruning plasticity with neuroregeneration. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021a.
- Shiwei Liu, Tianlong Chen, Zhenyu Zhang, Xuxi Chen, Tianjin Huang, Ajay Jaiswal, and Zhangyang Wang. Sparsity may cry: Let us fail (current) sparse neural networks together! *arXiv preprint arXiv:2303.02141*, 2023b.
- Zechun Liu, Barlas Oguz, Changsheng Zhao, Ernie Chang, Pierre Stock, Yashar Mehdad, Yangyang Shi, Raghuraman Krishnamoorthi, and Vikas Chandra. Llm-qat: Data-free quantization aware training for large language models. *arXiv preprint arXiv:2305.17888*, 2023c.
- Zhenhua Liu, Yunhe Wang, Kai Han, Wei Zhang, Siwei Ma, and Wen Gao. Post-training quantization for vision transformer. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.) *Advances in Neural Information Processing Systems*, 2021b. URL <https://openreview.net/forum?id=9TX5OsKJvM>.

- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. Chameleon: Plug-and-play compositional reasoning with large language models. arXiv preprint arXiv:2304.09842, 2023.
- Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models. arXiv preprint arXiv:2305.11627, 2023.
- Julieta Martinez, Jashan Shewakramani, Ting Liu, Ioan Andras, Wenyuan Zeng, and Raquel Urtasun. Permute, quantize, and re-tune: Efficient compression of neural networks. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 15694–15703, 2020. URL <https://api.semanticscholar.org/CorpusID:225103308>.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. arXiv preprint arXiv:1609.07843, 2016.
- Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. International Conference on Machine Learning, 2019.
- Dor Muhlgay, Ori Ram, Inbal Magar, Yoav Levine, Nir Ratner, Yonatan Belinkov, Omri Abend, Kevin Leyton-Brown, Amnon Shashua, and Yoav Shoham. Generating benchmarks for factuality evaluation of language models. arXiv preprint arXiv:2307.06908, 2023.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. Abstractive text summarization using sequence-to-sequence rnns and beyond. arXiv preprint arXiv:1602.06023, 2016.
- Nvidia. Nvidia a100 tensor core gpu architecture. <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper>, 2020.
- Chen Qian, Huayi Tang, Zhirui Yang, Hong Liang, and Yong Liu. Can large language models empower molecular property prediction? arXiv preprint arXiv:2307.07443, 2023.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiao Chen, Michihiro Yasunaga, and Diyi Yang. Is chatgpt a general-purpose natural language processing task solver? arXiv preprint arXiv:2302.06476, 2023.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. In-context retrieval-augmented language models. arXiv preprint arXiv:2302.00083, 2023.
- Victor Sanh, Thomas Wolf, and Alexander Rush. Movement pruning: Adaptive sparsity by re-tuning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds), Advances in Neural Information Processing Systems, volume 33, pp. 20378–20389. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/eaef15aabaa768ae4a5993a8a4f4fa6e4-Paper.pdf>.
- Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Krnias, John J Nay, Kshitij Gupta, and Aran Komatsuzaki. Arb: Advanced reasoning benchmark for large language models. arXiv preprint arXiv:2307.13692, 2023.
- Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Daniel Y Fu, Zhiqiang Xie, Beidi Chen, Clark Barrett, Joseph E Gonzalez, et al. High-throughput generative inference of large language models with a single gpu. arXiv preprint arXiv:2303.06865, 2023.
- Sidak Pal Singh and Dan Alistarh. Wood sher: Efficient second-order approximation for neural network compression. Advances in Neural Information Processing Systems, 18098–18109, 2020.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. A simple and effective pruning approach for large language models. arXiv preprint arXiv:2306.11695, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Roziere, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971, 2023.

- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Large language models still can't plan (a benchmark for llms on planning and reasoning about change). ArXiv, abs/2206.10498, 2022. URL <https://api.semanticscholar.org/CorpusID:249889477> .
- Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. VisionLLM: Large language model is also an open-ended decoder for vision-centric tasks. arXiv preprint arXiv:2305.11175, 2023.
- Dongkuan Xu, Ian EH Yen, Jinxi Zhao, and Zhibin Xiao. Rethinking network pruning—under the pre-train and re-tune paradigm. arXiv preprint arXiv:2104.08682, 2021.
- Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, and Yongfeng Zhang. Natural language is all a graph needs. arXiv preprint arXiv:2308.07134, 2023.
- Omer Zafrir, Ariel Larey, Guy Boudoukh, Haihao Shen, and Moshe Wasserblat. Prune once for all: Sparse pre-trained language models. arXiv preprint arXiv:2111.05754, 2021.
- Qingru Zhang, Simiao Zuo, Chen Liang, Alexander Bukharin, Pengcheng He, Weizhu Chen, and Tuo Zhao. Platon: Pruning large transformer models with upper confidence bound of weight importance. In International Conference on Machine Learning, pp. 26809–26823. PMLR, 2022.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori Hashimoto. Benchmarking large language models for news summarization. ArXiv, abs/2301.13848, 2023. URL <https://api.semanticscholar.org/CorpusID:256416014> .
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging LLM-as-a-judge with mt-bench and chatbot arena. arXiv preprint arXiv:2306.05685, 2023.
- Aojun Zhou, Yukun Ma, Junnan Zhu, Jianbo Liu, Zhijie Zhang, Kun Yuan, Wenxiu Sun, and Hongsheng Li. Learning fine-grained structured sparse neural networks from scratch. arXiv preprint arXiv:2102.04010, 2021.
- Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. arXiv preprint arXiv:1710.01878, 2017.
- Terry Yue Zhuo. Large language models are state-of-the-art evaluators of code generation. arXiv preprint arXiv:2304.14317, 2023.

A APPENDIX

A.1 RELATED WORKS

A.1.1 SPARSITY IN LARGE LANGUAGE MODELS

The advent of large-scale pre-trained models has led to the development of advanced post-training pruning methods, aiming to enhance the cost-effectiveness of these expansive models (Sanh et al., 2020; Chen et al., 2020a; Jaiswal et al., 2023b; Zafrir et al., 2021; Kurtic et al., 2022; Xu et al., 2021; Lagunas et al., 2021; Zhang et al., 2022; Frantar et al., 2021; Jaiswal et al., 2023a; Ma et al., 2023; Ji et al., 2023). Among them, Frantar et al. (2021) extend second-order pruning to the BERT-level scale, enabling the pruning of blocks of weights and achieving state-of-the-art results for sparse BERT. Frantar & Alistarh (2023) introduce SparseGPT for pruning large language models (LLMs) in a single shot without requiring re-training or re-tuning. They leverage column-wise second-order pruning, and successfully remove 100B weights from OPT-175B without a significant increase in perplexity. More recently, Sun et al. (2023) propose a straightforward pruning method that takes both weights and activations into account, demonstrating comparable performance to Frantar & Alistarh (2023). Li et al. (2022) reveal that activation sparsity is a prevalent phenomenon in Transformers (90% of intermediate output), yielding another opportunity for acceleration. Liu et al. (2023b) introduce a large-scale SMC-Bench, indicating that state-of-the-art magnitude- and/or gradient-based sparse algorithms fall short when applied out-of-the-box to larger-scale models and a selected of complex downstream tasks.

A.1.2 QUANTIZATION IN LARGE LANGUAGE MODELS

With the recent open-source releases of language models like BLOOM, Vicuna, LLaMa, OPT, etc., quantization has emerged as a widely embraced technique to alleviate the storage and computational overhead of deep learning models. Recent research endeavors have harnessed quantization to compress LLMs and they can be classified into the two mentioned approaches: Quantization-Aware Training (QAT), and Post-Training Quantization (PTQ). In QAT, the quantization objective is embedded into the LLM training process, enabling them to adapt to low-precision representations and handle precision loss caused by quantization. LLM-QAT (Liu et al., 2023c) proposes a data-free distillation method that leverages generations produced by the pre-trained model, preserving the original output distribution and allows quantizing LLaMa models independent of its training data. PEQA (Kim et al., 2023) operates through a dual-stage process: initially, the parameter matrix of each fully-connected layer undergoes quantization into a matrix of low-bit integers and a scalar vector; subsequently, re-tuning occurs on the scalar vector for each downstream task. QLoRA (Dettmers et al., 2023a) proposes an efficient re-tuning approach that reduces memory usage enough to re-tune a 65B parameter model on a single 48GB GPU while preserving full 16-bit re-tuning task performance by backpropagating gradients through a frozen, 4-bit quantized pre-trained language model into Low Rank Adapters (LoRA). PTQ involves quantizing the parameters of LLMs after the completion of the LLM's training phase. GPTQ (Frantar et al., 2022) proposes a novel layer-wise quantization technique based on approximate second-order information resulting a bitwidth reduction to 3 or 4 bits per weight, with minimal accuracy loss compared to the uncompressed version. AWQ (Lin et al., 2023a) based on the observation that weights are not equally important: protecting only 1% of salient weights can greatly reduce quantization error, employs an activation-aware approach by considering the significance of weight channels corresponding to larger activation magnitudes. SpQR (Dettmers et al., 2023c) works by identifying and isolating outlier weights, which cause particularly-large quantization errors, and storing them in higher precision, while compressing all other weights to 3-4 bits, and achieves relative accuracy losses of less than 1% in perplexity for highly-accurate LLaMA and Falcon LLMs.

A.1.3 LARGE LANGUAGE MODELS AND EVALUATION

Large language models (LLMs) are gaining increasing popularity in both academia and industry playing vital role in both research and daily use. With increasing popularity, several works (Li et al., 2023; Kaddour et al., 2023; Muhlgay et al., 2023; Zhang et al., 2023; Valmeekam et al., 2022; Liu et al., 2023a; Sawada et al., 2023; Qin et al., 2023; Zhuo, 2023; Lee et al., 2023) attempt to go beyond conventional perplexity to evaluate performance of LLMs across factuality, commonsense

reasoning, language understanding, reading comprehension, programming, instruction following abilities, etc. Muhlgay et al. (2023) propose a new metric FACTOR to understand factuality correct information in the LLM generated text. It found that although FACTOR accuracy and LMM perplexity tend to be highly correlated but sometimes induce different orderings between LMMs. They reported that pairs of models can share similar perplexity but differ significantly in terms of FACTOR accuracy. Lee et al. (2023) evaluate the performance and alignment of LLM distribution with humans using two different techniques: Monte Carlo Reconstruction (MCR) and Log Probability Reconstruction (LPR); and found LLMs exhibit limited ability in solving NLI tasks and simultaneously fail to capture human disagreement distribution. Zhang et al. (2023) attempt to investigate promise for automatic summarization with respect to human summary writers and found that LMM summaries are judged to be on par with human written summaries. Valmeekam et al. (2022) propose an extensible assessment framework to test the capabilities of LLMs on reasoning about actions and change, a central aspect of human intelligence and found that GPT-3 and BLOOM have dismal performance on these benchmarks. Despite these efforts to investigate the performance of dense LLMs comprehensively, it is surprising that no such efforts have been yet carried out for a more daunting case of compressed LLMs, which are derived from dense counterparts sharing significantly high similarity with them. Our work is first attempt to address this gap and encourage sparse community researchers to go beyond perplexity to evaluate the true merits and drawbacks of compression methods.

A.2 PROMPT DESIGN AND EXAMPLES FOR DIFFERENT TASK SETTINGS IN LLM-KICK

A.2.1 FACTOID-BASED QA

Prompt Design: Please give answer to this question. ~~QUESTION~~ The answer is

Example: Please give answer to this question. The film '10 things I hate about you' is based on which Shakespeare play? The answer is

Model Response: Please give answer to this question. The film '10 things I hate about you' is based on which Shakespeare play? The answer is *the taming of the shrew*.

A.2.2 MULTIPLE-CHOICE REASONING-BASED QA

Prompt Design: The following are multiple choice questions (with answers) about < SUBJECT NAME >. ~~QUESTION~~ nnA. < OPTION 1 > nnB. < OPTION 2 > nnC. < OPTION 3 > nnD. < OPTION 4 > nn Answer:

Example: The following are multiple choice questions (with answers) about algebra. Find the degree for the given field extension $Q(\sqrt{2}, \sqrt{3}, \sqrt{18})$ over Q . nnA. 0 nnB. 4 nnC. 2 nnD. 6 nn Answer:

Model Response: The following are multiple choice questions (with answers) about algebra. Find the degree for the given field extension $Q(\sqrt{2}, \sqrt{3}, \sqrt{18})$ over Q . nnA. 0 nnB. 4 nnC. 2 nnD. 6 nn Answer: B

A.2.3 IN-CONTEXT RETRIEVAL AUGMENTED QUESTION ANSWERING

① Closed Book Setting: For closed-book setting, we adopted the prompt from Touvron et al. (2023) as follows.

Prompt Design: Answer these questions. Q: <QUESTION> A:
Example: Answer these questions. Q: Who was the man behind The Chipmunks? A:
Model Response: Answer these questions. Q: Who was the man behind The Chipmunks? A: The man behind The Chipmunks was David Sarge, who was the founder of the Alphaville Virtual Real Estate Company.

② Open Book Setting: For open-book setting, we extend the above prompt as follows.

Prompt Design: <EVIDENCE> Answer these questions. Q: <QUESTION> A:
Example: ``Alvin and the Chipmunks (2007) - IMDb IMDb 17 January 2017 4:34 PM, UTC NEWS. A struggling songwriter named Dave Seville finds success ..." Answer these questions. Q: Who was the man behind The Chipmunks? A:
Model Response: ``Alvin and the Chipmunks (2007) - IMDb IMDb 17 January 2017 4:34 PM, UTC NEWS. A struggling songwriter named Dave Seville finds success ..." Answer these questions. Q: Who was the man behind The Chipmunks? A: Dave Seville .

A.2.4 IN-CONTEXT TEXT SUMMARIZATION

Prompt Design: A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Summarize the given story in less than 150 words while preserving high coherence, consistency, fluency, and relevance. <STORY>. ASSISTANT:
Example: A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Summarize the given story in less than 150 words while preserving high coherence, consistency, fluency, and relevance. <LIBYAN> Libyan and U.S. officials say the two governments held face-to-face talks in Tunisia ...have denied previous reports of talks with the government . ASSISTANT:

Model Response: The model response of one-shot magnitude pruned Vicuna-7B ASSISTANT is shown in Figure 8.

Figure 8: Output response of 10% compressed (unstructured one-shot) Vicuna-7b ASSISTANT.

A.2.5 MULTI-TURN CONVERSATION AND INSTRUCTION FOLLOWING

Prompt Design: A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: <QUESTION> ASSISTANT:
Example: A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Can I improve my time management skills? ASSISTANT:

Model Response: The model response of one-shot magnitude pruned Vicuna-7B ASSISTANT is shown in Figure 9.

Figure 9: Output response of 10% compressed (unstructured one-shot) Vicuna-7b ASSISTANT.

A.3 IN-CONTEXT SUMMARIZATION EVALUATION SETTINGS

For evaluating the performance of LLMs to generate high-quality in-context summarization, we focus on consistency, coherence, fluency, and relevance metrics. We prompt GPT-4 which has been recently identified to be highly effective as an automated evaluation framework for benchmark generation and performance assessments, to evaluate these metrics in comparison to the summaries generated by GPT-3.5. Examples of our prompts used for evaluating with GPT-4 Judge are shown in Figure 10. We also provide an example of GPT-4 Judge output in Figure 11.

Figure 10: Example of prompt used to evaluate the compressed LLM ASSISTANT using GPT-4 as Judge on consistency, coherence, fluency, and relevance of generated summaries.

Figure 11: GPT-4 Judge Evaluation of responses generated by GPT-3 (ASSISTANT 1) 10% compressed (unstructured one-shot) Vicuna-7b (ASSISTANT 2).

A.4 INSTRUCTION FOLLOWING ABILITY EVALUATION SETTING

For evaluating the responses generated by compressed LLMs, we closely follow the prompt design settings of MT-Bench (Zheng et al., 2023) using GPT-4 as judge. We prompt GPT-4 to rate the answers generated by compressed LLMs wrt. GPT-3.5 (text-davinci-003) model based on varying metrics (eg. correctness, helpfulness, logic, accuracy), on a scale of [0-10] and provides a

detailed explanation behind the score. Examples of our prompts used during evaluation for questions as well as GPT-4 Judge response are as shown in Figure 12, and 13, respectively.

Figure 12: Examples of prompts used for different categories to evaluate the compressed LLM ASSISTANT wrt. GPT-3.5 ASSISTANT using GPT-4 as a Judge.

Figure 13: GPT4-as-a-Judge evaluation of responses generated by GPT-3 (ASSISTANT) wrt. GPT-4 Judge on 10% compressed (unstructured one-shot) Vicuna-7b (ASSISTANT 2).

A.5 USEFUL LINKS FOR LLM-KICK

Table 1: Dataset and code link used in our work.

Method / Dataset	Download URL
FreebaseQA (Jiang et al., 2019)	https://huggingface.co/datasets/freebase_qa
MMLU Benchmark (Hendrycks et al., 2020)	https://huggingface.co/datasets/freebase_qa
TriviaQA (Joshi et al., 2017)	https://huggingface.co/datasets/trivia_qa
MT-Bench (Zheng et al., 2023)	https://huggingface.co/datasets/HuggingFaceH4/mt_bench_prompts
CNN/DailyMail Summarization (Nallapati et al., 2016)	https://cs.nyu.edu/~kcho/DMQA/
WikiText (Merity et al., 2016)	https://huggingface.co/datasets/HuggingFaceH4/mt_bench_prompts
Wanda (Sun et al., 2023)	https://github.com/locuslab/wanda
SparseGPT (Frantar & Alistarh, 2023)	https://github.com/IST-DASLab/sparsegpt
LLM-Judge (Zheng et al., 2023)	https://github.com/lm-sys/FastChat/tree/main/fastchat/llm_judge
GPTQ (Frantar et al., 2022)	https://github.com/qwopqwop200/GPTQ-for-LLaMa

A.6 COMPARISON WITH AWQ AND LLM-INT8

In this section, we considered evaluating AWQ (Lin et al., 2023b) and LLM.int8() (Dettmers et al., 2022) across our different task settings and we summarize our results on Vicuna-7B as in the following table. We observe that LLM.int8() despite its simplicity and ease-of-use, achieves better results than AWQ (8-bit), and GPTQ (8-bit) across all listed tasks.

Task	GPTQ	AWQ	LLM-int8()
Factoid-QA	60.14%	60.31%	61.02%
MCR-QA (MMLU)	47.10%	47.18%	47.82%
Retrieval Augmented QA	75.55%	75.89%	75.91%
Instruction Following (GPT4-Score)	9.74	9.72	9.81

Table 2: Performance comparison of AWQ and LLM-int8() on LLM-KICK.

A.7 UNDERSTANDING THE IMPACT OF K-SHOT FOR COMPRESSED LLM S

In this section, we aim to investigate how few-shot in-context learning examples can benefit SoTA pruning methods to preserve performance across various sparsity levels. Figure 14 illustrates the performance comparison of Vicuna-7B at varying sparsity ratios when augmented with k-shot in-context examples on MMLU benchmark. It is interesting to observe that k-shot in-context learning examples have marginal impact on dense network performance, while they significantly help in preserving the performance at high sparsity. More- Figure 14: k-shot results of over, we found 2-3 examples are sufficient to retain the performance, and supplementing additional examples doesn't necessarily provide further noticeable benefits.

A.8 SUMMARY OF VARIOUS PRUNING METHODS ON LLM-KICK

Task	Pruning Method	0%	10%	20%	30%	40%	50%
Factoid-QA	Magnitude	65.44	61.74	66.53	60.84	42.06	13.99
	SparseGPT	65.44	63.84	62.44	58.54	55.54	42.86
	Wanda	65.44	63.34	65.23	61.24	58.24	44.66
MCR-QA (MMLU)	Magnitude	0.471	0.466	0.455	0.422	0.339	0.050
	SparseGPT	0.471	0.470	0.460	0.437	0.395	0.308
	Wanda	0.471	0.469	0.460	0.455	0.425	0.386
In-context Retrieval (Long Story: Coherence)	Magnitude	5.883	6.112	5.855	5.567	4.329	1.233
	SparseGPT	5.883	6.033	5.533	6.067	5.567	5.067
	Wanda	5.883	6.0	5.783	5.933	5.267	5.033
Instruction Following (GPT-4 Score)	Magnitude	7.763	7.567	7.621	7.201	6.208	3.308
	SparseGPT	7.763	7.645	7.50	7.188	6.905	6.206
	Wanda	7.763	7.731	7.546	7.202	7.071	6.838

Table 3: Performance comparison of various pruning methods on Vicuna-7B with LLM-KICK.

A.9 ADDITIONAL RESULTS ON LLaMA-2

Figure 15: Compressed LLMs for Factoid-based QA. Performance comparison of compressed LLMs (LLaMa 1 & 2) on Factoid-QA task using FreebaseQA (Jiang et al., 2019). Results presented are for structured (N:M sparsity) and unstructured sparsity.

