

DVAGEN: Dynamic Vocabulary Augmented Generation

Wei Du^{1*}, Nuowei Liu^{1*}, Jie Wang^{1*}, Jiahao Kuang^{1*},
Tao Ji², Xiaoling Wang¹, Yuanbin Wu^{1†}

¹School of Computer Science and Technology, East China Normal University

²School of Computer Science, Fudan University

{weidu.cs, nwliu, jiewang.cs, jhkuang}@stu.ecnu.edu.cn

taoji@fudan.edu.cn, {xlwang, ybwu}@cs.ecnu.edu.cn

Abstract

Language models trained with a fixed vocabulary struggle to generalize to novel or out-of-vocabulary words, limiting their flexibility in handling diverse token combinations. Existing dynamic vocabulary approaches attempt to address this limitation but face challenges such as fragmented codebases, lack of support for modern LLMs, and limited inference scalability. To overcome these issues, we introduce DVAGEN, a fully open-source, unified framework designed for training, evaluation, and visualization of dynamic vocabulary-augmented language models. Our framework modularizes the pipeline for ease of customization, integrates seamlessly with open-source LLMs, and is the first to provide both CLI and WebUI tools for real-time result inspection. We validate the effectiveness of dynamic vocabulary methods on modern LLMs and demonstrate support for batch inference, significantly improving inference throughput.¹

1 Introduction

Most language models (Grattafiori et al., 2024; Yang et al., 2025) are typically trained on a fixed vocabulary, which constrains their ability to generalize beyond the training corpus. This limitation becomes apparent when encountering novel (out-of-vocabulary) words or when attempting to efficiently generate phrases composed of arbitrary token combinations.

Instead of relying solely on a pre-trained tokenizer with a static vocabulary, recent studies have adopted dynamic vocabulary approaches to augment the generation process, thereby enhancing modeling capabilities for natural languages (Lan et al., 2023; Cao et al., 2024; Liu et al., 2024) and even protein languages (Liu et al., 2025).

However, these prior works rely on diverse codebases, some of which are not yet fully open sourced, and lack a unified framework for training, inference, and visualization of results augmented by dynamic vocabulary. Furthermore, they primarily focus on GPT-2 based language models (Radford et al., 2019; Ferruz et al., 2022), without validation on Large Language Models (LLMs). Additionally, the absence of support for batch inference further increases inference latency, despite the methods demonstrating promising inference speed when processing a single input at a time.

To address these issues, we present DVAGEN and our key contributions are highlighted below.

A fully open-source, unified framework for training, evaluation, and visualization. We propose a unified framework that decouples individual modules, allowing users to customize each component. The framework supports the full training and inference pipeline (train, chat, and eval) for developing dynamic vocabulary-augmented applications. Furthermore, we are the first to offer both command-line tools and a WebUI interface for visualizing generation results.

Integration of existing open-source LLMs and their validation. Existing open-source LLMs can be easily integrated into DVAGEN in a plug-and-play manner. We further evaluate the language modeling performance of LLMs using our framework and validate the effectiveness of employing dynamic vocabulary methods on LLMs.

Supports batch inference to enhance inference throughput. Dynamic vocabulary methods demonstrate promising capabilities in reducing token usage during decoding. Owing to their sequence compression abilities, these methods exhibit low inference latency. Building on this foundation, DVAGEN supports batch inference, further maximizing inference speed.

* Equal Contribution

† Corresponding Author

¹Codes are publicly available at <https://github.com/AntNLP/DVAGEN>.

2 Background

Research on dynamic vocabulary generation addresses key limitations of standard language models that rely on **static tokenization**. Traditional approaches like Byte-Pair Encoding (BPE) and Word-Piece employ fixed vocabularies (Radford et al., 2019; Devlin et al., 2019), which struggle to incorporate domain-specific phrases or multi-token expressions during generation. Efforts to enhance static vocabularies, such as Multi-Word Tokenization (MWT) (Gee et al., 2023), expand vocabularies with frequent n-grams but remain inflexible post-training.

Retrieval-augmented methods have emerged as alternatives. Document-level approaches like RETRO (Borgeaud et al., 2022) integrate retrieved passages via cross-attention but incur high latency. The Copy-is-All-You-Need (CoG) framework (Lan et al., 2023) pioneered phrase-level retrieval through a two-stage pipeline: first retrieving relevant documents, then extracting candidate phrases. While CoG supports variable-length generation, it processes phrases as token sequences during input, creating inconsistency between input and output representations. CoG-2 (Cao et al., 2024) advanced this paradigm by enabling direct phrase retrieval without document pre-filtering. Using syntactic heuristics and self-reinforced training, CoG-2 improved knowledge-intensive task performance but maintains the output-only phrase handling limitation of its predecessor.

The **dynamic vocabulary** approach (Liu et al., 2024) constitutes a fundamental architectural shift. It introduces a causal Transformer-based phrase encoder that maps arbitrary text spans to atomic units, unifying their treatment in both input and output layers. This eliminates the token-phrase inconsistency seen in CoG variants. The method operates plug-and-play, requiring no model retraining when incorporating new phrases. The dynamic vocabulary’s end-to-end atomic processing provides a more unified solution for fluency, efficiency, and verifiability in text generation.

3 DVAGEN

In this section, we introduce the overall architecture of DVAGEN, as illustrated in Figure 1.

3.1 DVAModel

The DVA framework integrates a dynamic phrase encoder with a standard language model (LM)

through a projection layer, enabling on-the-fly vocabulary expansion. Formally, given a static vocabulary V and a phrase candidate set $P = \{p_1, \dots, p_m\}$ from PhraseSampler, the model operates as follows:

Phrase Encoding For each phrase $p_i \in P$, tokenize it into subwords $[w_1, \dots, w_s]$ using the Phrase Encoder’s static tokenizer (e.g., GPT-2’s BPE), then compute its embedding via the Phrase Encoder:

$$\mathbf{e}_{p_i} = \text{Projector}(\text{PhraseEncoder}(\mathbf{w}_{1:s})_s)$$

where $\mathbf{w}_{1:s}$ denotes subword embeddings of p_i , and the last token’s hidden state is taken as the vector representation of p_i . A Projector then maps this to the LM’s embedding space. Typically, the Phrase Encoder is a causal Transformer model, and the Projector is an MLP.

Vocabulary Expansion Concatenate phrase embeddings with LM’s original input/output embedding matrices:

$$\mathbf{E}'_{\text{in}} = [\mathbf{E}_{\text{in}}, \mathbf{E}_P], \quad \mathbf{E}'_{\text{out}} = [\mathbf{E}_{\text{out}}, \mathbf{E}_P]$$

where $\mathbf{E}_P = [\mathbf{e}_{p_1}, \dots, \mathbf{e}_{p_m}] \in \mathbb{R}^{d \times m}$. The expansion is dynamic, that is for each sentence, we can attach different \mathbf{E}_P for generation.

Generation The DVAModel then generate on the expanded dynamic vocabulary. At step t , we compute logits on both token vocabulary V and phrase candidates P , and choose the output $y_t \in V \cup P$ accordingly:

$$y_t \sim \text{Softmax}(h_t \mathbf{E}'_{\text{out}}) \in \mathbb{R}^{|V|+|P|}$$

where h_t is the last hidden state of the LM. We note that DVAGEN allows each sentence in a batch to have a different set of phrase candidates, in which case a phrase mask is incorporated in DVA Logits Processor to filter out a specific set in \mathbf{E}_P for that sentence.

3.2 PhraseSampler

The PhraseSampler \mathcal{S} extracts candidate phrases P from retrieved documents D using configurable strategies. DVAGEN provides a set of pre-defined strategies:

- *NToken*: Samples n -gram token sequences from tokenized text.

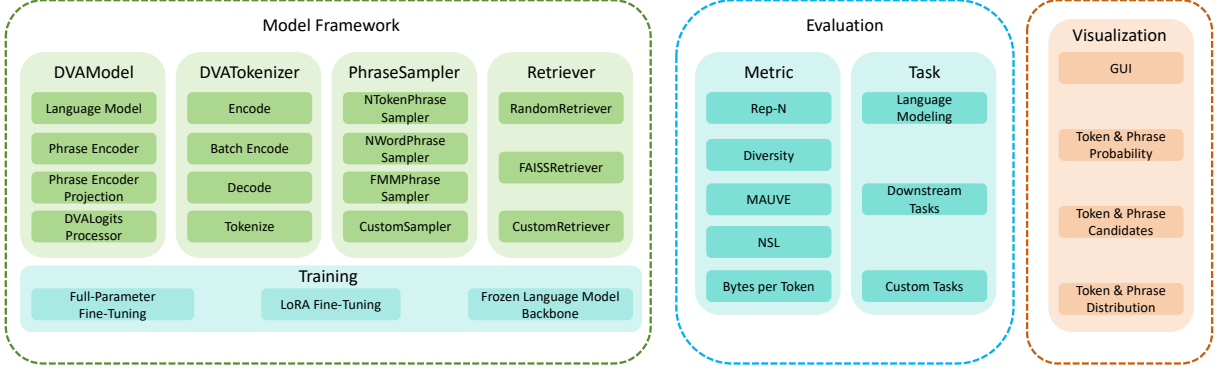


Figure 1: The overall architecture of DVAGEN. The light color boxes represent modules, while the dashed boxes indicate collections of modules with relevant functions.

- *NWord*: Extracts word-level n -grams via word tokenization (e.g., using NLTK).
- *FMM* (Forward Maximum Matching): Suppose the current position in the text is i from a corpus \mathcal{C} , then a segment $\mathcal{C}[i : i + m]$ is identified such that it appears in at least one sentence of the corpus \mathcal{C} , and m is the maximum value satisfying this condition. This segment will be extracted as a phrase, and the process continues from position $i + m$.

Custom samplers can be integrated by implementing the interface $P \leftarrow \mathcal{S}(D, \text{config})$.

3.3 DVATokenizer

Denote input text as $X = x_1, x_2, \dots, x_L$ where L is the length of input text. There are three main function in DVATokenizer.

Tokenize This function can split text to phrases and tokens according to P :

$$\text{Tokenize}(X, P) = \{x'_1, x'_2, \dots, x'_{L'}\}, L' \leq L$$

where $x'_i \in P \cup X$ is original token or candidate phrase.

Encode It can encode input text to input ids.

$$\text{Encode}(X, P) = \{\text{id}_1, \text{id}_2, \dots, \text{id}_{L'}\}, L' \leq L$$

$$\text{type}(\text{id}_i) = \begin{cases} \text{token}, & \text{if } \text{id}_i < |V| \\ \text{phrase}, & \text{otherwise} \end{cases}$$

If the id_i 's value exceeds the vocabulary size $|V|$, it is a token ID; otherwise, it is a phrase ID.

Decode This function is the inverse function of encode which can get the original text from the mixed ids.

$$\text{Decode}(\text{Encode}(X, P)) = X$$

3.4 Training

DVAGEN supports full-parameter fine-tuning, LoRA fine-tuning (Hu et al., 2022), and training with a frozen language model backbone. Different training strategies share a similar processing paradigm, as illustrated in Figure 2. Given a training dataset D , we first sample a batch containing m training samples. A PhraseSampler is then applied to sample phrases from the batch. With an arbitrary combination of tokens and phrases within each training sample, we employ a DVATokenizer to encode them into discrete token and phrase ids for subsequent training.




Figure 2: The overall training pipeline of DVAGEN.

3.5 Inference and Evaluation

Inference The inference paradigm differs from the training process by incorporating the retrieval of the top k most relevant supporting documents. The overall inference pipeline is illustrated in Figure 3. Given an input prefix, we employ a Retriever implemented using an embedding model with a FAISS backend (Douze et al., 2024) to enable efficient similarity search on dense vectors. Building on top of LangChain², we utilize the GPU version of FAISS (Johnson et al., 2019) to accelerate the retrieval process. The retrieved supporting documents are then used to sample phrases using

²<https://github.com/langchain-ai/langchain>

a **PhraseSampler**, and the sampled phrases serve as candidates for constructing the dynamic vocabulary. It is worth noting that for different input prefixes, the phrase candidates vary according to the retrieved supporting documents. Unlike previous studies (Lan et al., 2023; Liu et al., 2024) that do not support batch inference, we implement the **DVALogitsProcessor** to mask irrelevant logits (i.e., those corresponding to phrase ids associated with phrases unrelated to the current input). The positions to be masked are determined based on phrase candidates from other inputs within the same batch.




Figure 3: The overall inference pipeline of DVAGEN.

Evaluation Evaluation across various tasks are supported in DVAGEN, including basic language modeling and a range of downstream tasks. To ensure ease of use and fair comparisons, our framework includes implementations of a wide array of evaluation metrics. These metrics are categorized into three types: (1) generation quality (e.g., MAUVE (Pillutla et al., 2021), Rep-N and Diversity (Welleck et al., 2019), and Perplexity), (2) sequence compression (e.g., Normalized Sequence Length and Bytes per Token (Dagan et al., 2024)), and (3) downstream task performance (e.g., Recall, Precision, F1, and ROUGE-L (Lin, 2004)).

3.6 Visualization

Our web application is developed using the Django³, featuring a front-end interface built with Bulma⁴ to provide a clean and responsive user experience, and DVAGEN handles all text generation tasks in backend.

Figure 4 presents the input area of the front-end interface. The brown input box allows users to enter an initial query or prefix text, while the yellow phrase input area beneath it is specifically designed for defining dynamic vocabulary—users can specify phrases that the model should prioritize during generation. After completing the inputs,

users can click the "Generate" button to initiate the generation process.

Figure 5 displays the output area of the interface after generation is complete. The generated text appears in the grey output box. By clicking on any token or phrase in the box, users can view a scrollable list of candidate alternatives for that position, along with their associated probabilities. Users may then select a preferred alternative to replace the original token or phrase, thereby constructing a new prefix for controlled generation. To support more targeted interaction, the interface also offers filter options such as "Phrases", "Tokens & Phrases" and "Tokens", enabling users to refine the candidate list by category and focus on the specific category of alternatives.

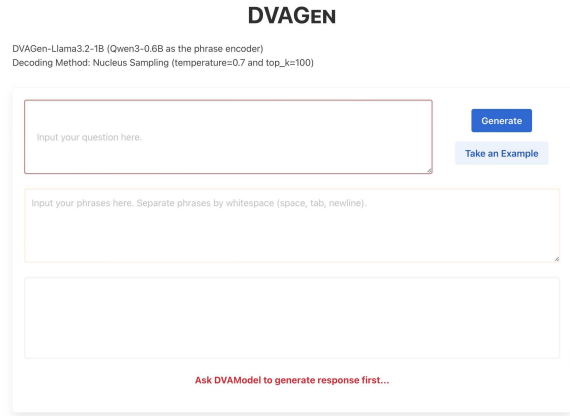


Figure 4: The screenshot of the DVAGEN WebUI interface.

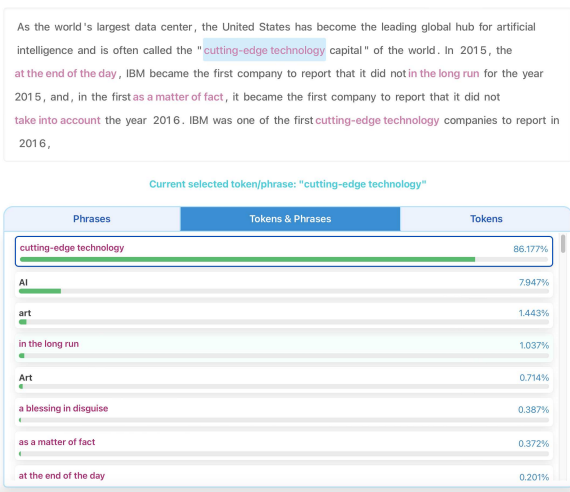


Figure 5: The screenshot displays the generation outputs. Phrases are highlighted in red, and the logit probabilities for each token and phrase can be viewed by clicking on the respective token or phrase.

³<https://www.djangoproject.com/>

⁴<https://bulma.io/>

To intuitively present the complex outputs generated by our model, we have developed a powerful visualization application.

Figure 6 illustrates the visualization generated by the application, which accurately captures the model’s interpretation of each token and phrase within the generated text. Tokens and phrases are distinguished using distinct color schemes—for instance, blue represents tokens, while purple-red denotes phrases. The shade of each color reflects the corresponding generation probability, with deeper shades indicating higher probabilities. A heatmap positioned at the top of the visualization clearly maps probability values to color intensities. Additionally, the application offers extensive customization options, enabling users to adjust token and phrase colors, font sizes, and image resolution to suit their specific analytical or presentation requirements.

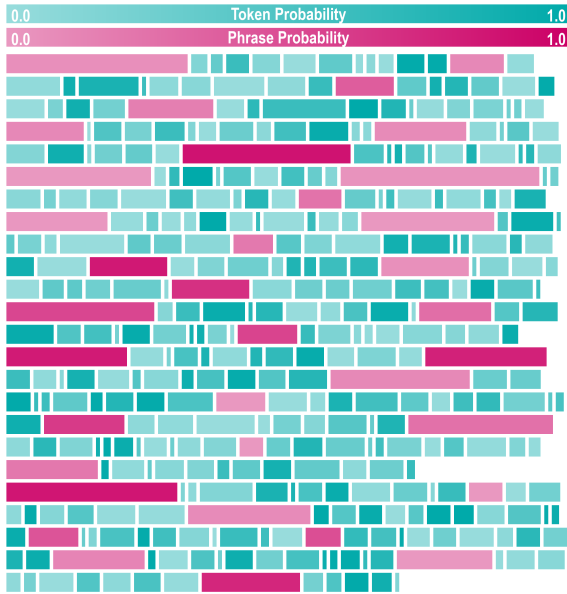


Figure 6: Visualization of generation outputs is represented in blue for token distribution and red for phrase distribution. The brightness of the colors reflects the logit probabilities, with darker colors indicating higher confidence.

4 Evaluation

4.1 Language Modeling

Setup We first show that DVAGEN can enhance the basic language modeling capabilities of LLMs. To evaluate this, we select the pre-trained base mod-

els Qwen3-0.6B⁵, Qwen3-1.7B⁶, Llama3.2-1B⁷, and Llama3.2-3B⁸. These models are fine-tuned on the WikiText-103 training set to serve as baselines. To ensure a fair comparison across different model families, a GPT-2 (Radford et al., 2019) tokenizer is employed to extract the first 32 tokens from each test sample as input. During inference with DVAGEN, we utilize the NTokenPhraseSampler to sample phrases from the given input prefixes. Qwen3-Embedding-0.6B⁹ is used as the embedding model for Retriever, and 32 documents are retrieved for each sample.

Results As shown in Table 1, our key findings can be summarized as follows:

(1) **DVAGEN improves generation quality while requiring fewer tokens to generate the same content.** DVAGEN helps maintain a lower Rep-N score and higher diversity, particularly for Llama models. The improvement in the MAUVE score indicates the generation of more natural and coherent language that closely matches human-written text. More importantly, training with DVAGEN significantly enhances sequence compression capability, demonstrating improved efficiency in reducing both tokens and decoding steps.

(2) **Freezing the language model during training is memory-efficient and achieves performance on par with that of a trainable backbone.** The results highlight that training with a frozen language model backbone yields performance on par with a trainable backbone while significantly reducing memory requirements. The evaluations offer insights for future research aimed at developing memory-efficient methods and reducing model size to enable more efficient training and inference.

4.2 Inference Performance

Setup To evaluate the generation performance of DVAGEN, we compare the tokens and UTF-8 bytes generated per second by Qwen3-0.6B and DVAGEN-Qwen3-0.6B. For a fair comparison, both models are evaluated on the WikiText-103 test set under identical settings. Each sample in a batch consists of 32 tokens as input prefix, and

⁵<https://huggingface.co/Qwen/Qwen3-0.6B-Base>

⁶<https://huggingface.co/Qwen/Qwen3-1.7B-Base>

⁷<https://huggingface.co/meta-llama/Llama-3.2-1B>

⁸<https://huggingface.co/meta-llama/Llama-3.2-3B>

⁹<https://huggingface.co/Qwen/Qwen3-Embedding-0.6B>

Models	MAUVE \uparrow	Rep-2 \downarrow	Rep-3 \downarrow	Rep-4 \downarrow	Diversity \uparrow	PPL \downarrow	NSL \downarrow	Bytes per Token \uparrow
Qwen3-0.6B	21.70	11.46	5.07	3.05	81.49	51.09	1.00	3.17
DVAGEN-Qwen3-0.6B ❄️	24.31	19.68	10.92	7.59	66.12	22.29	0.87	4.84
DVAGEN-Qwen3-0.6B 🔥	24.41	14.81	6.46	3.61	76.82	24.19	0.85	4.89
Qwen3-1.7B	22.74	10.21	4.14	2.19	84.18	57.55	1.00	3.27
DVAGEN-Qwen3-1.7B ❄️	24.92	18.64	10.26	7.21	67.76	21.73	0.86	4.88
DVAGEN-Qwen3-1.7B 🔥	23.57	16.05	7.88	4.99	73.47	19.93	0.99	4.63
Llama3.2-1B	24.74	17.28	8.73	5.29	71.50	28.65	1.00	4.00
DVAGEN-Llama3.2-1B 🔥	24.64	15.70	7.41	4.46	74.57	81.66	0.88	4.90
Llama3.2-3B	25.61	15.62	7.64	4.58	74.35	63.04	1.00	4.01
DVAGEN-Llama3.2-3B ❄️	25.26	16.85	8.50	5.41	71.96	43.98	0.87	4.97
DVAGEN-Llama3.2-3B 🔥	26.40	14.30	6.49	3.79	77.09	37.19	0.93	4.90

Table 1: Evaluation results on the WikiText-103 test set. Qwen3-0.6B is used as the phrase encoder and remains trainable across all settings. The 🔥 and ❄️ markers indicate whether the language model backbone is trainable or frozen, respectively.

the models are required to generate exactly 128 new tokens (or phrases when using DVAGEN) by setting both the minimum and maximum number of generated tokens to the same value. Note that, in alignment with previous work (Lan et al., 2023; Liu et al., 2024), the retrieval cost is not included, as the construction of phrase candidates can be conducted offline. The evaluations are conducted using a single NVIDIA RTX 4090 GPU.




Figure 7: Evaluation of performance in generating tokens and UTF-8 encoded bytes per second.

Results The evaluation results are presented in Figure 7, revealing two main findings:

(1) **Compared to the base model, DVAGEN achieves higher inference speed despite having more parameters.** DVAGEN-Qwen3-0.6B employs two Qwen3-0.6B models, serving as the language model and phrase encoder, respectively, resulting in a model twice the size of the base model. However, due to phrases typically consisting of multiple tokens and containing more UTF-8 bytes, DVAGEN achieves faster inference than the base model under the same settings.

(2) **DVAGEN supports batch inference and**

consistently maintains high inference speed as the batch size increases. The implementation of batch inference improves generation efficiency by approximately $7\times$ compared to processing one input at a time, as in the original implementations (Lan et al., 2023; Liu et al., 2024). Furthermore, the consistent inference performance with larger batch sizes demonstrates that DVAGEN effectively scales with workload, increases throughput, and optimizes GPU memory utilization without incurring additional latency.

4.3 Retrieval Performance

Setup In addition to the inference performance (excluding the retrieval process) discussed in Section 4.2, we also conduct experiments to compare retrieval latency during inference. The experimental setup is similar to the previous one, except that a single NVIDIA RTX 3090 GPU is used. Since many real-time applications require longer contexts and the Retriever requires substantial GPU memory during inference, we also evaluate its retrieval performance on CPU devices. This experiment is conducted on a server equipped with an Intel(R) Xeon(R) Silver 4214R CPU @ 2.40 GHz and 512 GB of RAM.

Results The comparisons of the proportions of different stages (primarily retrieval and generation) during inference are shown in Figure 8. The key observations are as follows:

(1) **When a GPU is utilized for inference, the generation stage consistently dominates the overall inference time, substantially exceeding the**




Figure 8: Evaluation of retrieval performance: (a) retrieval performed on a GPU device, (b) retrieval performed on a CPU device.

time spent on retrieval. As discussed in Section 3.5, DVAGEN retrieves relevant supporting documents on the GPU by default, with the retrieval process accounting for approximately 7% of the inference time when the batch size is set to 8. In addition, since retrieval within a batch is performed sequentially due to the independence of samples, retrieval time increases with larger batch sizes.

(2) Compared with GPU-based retrieval, CPU-based retrieval requires substantially more time. When the batch size is set to 8, CPU-based retrieval accounts for approximately half of the total inference time, slightly exceeding the duration of the generation stage (performed on a GPU device). This observation suggests that when using CPU-based retrieval to save GPU memory, the trade-off between batch size and inference latency should be carefully considered.

5 Conclusion

In this work, we present a fully open-source, modular framework that unifies training, evaluation, and visualization for dynamic vocabulary-augmented applications. The framework enables seamless customization of individual components and is the first to provide both CLI and WebUI tools for generation analysis. By supporting plug-and-play integration of existing open-source LLMs, we demonstrate the effectiveness of dynamic vocabulary methods in enhancing language modeling performance. Furthermore, with support for batch inference, our framework significantly improves inference throughput, offering a practical and scalable solution for deploying efficient LLM-based systems.

Acknowledgments

The authors wish to thank all reviewers for their helpful comments and suggestions. The corresponding author is Yuanbin Wu. This research was (partially) supported by National Key R&D

Program of China (2024YFC3308103).

References

- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, and 9 others. 2022. [Improving language models by retrieving from trillions of tokens](#). *Preprint*, arXiv:2112.04426.
- Bowen Cao, Deng Cai, Leyang Cui, Xuxin Cheng, Wei Bi, Yuexian Zou, and Shuming Shi. 2024. Retrieval is accurate generation. *arXiv preprint arXiv:2402.17532*.
- Gautier Dagan, Gabriel Synnaeve, and Baptiste Roziere. 2024. Getting the most out of your tokenizer for pre-training and domain adaptation. *arXiv preprint arXiv:2402.01035*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. [The faiss library](#).
- Noelia Ferruz, Steffen Schmidt, and Birte Höcker. 2022. Protgpt2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):4348.
- Leonidas Gee, Leonardo Rigutini, Marco Ernandes, and Andrea Zugarini. 2023. [Multi-word tokenization for sequence compression](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 612–621, Singapore. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.

- Tian Lan, Deng Cai, Yan Wang, Heyan Huang, and Xian-Ling Mao. 2023. [Copy is all you need](#). In *The Eleventh International Conference on Learning Representations*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Nuowei Liu, Jiahao Kuang, Yanting Liu, Changzhi Sun, Tao Ji, Yuanbin Wu, and Man Lan. 2025. Protein design with dynamic protein vocabulary. *arXiv preprint arXiv:2505.18966*.
- Yanting Liu, Tao Ji, Changzhi Sun, Yuanbin Wu, and Xiaoling Wang. 2024. [Generation with dynamic vocabulary](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18931–18948, Miami, Florida, USA. Association for Computational Linguistics.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. Mauve: Measuring the gap between neural text and human text using divergence frontiers. *Advances in Neural Information Processing Systems*, 34:4816–4828.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI*.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.