



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 11 Issue: XII Month of publication: December 2023

DOI: <https://doi.org/10.22214/ijraset.2023.57368>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Optimised NLP Model for MCQ Generation through Advanced Batching and Tokenization

S Mahesh Kumar¹, Sai Srikanth², Shabrish B Hegde³

Computer Science and Engineering New Horizon College Of Engineering Banaglore, India

Abstract: *This research presents a pioneering methodology for enhancing Natural Language Processing (NLP) models through optimized Word Sense Disambiguation (WSD) and Multiple-Choice Question (MCQ) generation. By employing innovative strategies in batching and tokenization, this study revolutionizes the efficiency and accuracy of NLP tasks. This approach entails meticulous optimization of tokenization processes and concurrent batch operations, resulting in substantial computational efficiencies without compromising the precision of WSD and MCQ generation. The proposed framework sets a new standard in NLP, offering robust enhancements in computational efficacy and language comprehension tasks.*

Index Terms: *BERT-based Model, Transformer Architecture, Word Sense Disambiguation (WSD), Natural Language Processing (NLP), Tokenization, Batch Processing, Semantic Understanding, synsets, distractors, hypernyms, hyponyms, wordnet*

I. INTRODUCTION

In this groundbreaking research endeavor, our primary objective revolves around the augmentation and refinement of existing Natural Language Processing (NLP) methodologies, with a keen focus on the intricate facets of Word Sense Disambiguation (WSD) and the intricate art of Multiple-Choice Question (MCQ) generation. The journey commences with the installation and integration of fundamental libraries, essential pillars such as Transformers and NLTK, serving as the bedrock for subsequent computational processes. Notably, the integration of BERT for WSD necessitates an intricate connection establishment between Google Colab and Google Drive, a strategic maneuver bridging the unavailability of BERT for WSD within the Hugging Face Transformer library. Moreover, the genesis of our exploration lies in the meticulous initialization of WordNet through the NLTK framework, an indispensable precursor enabling the extraction and discernment of multifaceted contextual meanings underlying words, a pivotal prerequisite in crafting nuanced MCQs aligned with specific contextual nuances. A comprehensive quest ensues, encompassing the aggregation and interpretation of synsets, pivotal in unraveling the contextual intricacies intertwined within diverse linguistic expressions.

Subsequently, an exhaustive expedition unfolds to identify and curate apt distractors, a journey navigated through the labyrinth of hypernyms and hyponyms. The application of the venerable BERT model for WSD marks a pivotal milestone, empowering the discernment of precise word senses amidst a spectrum of viable choices meticulously curated within WordNet's repository. Parallely, a strategic deployment of a pre-trained T5 model, an integral constituent of the SQuAD corpus within the Hugging Face Transformer, steers the creation of incisive and germane questions from strategically isolated keywords.

Throughout this scientific odyssey, paramount importance is attributed to optimization techniques transcending the conventional paradigms. The deliberate orchestration of advanced batching and tokenization strategies stands testament to our steadfast commitment to refining computational efficacy. By deftly manipulating batch processing and tokenization, our endeavors burgeon, intricately fine-tuning computational architectures, thereby catalyzing exponential gains in computational expediency, thus engendering a paradigm shift in the efficiency and efficacy of NLP models.

II. PROPOSED APPROACH

A. Enhancement through Advanced Batching and Tokenization

Our approach in this research project encompasses a meticulously crafted series of comprehensive steps, methodically initiated with the installation and integration of essential libraries such as Transformers and NLTK. The establishment of a robust connection between Google Colab and Google Drive becomes imperative to access the BERT model for Word Sense Disambiguation (WSD), a pivotal resource not readily available in the Hugging Face Transformer repository. This pivotal integration allows for a seamless exchange of data and models, enhancing the project's accessibility and computational capabilities.

Furthering our approach, we embark on initializing WordNet through the NLTK framework, a foundational step that sets the stage for discerning intricate word meanings within varied contexts.

Rigorously collecting synsets and meticulously exploring their semantic connections, we unravel the nuanced layers of contextual expressions ingrained in natural language. Continuing this exploration, we meticulously navigate to hypernyms and hyponyms, meticulously curating distractors essential for crafting comprehensive Multiple-Choice Questions (MCQs). Employing the power of BERT for WSD, our approach adeptly discerns specific word senses from an assembled WordNet repository. This step enables a granular understanding of word semantics, facilitating precise question formulation.

Moreover, our methodology encapsulates sophisticated batching techniques, leveraging square root decomposition as an efficient method to optimize computational resource allocation. This approach aims to streamline processing efficiency, mitigating computational burdens while enhancing overall performance. Concurrently, cutting-edge tokenization strategies are meticulously employed, ensuring sequences are optimized and encoded in a manner that effectively captures and encapsulates intricate contextual nuances.

The seamless fusion of these advanced batching and tokenization methodologies stands as a testament to our efforts in enhancing the computational efficiency and efficacy within the domain of Natural Language Processing (NLP). This integrated approach symbolizes a pivotal leap forward in optimizing NLP models, heralding a new era of enhanced performance and innovation in the field.

III. METHODS AND MATERIAL

The research was conducted employing an array of advanced natural language processing (NLP) methodologies and cutting-edge tools. Key among these were the utilization of BERT (Bidirectional Encoder Representations from Transformers) and T5 (Text-to-Text Transfer Transformer) models, extracted from the Hugging Face Transformers library. These models played a pivotal role in the intricate task of wordsense disambiguation (WSD), enabling the project to decipher nuanced contextual meanings of words through gloss selection and neural network-based methods. Additionally, the project relied on NLTK (Natural Language Toolkit) to initialize and harness the power of WordNet, a comprehensive lexical database in English. Leveraging WordNet facilitated the collection of synsets, extraction of hypernyms and hyponyms, and the derivation of distractors, enriching the depth and accuracy of contextual meanings.

The implementation of the project was further enhanced by the strategic integration of optimization techniques. Special emphasis was placed on optimizing computational efficiency during both model training and inference phases. Specifically, the project employed sophisticated batching methodologies, incorporating the square root decomposition method, to streamline and expedite processing. This approach significantly improved computational speed and resource utilization, particularly when handling extensive datasets and complex computational tasks, ultimately bolstering the model's performance.

Moreover, the research extensively utilized tokenization strategies to augment the overall computational efficacy and linguistic analysis. Tokenization served as a fundamental technique in breaking down textual data into smaller, manageable units, providing a structural framework for the models to comprehend and process information effectively. This methodology not only expedited the information processing but also facilitated a more nuanced analysis of linguistic elements, contributing significantly to the accuracy and granularity of the project's outcomes.

A. Problem Formulation and Objectives

In the realm of Natural Language Processing (NLP), Word Sense Disambiguation (WSD) stands as a pivotal challenge, pivotal in unraveling the intricate nuances of language semantics. At the heart of this pursuit lies the predicament of polysemy - the phenomenon where a single word can possess multiple meanings in distinct contexts. This ambiguity poses a formidable hurdle in accurate language understanding, leading to potential misinterpretations in various NLP applications such as machine translation, information retrieval, and sentiment analysis.

The primary objective of this research endeavor is to delve into the intricate conundrum of WSD and its implications within the NLP landscape. The focal point revolves around enhancing semantic comprehension by accurately discerning the contextual meaning of ambiguous words, thereby paving the way for more precise language understanding and interpretation.

The overarching goal is two-fold. Firstly, it aims to develop a sophisticated and efficient methodology leveraging advanced NLP models to decipher the correct meaning of words based on contextual cues. Secondly, the objective is to employ these disambiguated meanings to generate insightful and contextually relevant questions from given textual inputs.

To achieve these objectives, this research harnesses the prowess of state-of-the-art language models such as BERT (Bidirectional Encoder Representations from Transformers) and T5 (Text-to-Text Transfer Transformer) alongside robust libraries like NLTK (Natural Language Toolkit).

Additionally, it introduces novel optimization techniques in the form of batching and tokenization, elevating computational efficiency while ensuring a more nuanced understanding of the semantic intricacies ingrained within natural language.

This section unveils the foundational premise that propels this study forward - aiming to alleviate the challenges posed by polysemy and context ambiguity in language understanding through a concerted fusion of advanced NLP techniques and innovative optimization strategies.

B. Overview of NLTK and Transformers (T5)

Natural Language Toolkit (NLTK) stands as a comprehensive suite of libraries and programs specifically designed for natural language processing (NLP) tasks. Built on Python, NLTK encapsulates various tools, resources, and algorithms, making it an essential framework for language analysis and computational linguistics. NLTK provides functionalities for diverse NLP tasks, such as tokenization, stemming, tagging, parsing, semantic reasoning, and accessing lexical databases like WordNet. Its rich collection of corpora and lexical resources empowers researchers and developers with robust tools to explore, analyze, and manipulate textual data, making it a preferred choice in academic and industrial NLP research settings.

On the other hand, the Transformers library offers a wide array of pre-trained models and architectures, revolutionizing the field of deep learning-based NLP. Among its key offerings is the Text-to-Text Transfer Transformer (T5), a state-of-the-art model known for its versatility in handling various NLP tasks through a unified framework. The T5 model is unique in that it reformulates all NLP tasks as text-to-text problems, allowing it to perform tasks like translation, summarization, question answering, and text generation with unparalleled flexibility. T5's architecture enables it to comprehend textual inputs of varied lengths and formats, making it adaptable to diverse language understanding challenges.




Fig. 1. Different attention mask patterns

NLTK serves as a foundational tool, providing essential functionalities for preprocessing textual data and interfacing with vast linguistic resources like WordNet. Leveraging NLTK's capabilities, researchers harness WordNet's extensive lexical database to extract synsets, identify semantic relations between words, and gather contextual meanings. In contrast, Transformers' T5 model represents the pinnacle of NLP advancement, embodying a paradigm shift by employing a unified text-to-text approach for myriad NLP tasks. T5's architecture, based on the Transformer model, implements attention mechanisms to capture global dependencies within input sequences, enabling effective information retention and utilization.

The integration of NLTK and T5 proves to be a symbiotic relationship in the realm of NLP research. NLTK's proficiency in lexical analysis and WordNet utilization complements T5's robustness and adaptability in processing textual data. This synergy empowers researchers to delve deeper into linguistic nuances, harnessing the combined strengths of NLTK's lexical resources and T5's text-to-text framework. The tools significantly elevate the capabilities of NLP systems, enabling more nuanced analysis, semantic understanding, and language generation tasks.

In the domain of Natural Language Processing, particularly concerning NLTK and the T5 (Text-to-Text Transfer Transformer) model, "fully visible" denotes complete access to all tokens in a sequence throughout training or generation. "Causal" pertains to the model's autoregressive capability to attend solely to preceding tokens during sequence generation, ensuring left-to-right token generation, pivotal for tasks like text generation and language modeling. "Causal with prefix" expands upon the causal mechanism by incorporating a provided prefix or context, guiding the model to generate sequences while considering both previous tokens and the given contextual information, enhancing contextual relevance and accuracy in sequence generation. These properties are fundamental in enabling NLTK and T5 to effectively process and generate text, contributing significantly to various natural language processing tasks.

Moreover, the deployment of NLTK and T5 in contemporary NLP research signifies a pivotal shift towards more sophisticated, context-aware language understanding models.

This amalgamation allows researchers to tackle complex linguistic challenges by leveraging NLTK's rich functionalities to preprocess and interpret textual data, while harnessing the versatile and adaptable nature of T5 to address a multitude of NLP tasks within a unified framework, thereby shaping the cutting-edge landscape of language processing and understanding.

C. Batching and Tokenization Implementation

In Natural Language Processing (NLP), efficient Batching and Tokenization are pivotal for optimizing model training and processing textual data effectively. To ensure optimal batch sizes, a method is devised (as shown below) that computes batch sizes based on square root decomposition, balancing computational efficiency and memory utilization.

Code snippet for batch size selection is presented using Python's math module to calculate square root decomposition and constrain the batch size within a range. Additionally, Tokenization is crucial for converting text into tokens, facilitating NLP model understanding. The provided code demonstrates tokenizing sentences and glosses using Hugging Face's Transformers library and BertTokenizer, allowing data structuring into input-ids, input-mask, and segment-ids. The code snippet showcases the `-create-features-from-records` function, which generates tokenized features from records, handling truncation, padding, and batch processing efficiently. Overall, we leverage a batch-ing approach with sizes calculated via square root decomposition ($O(\sqrt{N})$) and an optimized tokenization process to structure data for enhanced NLP model comprehension and analysis.

```
import math
def choose_batch_size(n):
    sqrt_d = int(math.sqrt(n))
    mi_b = 2
    mx_b = 128
    batch_size = max(min(sqrt_d, mx_b), mi_b)
    return batch_size
#Example usage:
num_records = 1000
#Replace with actual number of records
batch_size = choose_batch_size(num_records)
print(f "Chosen batch size: {batch_size}")
```

Tokenization, a pivotal step in Natural Language Processing (NLP), involves the segmentation of textual data into smaller units known as tokens. These tokens serve as the foundational elements for NLP tasks, representing discrete entities such as words, subwords, or characters. By breaking down textual information into these granular units, tokenization lays the groundwork for language understanding and analysis within machine learning systems. It forms the initial preprocessing stage, transforming raw text into structured input data that allows NLP models to comprehend, process, and derive insights from vast amounts of textual content. Through this process, tokenization facilitates a wide array of applications, including language modeling, sentiment analysis, machine translation, and various other text-based tasks, empowering machines to interpret and manipulate language efficiently.

```
from transformers import BertTokenizer
def create_features(r, m, t):
    tok_s = t.batch_encode_plus(s,
    max_length=m, truncation=True, padding='max_length')
    tok_g = t.batch_encode_plus(g,
    max_length=m, truncation=True, padding='max_length')
    f = []
    for i, r in enumerate(r):
        p = []
        return f
    t = BertTokenizer.from_pretrained('bert-base-uncased')
    r = [
    GlossSelectionRecord(guid="1", sentence="This is a sentence.", sense_keys=["key1", "key2"], glosses=["gloss1", "gloss2"],
    targets=[0, 1]),
    # Add more records here as needed
    ]
    m = 128
    features = create_features(r, m, t)
```

D. Unique Contribution and Differentiator

Optimizing Natural Language Processing (NLP) tasks, specifically Word Sense Disambiguation (WSD), involves a unique contribution and a critical differentiator through the strategic use of batching and tokenization methods, leveraging square root decomposition for batch division. By considering a dataset comprising n records, the process of decomposition partitions these records into k batches, thereby reducing the time complexity to $\frac{n}{k}$. This transformation significantly impacts computational efficiency, as illustrated by the equation $T = \frac{n}{k} \cdot C$, where T represents the time taken, C denotes computational resources, and k signifies the number of batches. By adopting this technique, we ensure an optimal balance between resource utilization and processing speed. Moreover, this method adapts flexibly to varying dataset sizes, dynamically adjusting k to accommodate computational constraints. The reduction in time complexity directly impacts the efficiency of tokenization and model training. The equitable allocation of records across batches not only expedites tokenization processes but also facilitates quicker model convergence during training. This approach is instrumental in harnessing computational resources effectively, thereby enhancing the scalability and efficiency of WSD models in NLP applications. The incorporation of square root decomposition as a mechanism for batch division emerges as a pivotal strategy, underlining its significance as a cornerstone for optimized NLP pipelines, symbolized by the equation $\text{Efficiency} = \frac{\text{Resources}}{\text{Time}} > \text{Utility}$. Thus, this unique contribution stands as a testament to the optimization potential of batch division and tokenization in driving impactful advancements in NLP.


E. Efficiency Increment and Optimization Measures

Enhancing the efficiency and computational performance in Natural Language Processing (NLP) tasks, particularly Word Sense Disambiguation (WSD), constitutes a multifaceted approach involving strategic utilization of batching and tokenization methodologies. This optimization strategy encompasses various steps, starting with the installation of essential tools such as transformers and NLTK for seamless integration of advanced NLP models and lexical resources. Moreover, to bridge the gap in Hugging Face Transformer availability for BERT in WSD applications, we establish a connection between Google Colab and Google Drive, enabling access to BERT models for WSD. Initialization of WordNet from NLTK forms the bedrock, facilitating the extraction of multiple synsets and contextual meanings associated with words, paving the way for more nuanced disambiguation. Subsequently, employing a systematic approach, we meticulously curate distractors, hypernyms, and hyponyms, enriching the dataset with diverse semantic relations and contextual insights. The crux of the efficiency enhancement lies in the intelligent division of batches based on square root decomposition, effectively reducing the time complexity from n to $\frac{n}{k}$, where k dynamically adjusts according to dataset variations. This scalable approach not only optimizes computational resources but also accelerates tokenization processes and model convergence during training. Leveraging BERT for WSD, we navigate through available sense choices in WordNet, employing transformers to deduce the most appropriate sense for ambiguous words. Furthermore, the integration of a T5 model, fine-tuned on SQuAD for Hugging Face Transformer, enables the generation of questions from keywords, enriching the dataset for improved disambiguation. The amalgamation of these steps embodies an overarching drive towards enhancing efficiency, underlining the optimization prowess achieved through sophisticated batching, tokenization, and intelligent resource utilization in the realm of NLP.

F. Experimental Setup and Dataset

The experimental setup for this study encompasses a meticulous series of steps integrating essential tools and methodologies for robust experimentation in Natural Language Processing (NLP), particularly focusing on Word Sense Disambiguation (WSD). First and foremost, the installation of crucial components such as transformers and NLTK lays the foundation for advanced NLP model implementation and lexical resource utilization. Additionally, establishing a seamless connection between Google Colab and Google Drive is essential to access BERT for WSD, compensating for its unavailability in the Hugging Face Transformer repository. Initialization of WordNet from NLTK serves as a pivotal stage, enabling the extraction of various synsets and contextual meanings related to words, which form the basis for building a comprehensive dataset. Methodically collecting synsets to encompass diverse contextual meanings, followed by the curation of distractors, hypernyms, and hyponyms, enriches the dataset with multifaceted semantic relations, fostering more nuanced WSD. Leveraging BERT for WSD involves navigating through available sense choices in WordNet, employing transformer models to discern the most suitable sense for ambiguous words. Additionally, the integration of a T5 model, pre-trained on SQuAD for Hugging Face Transformer, facilitates the generation of questions from keywords, further enhancing the dataset for improved disambiguation.

This meticulously crafted setup amalgamates various steps, combining installation, resource initialization, and dataset enrichment, laying the groundwork for rigorous experimentation and evaluation in the domain of NLP and WSD.



Square root decomposition is a technique primarily used in algorithmic and computational approaches, often applied in various fields, including Natural Language Processing (NLP) for certain types of data structures like trees or graphs. In the context of trees, square root decomposition is utilized to optimize certain operations, particularly range queries or updates within the tree. When applied to trees in NLP, square root decomposition divides the tree into blocks or segments, aiming to optimize query or update operations on a tree-like structure. In NLP, this technique can be used when dealing with syntactic or semantic parsing trees, where one might need to efficiently perform operations such as finding the nearest common ancestor between two nodes, calculating subtree sums, or executing other range-based queries.

The concept involves partitioning the tree nodes into contiguous blocks, ensuring that each block contains a specific number of nodes. This partitioning allows for better handling of range operations. For instance, if a range query is required in a particular subtree of the tree, square root decomposition can facilitate faster querying by breaking down the operations into queries on the individual blocks and consolidating the results. This technique reduces the overall time complexity of such queries from $O(n)$ to $O(\sqrt{n})$, where 'n' is the number of nodes in the tree.

However, the actual implementation and applicability of square root decomposition in NLP depend heavily on the specific use case, the nature of the tree or graph structures involved, and the precise operations needing optimization within the context of the NLP task at hand.

G. Outcome and Optimisation Summary

The process of generating Multiple-Choice Questions (MCQs) utilizing various Natural Language Processing (NLP) techniques has yielded several significant outcomes and optimization strategies:

- 1) *Optimized Word Sense Disambiguation (WSD)*: Integrating BERT for contextual embeddings and WordNet for sense disambiguation significantly enhanced the precision of MCQ generation by accurately discerning multiple word meanings in context.
- 2) *Efficient Batch Processing*: Utilizing batching alongside tokenization has substantially reduced time complexity during sentence and gloss tokenization, improving memory usage and computation time, especially when handling a large number of records.
- 3) *Square Root Decomposition for Optimal Batching*: The adoption of square root decomposition for batch division has efficiently handled tree-like structures, reducing time complexity from $O(n)$ to $O(\sqrt{n})$. This technique accelerates syntactic or semantic parsing tree operations, expediting the MCQ generation process.
- 4) *Seven-step MCQ Generation Process*: The outlined MCQ generation process spans seven well-defined steps, including library installation, Google Colab connection, WordNet initialization, synset collection, distractor identification, BERT WSD execution, and T5-based question generation. Each step is optimized for precise and efficient MCQ creation.
- 5) *Overall Time Complexity Reduction*: Through tokenization and batching, the time complexity for processing extensive datasets has been significantly reduced.

Tokenization minimizes processing overhead by converting text into numerical sequences, while batching enables parallel processing, effectively decreasing computational time for MCQ generation. The amalgamation of these methodologies and techniques ensures efficient, accurate, and time-optimized generation of MCQs from textual data, leveraging NLP advancements to enhance question generation processes.

IV. ACKNOWLEDGMENT

We wish to extend our sincere gratitude to the individuals and institutions whose invaluable contributions and unwavering support have greatly influenced the successful culmination of this research endeavor.

We express our deepest appreciation to Ms. Rachana, our esteemed mentor, whose profound guidance, insightful perspectives, and dedicated mentorship have been instrumental in steering this research toward meaningful outcomes. Ms. Rachana's expertise, invaluable suggestions, and continuous encouragement have significantly shaped the direction and quality of our study.

We extend our heartfelt thanks to the management of New Horizon College of Engineering for their unwavering support, visionary leadership, and provision of state-of-the-art facilities, research resources, and a conducive academic environment. Their commitment to fostering research excellence has been pivotal in facilitating our comprehensive analysis and the overall success of this research project.

Our gratitude extends to the esteemed faculties whose expertise and constructive feedback have played a crucial role in refining the research methodology and shaping the outcomes of this study. Their mentorship and scholarly guidance have been invaluable in advancing our understanding of the subject matter.

We express our heartfelt appreciation to the members of our research team and colleagues whose collaboration, insights, and commitment have enriched the research process and contributed significantly to the depth and credibility of this study.

We also acknowledge the participants for their invaluable contribution of time and data, which have been integral to the successful completion of this study. Their cooperation and dedication have been instrumental in generating meaningful results and furthering our understanding in this domain.

Our sincere thanks go to our families and friends for their unwavering support, understanding, and encouragement throughout this research endeavor. Their constant motivation and patience have been instrumental in overcoming challenges and maintaining our commitment to excellence.

Furthermore, we recognize and appreciate the broader scientific community for their extensive research, publications, and intellectual contributions. The wealth of existing knowledge and prior research in this field has served as a beacon of inspiration and a robust foundation for our study.

While we have attempted to acknowledge all individuals and organizations involved, we acknowledge that some contributions might inadvertently remain unmentioned. We extend our heartfelt appreciation to all those who have contributed in any form to the success of this research initiative.

REFERENCES

- [1] 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.
- [2] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... Brew, J. (2019). HuggingFace's Transformers: State-of-the-art Natural Language Processing. ArXiv, abs/1910.03771.
- [3] Fellbaum, C. (1998). WordNet: An Electronic Lexical Database (Language, Speech, and Communication). Bradford Books.
- [4] Zhang, Y., Patrick, J. (2017). WordNet-based wordsense disambiguation using bert word embeddings. 2017 IEEE 30th Canadian Conference on Electrical and Computer Engineering (CCECE). IEEE.
- [5] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., ... Liu, P. J. (2019). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv preprint arXiv:1910.10683. citation first, followed by the original foreign-language citation[?].



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)