

Brevity is the soul of sustainability: Characterizing LLM response lengths

Soham Poddar¹, Paramita Koley², Janardan Misra³, Niloy Ganguly¹, Saptarshi Ghosh¹

¹ Indian Institute of Technology, Kharagpur, India

² Indian Statistical Institute, Kolkata, India

³ Accenture Labs, Bangalore, India

Abstract

A significant portion of the energy consumed by Large Language Models (LLMs) arises from their inference processes; hence developing energy-efficient methods for inference is crucial. While several techniques exist for inference optimization, output compression remains relatively unexplored, with only a few preliminary efforts addressing this aspect. In this work, we first benchmark 12 decoder-only LLMs across 5 datasets, revealing that these models often produce responses that are substantially longer than necessary. We then conduct a comprehensive quality assessment of LLM responses, formally defining *six* information categories present in LLM responses. We show that LLMs often tend to include redundant or additional information besides the minimal answer. To address this issue of long responses by LLMs, we explore several simple and intuitive prompt-engineering strategies. Empirical evaluation shows that appropriate prompts targeting length reduction and controlling information content can achieve significant energy optimization between 25-60% by reducing the response length while preserving the quality of LLM responses.

1 Introduction

Large Language Models (LLMs) have seen rapid adoption in different applications, including virtual assistants for daily life, and in scientific research. However, the deployment of these models require massive computational resources, leading to significant energy consumption and carbon emissions. Although a single inference operation may not consume much energy, such operations are performed enormously many times under deployment. For example, OpenAI’s ChatGPT processes more than 1B queries per day, which consumes an estimated 0.5GWh of energy. In recent times, ninety percent of the total cloud computing demand for Amazon Web Services was for model inference

purposes (Barr, 2019). Hence, LLM inferences contribute to a significant energy impact (Wu et al., 2022; Patterson et al., 2022) and merit further investigation.

Our prior work established that the energy consumption during inference is closely correlated with the length of the response generated by LLMs (Poddar et al., 2025). Notably, output length has a greater impact on inference energy usage than input length because inputs may benefit from parallel processing and caching on GPUs, whereas output generation is predominantly sequential. Additionally, we demonstrated in (Poddar et al., 2025) that the energy expended during inference is largely independent of task complexity or type, depending largely on the output size. Together, these findings highlight the potential for optimizing inference energy consumption by limiting the output length.

A key consideration in this context is how popular LLMs behave - specifically, *whether LLMs tend to generate more text than necessary for certain tasks*. To investigate this, in this work, we benchmark 12 LLMs across 5 tasks/datasets spanning different use cases. We focus on factual question-answering datasets primarily consisting of factual and procedural questions, deliberately excluding opinion-based questions, as determining their minimal response length can be challenging. Our primary finding is that, *for factual questions, LLM-generated responses are significantly longer than the target responses*, regardless of whether the targets are human-written or machine-generated.

It can be argued that longer answers are more useful to the users, since the additional information beyond the core information need – such as explanations, summaries, or conversational enhancements – can enhance user experience. But it is important to recognize that such enhancements come at the cost of increased energy consumption and environmental costs. A useful analogy is air travel via economy vs. business class, where the core need

(timely travel) is distinct from optional enhancements that improve the experience but come at an additional cost. This is the first work to bring forward the trade-off between energy efficiency and user satisfaction, a pivotal contribution that unveils an entirely new perspective on LLM responses.

Next, we categorize the LLM-generated responses by identifying six distinct information categories through a comprehensive manual survey. This categorization reveals that LLM responses often include, along with the minimal answers, explanations, examples, and reasoning information, as well as redundant or irrelevant text. In fact, our study reveals that, on average, the core answer constitutes only about 42% of the total responses. While the extra elements can enhance readability and potentially increase user satisfaction (for instance, users might appreciate supplementary context or a polite add-on), they also incur higher inference energy costs. Consequently, we argue that LLMs should be aligned to generate minimal answers by default, incorporating only a limited amount of additional content and excluding extraneous or conversational tokens.

Finally, we explore several prompt-based strategies to generate more concise responses from LLMs. These include simple directives (such as “Answer briefly”), providing in-context examples to guide the LLM toward optimal response length, specifying a desired output length, and reducing non-essential content such as conversational and irrelevant information. We evaluate these prompt engineering strategies in terms of how they affect the length, energy, and information content of the LLM responses. We also check how they affect the quality of LLM responses in terms of ROUGE-L F1 scores with the target (gold standard) answers. Empirical evaluation shows that some of our prompts achieve significant energy optimization by reducing the response lengths, while preserving (or sometimes even improving) the quality of LLM responses. For instance, one of our best strategies is simply where we direct the LLM to provide only the minimal answer; this strategy achieves up to 88% reduction in LLM responses, which leads up to 58% reduction in energy consumption, while actually improving the response quality.

2 Related Works

The optimization of energy consumption in large language models (LLMs) has been extensively

Dataset	Description
DOLLY (Conover et al., 2023)	instruction following records with human-curated answers in several QA and classification categories
GOOQA (Khashabi et al., 2021)	open-domain dataset with factual questions from Google searches and answers mined from Google’s responses.
MS-MARCO (Nguyen et al., 2016)	Open-domain QA dataset, with questions from Bing queries and crowd-sourced answers.
NARRATIVEQA (Kočíský et al., 2018)	QA dataset where Wikipedia summaries of novels were given to humans to frame question-answer pairs.
TWEETQA (Xiong et al., 2019)	factual QA dataset, curated by humans based on context from tweets used by journalists.

Table 1: List of datasets used in our experiments.

studied in the literature. Notable white-box approaches include pruning (Kurtić et al., 2024), quantization (Xiao et al., 2023; Frantar et al., 2022), early exit strategies (Schuster et al., 2022), and embedding recycling (Saad-Falcon et al., 2022). Hardware or OS-level optimizations focus on optimal data transfer and device utilization, for example, caching (Ramírez et al., 2023), scheduling (Yu et al., 2022), etc. Black-box approaches include cascading (Chen et al., 2023), speculative decoding (Leviathan et al., 2023), input compression (Jiang et al., 2023a,b), among others. A detailed literature survey has been given in Appendix A.

However, the area of output compression remains relatively underexplored, with only a few preliminary efforts in the literature, namely prompt augmentation with generation directives (Li et al., 2024). In this work, our primary focus is to comprehensively explore the underexplored avenue of output compression for energy-efficient inference.

3 Do LLMs generate too much text?

In this section, we investigate the question – *Are LLMs generating too much text, more than necessary for certain types of tasks?* We primarily focus on objective factual queries where concise and precise responses are often sufficient and compressing the response can improve energy efficiency trading-off with user satisfaction. Note that output compression may not be suitable for long-form tasks.

3.1 Experimental Setup

Datasets: We choose 5 datasets, namely DOLLY, GOOQA, MS-MARCO, NARRATIVEQA, and

Dataset	tgt len	GPT 3.5t	GPT 4om	gemma -2-2b	gemma -2-9b	Llama -2-13b	Mistral -7B	Llama -3.2-3B	Llama -3.1-8B	Phi-3 -mini	Phi-3 -small	falcon -11B	vicuna -13b
DOLLY	68.3	78.2	123.3	191.1	152.3	176.3	150.4	229.0	539.8	1445.1	477.5	168.3	84.2
GOOAQ	27.4	63.1	142.3	212.5	162.4	219.3	161.8	311.9	535.2	1333.9	418.5	177.5	89.0
MS-MARCO	22.8	40.3	127.8	124.4	138.8	169.3	139.4	280.6	683.3	1620.0	1333.2	63.9	94.8
NARRATIVEQA	7.1	23.4	44.7	35.0	25.1	36.9	37.1	246.4	671.8	1215.6	1462.6	30.9	38.0
TWEETQA	4.0	12.7	22.8	89.6	44.8	71.9	111.6	182.6	395.4	1671.7	1108.2	80.8	49.1

Table 2: Average target response length (second column) vs response lengths for different LLMs on the 5 datasets. In almost all cases, LLM responses are substantially longer than the target response lengths, across all datasets.

TWEETQA for LLM evaluations, for reasons described below. The datasets are described in Table 1 (refer Appendix F for examples).

We selected these 5 datasets for the following reasons: **(1) Diverse tasks:** The datasets cover a diverse range of tasks, including logical reasoning (DOLLY, GOOAQ), instruction following (DOLLY), knowledge-dependent closed-QA (DOLLY, GOOAQ), knowledge-dependent short contextual QA (TWEETQA) and long contextual open-QA (NARRATIVEQA, MS-MARCO). **(2) Diverse domains:** These datasets also cover diverse domains, ranging from tweet-based queries in TWEETQA, regular factual and procedural queries from different categories in DOLLY and GOOAQ, and curated queries from fiction in NARRATIVEQA. GOOAQ and MS-MARCO also present several factoid-based questions of practical interest, sampled from google searches and bing searches. **(3) Diverse answer types:** Target answers in our selected datasets include human-generated, human-annotated, and machine-generated responses. To ascertain whether LLMs are generating more text than is required, we need to compare LLM responses with minimal and sufficient answer lengths. Human-written answers serve as the proxy of answers of reasonable lengths in DOLLY, MS-MARCO, NARRATIVEQA, and TWEETQA, while GOOAQ contains answers to short fact-based and longer procedural questions automatically scraped from Google search answers. This results in diverse answer types, and a rich space of answers.

From each of the datasets, we select 1,024 samples from the validation/test splits for evaluating the LLMs. 5K samples are also extracted from the train splits to be used in Section 5. To test LLMs generalizability, we use this pool of 25K training samples for different strategies, rather than using individual datasets for training and testing the models separately.

Large Language Models: In this work, we study the generation of 10 open-source LLMs from different families – Meta Llama-2/3, Microsoft Phi-3,

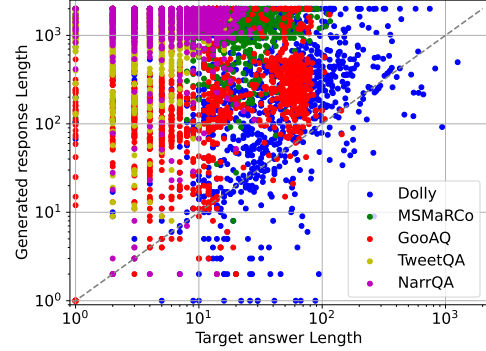


Figure 1: Target answer length vs LLM generated response length by Phi-3-small on different datasets

Google Gemma-2, Mistral, Falcon, Vicuna. We also test the closed-source OpenAI GPT-3.5-turbo and GPT-4o-mini LLMs. These models are listed in Table 2 and detailed in App:B.

For reproducibility, models were loaded in standard half-precision and made to generate greedily up to a hard limit of 2K tokens. Note that this limit does *not* affect the response generated, and is rather just a cut-off if the model keeps on generating tokens. The repetition penalty parameter was set to 1.2 (Keskar et al., 2019) to restrict models from generating repetitive text. Details of the hardware/software setup are reported in App: C.

Measuring the length of LLM responses: We measure the length of the response by counting the number of tokens present in the text, where the tokenizer of the corresponding model has been used to tokenize the model output as well as the gold answer. For the closed-source GPT-3.5/4 models from OpenAI, we use the SpaCy library to tokenize the responses and gold answers. We refer to the length of gold-standard answers from the datasets (mostly human-written responses) as “**target length**” and the length of LLM-generated response as “**generated length**”.

3.2 Quantitative Analysis

Table 2 presents the average target response length and generated response length in terms of tokens

	DS-R1-Llama-8B	DS-R1-Qwen-14B	Gemma-2-9B	Target Answer
DOLLY	586.2 (64.8%)	669.1 (73.8%)	152.2	68.3
GooAQ	552.6 (52.5%)	759.5 (73.1%)	162.4	27.4

Table 3: Average response lengths of Deepseek reasoning models on two datasets (% of thinking/reasoning tokens inside parenthesis).

for all dataset-LLM pairs. In almost all cases, LLMs generate responses substantially longer than target responses. Even for extremely brief target responses (avg < 10 tokens) in NARRATIVEQA and TWEETQA, LLM responses are relatively long (20–1.5K tokens).

We can classify the models based on the ratio of their average response length to the average target answer length into three categories – (1) **Moderate (mostly within 1-3 times)**: models from GPT-3 family generate responses within 1-3 times of the target answer length, (2) **Long (mostly between 3-10 times)**: models from GPT-4, vicuna, gemma-2, llama-2, and mistral generate moderately long responses, and (3) **Very Long (often more than 10 times)**: responses of llama-3.1, and phi-3 family are generally very long.

Figure 1 plots the distribution of generated response length with target answer length across all datasets for phi-3-small; the presence of datapoints mostly above the diagonal signifies that LLMs generate overly long responses.

Reasoning LLMs are gaining popularity as they perform well in complex scientific tasks. They do this by generating long reasoning traces for optimal final answers, requiring higher time and energy for inference. To further analyze the behavior of reasoning LLMs, Table 3 reports average response length for DeepSeek-R1-Distill-Llama-8B and DeepSeek-R1-Distill-Qwen-14B on the Dolly and GooAQ datasets; we observe that these models generate very long outputs on average, with reasoning tokens constituting the majority of them.

Note that our selected datasets include examples with both human-written responses and standard objective responses to serve as a proxy for minimal/sufficient response length. Thus, we can conclude that *most LLMs usually generate responses substantially longer than desired for the use cases under consideration in this study.*

3.3 Categorizing information in responses

To gain further insight on overly long responses of LLMs, we categorize the information present in the

responses. Low et al. (1973) suggests a taxonomy of irrelevant, relevant, and redundant information for general concept transfer. Following this line of work and our inspection of 300 responses by different LLMs, we develop the following classification scheme for the tokens contained in LLM responses.

1. Minimal answer (MINANS): These tokens provide a concise and to-the-point answer to the query/prompt. Note that the minimal answer can be correct or wrong; we are *not* differentiating between correct and wrong answers in this analysis.

2. Additional helpful information (ADDINFO): These include information that adds valuable context and aims to expand the knowledge of the user. Example: User: “What is capital of France?” LLM: “Paris. *It is a center of art and culture*”.

3. Reasoning/Explanations (EXPLAIN): Here the LLM states its thought process, logical steps that it followed to answer the query. Example: User: “What is 2 + 2?” LLM: “*First, you take 2 and add another 2, which equals 4.* Final answer: 4”.

4. Conversational Enhancements (CONVERS): Tokens that are included to maintain politeness, enhance the conversational tone, or ask users if they require further information on some aspects. Examples include “*Let me know if you need anything else*”, “*I hope that helps!*”.

5. Redundant Information (REDINFO): In some longer responses, LLMs provide a summary of the previous content at the end. In some other cases, the responses contain rephrased versions of information already given (in the same response). Example: “Netherlands is the official name of the country as a whole, including Holland and other regions like Friesland, Gelderland. *So, while Holland is part of the Netherlands, Netherlands is the entire country.*”

6. Irrelevant Information (IRREL): At times, LLMs generate tokens unrelated to the user’s query, e.g., when the model hallucinates, or produces repeating patterns of tokens. Example: User: “Who discovered America?” LLM: “Columbus! *By the way, I recently read about American recipes.*”

3.4 Analysis of LLM responses

Two of the authors annotated an additional 500 LLM responses, sampling uniformly from the 5 datasets and 10 LLMs, marking spans from the responses into the categories defined above. The first 50 were marked by both authors and achieved pairwise F-measure (Deleger et al., 2012) of 0.764,

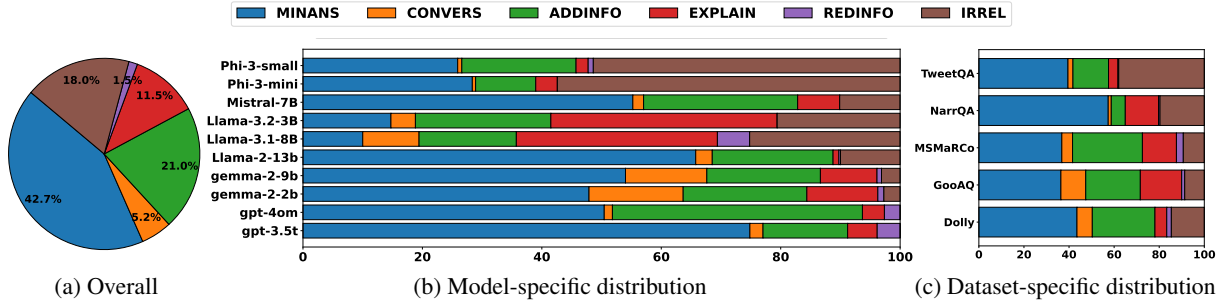


Figure 2: Distribution of information categories in LLM responses.

macro-averaged across different categories, which shows reasonable agreement. The major disagreements were in the REDINFO class, particularly in cases where parts of minimal answer, additional and redundant information were intertwined. Some example annotations are shown in Appendix E. We are releasing this annotated dataset to promote further research on LLM responses.¹

The distribution of the categories across all responses is given in Figure 2a, whereas Figure 2b and Figure 2c present model-specific and dataset-specific category distributions. Figure 2a shows that though the answers primarily contain minimal answer (MINANS) and additional information (ADDINFO), *irrelevant information* ($\approx 18\%$) and *conversation enhancers* ($\approx 5.2\%$) form a significant portion of the responses. Reducing these can immediately lead to much lower consumption of energy. Further, restricting additional information ($\approx 21\%$) and explanations ($\approx 11.5\%$) can further compress the answers without degrading the quality, where concise answers are sufficient.

Figure 2b shows the distribution of information categories in the responses by different LLMs. We see that, for some LLMs such as GPT-3,4, Mistral-7B, gemma-2, Llama-3.2-3B, most of the generated tokens belong to ‘minimal answer’ with a moderate portion of additional information. On the other hand, llama-3 generates a lot of explanation (CoT reasoning) and irrelevant information, and phi-3 generates a high amount of irrelevant text, with minimal response being rather small parts of the generations. These model-specific trends in responses seem to be consistent across a model family, suggesting that the training setup of an LLM is the primary factor in determining the distribution of content categories within a response in comparison to the relative size of the model. Interestingly, *older LLMs (Llama-2.1, GPT-3.5) give more min-*

imal answers, whereas newer LLMs (Llama-3.1, Llama-3.2, GPT-4om), with updated pre-training data and strategies, generate more of additional information and explanations.

For dataset-specific trends (Figure 2c), queries in NARRATIVEQA and TWEETQA are objective and factual, based on a given context, resulting in responses being mostly minimal answers with low amount of additional information. In MS-MARCO, DOLLY, and GOQA, queries are open-domain questions, containing both factual and reasoning queries, resulting in decreasing fraction of minimal responses and increasing additional info and conversational enhancements.

Based on our observations above, we now explore several strategies to reduce the LLM response lengths, while still retaining the response quality.

4 Some common strategies for shortening LLM responses

We now focus on some common strategies for shortening LLM responses to achieve energy-efficient inference. We explore simple, intuitive white-box approaches, including response truncation and fine-tuning of LLMs, as well as black-box strategies based on prompt engineering.

Truncating responses refers to stopping the generation early, e.g., just after the minimal answer is obtained. However, we notice that in about 19% of the responses studied in Section 3.4, the minimal answer is *not* at the beginning of the LLM response. In most of these cases, the LLMs output conversational enhancers first to set up the context or for readability, followed by the minimal answer, thus limiting the scope of strategies for truncating responses.

Fine-tuning LLMs: Next we explore whether simple fine-tuning can help in controlling the response length. Toward that, we performed some basic fine-

¹<https://github.com/sohampoddar26/LLM-brevity>

tuning experiments. We fine-tuned Phi-3-mini (one of the LLMs that generate very long responses, as stated earlier) with a pool of human-written question-answers from the DOLLY and GOOAQ datasets, using LoRA with $r = 16$, for 100 epochs. The models were trained on the target answer tokens with the hypothesis that it would enable the models to generate responses similar to humans on these particular datasets. We compare the average ratio of generated and target response lengths for the DEFAULT model and the fine-tuned model. We observe that *simple fine-tuning is leading to 1.24 times higher ratio for DOLLY and 1.26 times higher ratio for GOOAQ, in comparison with the DEFAULT model*. This is an unexpected result, where fine-tuning is (slightly) increasing the response length instead of shortening the length. This behavior could stem from extensive pretraining on large-scale data, making small-scale fine-tuning insufficient to alter its learned patterns. Additionally, a model’s intrinsic stopping mechanism may not be functioning as intended after the fine-tuning, causing it to overestimate the length for the responses.

Therefore, it is prudent to explore prompt engineering strategies leveraging LLM’s zero-shot capabilities, drawing on our observations in Section 3. We explore this direction in the next section.

5 Prompt engineering strategies for controlling LLM responses

Here we investigate various prompt engineering strategies aimed at generating concise LLM responses, thereby reducing inference energy.

5.1 Prompt engineering strategies

We first describe the different strategies we used to control response length. The exact prompts used are given in Table 4.

BRIEF: Following (Li et al., 2024), the default query is modified by appending “Answer briefly” at the end of the query.

BM25-INCONTEXT: In this method, we provide a set of in-context examples to help LLM in deciding the ideal response length of the given query. Specifically, we provide a sequence of 10 sample queries and their respective gold-answers as in-context examples, followed by the given query. These samples were chosen from the 25K train set based on the similarity with the given query, as predicted by the BM25+ retrieval model (Trotman et al., 2014).

LIMIT-LEN: In this approach, we specify an estimated length of the response and instruct the LLM to restrict the generation length within that limit. Specifically, we add “Answer within \mathcal{X} words” at the end of the prompt, where \mathcal{X} is the limit of generation that we obtain through three ways: (1) **BM25-LENGTH:** Top-k similar samples are selected from the pool of samples using BM25+ model, and \mathcal{X} is set to the median of their target lengths. (2) **GOLDRESLEN:** To serve as an oracle baseline, \mathcal{X} is set to the actual gold-answer length corresponding to the query. (3) **PREDRESLEN:** Here we employ a supervised framework for learning the ideal length of a given query. Specifically, we train an encoder-only transformer-based regression model DEBERTA-V3-LARGE (He et al., 2021) with scalar output to predict \mathcal{X} of the given query. The model is trained with queries and the length of their target (gold standard) responses from the training set.

LIMIT-CAT: Based on our observations regarding the distribution of various categories in LLM responses (see Figure 2a), we hypothesize that responses can be made shorter and energy-efficient while preserving response quality, by reducing the presence of CONVERS, REDINFO, and IRREL in the responses. Accordingly, we design LIMIT-CAT, a set of prompt directives to control the presence of various categories in the response. We tried several combinations of prompts to restrict the category distribution in the responses. Out of these we report results for the following two strategies which performed the best: (1) **MINANS:** Model is asked to provide the minimal answer only. (2) **MADDNORED:** Model is asked to provide the minimal answer, along with some additional information if applicable. It is also explicitly asked *not* to produce redundant text or conversational enhancements (exact prompts in Table 4).

5.2 Results

We empirically validate the efficacy of the prompt engineering strategies against the default prompt, by investigating these questions: (1) Can the prompt engineering strategies achieve a reduction in LLM response length? (2) Does the length reduction degrade the response quality? (3) Are these approaches able to achieve notable reduction in terms of energy? (4) How do these methods influence the category distribution in the LLM responses?

Effects on response length: Fig. 3 reports the ratio of generated response length to the target length

query	<Context if present> \n
prompt	### Question: <Query text> \n <query> <response directive> ### Response:
Response directives	
BRIEF	Answer briefly.
MINANS	Only provide the minimal answer.
MADDNoRED	Provide the minimal answer. Provide some additional information where applicable, but do not produce redundant text or politeness and conversational enhancements.
LIMIT-LEN	Answer within X words.
BM25-INCONTEXT	<10 examples> <default prompt>

Table 4: Prompt directives to shorten LLM responses.

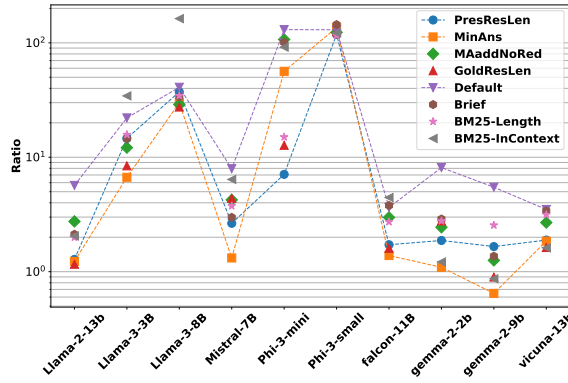


Figure 3: Ratio of generated response length to target length for different LLMs using each prompting strategy, averaged over all datasets.

for each of the above-mentioned strategies, averaged across all datasets. We observe that prompt directives are generally able to reduce the response length for almost all models relative to the DEFAULT setup.

MINANS achieves the highest length reduction across all models with about 60% decrease, followed by PREDRESLEN at 53%. The supervised length estimator PREDRESLEN performs at par with the oracle (GOLDRESLEN) with superior length reduction for some models, namely mistral (66%), llama-2 (77%), and gemma-2-9b (69%), showing its efficacy in accurate estimation of ideal target length. BM25-LENGTH and BRIEF perform comparably while achieving lower compression of about 38%. on average Length reduction is minimal for BM25-INCONTEXT, as some models fail to understand the desired output length from examples. MADDNoRED produces relatively less compression too as we specifically ask the models to generate additional information. However, some LLMs, like phi-3-small or llama-3-8B fail to

respond to such prompts appropriately.

Effect on response quality: While response quality is abstract and difficult to quantify, we wanted to check if the user’s *basic information need* is being fulfilled in the LLM responses. To this end, we compute the Rouge-L F1 scores of the LLM responses with respect to the target answers. While this is no substitute for a human evaluation, this will give us some insights into the quality of responses using different strategies.

Fig. 4 shows the Rouge-L-F1 score of each strategy for different models, averaged across all datasets. We observe that employing the above-mentioned prompting strategies generally improves the F1-scores of the responses upon the default setup in most cases, since the models are instructed to generate more information-dense answers, that better match the target answers. Specifically, we find that *shorter responses lead to substantial improvements in ROUGE precision, with marginal drop in ROUGE recall, leading to increased F1 scores* (refer to Appendix G for the precision, recall and F1-scores for all the methods).

MINANS emerges as the most effective technique with the highest gain in response quality, followed by GOLDRESLEN, PREDRESLEN and BRIEF. These results demonstrate that both simple directives for shortening responses (MINANS, BRIEF) as well as providing accurate length estimates (PREDRESLEN) can achieve significant length reduction while preserving the response quality. MADDNoRED and BM25-LENGTH shows a smaller performance gain than the default, which can be attributed to longer response with lesser precision. BM25-INCONTEXT generates even larger responses leading to drop in performance.

Effects on Energy Consumption: To understand the effect of prompt engineering on inference energy, we report inference energy for different models for different strategies, averaged across all datasets in Figure 5. The inference energy consumed per sample is measured in milli-watt-hour (mWh) using the popular CodeCarbon library (Courty et al., 2024) (see Appendix C).

Most methods demonstrate improved energy efficiency compared to the default approach, except for BM25-INCONTEXT, which fails to reduce energy usage across all models, attributing to both significantly longer input and its ineffectiveness in achieving length reduction. Among the evalu-

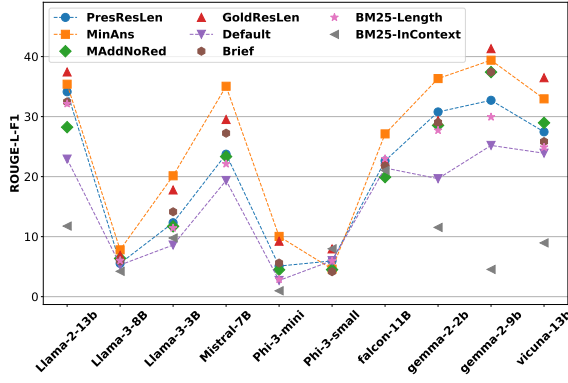


Figure 4: ROUGE-L-F1 scores for different models using different strategies, averaged across all datasets.

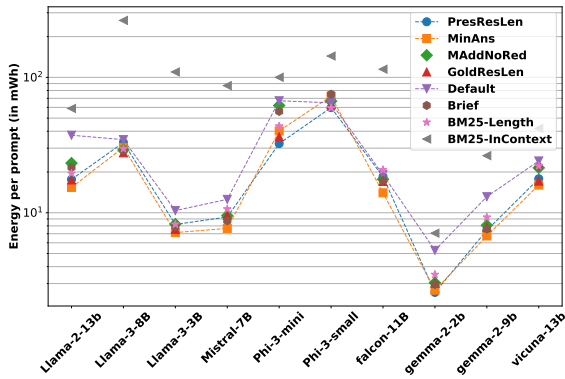


Figure 5: Inference Energy for different models using different strategies, averaged over all datasets

ated techniques, MINANS emerges as most energy-efficient method with 28%, attributing to highest length reduction. MINANS is closely followed by PREDRESLEN (comparable with oracle GOLDRSLEN) with about 26% reduction, highlighting the importance of advanced energy-efficient supervised approaches for estimating optimal response lengths. Methods like BM25-LENGTH achieves lower energy efficiency, due to its compression inefficiency and computational overhead.

Effects on information categories in responses:

Finally, to understand how these prompts affect the information category distribution in the LLM responses, we perform the following experiment. We select a subset of the category-annotated queries for which the responses were manually annotated to identify information categories (Sec. 3.4). For each of the queries, we employ the following prompts – BRIEF among the general prompts, PREDRESLEN among the length-specific prompts, and MINANS, and MADDNORED among the category-specific prompts. The responses of the

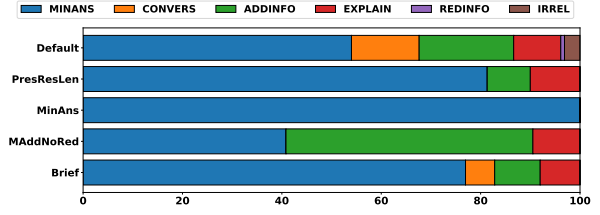


Figure 6: Distribution of information categories in responses of Gemma-2-9B after incorporating 4 prompt strategies, compared to DEFAULT.

two models on these four prompts are further annotated by the authors into the same information categories.

Figure 6 presents a representative result for the gemma-2-9B LLM, showing the category distribution for the generated responses, averaged across all queries for each model. Gemma-2 is seen to give almost entirely minimal answers for MINANS, and a good balance of minimal and additional information with the MADDNORED strategy. BRIEF and PREDRESLEN give relatively short answers containing mostly minimal answers, with BRIEF having some amounts of all desirable classes.

Table 5 shows a few examples of responses by Gemma-2-9B before and after introducing the prompt strategies. It can be seen that MINANS leads the Gemma-2-9B model to generate just the minimal answer and correctly remove the explanations. We also see examples for BRIEF and MADDNORED which shorten the responses hugely, while still retaining some useful extra information. However, we also observed cases where the prompts are making unnecessary cuts. For example, as stated earlier, we experimented with different LIMIT-CAT strategies; we observed that, if we ask the Gemma-2-9B model to remove the conversation elements, there are cases where it also removes explanations and additional info as well, leading to a more minimal answer that may affect user satisfaction. This suggests that further research is needed to systematically identify prompts that lead to the desired behaviour (with respect to both response length and quality) in various LLMs.

6 Concluding Discussion

In this work, we comprehensively explore output compression for energy-efficient inference. Our specific contributions are as follows. (1) We are the first to systematically establish that large language models (LLMs) often generate very long responses which lead to high energy consumptions; We also

Strategy	Gemma-2-9B Response
Query: 29 cm is equal to how many m?	
DEFAULT	Here's how to convert centimeters (cm) to meters (m): **Remember the conversion:** 1 meter = 100 centimeters. **Divide by 100:** To find out how many meters are in 29 cm, divide 29 by 100. 29 cm / 100 cm/meter = 0.29 meters Therefore, 29 cm is equal to 0.29 meters
MINANS	0.29
BRIEF	0.29 m
MADDNoRED	0.29 **Explanation:** There are 100 centimeters (cm) in one meter (m). To convert from centimeters to meters, divide by 100."
Query: <INFORMATION from different sources> ### Question: what is a necromonger in riddick?	
DEFAULT	In the Chronicles of Riddick film franchise, Necromongers are a fanatical warrior-religious order ... <LONG ADDITIONAL INFO> Let me know if you have any other questions about the Chronicles of Riddick!
MINANS	A fanatical religious warrior race seeking universal conversion or death.
BRIEF	A Necromonger is a member of a fanatical religious empire in the *Riddick* franchise. They believe life is an enemy of the universe and seek to conquer and convert all others through violence.
MADDNoRED	A Necromonger is a member of a fanatical religious empire in the *Chronicles of Riddick* franchise. They believe life is an affront to the universe and seek to eradicate it through conversion or death. Their society is highly structured with strict rules and rituals.
Query: How to connect iphone to LG TV to play music?	
DEFAULT	Here's how to connect your iPhone to an LG TV for screen mirroring, along with some important things to keep in mind: <LONG ADDITIONAL INFO> Let me know if you have any other questions!
MINANS	Use AirPlay 2 on your iPhone and ensure your LG TV supports it. Then, select ""Screen Mirroring"" in Control Center and choose your LG TV from the list.
BRIEF	1. **Enable Miracast on your LG TV.** (Settings > General > Screen Share) 2. **Turn on AirPlay Mirroring on your iPhone.** (Control Center > Screen Mirroring). 3. **Select your LG TV from the list of available devices.** Let me know if you need detailed instructions!
MADDNoRED	1. **Enable Miracast on your LG TV.** (Settings > General > Screen Share) 2. **Turn on AirPlay Mirroring on your iPhone.** (Control Center). Your iPhone should now appear as a connection option on your LG TV. Select it to mirror your screen.

Table 5: Sample queries and responses of Gemma-2-9B after introducing prompt strategies.

categorize LLMs based on their generation lengths. (2) We present the first comprehensive information categorization of the LLM responses, identifying 6 information classes; (3) Based on our findings, we develop various length-based and category-based prompts and evaluate them extensively in terms of how they affect the length, energy and quality of LLM responses. We observe that appropriate prompts can achieve significant energy reduction by reducing response length, alongside improving response quality. Overall, we think the present work takes the first important step toward exploring the utility of output compression in reducing the energy consumption of LLMs.

Our work opens up several future directions for further work (more details in Appendix H). (1) **What constitutes an ideal response length** under different scenarios/applications should be studied more, and various factors influencing that decision should be identified. Measuring user satisfaction would likely require different approaches, including human surveys and automated scoring metrics and need a study on its own. (2) **User personalisation:** Another possible direction is to consider user-preferences while optimizing response length, since different users can prefer responses of different lengths for the same question/task. Such personalization is out of scope for the current work. However, it presents a promising line of future research. (3) **Automatic information categorization:** We attempted to automate the categorization of LLM responses. To this end, we employed OpenAI GPT-4o with a detailed prompt describing all the information categories. We achieved an average

F1 score of 0.391 over all tokens. Hence, efficient methods for distinguishing categories in LLM responses should be investigated. (4) **Pre-training:** Model families seem to generate responses with similar category distributions, indicating the pre-training of the LLMs should be revisited to have better control over what information categories should be included in the response. (5) **Model alignment:** Efforts should be made to align LLMs towards generating responses of desirable lengths. (6) **Efficient decoding algorithms** may be employed to control category information within the generated responses in the real-time, to achieve a better compromise between readability and energy-optimality.

Acknowledgments

The authors thank the anonymous reviewers whose suggestions helped to improve the work. The research is partially supported by a research grant from Accenture Corporation. Paramita Koley is supported by SERB NPDF Fellowship.

Limitations

While this study offers a comprehensive analysis and valuable insights, some limitations should be noted: (1) Limited applicability to long-answer tasks: The potential for energy optimization through output compression is limited in tasks requiring detailed and lengthy responses, such as creative writing or critical decision-making tasks with explanations. Our approaches are primarily effective for objective, factual queries where concise

and brief responses are both sufficient and desirable. (2) Subjectivity of Ideal Response Length: The notion of an ideal response length is subjective and varies significantly among individuals. As a result, a one-size-fits-all approach may not effectively meet the diverse needs of users. Here we argue that the default behaviour of LLMs can be tuned so that they generate shorter responses, so that energy consumption is minimized. A user can ask for more details or longer responses if she desires. (3) Use of standard datasets: Our experiments have been performed on publicly available rich data for testing LLMs. However, the LLMs we used might have been pre-trained on parts of these standard datasets, causing them to behave in a certain way. This may lead to slight variations in results for real-world use cases, but the overall findings should be generalizable enough.

References

- Jeff Barr. 2019. Amazon ec2 update—inf1 instances with aws inferentia chips for high performance cost-effective inferencing.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Lucía Bouza, Aurélie Bugeau, and Loïc Lannelongue. 2023. How to estimate carbon footprint when training deep learning models? a guide and review. *Environmental Research Communications*, 5(11):115014.
- Lingjiao Chen, Matei Zaharia, and James Zou. 2023. Frugalgpt: How to use large language models while reducing cost and improving performance. *arXiv preprint arXiv:2305.05176*.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#).
- Benoit Courty, Victor Schmidt, Sasha Luccioni, Goyal-Kamal, MarionCoutarel, Boris Feld, Jérémy Lecourt, LiamConnell, Amine Saboni, Inimaz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Hugues de Lavoreille, Niko Laskaris, Edoardo Abati, Douglas Blank, Ziyao Wang, Armin Catovic, Marc Alencon, Michał Stęchły, Christian Bauer, Lucas-Otavio, JPW, and MinervaBooks. 2024. [mlco2/codecarbon: v2.4.1](#).
- Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and unfairness in information retrieval systems: New challenges in the llm era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6437–6447.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359.
- Louise Deleger, Qi Li, Todd Lingren, Megan Kaiser, Katalin Molnar, Laura Stoutenborough, Michal Kouril, Keith Marsolo, Imre Solti, et al. 2012. Building gold standard corpora for medical natural language processing tasks. In *AMIA Annual Symposium Proceedings*, volume 2012, page 144. American Medical Informatics Association.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Mathilde Jay, Vladimir Ostapenco, Laurent Lefèvre, Denis Trystram, Anne-Cécile Orgerie, and Benjamin Fichel. 2023. An experimental comparison of software-based power meters: focus on cpu and gpu. In *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, pages 106–118. IEEE.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. 2023. Towards mitigating llm hallucination via self reflection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1827–1843.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023a. Llmllingua: Compressing prompts for accelerated inference of large language models. *arXiv preprint arXiv:2310.05736*.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023b. Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression. *arXiv preprint arXiv:2310.06839*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Daniel Khashabi, Amos Ng, Tushar Khot, Ashish Sabharwal, Hannaneh Hajishirzi, and Chris Callison-Burch. 2021. Gooaq: Open question answering with diverse answer types. *arXiv preprint*.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.

- Eldar Kurtić, Elias Frantar, and Dan Alistarh. 2024. Zi-plm: Inference-aware structured pruning of language models. *Advances in Neural Information Processing Systems*, 36.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- Baolin Li, Yankai Jiang, Vijay Gadepally, and Devesh Tiwari. 2024. Toward sustainable genai using generation directives for carbon-friendly large language model inference. *arXiv preprint arXiv:2403.12900*.
- Lorraine A Low, Jacqueline Fortier, William Mickalide, David Page, and Diane Troumpalos. 1973. Role of relevant, irrelevant, and redundant information in simple concept transfer. *Bulletin of the Psychonomic Society*, 1:267–269.
- Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power hungry processing: Watts driving the cost of ai deployment? In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 85–99.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A human generated machine reading comprehension dataset](#). *CoRR*, abs/1611.09268.
- David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R So, Maud Texier, and Jeff Dean. 2022. The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7):18–28.
- Soham Poddar, Paramita Koley, Janardan Misra, Niloy Ganguly, and Saptarshi Ghosh. 2025. Towards sustainable nlp: Insights from benchmarking inference energy in large language models. *arXiv preprint arXiv:2502.05610*.
- Guillem Ramírez, Matthias Lindemann, Alexandra Birch, and Ivan Titov. 2023. Cache & distil: Optimising api calls to large language models. *arXiv preprint arXiv:2310.13561*.
- Jon Saad-Falcon, Amanpreet Singh, Luca Soldaini, Mike D’Arcy, Arman Cohan, and Doug Downey. 2022. Embedding recycling for language models. *arXiv preprint arXiv:2207.04993*.
- Tal Schuster, Adam Fisch, Jai Gupta, Mostafa Dehghani, Dara Bahri, Vinh Tran, Yi Tay, and Donald Metzler. 2022. Confident adaptive language modeling. *Advances in Neural Information Processing Systems*, 35:17456–17472.
- Andrew Trotman, Antti Puurula, and Blake Burgess. 2014. Improvements to bm25 and language models examined. In *Proceedings of the 19th Australasian Document Computing Symposium*, pages 58–65.
- Zhichao Wang, Bin Bi, Shiva Kumar Penttala, Kiran Ramnath, Sougata Chaudhuri, Shubham Mehrotra, Xiang-Bo Mao, Sitaram Asur, et al. 2024. A comprehensive survey of llm alignment techniques: Rlhf, rlaf, ppo, dpo and more. *arXiv preprint arXiv:2407.16216*.
- Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. 2022. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems*, 4:795–813.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pages 38087–38099. PMLR.
- Wenhan Xiong, Jiawei Wu, Hong Wang, Vivek Kulkarni, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. [TWEETQA: A social media focused question answering dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5020–5031, Florence, Italy. Association for Computational Linguistics.
- Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. 2022. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 521–538.

Appendix

A Related Works

The optimization of energy consumption in large language models (LLMs) has been extensively studied in the literature.

White-box approaches: These techniques focus on architectural modifications to LLMs and represent a well-established line of research. Notable methods include pruning (Kurtić et al., 2024), quantization (Xiao et al., 2023; Frantar et al., 2022), early exit strategies (Schuster et al., 2022), and embedding recycling (Saad-Falcon et al., 2022), among others.

Hardware- or OS-level optimizations focus on improving energy efficiency by leveraging hardware innovations or operating system enhancements. Yu et al. (Yu et al., 2022) propose an optimal scheduling strategy by making it iteration-level from user-request-level. Flash-attention (Dao et al., 2022) effectively reduces latency by optimizing the number of memory read and write operations required in calculating model states. Ramirez et al. (Ramírez et al., 2023) employs a student model along with the larger LLM, which gradually gains proficiency in independently handling increasing user requests, mimicking the popular caching technique in OS.

Black-box approaches avoid altering the underlying model, unlike the white-box models. Instead, they employ multi-level frameworks that incorporate smaller LLMs for specific tasks. Chen et al. (Chen et al., 2023) highlight the efficacy of cascading and multiplexing in reducing inference energy and cost. Leviathan et al. (Leviathan et al., 2023) introduce speculative decoding that accelerates the decoding by running Larger LLMs on the output of approximately smaller models.

Input compression techniques apply for large-input tasks like legal question answering, multi-document summarization, multi-document retrieval, etc, where the input is excessively high. Jiang et al. (Jiang et al., 2023a) propose LLMLingua, a coarse-to-fine prompt compression method that involves a budget controller to maintain semantic integrity under high compression ratios. LongLLMLingua (Jiang et al., 2023b) further refine the method, resulting in achieving higher compression.

As already shown in literature (Luccioni et al., 2024; Poddar et al., 2025), inference energy is also

closely tied to output length, presenting a scope for reducing energy consumption by **Output compression**. However, this area remains relatively underexplored, with only a few preliminary efforts in the literature. For instance, prompt augmentation with generation directives has been investigated as an initial step (Li et al., 2024). In this work, we aim to comprehensively explore the underexplored avenue of output compression, advancing the state of the art in energy-efficient inference.

B Details of LLMs used

Table 6 states the links to the specific LLM variants that we used in our experiments.

C Details of Hardware/Software setup

We perform our experiments on a single NVIDIA A6000 GPU with 48GB VRAM hosted in a local server with Intel Xeon Silver 4210R processor and 128GB RAM, running Ubuntu 20.04-LTS. We use Pytorch version 2.3 (with CUDA 12.1). We employ vLLM (Kwon et al., 2023) for all our experiments, which is a fast and easy-to-use library for LLM inference and serving.

To measure energy consumption of LLMs in offline mode, we used the CodeCarbon (Courtney et al., 2024) library with default settings. The library samples every 15 seconds the GPU power usage from the *nvidia-smi* app, CPU-power from Intel RAPL files and an approximate RAM power. These power values are integrated over time to get the total energy usage. Note that CodeCarbon is a popular and reliable tool for measuring energy consumption (Jay et al., 2023; Bouza et al., 2023) and has been used in prior works to measure the energy consumption of LLMs (Poddar et al., 2025).

D More Length ratio plots

Figure 7 shows scatter plots comparing the target answer lengths and the generated response lengths for some more LLMs. In all cases, we see that LLMs generate substantially longer responses than the target answer lengths, in most cases.

E Annotation Examples

Figure 8 shows screenshots of the annotation of LLM responses with information categories.

F Examples from Datasets

Table 7 demonstrates a few illustrative examples from each dataset.

Model	Model description link
GPT-3.5-turbo & GPT-4o-mini	https://platform.openai.com/
Llama-2-13B (13B params)	https://huggingface.co/meta-llama/Llama-2-13b-chat-hf
Llama-3-3B (3B params)	https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct
Llama-3-8B (8B params)	https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct
Gemma-2-2B (2.6B params)	https://huggingface.co/google/gemma-2-2b-it
Gemma-2-9B (9B params)	https://huggingface.co/google/gemma-2-9b-it
Phi-3-mini (4B params)	https://huggingface.co/microsoft/Phi-3-mini-4k-instruct
Phi-3-small (8B params)	https://huggingface.co/microsoft/Phi-3-small-8k-instruct
Mistral-7B (7B params)	https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3
Falcon-11B (11B params)	https://huggingface.co/tiiuae/falcon-11B
Vicuna-13B (13B params)	https://huggingface.co/lmsys/vicuna-13b-v1.5

Table 6: Links to specific models versions we used in our experiments

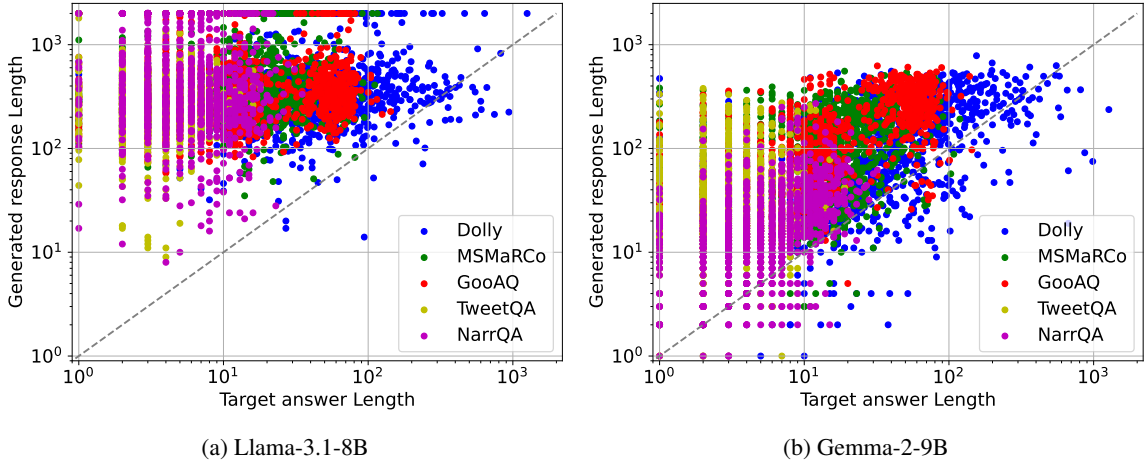


Figure 7: Target answer length vs LLM generated response length on different datasets

G Detailed ROUGE-L metrics for prompt engineering strategies

Table 8 gives the average ROUGE-L Precision scores for all the prompt engineering strategies. Table 9 gives the average ROUGE-L Recall scores for all the prompt engineering strategies. Table 10 gives the average ROUGE-L F1 scores for all the prompt engineering strategies.

H Detailed discussion about the scope and potential future work

In this work, we explore the question whether LLMs are generating too much text, which add up the energy consumption and environmental impact of the models. To this end, we benchmark the responses of various open-source LLMs for factual objective queries. We find that **LLMs are generating long responses**, often orders of magnitude higher than human-written answers. Revisiting decoding and training strategies with the insights from our experiments can be a promising direction for future research toward devising solutions that

produce concise and accurate answers.

Response categorization can play a pivotal role in devising techniques for forcing the models to generate brief (as desired) and accurate answers. Specifically, we established that the minimal answers contain enough information to achieve accurate responses, while redundant information could be removed completely in most cases. **What constitutes an ideal response** under different scenarios/applications should be studied more, and various factors influencing that decision should be identified. We argue that the composition of the responses should be based on users preferences too, whether they want concise or explanatory or creative answers. However, we argue that, from an energy optimality perspective, the LLM should, by default, generate a brief and succinct response for factual objective user queries.

We employed OpenAI GPT-4o for categorizing the responses from the set of 500 annotated examples, which is found incapable to differentiate between certain pair of categories like irrelevant

and redundant, reasoning and redundant, highlighting the need for more sophisticated, focused approaches. Overall it only achieved an average F1 score of 0.391 over all tokens. **Effective Methods for distinguishing categories**, between necessary (minimal answer and strongly related additional information), redundant (polite, summary, repetition, reasoning), and irrelevant content should be investigated. Towards that, we make our annotated LLM responses publicly available. Researchers can employ this dataset to develop such supervised framework (e.g. using Longformer (Beltagy et al., 2020) based models), that can be used to automatically classify LLM responses into different categories, which will remove the human annotation dependency and enable large-scale evaluation of LLM responses.

We found that Minimal answers are sometimes preceded by other categories (polite and conversational enhancements, additional information), limiting the scope for simple early stopping strategies such as stopping after 'k' tokens and discarding incomplete sentences, for getting the ideal minimal answer. Simple fine-tuning on these open-domain QA datasets proved ineffective, though it may prove beneficial for certain controlled environments. Fine-tuning with a length penalty may allow models to capture desired lengths better. Model families seem to generate responses with similar category distributions, indicating the **pre-training of the LLMs should be revisited** to have better control over the information categories in the response.

We also found irrelevant information are present in nearly 20% of responses across all models. This happens when the model is unsure of what to predict next. Generation can be shortened greatly if the models are stopped from generating irrelevant information, towards which **hallucination detection and prevention** (Manakul et al., 2023; Ji et al., 2023) are promising starting steps. **Model alignment** have been used to mitigate harmful content/biases in responses through re-training (Wang et al., 2024; Dai et al., 2024). Efforts should also be made to align LLMs towards generating responses of lengths that are desirable for a given situation.

Prompt engineering strategies show promise in reducing the response length and energy consumption, while preserving both precision and recall of the response for most of the models. Specifically, length-based prompts with sophisticated length estimator (PREDRESLEN), along with

certain category-specific prompts (MINANS and MADDNORED) are quite effective in achieving energy-aware inferences across different use cases. However, they are not enough to effectively control the desired proportions of the information categories. **Efficient decoding time algorithms** may be employed to control category information within the generated response in the real-time, enabling us to achieve a better compromise between readability and energy-optimality. MLP or RNN-based methods may be trained on our annotated response category data and incorporated into the final language modelling. That will also allow us a finer control for modulating the category distribution in the generated response, enhancing its scope and applicability.

The answer is "Five." The famous painting, created by Cassius Marcellus Coolidge in the early 20th century, depicts five dogs playing poker. This
 •REDINFO •MANS •ADDINFO

iconic artwork has become synonymous with canine card games and continues to be widely recognized today. ### Step-by-step reasoning process:

•REASON

To solve this question, one needs to identify the specific details mentioned - that it's an animal-themed painting featuring poker players. Given these clues, we can narrow down our search for the correct answer.

1. Identify key elements of the problem: Animal theme, Poker game, Painting.
2. Recall common depictions of animals engaging in human activities like gaming or leisure pursuits.
3. Recognize the association between paintings of animals and popular culture references (e.g., movies, literature).
4. Connect the dots between known artworks featuring animals engaged in various activities and those specifically related to gambling or entertainment.

By following these steps, one should arrive at the conclusion that the most fitting response would indeed be 'five', as it directly relates to the title of the well-known painting depicting five dogs playing cards.

This type of logical deduction requires attention to detail, knowledge of cultural icons, and the ability to connect seemingly disparate pieces of information into a coherent narrative. It also highlights the importance of understanding context when approaching questions about art, history, or pop culture.

•POLITE

10 pm EST is equivalent to midnight in AEDT (Australian Eastern Daylight Time). This means that if it's already past your bedtime at home, you can go straight into sleep mode. It will be the next day when time changes for those living on Australian East Coast and New Zealand
 •MANS •IRREL •ADDINFO

due east of Greenwich Mean Time zone where they are observing DST or Summer Saving hours which ends this Sunday November first weekend morning after sunrise local standard times change back by one hour from UTC+13hours/AET-4hr = UCT +9 hrs /AEST -5 hr . So now we have two different scenarios depending upon whether someone lives within these regions who still follow summer saving rules until end Nov : If yes then their clock would read as follows :- Sunset Oct 8th @6pm EDT , Mid Night October nineteenth(Sun) & Dawn September twenty thirdrd@7am PDT ; While For Those Who Don t Follow These Rules Their Clock Will Show The Following Times:- Evening Sept Twenty First At Five PM CDT; Morning Sep Thirtieth From Six AM MDT To Early Afternoon On Same Date With No Change In Local Standard Or Civil TIME ZONE! Now let us see how much difference there really was between our respective locations during each phase mentioned above using some simple calculations based off data provided earlier today regarding current weather conditions across North America along with historical averages over several decades ago We know what happened yesterday evening because everyone knows about Halloween night right?? But did anyone notice anything unusual happening around here last Friday afternoon before dark came down again ? Well maybe not exactly but certainly something worth mentioning since its been quite cold lately especially compared

•IRREL

Figure 8: Annotation examples depicting the information categories

Dataset	Example
DOLLY (Conover et al., 2023)	<p>Instruction: Instead of making a peanut butter and jelly sandwich, what else could I combine peanut butter with in a sandwich? Give five ideas.</p> <p>Answer: Instead of jelly, try one of the following with peanut butter in a sandwich: Honey Raisins Fruit preserves Bananas Sliced apples</p>
GOOQA (Khashabi et al., 2021)	<p>Question: what are the steps for decision making?</p> <p>Answer(collection): Step 1: Identify the decision You realize that you need to make a decision. , Step 2: Gather relevant information. Step 3: Identify the alternatives. Step 4: Weigh the evidence. Step 5: Choose among alternatives. Step 6: Take action. Step 7: Review your decision and its consequences.</p>
MS-MARCO (Nguyen et al., 2016)	<p>Open-domain question-answering dataset with questions sampled from Bing user query logs and answers generated by humans.</p>
NARRATIVEQA (Kočíský et al., 2018)	<p>Title: Ghostbusters II Summary snippet: . . . Peter’s former girlfriend Dana Barrett has had a son, Oscar. . . Story snippet: DANA (setting the wheel brakes on the buggy) Thank you, Frank. I’ll get the hang of this eventually. She continues digging in her purse while Frank leans over the buggy and makes funny faces at the baby, OSCAR, a very cute nine-month old boy. FRANK (to the baby) Hiya, Oscar. What do you say, slugger? FRANK (to Dana) That’s a good-looking kid you got there, Ms. Barrett Question: How is Oscar related to Dana? Answer: her son</p>
TWEETQA (Xiong et al., 2019)	<p>Passage: Oh man just read about Paul Walkers death. So young. Ugggh makes me sick especially when it’s caused by an accident. God bless his soul. – Jay Sean (@jaysean) December 1, 2013 Question: why is sean torn over the actor’s death? Answer: walker was young</p>

Table 7: List of datasets with illustrative examples.

Dataset	gemma -2-2b	gemma -2-9b	Llama -2-13b	Mistral -7B	Llama -3.2-3B	Llama -3.1-8B	Phi-3 -mini	Phi-3 -small	falcon -11B	vicuna -13b
Default	16.0	21.3	19.5	16.1	5.8	3.2	2.3	5.3	24.0	22.4
Brief	30.1	41.5	31.7	24.8	12.3	3.8	5.4	4.0	21.8	22.3
BM25-InContext	12.6	5.0	12.6	50.3	7.9	2.6	0.9	8.0	27.0	10.4
GoldResLen	31.7	45.2	41.0	31.6	15.8	4.5	8.9	7.0	24.1	36.3
MinAns	46.5	55.9	42.9	40.5	22.2	5.8	11.7	5.9	32.8	35.6
MAddNoRed	30.7	43.2	26.7	20.1	8.6	4.0	4.4	3.5	18.5	28.2
BM25-Length	25.0	26.6	30.7	18.3	8.0	3.7	2.3	4.9	23.5	22.2
PresResLen	33.3	35.9	38.3	24.0	9.6	3.5	5.3	5.6	25.6	29.2

Table 8: Average ROUGE-L Precision scores.

Dataset	gemma -2-2b	gemma -2-9b	Llama -2-13b	Mistral -7B	Llama -3.2-3B	Llama -3.1-8B	Phi-3 -mini	Phi-3 -small	falcon -11B	vicuna -13b
Default	55.1	57.0	56.9	56.0	60.7	64.4	26.7	55.0	41.0	47.9
Brief	46.5	46.0	50.6	52.2	55.7	62.2	27.4	56.9	41.3	51.4
BM25-InContext	12.1	5.4	15.2	56.8	61.0	67.9	5.2	34.8	22.8	10.0
GoldResLen	41.8	42.3	39.9	43.4	53.2	60.7	23.1	51.6	31.1	44.6
MinAns	38.7	36.6	41.0	44.7	48.5	60.5	25.9	51.1	34.2	45.0
MAddNoRed	43.4	44.1	50.5	52.7	55.4	61.6	28.9	58.2	41.5	48.9
BM25-Length	48.1	51.2	48.6	51.3	58.4	64.2	19.3	54.0	39.5	49.0
PresResLen	43.1	44.8	41.9	43.4	55.5	62.6	18.8	50.2	35.4	40.6

Table 9: Average ROUGE-L Recall scores.

Dataset	gemma -2-2b	gemma -2-9b	Llama -2-13b	Mistral -7B	Llama -3.2-3B	Llama -3.1-8B	Phi-3 -mini	Phi-3 -small	falcon -11B	vicuna -13b
Default	19.7	25.2	22.9	19.4	8.6	5.3	2.7	6.0	21.4	23.9
Brief	29.0	37.4	32.5	27.2	14.1	5.9	5.6	4.2	21.7	25.8
BM25-InContext	11.5	4.5	11.8	46.8	9.8	4.2	1.0	8.0	21.0	9.0
GoldResLen	29.3	41.3	37.4	29.5	17.8	6.9	9.2	8.0	22.5	36.5
MinAns	36.3	39.4	35.4	35.1	20.1	7.8	10.0	4.6	27.1	33.0
MAddNoRed	28.5	37.4	28.2	23.3	11.7	6.4	4.5	4.5	19.9	29.0
BM25-Length	27.7	29.9	32.1	22.1	11.4	6.0	2.9	5.9	23.0	24.9
PresResLen	30.8	32.7	34.1	23.7	12.3	5.5	5.1	6.0	22.6	27.5

Table 10: Average ROUGE-L F1 scores.