

# Unused information in token probability distribution of generative LLM: improving LLM reading comprehension through calculation of expected values.

Krystian Zawistowski  
0009-0002-9589-9030  
Samsung Research Poland  
Warsaw

Email: k.zawistowsk@samsung.com, krystian.zawistowski@zoho.com

**Abstract**—LLM text decoding is key component for perceived LLMs quality. We demonstrate two experiments showing that decoding methods could be improved by manipulation of token probabilities. First, we test few LLM on SummEval summary scoring dataset, to measure reading comprehension. We compare scores from greedy decoding to expected values over the next token distribution<sup>1</sup>. We scale logits by large temperature to increase the entropy of scores. This allows strong improvement of performance on SummEval (in terms of correlations to human judgement). We see improvement from 6-8% to 13-28% for 7B Mistral and from 20%-46% to 37%-56% for Mixtral, beating GPT 4 0314 result on two metrics. Part of the gain seems related to positional bias. Secondly, we use probability-based tree sampling algorithm, to examine all most probable generations for given prompt.

## I. INTRODUCTION

GENERATIVE LLMs are trained on large text corpora as estimators of next token probability conditional on prior text. Then sampling from such probability distribution is performed, or token with largest probability is chosen (greedy decoding). Typically, one introduces parameter  $T$  – called temperature. Let  $l_i$  be the logit for  $i$ -th token. Then token probability is as follows:

$$p_i = \frac{e^{\frac{l_i}{T}}}{\sum_j e^{\frac{l_j}{T}}} \quad (1)$$

Question arises: what the  $T$  should be? Typically  $T \in (0, 1]$ , with greedy decoding as limit in 0 and larger values corresponding to greater diversity (but also greater randomness). Research [1] shows that human generated text often does not correspond to modelled highest probability. Human choice of words is not guided by greatest probability, as “humans optimize against stating the obvious”. The author of this observation, H. P. Grice in [2] gave following example: suppose that I meet a man with a gas tank asking me to sell some gas and I answer “There’s gas station over the corner”. I said only

a bit of information that doesn’t answer directly, while there’s lots of implicit information not being said (gas station is open, it sells gasoline and has it available and you can go there buy some). Such concepts might be pretty foreign to LLMs: when we asked Mixtral Instruct “How to get gasoline in Fresno”, it gave us long instruction on finding gas station on a map, choosing best gas station, operating the pump safely, paying and so on, despite the fact that most of it is irrelevant to the problem of getting gasoline in Fresno specifically.

Thus, a tradeoff arises. Probability maximization with small temperature doesn’t give us natural, relevant responses. Large  $T$  is not ideal either, introducing more randomness, as low probability token might be either very informative or very wrong. Is however a single fixed value of  $T$  sufficient, even for specific use case?

We conjecture that decoding should be more dynamically controlled to more fully utilize the information in the distribution.

## II. SUMMARY EVALUATION WITH EXPECTED VALUE DECODING.

### A. Expected value decoding.

Currently, greedy decoding is often used for label based QA<sup>2</sup>. We want to find whether relative probabilities of few most probable tokens are informative. We test whether greedy decoding can be outperformed by calculation of expected value. We evaluate our approach on SummEval [4] dataset. It contains 1600 article summaries with human annotations for relevance, fluency, coherence and consistency of a summary. We compare our result against known LLM-based evaluations ([5], [6]). LLM is asked to evaluate relevance (or other feature) on Likert scale (from 1 to 5). We use MCQ prompts from [6], where LLM answer is A, B, C, D or E (A is 1 – worst, E is 5

<sup>1</sup>Source code released here: [https://github.com/kzawisto/unused\\_information\\_llm](https://github.com/kzawisto/unused_information_llm)

<sup>2</sup>For example in LM Evaluation Harness, standard set of LLM benchmarks [3]

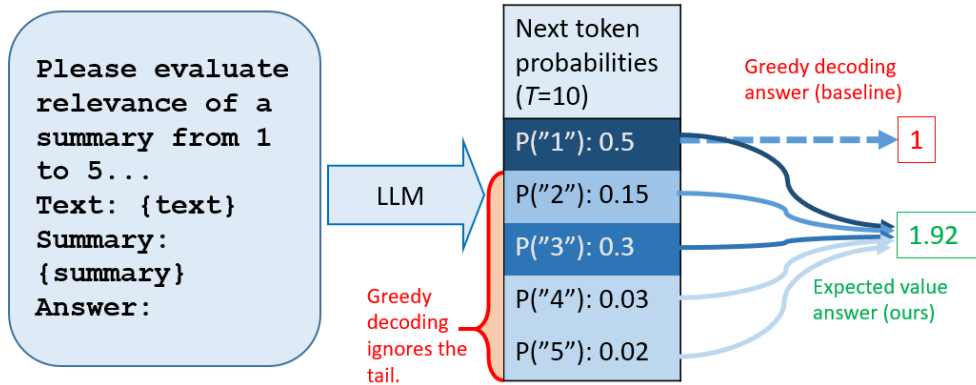


Fig. 1. Conceptual diagram of presented approach: instead of answering with most probable token, we calculate expected value for temperature  $T = 10$  to utilize residual information in next-token distribution.

– best). Let  $p(A), p(B) \dots$  be probability of token “A”, “B”... Expected value score (that we use) is calculated as follows:

$$E(s) = p(A) + 2p(B) + 3p(C) + 4p(D) + 5p(E), \quad (2)$$

while greedy score (a standard baseline) is:

$$s_{max} = F \left( \arg \max_{t \in \{A, B, C, D, E\}} p(t) \right), \quad (3)$$

where  $F$  is mapping  $\{A \rightarrow 1, B \rightarrow 2 \dots\}$ .  $E(s)$  depends on  $T$ , as (1) shows, while  $s_{max}$  does not. However, our initial experiments show that for  $T \in (0, 1]$  these values are very close to each other,  $E(s) \approx s_{max}$ . LLM is spuriously certain about its answer and assigns near 100% probability to selected answer. To fix that we increased entropy of score distribution by setting very high  $T = 10$  (i.e. we want for these scores to have smooth continuous distribution over the  $[1, 5]$  interval). Conceptual diagram is presented on Figure 1.

### B. Results

We evaluated Pearson correlation of scores to human judgements, the scores being calculate either with greedy method (3) or expected value method 2 with  $T = 10$ . We saw strong improvement in metric correlation to human judgement. In addition,  $E(s)$  scores from Mixtral 8x7B Instruct [7] beat GPT3.5 and nearly match GPT4 results from [6], see Table I.

Furthermore, strong improvements were produced for small and quantized LLM too. We evaluated 3 LLMs, from 7B to 47B parameters and all metrics are consistently improved. We compare scores for quantized and float16 Mixtral Instruct in Table II. Surprisingly, quantized Mixtral performance is only slightly worse than float16 version.

Up to 4.4 times improvement is achieved for Mistral v0.2 Instruct 7B from [8] (from 6.4% to 28.4% on relevance). We show these results in Table III. In Table IV we show result for SOLAR 10.7B Instruct [9]. While authors of the model reported it to outperform much larger Mistral 7x8B on some benchmarks, we see nothing similar for SummEval. Consistently with Mistral results, for SOLAR the largest gain

was observed for relevance evaluation (from 19% to 43%). For every model we show results for float16 inference and also for model checkpoints quantized with use of GPTQ [10].

We see that summarization metrics, being relevant automated metrics for reading comprehension, strongly improve with the number of parameters. Also, quantized 4-bit LLMs are very strong performers proportionally to their size and outperform similar size float16 models (for instance, quantized Mixtral, having 24GB in parameters and Mixture-of-Experts architecture strongly outperforms float16 SOLAR with 21 GB of parameters). For this reason quantized LLMs might be viable, cost-effective option for RAG and other similar use cases. This phenomenon is similar to emergent abilities of LLMs [11] where larger sizes lead to qualitative improvement in LLM performance. In this case too, quantized LLMs [12] retain large portion of their emergent capabilities.

Gains are particularly strong for relevance and consistency: this is important for systems that rely on reading comprehension, like RAG expert systems. We used vLLM [13] and Transformers [14] for implementation.

### C. Positional bias.

Previously it was reported that LLM preference for candidate responses might be altered [15] by simply reordering the responses in the prompt. This effect is called positional bias. Our experimental setup might be affected by it, as we use multi choice question answering prompts from [6].

We modified our approach as follows: we evaluate our score for two nearly identical MCQ prompts that differ by the order of answer candidates. One prompt has answer candidates in A, B, C... order, the other in E, D, C... order. Having done that, we average the scores for two prompts, doing that for every example we evaluate.

We performed this experiment for Mistral 7B for relevance evaluation and results can be found in Table V. Quite interestingly averaging out positional bias produces strong improvement for greedy decoding, while there’s no big difference for  $E(s)$  decoding. Furthermore stronger improvement is found

TABLE I  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL: MIXTRAL-INSTRUCT AND PRIOR RESULTS FOR OPENAI MODELS.

Metric	GPT3.5 0301 [6]	GPT4 0314 [6]	ChatGPT [5]	Mixtral E(s), T=10	Mixtral (greedy)
Fluency	0.431	<b>0.600</b>	0.384	0.392	0.222
Relev.	0.395	0.461	0.459	<b>0.555</b>	0.457
Consist.	0.484	<b>0.618</b>	0.516	0.506	0.397
Coher.	0.416	0.428	0.438	<b>0.485</b>	0.427

TABLE II  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL FOR MIXTRAL 8x7B v 0.1 INSTRUCT.

Metric	Fp16, greedy	Fp16, E(s)	Int4, greedy	Int4, E(S)
Fluency	0.222	0.392	0.235	<b>0.405</b>
Relev.	0.457	0.555	0.464	<b>0.564</b>
Consist.	0.379	<b>0.506</b>	0.293	0.470
Coher.	0.428	<b>0.485</b>	0.388	0.438

TABLE III  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL FOR MISTRAL 7B v 0.2 INSTRUCT.

Metric	Fp16, greedy	Fp16, E(s)	Int4, greedy	Int4, E(S)
Fluency	0.060	<b>0.134</b>	-0.045	0.061
Relev.	0.064	<b>0.284</b>	0.074	0.264
Consist.	0.061	<b>0.252</b>	0.076	0.249
Coher.	0.084	<b>0.199</b>	0.042	0.176

TABLE IV  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL FOR SOLAR 10.7B INSTRUCT.

Metric	Fp16, greedy	Fp16, E(s)	Int4, greedy	Int4, E(S)
Fluency	0.187	0.240	0.187	<b>0.251</b>
Relev.	0.192	<b>0.427</b>	0.165	0.364
Consist.	0.298	0.331	0.156	0.194
Coher.	0.305	<b>0.362</b>	0.200	0.267

for Fp16 model, than for Int4 model. Table VI shows few more experiments for Nous Hermes DPO Mistral 7B model, where either prompt candidates put in ascending (A, B, C... order), reversed (E, D, C... order) or random order. NaN correlations indicate that the model had predicted identical result for all test examples. One new result here is that random order is remarkably bad, with many correlations dropping by 30% or more. From human point of view prompt says exact same thing, but this is not the same for LLMs, which cannot generalize when the structure is altered. Also, NousHermes Mistral, undergoing more extensive finetuning and alignment, outperforms Mistral Instruct on metrics related to logical reasoning, but underperforms on fluency and coherence.

This suggests that gains from  $E(s)$  method might be related to positional bias, but details of it are not clear without further research.

This looks related to spurious certainty of LLM we already mentioned, our conjecture of temperature misconfiguration and improved results for very high  $T = 10$ . LLM, when having no good candidate hypothesis, seems to overreact to weak signals - instance of this problem is positional bias. While averaging out provides specific solution to positional bias, setting large temperature provides general solution: as LLM might take into account more candidate hypotheses, which presently are dominated by overreaction to spurious signal.

These problems could be related to the use of softmax

function in attention heads. Neural net limitations with respect to softmax and the rank of matrix under it were brought to attention by [16] (which proposes high rank RNN). Similar problems might reemerge in case of transformers and attention, which use relatively small matrices for attention heads. For softmax it does not matter whether signal is weak or strong, only whether it is the strongest among provided candidates. It is also true however, that weak signal suppression can be learned by the attention head in the pretraining process, especially when bias matrices are added to  $Q$  and  $K$  - so it is impossible to tell more without further study.

D. Statistical analysis.

We evaluate our results on 1600 samples from SummEval dataset, calculating Pearson correlation to human judgement evaluations on four metrics: fluency, relevance, consistency, coherence. We recalculate correlations for 1600 ChatGPT-evaluated samples provided by [5] and 1200 samples evaluated by GPT3.5 0301 and GPT4 0314, provided by [6].

We evaluate statistical significance with use of bootstrap method. We randomly shuffle series of human evaluation metric  $x_i$  and we do it 10000 times. For every random shuffle  $\hat{x}_i$  and Pearson correlation coefficient  $r$  we calculate  $\bar{x}_i = \frac{r}{\sqrt{1-r^2}}\hat{x}_i + x_i$ . Clearly for large sample size Pearson correlation  $corr(x_i, \bar{x}_i) \rightarrow r$ . We examine the empirical cumulative distribution  $P(\rho|r, x_i)$  of  $corr(x_i, \bar{x}_i)$ .

TABLE V  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL FOR MISTRAL 7B V 0.2 INSTRUCT - POSITIONAL BIAS ANALYSIS.

Metric	Fp16, greedy	Fp16, E(s)	Int4, greedy	Int4, E(S)
Relevance - standard.	0.064	0.284	0.074	0.264
Relevance - average.	0.245	0.295	0.161	0.307

TABLE VI  
PEARSON CORRELATIONS TO HUMAN JUDGEMENT ON SUMMEVAL FOR MISTRAL 7B NOUS HERMES DPO FP16 - POSITIONAL BIAS ANALYSIS.

Metric	Greedy, ascending	Greedy, reverse	Greedy, random	E(S), ascending	E(S), reversed	E(S), random
Relevance	0.32	0.36	0.092	0.44	0.42	0.2
Consistency	0.38	0.33	0.17	0.5	0.42	0.3
Fluency	0.097	NaN	0	0.12	0.0001	0.015
Coherence	0.16	NaN	0.04	0.24	0.06	0.061

With this we seek to evaluate, whether the difference of two sample correlations  $r_1$  for sample 1 and  $r_2$  for sample 2 is statistically significant. Significant difference of  $r_1$  and  $r_2$  would correspond to  $r_2$  being unlikely result if real correlation for sample 2 was  $r_1$ :

$$P(\rho < r_2 | r_1, x_i) < 5\%$$

and

$$1 - P(\rho > r_2 | r_1, x_i) < 5\%,$$

according to  $p$ -value testing methods.

Our estimates suggests that our results in Table I for consistency evaluation and relevance evaluation with Mixtral are significantly better than GPT model results, as far as statistical significance is concerned.

In addition, almost all improvements of  $E(s)$  method over greedy method provide significant difference in correlation. Only exception is consistency evaluation for SOLAR for Fp16 model in Table IV, which is not statistically significant (the difference is 3.3% while significance threshold corresponds to 3.4%).

### III. TREE-BASED SAMPLING

To further develop our hypothesis we propose an LLM inference analysis method that, for a given prompt, seeks to find all probable completions that could be generated by nucleus sampling - to give complete, nearly deterministic picture, what LLM outputs could be for given prompt. As a foundation we use tree-search based sampling algorithm. We use priority queue mechanism, where most probable completions are evaluated first (like in Dijkstra algorithm). Tree sampling (a.k.a beam search) is broadly implemented approach<sup>3</sup> in generative language models. Recently we saw very similar algorithm to ours [17] applied to compiler optimization (highest probability output produces superior compiler parametrization). Other controlled beam search techniques used for improved natural language generation can be found in [18], [19].

We utilize priority based tree sampling to find all possible or most probable completions for nucleus sampling. Algorithm 1 shows this procedure in pseudocode. This algorithm has exponential asymptotic complexity: every iteration produces  $N$  new sequences without fixed lower bound for  $N$  (LLM

tokenizers have tens of thousands of tokens), leading to exponential divergence  $K^N$  for  $K$  new tokens. One could decrease  $N$  by adjusting  $T$  and  $\hat{p}$ . We notice that for some prompts  $N$  is small number and large values of  $N$  indicate a qualitative change in the text generation (such as going from direct answer to user query, to additional not needed remarks). As an example of this we evaluate following prompt for Mixtral instruct<sup>4</sup>:

```
<s> [INST]Please provide one original,
creative paraphrase for sentence
"My name is John Kennedy"
and write new line after it[/INST]
Answer:\n\n"
```

Outputs with their evaluated probabilities can be found in Table VII. We used nucleus sampling threshold  $\hat{p} = 0.9$  and temperature  $T = 2$  and we show outputs with  $p > 0.1\%$ . Temperature is large, and reason is that for smaller temperatures only first, most probable output would be generated, while now it is generated with 73% probability. Output distribution is uneven, with top output 41 times more probable than second most probable output and very fat tailed with about 23% of probability mass distributed among very unlikely outputs ( $p < 0.1\%$ ).

We did not get any diversity of the paraphrase, despite asking for it explicitly, we got only two options: "I go by the name of Kennedy John" and "I go by the name of JFK". At the same time LLM becomes unpredictable in the tail of the distribution as various additional comments follow after requested text. Output is thus not diverse and diversity we get provides little benefit: it might be a problem, when long unpredictable output follows the answer, evading usual stopping mechanisms of the inference (here we stop inference on two new lines in a row.)

This type of behavior can be easily explained with reference to the content of training corpora for LLM. Specific tasks like "paraphrase this sentence" or specific sentences like "My name is John Kennedy" are likely rare in the large internet crawl corpora. At the same time, casual conversation is more frequent, so LLM can generate diverse full sentence answers (but this is not very useful for instruction-following tasks).

<sup>3</sup>Available in popular library Transformers [14].

<sup>4</sup>The prompt contains no new lines, but line wrapping was added for clarity.

---

**Algorithm 1** Tree-crawling topP algorithm

---

**Require:**

$t_1 \dots t_n$  {Prompt input sequence.}  
 $\phi(t_1 \dots t_n) \rightarrow l_i$  {LLM that maps token sequence to next-token log-probabilities}  
 $\hat{p} \in (0, 1)$  {TopP probability threshold}  
 $\alpha$  {Minimum loglikelihood of completion.}  
 $StopTokens$  {Tokens that terminate inference, such as newline or end-of-sentence.}  
 $MaxSteps$  {Max number of LLM evaluations}

**Ensure:**

$Queue \leftarrow [(t_1 \dots t_n, 0)]$  {Priority queue ordered by second argument}  
 $Complete \leftarrow []$  {Generated sequences terminated on stop tokens.}  
 $Incomplete \leftarrow []$  {Generated sequences terminated on minimum logprobability  $\alpha$ }  
 $i \leftarrow 0$

**while**  $Queue$  not empty and  $i < MaxSteps$  **do**

$seq_i, lproba \leftarrow Queue.pop()$

$l_i \leftarrow \phi(seq)$

**for all**  $token, l \in TopPCandidates(l_i, \hat{p})$  **do**

$el \leftarrow (concat(seq, token), lproba + l)$

**if**  $token \in StopTokens$  **then**

$Complete.append(el)$

**else if**  $l + x \geq \alpha$  **then**

$Queue.append(el)$

**else**

$Incomplete.append(el)$

**end if**

**end for**

$i \leftarrow i + 1$

**end while**

**return**  $Complete, Incomplete$

---

TABLE VII  
 TREE SAMPLING ALGORITHM OUTPUT FOR MIXTRAL INSTRUCT, COMPLETED SENTENCES.

Probability	LLM output
0.73	I go by the name of Kennedy John."
0.018	I go by the name of JFK"
0.005	I go by the name of JFK" is a creative paraphrase for "My name is John Kennedy."
0.005	I go by the name of JFK" is a creative paraphrase for the sentence "My name is John Kennedy."
0.003	I go by the name of JFK" is a creative paraphrase of "My name is John Kennedy."
0.002	I go by the name of JFK" is a possible creative paraphrase for "My name is John Kennedy."
0.001	I go by the name of JFK" is a possible paraphrase for "My name is John Kennedy."
0.001	I go by the name of JFK" followed by, "What an honor to make your acquaintance!"

Results point to causes of few problems of LLMs that we believe to be fairly widespread in applications based on prompting LLM and parsing their output.

- 1) Repetitive output – LLM output might be often identical, as there is single completion with very large probability. LLM might sometimes provide little advantage compared to retrieval based or rule based system.
- 2) Stopping instability – when LLM completes desired output, many different unwanted follow-up comments might be produced, disturbing rule based inference stopping mechanism.
- 3) Uncontrollability – when LLM is asked to do something, it can ignore instruction.

- 4) Hallucination - false or otherwise unwanted outputs might be produced by LLM in rare cases, while being undetected in tests using standard decoding.

These issues are rarely detected by commonly used accuracy-based benchmarks. There are some generic automated metrics such as MAUVE for diversity [20], but goal of our method is to analyze these issues in specific use cases and provide guidelines on configuration and further analysis. Our algorithm allows to analyze influence of modified prompts and system configuration on output probabilities, allowing greater degree of reliability and objectivity in the development, as different prompts, LLMs or sampling algorithm can be compared. On more general terms it seems that decoding heuristics like

TopP seem to fail our expectations, where entropy is extremely small or extremely large. Instead other approaches might be investigated: a) Scaling  $T$  for entropy extremes of next-token probability distribution. b) Taboo sampling – tree sampling with penalty for token and substring repetition. c) Stopping generation for large entropy spikes (as that would likely result in unwanted output). This might allow new improvements in few important KPI for LLM, such as output diversity, controllability and safety.

Among related work we may mention Mirostat [21], a sampling algorithm similar to TopP. Fixed perplexity objective allows to avoid both incoherence for large  $p$  and repetition for small  $p$ , while similar approach of [22] seeks to generate text with locally constant information content. Also vLLM [13] implements repetition penalty to stabilize low  $p$  inference. Another tradeoff however remains, as high Shannon information makes no difference between highly informative word in human terms and unwanted token overrepresented in the training dataset (as Shannon information is defined as the inverse of probability). In addition, prioritizing largest possible probability is very useful for some use cases, such as multi-choice QA [3] or LLM for source code generation (see Fig. 7 in [23]) or compiler configuration [17]. Balancing this tradeoff for humal language is a problem with no general solution yet known to us. We thus seek to provide a tool for analyzing specific problems, such as prompting, interaction between LLM and rule based scripts or stopping LLM inference reliably.

#### IV. CONCLUSIONS

We show LLM decoding method that improves performance for answers given on quantitative scale: such as “evaluate relevance of summary on a scale from 1 to 5”. On SummEval summary evaluation dataset the method produces strong improvements, with open source LLM nearly matching much larger GPT3.5 and GPT4, with GPT4 0314 outperformed on relevance and coherence. Such improvement supports our hypothesis that the temperature might be not optimally configured in standard LLM decoding, as token probabilities do not reflect real world probabilities and small and large temperatures serve different purposes.

We demonstrate new LLM analysis method using priority based tree sampling algorithm, useful for study of some niche problems in LLM, such as the diversity and controllability of the output.

We show reading comprehension metrics for few different LLM with sizes 7B, 10.7B and 47B with float16 (half precision) inference and 4 bit GPTQ quantization. Summarization metrics strongly improve with the number of parameters, and quantized 4-bit LLM are effective in proportion to their size (which is of interest for RAG on low-powered systems).

#### REFERENCES

- [1] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, “The curious case of neural text degeneration,” in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. doi: 10.48550/arXiv.1904.09751. [Online]. Available: <https://openreview.net/forum?id=rygGQyrFvH>
- [2] H. P. Grice, “Logic and conversation,” in *Syntax and Semantics: Vol. 3: Speech Acts*, P. Cole and J. L. Morgan, Eds. New York: Academic Press, 1975. doi: 10.1163/9789004368811\_003 pp. 41–58.
- [3] L. Gao, J. Tow, B. Abbasi, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, A. Le Noac’h, H. Li, K. McDonell, N. Muennighoff, C. Ociepa, J. Phang, L. Reynolds, H. Schoelkopf, A. Skowron, L. Sutawika, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou, “A framework for few-shot language model evaluation,” 12 2023. doi: 10.5281/zenodo.10256836. [Online]. Available: <https://zenodo.org/records/10256836>
- [4] A. R. Fabbri, W. Kryściński, B. McCann, C. Xiong, R. Socher, and D. Radev, “SummEval: Re-evaluating summarization evaluation,” *Transactions of the Association for Computational Linguistics*, vol. 9, 2021. doi: 10.1162/tacl\_a\_00373. [Online]. Available: <https://aclanthology.org/2021.tacl-1.24>
- [5] J. Wang, Y. Liang, F. Meng, Z. Sun, H. Shi, Z. Li, J. Xu, J. Qu, and J. Zhou, “Is ChatGPT a good NLG evaluator? a preliminary study,” in *Proceedings of the 4th New Frontiers in Summarization Workshop*, Y. Dong, W. Xiao, L. Wang, F. Liu, and G. Carenini, Eds. Singapore: Association for Computational Linguistics, Dec. 2023. doi: 10.18653/v1/2023.newsum-1.1 pp. 1–11. [Online]. Available: <https://aclanthology.org/2023.newsum-1.1>
- [6] C. Shen, L. Cheng, X.-P. Nguyen, Y. You, and L. Bing, “Large language models are not yet human-level evaluators for abstractive summarization,” in *Findings of the Association for Computational Linguistics: EMNLP 2023*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023. doi: 10.18653/v1/2023.findings-emnlp.278 pp. 4215–4233. [Online]. Available: <https://aclanthology.org/2023.findings-emnlp.278>
- [7] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed, “Mixtral of experts,” 2024. doi: 10.48550/arXiv.2401.04088
- [8] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed, “Mistral 7b,” 2023. doi: 10.48550/arXiv.2310.06825
- [9] D. Kim, C. Park, S. Kim, W. Lee, W. Song, Y. Kim, H. Kim, Y. Kim, H. Lee, J. Kim, C. Ahn, S. Yang, S. Lee, H. Park, G. Gim, M. Cha, H. Lee, and S. Kim, “Solar 10.7b: Scaling large language models

- with simple yet effective depth up-scaling,” 2024. doi: 10.48550/arXiv.2312.15166
- [10] E. Frantar, S. Ashkboos, T. Hoefler, and D. Alistarh, “Gptq: Accurate post-training quantization for generative pre-trained transformers,” 2023. doi: 10.48550/arXiv.2210.17323
- [11] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus, “Emergent abilities of large language models,” 2022. doi: 10.48550/arXiv.2206.07682
- [12] P. Liu, Z. Liu, Z.-F. Gao, D. Gao, W. X. Zhao, Y. Li, B. Ding, and J.-R. Wen, “Do emergent abilities exist in quantized large language models: An empirical study,” 2023. doi: 10.48550/arXiv.2307.08072
- [13] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. Gonzalez, H. Zhang, and I. Stoica, “Efficient memory management for large language model serving with pagedattention,” in *Proceedings of the 29th Symposium on Operating Systems Principles*, ser. SOSP ’23. New York, NY, USA: Association for Computing Machinery, 2023. doi: 10.1145/3600006.3613165. ISBN 9798400702297 p. 611–626. [Online]. Available: <https://doi.org/10.1145/3600006.3613165>
- [14] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush, “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Q. Liu and D. Schlangen, Eds. Online: Association for Computational Linguistics, Oct. 2020. doi: 10.18653/v1/2020.emnlp-demos.6 pp. 38–45. [Online]. Available: <https://aclanthology.org/2020.emnlp-demos.6>
- [15] C. Shen, L. Cheng, X.-P. Nguyen, Y. You, and L. Bing, “Large language models are not yet human-level evaluators for abstractive summarization,” in *Findings of the Association for Computational Linguistics: EMNLP 2023*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023. doi: 10.18653/v1/2023.findings-emnlp.278 pp. 4215–4233. [Online]. Available: <https://aclanthology.org/2023.findings-emnlp.278>
- [16] Z. Yang, Z. Dai, R. Salakhutdinov, and W. W. Cohen, “Breaking the softmax bottleneck: A high-rank rnn language model,” 2018. doi: 10.48550/arXiv.1711.03953
- [17] D. Grubisic, C. Cummins, V. Seeker, and H. Leather, “Priority sampling of large language models for compilers,” 2024. [Online]. Available: <http://arxiv.org/abs/2402.18734>
- [18] C. Meister, M. Forster, and R. Cotterell, “Determinantal beam search,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, C. Zong, F. Xia, W. Li, and R. Navigli, Eds. Online: Association for Computational Linguistics, Aug. 2021. doi: 10.18653/v1/2021.acl-long.512 pp. 6551–6562. [Online]. Available: <https://aclanthology.org/2021.acl-long.512>
- [19] L. Vilnis, Y. Zemlyanskiy, P. Murray, A. Passos, and S. Sanghai, “Arithmetic sampling: parallel diverse decoding for large language models,” in *Proceedings of the 40th International Conference on Machine Learning*, ser. ICML’23. JMLR.org, 2023. doi: 10.5555/3618408.3619870
- [20] K. Pillutla, L. Liu, J. Thickstun, S. Welleck, S. Swayamdipta, R. Zellers, S. Oh, Y. Choi, and Z. Harchaoui, “Mauve scores for generative models: Theory and practice,” 2023. doi: 10.5555/3648699.3649055
- [21] S. Basu, G. S. Ramachandran, N. S. Keskar, and L. R. Varshney, “{MIROSTAT}: A {neural} {text} {decoding} {algorithm} {that} {directly} {controls} {perplexity},” in *International Conference on Learning Representations*, 2021. doi: 10.48550/arXiv.2007.14966
- [22] C. Meister, T. Pimentel, G. Wiher, and R. Cotterell, “Locally typical sampling,” *Transactions of the Association for Computational Linguistics*, vol. 11, pp. 102–121, 2022. doi: 10.48550/arXiv.2202.00666. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252918666>
- [23] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba, “Evaluating large language models trained on code,” 2021. doi: 10.48550/arXiv.2107.03374