World Journal of Innovation and Modern Technology, Vol. 7, Issue 4, (Aug)
ISSN 2682-5910
2024

DOI: 10.53469/wjimt.2024.07(04).10

Leveraging Large Language Models for Information Retrieval from NEPA Documents

Zibu Wei[1], Yang Luo[2], Guokun Xu[3], Zhengning Li[4], Yibo Yin[5], Ying Xie[6]

[1]Computer Science, University of California, Los Angeles, Los Angeles, USA
[2]Computer Science, China CITIC Bank Software Development Center, Beijing, China
[3]Computer Science, Beijing Foreign Studies University, Beijing, China
[4]Computer Science, Georgetown University, Washington, D.C. USA
[5]Computer Science, Contemporary Amperex Technology USA Inc, Auburn Hills, USA
[6]Computer Science, San Francisco Bay University, Fremont, USA
[1]zibuwei@ucla.edu, [2]luoyangdxx@163.com, [3]nicolashsu@hotmail.com, [4]zl132@georgetown.edu,
[5]epark00007@gmail.com, [6]floraxlr999@gmail.com

Abstract: ***This paper explores the application of large language models (LLMs) to efficiently and accurately extract relevant information from National Environmental Policy Act (NEPA) documents, specifically focusing on environmental impact statements (EIS). NEPA mandates federal agencies to evaluate the environmental effects of their proposed actions, and EIS documents are essential for this process. However, these documents are often lengthy and complex, making manual information extraction time-consuming and error-prone. We address this challenge by leveraging advanced natural language processing techniques and the newly introduced NEPAQuAD1.0 dataset, which contains 1,450 question-answer pairs generated under human supervision. Our approach involves fine-tuning the Meta-Llama-3.1-8B-Instruct model on this dataset. The results demonstrate significant improvements in retrieval accuracy and efficiency compared to baseline models, highlighting the potential of LLMs to streamline the environmental review process and provide valuable insights for environmental policy analysis. This work contributes to the broader field of natural language processing by offering a robust method for handling complex, domain-specific information retrieval tasks.***

Keywords: Large Language Models, NEPA, Information Retrieval, Environmental Impact Statements, Natural Language Processing, NEPAQuAD1.0

1. INTRODUCTION

The National Environmental Policy Act (NEPA) mandates that federal agencies evaluate the environmental effects of their proposed actions, typically documented in Environmental Impact Statements (EIS). These comprehensive reports are crucial for integrating environmental considerations into federal decision-making. However, the complexity and volume of EIS documents pose significant challenges for manual review and information extraction.

Large language models (LLMs), especially those based on transformer architecture like GPT-4, offer transformative potential for enhancing information retrieval from these extensive documents. LLMs can process and understand vast amounts of textual data, making them well-suited for handling the intricacies of EIS documents.

This paper leverages LLM capabilities to develop a robust system for accurately retrieving relevant information from NEPA documents. We fine-tune the Meta-Llama-3.1-8B-Instruct model on the NEPAQuAD1.0 dataset, which consists of 1,450 question-answer pairs generated under human supervision. This dataset addresses the specific complexities of environmental impact assessments.

Our methodology involves several technical steps. First, we preprocess the NEPA documents to ensure the text is suitable for model training and evaluation, including tokenization, normalization, and segmentation. Second, we fine-tune the Meta-Llama-3.1-8B-Instruct model on this dataset, optimizing it for question answering based on EIS documents. Third, we implement an attention mechanism to improve the model's focus on relevant parts of the input text.

We evaluate the model's performance using BLEU scores, BERTScore, and a custom scoring system with GPT-4

Turbo for factual accuracy. These metrics comprehensively assess the model's ability to retrieve accurate and relevant information from the documents. Our results show significant improvements in retrieval accuracy and efficiency compared to baseline models, demonstrating the potential of LLMs to streamline the environmental review process.

In summary, this paper presents a novel approach to using LLMs for information retrieval from NEPA documents, addressing a critical need in environmental policy analysis. Our findings offer valuable insights into applying advanced natural language processing techniques to complex, domain-specific tasks, contributing to the fields of machine learning and deep learning.

2. RELATED WORK

The application of large language models (LLMs) for information retrieval and natural language understanding has garnered significant attention in recent years, particularly with the development of transformer-based architectures. This section reviews the key advancements and methodologies relevant to our study, drawing on recent literature. The transformer architecture, introduced by Vaswani et al. [1], has become foundational in natural language processing (NLP) due to its ability to handle long-range dependencies in text effectively through self-attention mechanisms . Building on this, Devlin et al. [2] developed BERT (Bidirectional Encoder Representations from Transformers), which pre-trains deep bidirectional representations by jointly conditioning on both left and right context in all layers . This innovation significantly improved performance on a range of NLP tasks.

Further extending the capabilities of transformers, Brown et al.[3] advanced this approach with the development of GPT-3, demonstrating that LLMs could perform few-shot learning, thereby reducing the need for extensive task-specific data .
Radford et al. [4] introduced the concept of unsupervised multitask learning with language models, emphasizing the versatility of LLMs in various text-based applications .

The exploration of transfer learning limits was conducted by Raffel et al. [5] with the T5 model, which unified text-to-text frameworks for different NLP tasks, emphasizing the adaptability of pre-trained models across various domains . Similarly, the robustness of BERT was enhanced by Liu et al. [6] through the development of RoBERTa, which optimized pre-training techniques to further improve model performance .

In the context of question answering and comprehension, several studies have introduced innovative models and benchmarks. Wang et al. [7] created the GLUE benchmark, providing a platform for evaluating the performance of NLP models across multiple tasks . Kwiatkowski et al. [8] developed the Natural Questions benchmark, specifically designed to assess models' abilities to retrieve and answer questions based on real-world data .

The introduction of models like XLNet by Yang et al. [9], which generalized autoregressive pre-training methods, offered an alternative to traditional masked language models by capturing bidirectional contexts without the limitations of independence assumption . Similarly, Lewis et al. [10] introduced BART, a denoising autoencoder that combined the strengths of bidirectional and autoregressive models for text generation and comprehension tasks .

Recent advancements also include the development of ALBERT by Lan et al. [11] , which reduced model size while maintaining performance through parameter-sharing techniques . Clark et al. [12] proposed ELECTRA, a model that pre-trains text encoders as discriminators rather than generators, providing a more efficient pre-training approach .

In the realm of conversational AI, Zhang et al. [13] introduced DialoGPT, a large-scale generative pre-trained model specifically designed for dialogue generation, highlighting the application of LLMs in interactive systems . Additionally, Zhang et al. [14] presented SG-Net, a syntax-guided machine reading comprehension model that integrated syntactic information to enhance understanding and retrieval capabilities .

Gao et al. [15] proposed SimCSE, a contrastive learning approach for sentence embeddings that simplified the training process while achieving state-of-the-art performance on sentence-level tasks . These advancements collectively highlight the ongoing evolution and refinement of LLMs and their applications in NLP, providing a strong foundation for our study.

In addressing the challenge of efficiently extracting information from NEPA documents, we leverage the insights from Smith et al. [16] , which highlight the use of advanced data processing techniques in large-scale data environments. Their work on comparative usage patterns in cloud computing and hyperscale data centers informs our approach to handling the substantial volume of data found in EIS documents, ensuring scalability and efficiency in our extraction process.

Additionally, the methodology employed in Johnson et al. [17] provides a foundational basis for our model fine-tuning approach. Their dual-augmentor framework for domain generalization is pivotal in enhancing the robustness of our Meta-Llama-3.1-8B-Instruct model, allowing it to better generalize across the diverse and complex information found in NEPAQuAD1.0 dataset, ultimately improving the accuracy of information retrieval. Peichen et al.[18] underscore the relevance and applicability of LLMs in handling complex environmental documents.

By building on these methodologies, our research leverages the strengths of fine-tuned LLMs, particularly the Meta-Llama-3.1-8B-Instruct model, to address the specific challenges of extracting relevant information from NEPA documents. This work aims to contribute to the existing body of knowledge by demonstrating the efficacy of LLMs in a domain-specific context, emphasizing their potential to streamline the environmental review process.

3.ALGORITHM AND MODEL

3.1 MODEL NETWORK

Our model is built upon the Meta-Llama-3.1-8B-Instruct architecture, which represents a state-of-the-art language model designed for complex natural language understanding and generation tasks. This architecture is particularly well-suited for tasks involving large-scale document analysis, such as processing NEPA documents to extract relevant information efficiently. The model’s architecture consists of several integral components that work in tandem to achieve high accuracy and contextual understanding.

3.1.1 Encoder

The encoder is a fundamental component of the model, tasked with processing the input context χ and transforming it into a hidden representation h, as captured by the following equation:

$$h = Encoder(\chi) \tag{1}$$

This transformation is pivotal in the model's ability to understand and represent the semantic content of the input text. The encoder is architected with multiple layers of multi-head self-attention and feed-forward neural networks. Each of these layers plays a crucial role in enabling the model to capture intricate patterns, dependencies, and hierarchical structures present within the input text. This architecture is particularly effective for handling the complex, technical, and variable language found in NEPA (National Environmental Policy Act) documents.

The self-attention mechanism within the encoder is key to its ability to dynamically focus on different parts of the input sequence. This mechanism computes a weighted sum of the input representations, allowing the model to prioritize certain words or phrases depending on their relevance to the context. This dynamic focus is essential when processing long and detailed NEPA documents, where certain sections may carry more significance than others in response to specific queries.

Mathematically, the encoder's self-attention mechanism can be described as:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{2}$$

In this equation, Q, K, and V represent the query, key, and value matrices, respectively, derived from the input text. The dimension of the key vectors is denoted by d_k. The operation involves computing the dot product of Q and K, scaling by $\sqrt{d_k}$, and then applying a softmax function to normalize the result. This produces the attention weights, which are subsequently used to compute a weighted sum of the value matrix V. The softmax function ensures that the sum of attention weights is one, thus highlighting the most relevant parts of the input sequence.

The encoder's self-attention mechanism not only allows the model to focus on the most relevant parts of the input but also to maintain a global understanding of the entire sequence. This balance between local focus and global awareness is particularly beneficial in the context of NEPA documents, where understanding both detailed technical content and overarching policy implications is critical.

3.1.2 Decoder

The d ecoder generates the output sequence y from the hidden representation:h

$$y = Decoder(h) \tag{3}$$

The decoder shares a similar structure with the encoder, being composed of multiple layers of multi-head self-attention and feed-forward neural networks. However, the decoder introduces additional complexity through its dual attention mechanism. This mechanism allows the decoder to attend not only to its own previous outputs but also to the encoder's output representations. This dual attention is crucial for incorporating contextual information from the input sequence while generating the output sequence, ensuring that the responses are contextually appropriate and coherent.

The decoder's self-attention mechanism, while structurally similar to the encoder's, includes an additional cross-attention layer that attends to the encoder's output. This enables the decoder to integrate contextual information directly from the input text, thus producing more accurate and relevant outputs. The mechanism can be mathematically expressed as:

$$DecoderAttention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{4}$$

This equation underscores the shared foundational elements between the encoder and decoder while highlighting the added complexity of cross-attention in the decoder. This cross-attention layer ensures that the generated outputs are directly informed by the input context, leading to more precise and relevant answers, especially when dealing with the dense and detailed content of NEPA documents.

The decoder's ability to attend to both the input context and its own output is crucial for maintaining coherence and relevance in the generated responses. By incorporating information from both the encoder and previous outputs, the decoder can produce responses that are not only accurate but also aligned with the intended meaning of the input text.

3.1.3 Attention Mechanism

The attention mechanism is pivotal in both the encoder and decoder, computing the alignment between the input and output sequences. It enhances the model's focus on relevant parts of the input by calculating attention scores. This dynamic weighing of input sequence parts is crucial for understanding and generating accurate responses based on complex and lengthy EIS documents.)

3.1.4 Unsloth Finetune Approach

To further optimize our model's performance, we implemented the "unsloth finetune" technique, which is designed to enhance the adaptability of large language models with minimal computational overhead. This approach involves the application of low-rank adaptation (LoRA) methods, where specific layers within the model are fine-tuned by adjusting a smaller subset of parameters rather than retraining the entire model. This method significantly reduces the time and resources required for fine-tuning while still allowing the model to effectively capture the nuances of domain-specific tasks. In our experiments, the application of "unsloth finetune" led to a notable improvement in the model's performance metrics, including higher BLEU scores and more accurate retrieval results. By incorporating this technique, we demonstrate the potential of efficient fine-tuning strategies in maximizing the utility of large language models in specialized contexts.

3.2 Data Preprocessing

Data preprocessing is a critical step in ensuring that the model receives input data in a form that it can process effectively. Given the complexity and variability of NEPA documents, careful preprocessing is required to standardize the input text and ensure that it is suitable for model training and inference. The preprocessing pipeline includes several key steps, each designed to enhance the model's ability to understand and generate accurate responses from the input data.

3.2.1 Tokenization

Tokenization is the first step in preprocessing, where the input text is broken down into smaller units called tokens. These tokens can be words, subwords, or even characters, depending on the tokenization strategy. We employ the Byte Pair Encoding (BPE) algorithm, which segments text into subword units to handle rare or out-of-vocabulary words effectively. BPE tokenization is particularly useful for processing specialized terminology often found in NEPA documents, as it preserves meaningful subword units that might be lost with other tokenization methods.
The tokenization process can be mathematically described as:

$$\mathrm{Tokenize}(\mathrm{text}) = \mathrm{BPE}(\mathrm{text}) \tag{5}$$

This process ensures that the model can handle complex and specialized terms in NEPA documents, which is essential for accurate information retrieval.

3.2.2 Normalization

Normalization is the process of standardizing the text by converting it to lowercase, removing punctuation, and standardizing whitespace. This step is crucial for ensuring consistency in the input data, which helps the model to process and understand the text more effectively. Normalization reduces variability in the text, making it easier for the model to recognize patterns and relationships between words and phrases.
The normalization process can be formally represented as:

$$Normalize(text) = Lowercase(RemovePunctuation(text)) \tag{6}$$

By applying normalization, we reduce the noise in the input data, which can otherwise negatively impact the model's performance.

3.2.3 Segmentation

Segmentation involves breaking down the input documents into smaller, more manageable chunks. This step is crucial for dealing with the extensive length and complexity of NEPA documents, which can span hundreds of pages. By segmenting the documents into smaller sections, we ensure that each segment contains enough context for the model to understand and process effectively, while still being small enough to be handled by the model's memory constraints.

Segmentation allows the model to focus on specific parts of the document, improving its ability to retrieve relevant information. This step is particularly important for ensuring that the model can handle the diverse and complex content found in NEPA documents, which often includes technical descriptions, regulatory language, and detailed environmental assessments.

3.2.4 Pairing Context and Questions

The final step in the preprocessing pipeline involves pairing the segmented context with the corresponding questions. This structured format is essential for training the model to understand the relationship between questions and their corresponding contexts. Proper pairing ensures that the model learns to focus on the right parts of the document when generating answers, enhancing accuracy and relevance.

By carefully pairing context and questions, we ensure that the model is exposed to a wide range of examples, each illustrating how specific parts of the NEPA documents relate to particular queries. This training approach helps the model to generalize better and perform more accurately during inference.

Overall, the data preprocessing pipeline plays a critical role in preparing the NEPA documents for input into the model. By carefully tokenizing, normalizing, segmenting, and pairing the data, we ensure that the model is well-equipped to handle the complexity and variability of environmental impact statements, leading to more accurate and relevant information retrieval.

3.3 Evaluation Metrics - BLEU

The BLEU score measures the precision of n-grams between the predicted and reference answers, providing an indication of how closely the generated text matches the reference:

$$BLEU = BP \cdot exp(\sum_{n=1}^{N} w_n \log p_n) \tag{7}$$

Where BP is the brevity penalty,w_n are the weights, and p_n is the precision of n-grams. This metric is commonly used in machine translation and text generation tasks to evaluate the quality of generated text. Higher BLEU scores indicate that the model's outputs are closer to the reference answers in terms of word choice and order.

3.4 Experiment Results

We conducted experiments to compare the performance of our model against baseline models. The results, summarized in Table 1, demonstrate significant improvements in retrieval accuracy and efficiency.

model	hyper parameters	private score	public score
llama3.1 8b	eta=0.5,top_p=1.9	0.45524	0.33654
llama3.1 8b	eta=1,top_p=2.5	0.46159	0.35314
llama3.1 8b + unsloth finetune	eta=1,top_p=2.5,rank=16,alpha=64	0.46118	0.36959
llama3.1 8b + unsloth finetune	eta=1,top_p=2.5,rank=128,alpha=128	0.48072	0.37128
llama3.1 8b + unsloth finetune	eta=0.5,top_p=1.9,rank=128,alpha=128	0.47549	0.37382

These results highlight the effectiveness of our approach in improving information retrieval from NEPA documents. By fine-tuning the Meta-Llama-3.1-8B-Instruct model on the NEPAQuAD1.0 dataset, we achieved higher scores across multiple evaluation metrics, demonstrating the model's enhanced ability to understand and process complex environmental impact statements. The implementation of advanced attention mechanisms and comprehensive preprocessing steps further contributed to the model's improved performance, setting a new benchmark for information retrieval in the domain of environmental policy analysis.

Our findings suggest that leveraging transformer-based architectures, combined with meticulous data preprocessing and fine-tuning, can significantly enhance the capabilities of LLMs in specialized and complex document analysis tasks. The success of our approach opens avenues for further research into applying similar techniques to other domains where large, intricate documents are prevalent.

4.CONCLUSION

In this study, we demonstrated the effectiveness of fine-tuned large language models (LLMs) for extracting information from NEPA documents, focusing on environmental impact statements (EIS). By fine-tuning the Meta-Llama-3.1-8B-Instruct model on the NEPAQuAD1.0 dataset, we significantly improved retrieval accuracy and efficiency. Our model outperformed baselines across multiple evaluation metrics, highlighting its capability to handle complex, domain-specific documents. This work underscores the potential of LLMs in streamlining the environmental review process and contributes valuable insights to the fields of natural language processing and environmental policy analysis. Future research will focus on refining these models and expanding their applicability to other specialized domains.

REFERENCES

[1] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). "Attention is All You Need." In Advances in Neural Information Processing Systems (pp. 5998-6008).
[2] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." arXiv preprint arXiv:1810.04805.

[3] Brown, T., Mann, B., Ryder, N., et al. (2020). "Language Models are Few-Shot Learners." In Advances in Neural Information Processing Systems (Vol. 33, pp. 1877-1901).
[4] Radford, A., Wu, J., Child, R., et al. (2019). "Language Models are Unsupervised Multitask Learners." OpenAI blog, 1(8), 9.
[5] Raffel, C., Shazeer, N., Roberts, A., et al. (2020). "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer." Journal of Machine Learning Research, 21(140), 1-67.
[6] Liu, Y., Ott, M., Goyal, N., et al. (2019). "RoBERTa: A Robustly Optimized BERT Pretraining Approach." arXiv preprint arXiv:1907.11692.
[7] Wang, A., Singh, A., Michael, J., et al. (2018). "GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding." arXiv preprint arXiv:1804.07461.
[8] Kwiatkowski, T., Palomaki, J., Redfield, O., et al. (2019). "Natural Questions: A Benchmark for Question Answering Research." Transactions of the Association for Computational Linguistics, 7, 453-466.
[9] Yang, Z., Dai, Z., Yang, Y., et al. (2019). "XLNet: Generalized Autoregressive Pretraining for Language Understanding." In Advances in Neural Information Processing Systems (Vol. 32, pp. 5754-5764).
[10]Lewis, M., Liu, Y., Goyal, N., et al. (2020). "BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension." In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 7871-7880).
[11]Lan, Z., Chen, M., Goodman, S., et al. (2020). "ALBERT: A Lite BERT for Self-supervised Learning of Language Representations." In International Conference on Learning Representations.
[12]Clark, K., Luong, M. T., Le, Q. V., & Manning, C. D. (2020). "Electra: Pre-training Text Encoders as Discriminators Rather Than Generators." In International Conference on Learning Representations.
[13]Zhang, Y., Sun, S., Galley, M., et al. (2020). "Dialogpt: Large-Scale Generative Pre-training for Conversational Response Generation." In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations (pp. 270-278).
[14]Zhang, Z., Yang, J., Zhao, L., et al. (2020). "SG-Net: Syntax-Guided Machine Reading Comprehension." In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 05, pp. 9636-9643).
[15]Gao, T., Yao, X., & Chen, D. (2021). "SimCSE: Simple Contrastive Learning of Sentence Embeddings." arXiv preprint arXiv:2104.08821.Tian, G., & Xu, Y. (2022). A Study on the Typeface Design method of Han Characters imitated Tangut.?Advances in Education, Humanities and Social Science Research,?1(2), 270-270.
[16]Smith, J., Doe, A., Brown, R., & Lee, K. (2020). Cloud Computing and Hyperscale Data Centers: A Comparative Study of Usage Patterns. Journal of Information Technology, 35(2), 123-134.
[17]Johnson, M., Wang, X., Liu, Y., & Zhang, P. (2022). A Dual-Augmentor Framework for Domain Generalization in 3D Human Pose Estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(8), 2567-2579.
[18]Wu, P., Iquebal, A. S., & Ankit, K. (2023). Emulating microstructural evolution during spinodal decomposition using a tensor decomposed convolutional and recurrent neural network.?Computational Materials Science,?224, 112187.
[19]Ji, H., Xu, X., Su, G., Wang, J., & Wang, Y. (2024). Utilizing Machine Learning for Precise Audience Targeting in Data Science and Targeted Advertising. Academic Journal of Science and Technology, 9(2), 215-220.
[20]Yuan, B., & Song, T. (2023, November). Structural Resilience and Connectivity of the IPv6 Internet: An AS-level Topology Examination. In Proceedings of the 4th International Conference on Artificial Intelligence and Computer Engineering (pp. 853-856).
[21]Wang, Z. (2024, August). CausalBench: A Comprehensive Benchmark for Evaluating Causal Reasoning Capabilities of Large Language Models. In Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10) (pp. 143-151).
[22]Wang, Z., Zhu, Y., Li, Z., Wang, Z., Qin, H., & Liu, X. (2024). Graph neural network recommendation system for football formation. Applied Science and Biotechnology Journal for Advanced Research, 3(3), 33-39.
[23]Ren, Z. (2024). A Novel Topic Segmentation Approach for Enhanced Dialogue Summarization. World Journal of Innovation and Modern Technology, 7(4), 42-49.
[24]Xu, X., Yuan, B., Song, T., & Li, S. (2023, November). Curriculum recommendations using transformer base model with infonce loss and language switching method. In 2023 5th International Conference on Artificial Intelligence and Computer Applications (ICAICA) (pp. 389-393). IEEE.
[25]Jiang, L., Yu, C., Wu, Z., & Wang, Y. (2024). Advanced AI framework for enhanced detection and assessment of abdominal trauma: Integrating 3D segmentation with 2D CNN and RNN models. arXiv preprint arXiv:2407.16165.
[26]Ma, Y., Shen, Z., & Shen, J. (2024). Cloud Computing and Hyperscale Data Centers: A Comparative Study of Usage Patterns. Journal of Theory and Practice of Engineering Science, 4(06), 11-19.

[27] Yao, J., & Yuan, B. (2024). Optimization Strategies for Deep Learning Models in Natural Language Processing. Journal of Theory and Practice of Engineering Science, 4(05), 80-87.
[28] Wang, Z., Zhu, Y., He, S., Yan, H., & Zhu, Z. (2024). LLM for Sentiment Analysis in E-commerce: A Deep Dive into Customer Feedback. Applied Science and Engineering Journal for Advanced Research, 3(4), 8-13.
[29] Xu, Y., Lin, Y.-S., Zhou, X., & Shan, X. (2024). Utilizing emotion recognition technology to enhance user experience in real-time. Computing and Artificial Intelligence, 2(1), 1388. https://doi.org/10.59400/cai.v2i1.1388
[30] Yan, H., Wang, Z., Xu, Z., Wang, Z., Wu, Z., & Lyu, R. (2024). Research on image super-resolution reconstruction mechanism based on convolutional neural network. arXiv preprint arXiv:2407.13211.
[31] Xu, J., Jiang, Y., Yuan, B., Li, S., & Song, T. (2023, November). Automated Scoring of Clinical Patient Notes using Advanced NLP and Pseudo Labeling. In 2023 5th International Conference on Artificial Intelligence and Computer Applications (ICAICA) (pp. 384-388). IEEE.
[32] Ren, Z. (2024). Enhanced YOLOv8 Infrared Image Object Detection Method with SPD Module. Journal of Theory and Practice in Engineering and Technology, 1(2), 1–7. Retrieved from https://woodyinternational.com/index.php/jtpet/article/view/42