

An Efficient Deep Learning Model for Sequential Sentence Classification of Medical Abstracts

Lalan Kumar, Student, Hyderabad Institute of Technology and Management, Hyderabad, India,
lalan.k8977@gmail.com

Bhaskar Das, Associate Professor, Hyderabad Institute of Technology and Management,
Hyderabad, India, bhaskardas.cse@hitam.org

K. Chaitanya Venkat, Student, Hyderabad Institute of Technology and Management, Hyderabad,
India, venkatchaitu477@gmail.com

Abstract In medical research, the rising volume of randomized controlled trial (RCT) papers makes it challenging for researchers to quickly find and extract key information. Many abstracts are unstructured, which complicates skimming and synthesis of critical content. This project aims to address this issue by developing a natural language processing (NLP) model to classify sentences in medical abstracts by their functional roles—such as objective, methods, results, and conclusions. Using core NLP techniques like tokenization and word embeddings, this model organizes unstructured text into clear, labeled sections, improving navigation and comprehension. For example, abstracts can be segmented into easily identifiable parts, enabling quicker understanding of the study's focus and findings. By converting dense, complex abstracts into structured formats, SkimLit seeks to streamline the research process, allowing scholars to efficiently locate and assess pertinent information.

Keywords — *Natural Language Processing (NLP), Randomized Controlled Trials (RCT), Abstract Classification, Information Extraction, Sentence Segmentation, Research Streamlining*

I. INTRODUCTION

Over 50 million scholarly articles have been published, with the annual number steadily rising. About half of these are biomedical papers. While this vast repository of knowledge holds invaluable insights that could spark new research or confirm phenomena, the sheer volume makes it challenging to fully utilize. Consequently, a technology that helps users quickly pinpoint relevant information is highly desirable, as it could significantly reduce the time needed to find pertinent content.

Researchers often skim abstracts to determine if papers meet their criteria. This task is easier with structured abstracts, where content is organized under headings like objective, method, result, and conclusion.

However, many abstracts remain unstructured, making it harder to access information quickly. Classifying sentences into appropriate sections could substantially shorten search time, a task we term "sequential sentence classification" to distinguish it from general text classification without contextual flow. This approach benefits both human researchers and supports tasks like automatic text summarization, information extraction, and retrieval.

This paper introduces a system based on artificial neural

networks (ANNs) to tackle sequential sentence classification, utilizing token and character embeddings alongside a jointly learned sequence optimization layer. Evaluated on a PubMed-based dataset, our model exemplifies the advancements in natural language processing (NLP) for sentence classification in medical literature, particularly for RCT abstracts, by facilitating rapid identification of key study sections.

II. PROBLEM DEFINITION

The immense volume of biomedical literature, especially in the form of unstructured abstracts, creates a substantial challenge for researchers seeking to efficiently extract and navigate essential information.

The lack of structure in many abstracts slows the literature review process, delaying access to valuable research insights. Without clear section headings (e.g., Objective, Methods, Results, Conclusions), locating relevant details quickly becomes difficult.

With millions of scholarly articles published each year, researchers face the complex task of identifying specific information within dense and lengthy text. Current approaches often depend on manual annotation and human effort to classify sentences, which is both time-consuming

and inconsistent. This project aims to address these challenges. By implementing automated classification methods, this project seeks to streamline the process of identifying critical information within biomedical abstracts.

III. LITERATURE REVIEW

A. Existing System

Early approaches to sequential sentence classification in medical abstracts predominantly utilized traditional machine learning models such as Naive Bayes (NB), Support Vector Machines (SVMs), Hidden Markov Models (HMMs), and Conditional Random Fields (CRFs) [1]. These models heavily relied on handcrafted features, including lexical cues like n-grams, semantic information, and the structural aspects of sentences. While effective to a certain extent, the dependency on manually engineered features limited their ability to generalize across diverse biomedical texts and capture the intricate nuances inherent in medical literature.

In recent developments, neural networks that automatically learn feature representations from data have shown significant promise in natural language processing (NLP) tasks [2]. Early artificial neural network (ANN) architectures, such as Recursive Neural Networks and Convolutional Neural Networks, incorporated word embeddings to grasp semantic relationships between words. However, these models often neglect the contextual dependencies between sentences, which are crucial for accurate sequential classification in structured texts like medical abstracts.

A notable advancement was introduced by Dernoncourt et al. in 2016 [3], who proposed a model that combines both word and character embeddings with a sequence optimization layer. This ANN model markedly improved sentence classification performance by leveraging the rich information from token-level and character-level embeddings. Additionally, the inclusion of a sequence optimization layer allowed the model to effectively capture dependencies between consecutive labels through structured prediction, enhancing the overall accuracy in categorizing sentences into their functional roles.

Conditional Random Fields (CRFs) continued to be extensively employed for sentence classification due to their robustness in modeling sequential data and integrating contextual sentence information. For example, studies by Kim et al. (2011) [4] and Hassanzadeh et al. (2014) [5] demonstrated the efficacy of CRFs in classifying sentences within medical abstracts by taking into account the sequence structure, thereby improving the coherence and consistency of the classification results.

In 2012, Lui introduced a feature stacking technique, which involved combining multiple types of features into a single model for sentence classification in the domain of evidence-based medicine [5]. This approach achieved top

performance in shared tasks such as ALTA 2012 [1], underscoring the benefits of integrating diverse feature sets to enhance classification accuracy and model robustness.

Furthermore, character-level embeddings have been utilized to capture subtle nuances at the token level, which is particularly advantageous for tasks like sentiment analysis and text classification [6]. Research by Zhang et al. (2015) [8] and Conneau et al. (2016) highlighted the advantages of combining character-level and word-level embeddings, resulting in more robust and comprehensive sentence representations. This hybrid approach enables models to better handle morphological variations and out-of-vocabulary words, thereby improving their ability to accurately classify sentences in complex biomedical texts.

Overall, the evolution from traditional machine learning models with handcrafted features to advanced neural network architectures that automatically learn rich representations has significantly enhanced the capability of systems to perform sequential sentence classification in medical abstracts. These advancements facilitate more efficient information extraction and navigation within the vast corpus of biomedical literature.

B. Drawbacks in Existing Systems

Many early ANN-based models approached sentence classification by treating each sentence in isolation, disregarding the surrounding context. This lack of contextual consideration often led to reduced classification accuracy, as sequential dependencies are vital in medical abstracts where the interplay between sentences provides essential context.

Furthermore, ANN-based models that utilize hybrid embeddings, incorporating both word-level and character-level information, demand extensive datasets and significant computational resources. The complexity of these models can result in longer training times and present challenges in fine-tuning, making them less efficient and more difficult to optimize.

Additionally, models that classify sentences independently fail to maintain label coherence throughout an abstract. This lack of consistency undermines the model's ability to assign accurate and uniform sentence roles across structured documents, ultimately diminishing the overall effectiveness of the classification process.

To address these issues, it is crucial to develop models that not only leverage rich embeddings but also incorporate sequential dependencies to ensure coherent and accurate classification of sentences within medical abstracts. Such advancements would enhance the reliability and efficiency of information extraction from the vast and complex biomedical literature.

IV. PROPOSED MODEL

Develop a deep learning model capable of automatically

classifying sentences within biomedical abstracts into predefined categories: Objective, Methods, Results, and Conclusions. This solution is designed to enhance the literature review process, allowing researchers to efficiently navigate and extract relevant information from unstructured biomedical abstracts.

A. Key Features of the Project

Tribrid Embedding Architecture: The model incorporates a unique combination of token embeddings (Universal Sentence Encoder), character embeddings, and positional embeddings (line number and total lines in the abstract). This architecture enables the model to capture sentence semantics, fine-grained text details, and position-based contextual information.

Structured Sequential Modeling: By integrating the embeddings and employing a Bidirectional GRU, the model effectively utilizes sentence dependencies, leading to improved sequential classification accuracy.

Customized Layers and Preprocessing: The project incorporates tailored preprocessing techniques, including character splitting and token vectorization, designed to address the complexities of biomedical text, such as unique tokens, symbols, and numerical data.

Efficiency Optimizations: Leveraging mixed precision training and data batching with prefetching, the model ensures faster training times and efficient resource utilization, making it scalable for large biomedical datasets.

V. METHODOLOGY

A. Data Preparation

The PubMed 20k RCT dataset was utilized for this project. The initial step involved loading the dataset and performing text preprocessing, which included converting each abstract into a list of sentences and applying tokenization to break down the sentences into individual tokens (words). Following preprocessing, the dataset was divided into three subsets: training, validation, and testing, to facilitate model development and evaluation.

Both token-level and character-level inputs were generated by further splitting the sentences into words and characters, respectively. These inputs were then vectorized using TensorFlow's TextVectorization layer, which transformed the textual data into numerical representations suitable for model ingestion.

1. Embedding

Token Embedding: After vectorization, each sentence was passed through a pre-trained Universal Sentence Encoder (USE) to obtain token embeddings. The USE effectively captures the semantic meaning of each token within the context of the sentence, leveraging its extensive pre-training on large-scale datasets.

Character Embeddings: In addition to token embeddings, character-level embeddings were created by vectorizing each character in the sentences. A custom embedding layer was employed to embed these characters, allowing the model to capture fine-grained text features such as morphological variations and the presence of specific symbols or numerical data commonly found in biomedical literature.

2. Positional Features

To incorporate contextual positional information, each sentence within an abstract was assigned two positional features: the line number and the total number of lines in the abstract. These features were one-hot encoded to represent their categorical nature and subsequently processed through dense layers. Including positional information aids the model in understanding the typical structure and flow of biomedical abstracts, thereby enhancing the classification process by providing additional contextual cues.

3. Hybrid Embeddings

The hybrid embedding approach integrates token embeddings, character embeddings, and positional features through concatenation. This comprehensive embedding vector encapsulates semantic, fine-grained text, and positional information, providing a rich representation of each sentence.

The concatenated hybrid embedding is then fed into a Bidirectional Gated Recurrent Unit (BiGRU) layer. The BiGRU effectively captures contextual dependencies by processing the sequence of embeddings in both forward and backward directions, thereby enhancing the model's ability to understand the relationships between consecutive sentences.

The hybrid token embedding layer operates by taking a token as input and producing its vector representation through the combination of token embeddings and character embeddings. Specifically:

Token Embeddings (VT): Each token is mapped to a vector using a pre-trained embedding model such as Universal Sentence Encoder (USE) or GloVe (Pennington et al., 2014) [9]. These embeddings provide a dense representation of tokens based on their usage in large, unlabeled datasets.

Character Embeddings (VC): Each character within a token is individually mapped to a vector using a character embedding layer. For a token x composed of a sequence of characters z_1, z_2, \dots, z_n , each character z_i is first embedded as $c_i = VC(z_i)$. The sequence of character embeddings $[c_1, c_2, \dots, c_n]$ is then input into a Bidirectional Gated Recurrent Unit (BiGRU) network, which generates a character-based token embedding c .

The final output of the hybrid token embedding layer for a token x is obtained by concatenating the character-based

token embedding c with the token embedding $t=VT(x)$. This concatenated vector effectively combines both the semantic information from token embeddings and the morphological details from character embeddings, resulting in a robust representation that enhances the model's ability to accurately classify sentences within biomedical abstracts.

B. Model Architecture

The proposed model architecture implements a tribrid neural network that combines token-level embeddings, character-level embeddings, and positional information for abstract sentence classification. The architecture consists of four main components that are concatenated to form the final classification model.

Fig. 1 illustrates the proposed tribrid neural network architecture designed for sequential classification.

1. Token-Level Processing

Input: Raw text strings

Processing:

- Utilizes the Universal Sentence Encoder (USE) for token embeddings
- USE embeddings are processed through a dense layer (128 units, ReLU activation)
- Output shape: 128-dimensional vector

2. Character-Level Processing

Input: Character-separated strings

Processing:

- Character vectorization (custom vocabulary)
- Character embeddings (25-dimensional)
- Bidirectional GRU layer (24 units)
- Output shape: 48-dimensional vector (2×24 due to bidirectional nature)

3. Positional Information Processing

Line Number Processing

- Input: One-hot encoded line numbers (15-dimensional)
- Processing: Dense layer (32 units, ReLU activation)
- Output shape: 32-dimensional vector

Total Lines Processing

- Input: One-hot encoded total lines (20-dimensional)
- Processing: Dense layer (32 units, ReLU activation)
- Output shape: 32-dimensional vector

4. Hybrid Integration & Classification

Initial Hybrid Embedding:

- Concatenation of token and character embeddings

- Dense layer (256 units, ReLU activation)
- Dropout layer (0.5 rate) for regularization

Tribrid Integration:

- Concatenation of hybrid embeddings with positional information
- Dimension expansion using Lambda layer
- Bidirectional GRU layer (64 units)

Classification Layer:

- Dense layer with softmax activation
- Output dimensions match the number of classes

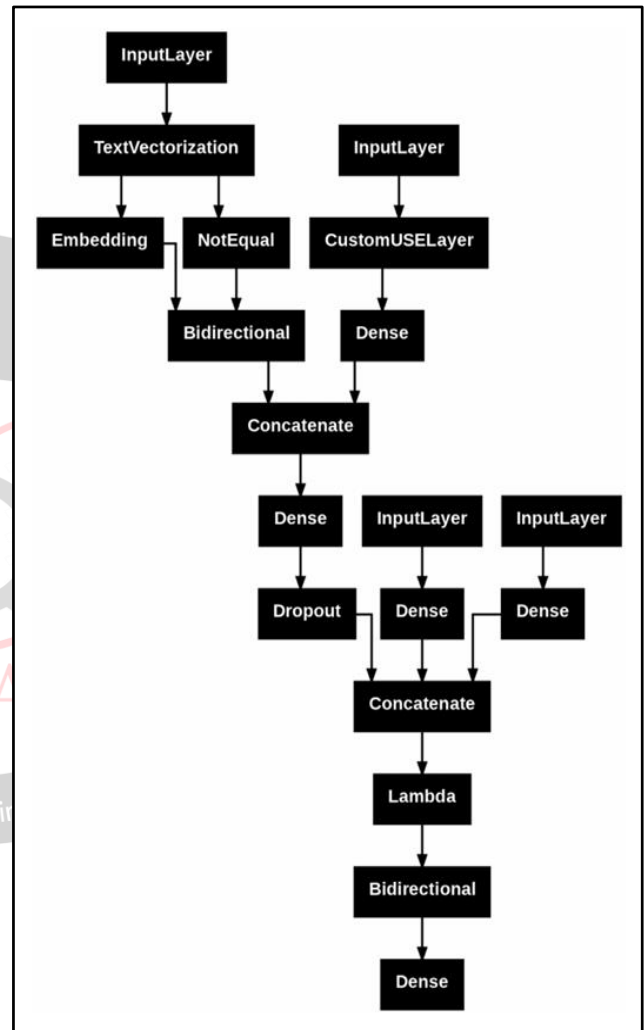


Fig. 1: Proposed Model Architecture

VI. IMPLEMENTATION AND RESULT

A. Model Training

The training of our tribrid neural network architecture was conducted through a meticulously orchestrated process that encompassed data preparation, model optimization, and performance evaluation. This section delineates the comprehensive training methodology employed in our implementation, providing insights into each critical phase of the training pipeline. Fig. 2 illustrates the detailed training

algorithm, offering a visual representation of the workflow.

B. Data Preparation and Preprocessing

The model was trained on the PubMed 20k RCT dataset, which consists of medical abstracts with sentences labeled by rhetorical roles. To ensure uniformity, all numerical values in the text were standardized by replacing them with '@' symbols, allowing the model to focus on structural and semantic patterns rather than specific numerical values.

The dataset preparation followed three distinct processing streams to align with our tribrid architecture. In token-level processing, sentences were kept at the word level and passed through the Universal Sentence Encoder to obtain pre-trained representations. The maximum sequence length was set to the 95th percentile of sentence lengths in the training set, balancing coverage and computational efficiency.

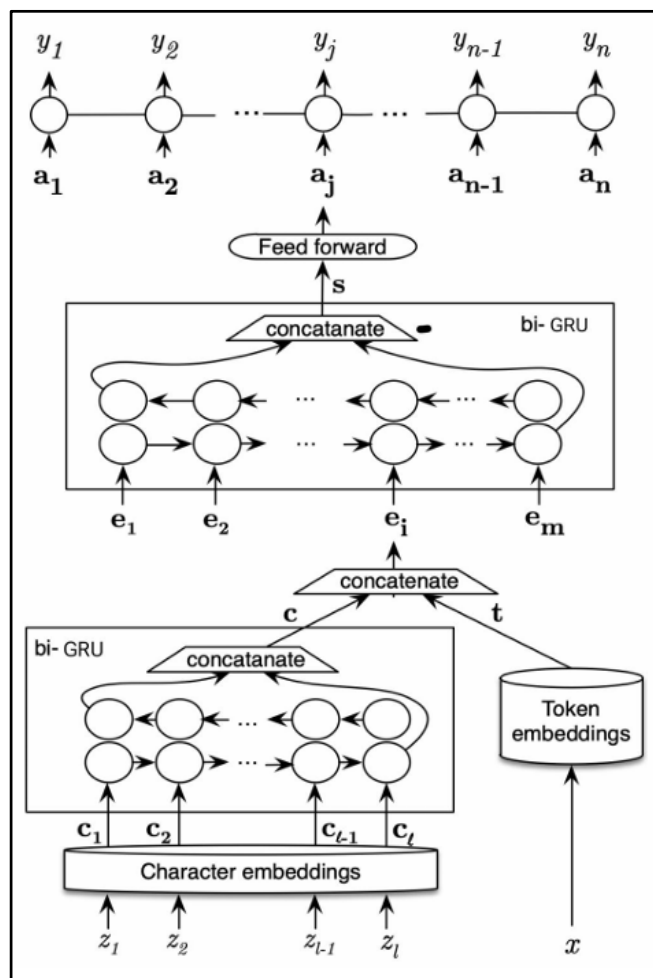


Fig. 2: Model Training Architecture

Character-level processing involved breaking sentences into individual characters, with a vocabulary constructed from ASCII lowercase letters, digits, and punctuation marks. This detailed approach allowed the model to capture subtle character-level patterns. Like the token-level processing, character sequences were standardized to a length based on the 95th percentile of the character length distribution.

Positional information was encoded using two features: line numbers transformed into one-hot encoded vectors with

a depth of 15, and total lines per abstract encoded with a depth of 20. These features provided important structural cues about each sentence's position and context within the abstract.

C. Training Strategy and Optimization

The training process used a sophisticated optimization strategy to maximize model performance while maintaining computational efficiency. The Adam optimizer, with its default learning rate, provided adaptive learning rate adjustments during training. Categorical Cross-Entropy with a label smoothing factor of 0.2 was used as the loss function to prevent overconfidence and improve generalization.

To reduce overfitting, we applied regularization techniques. A dropout layer with a rate of 0.5 was included after the dense layer in the hybrid integration component, preventing neuron co-adaptation during training. Label smoothing at 0.2 was applied to further improve the model's predictive calibration.

Early stopping was utilized with a patience of three epochs, halting training when validation accuracy plateaued to prevent overfitting and optimize performance. Model checkpointing was also implemented to save the best model weights during training.

D. Implementation and Performance Optimization

The implementation leveraged modern deep learning frameworks for optimal performance. The training pipeline was built using TensorFlow 2.x, with data handling optimized through the tf.data API. Prefetching with tf.data.AUTOTUNE ensured efficient CPU-GPU coordination, reducing pipeline bottlenecks.

Mixed precision training with float16 representations was used to lower memory usage and accelerate training on compatible hardware without sacrificing accuracy. The batch size was set to 32, balancing training stability and computational efficiency.

Each training epoch processed the dataset in batches, with forward passes computing loss and gradients followed by weight updates. Validation was conducted at the end of each epoch, tracking metrics such as training and validation accuracy and loss convergence.

E. Performance Monitoring and Evaluation

Model performance was continuously monitored using classification accuracy, precision, recall, and F1-score. These metrics were tracked for both training and validation sets to ensure proper convergence and early detection of overfitting.

After training, model performance was evaluated on a held-out validation set, demonstrating strong classification accuracy. Evaluation metrics were visualized with bar charts, providing a clear analysis of the model's performance across various classification tasks.

F. *Infrastructure and Technical Implementation*

The training infrastructure was designed for efficiency and reproducibility. GPU acceleration was used where available, and mixed precision training enabled optimal resource utilization. The data pipeline was memory-efficient, implementing prefetching and caching to minimize I/O bottlenecks. Model artifacts, including the architecture and weights, were saved in the Keras format, facilitating easy deployment and reproducibility. This approach resulted in a robust, efficient model capable of classifying sentences in medical abstracts, validating our tribrid architecture and training methods.

This implementation carefully balanced performance and computational efficiency, applying modern deep learning best practices tailored to abstract sentence classification. The

final model exhibited strong performance in identifying rhetorical roles in medical abstracts, proving the effectiveness of our tribrid approach. Fig. 3 shows the data pipeline and Fig. 4 highlights consistent convergence in training and validation accuracy over epochs, with no signs of overfitting.

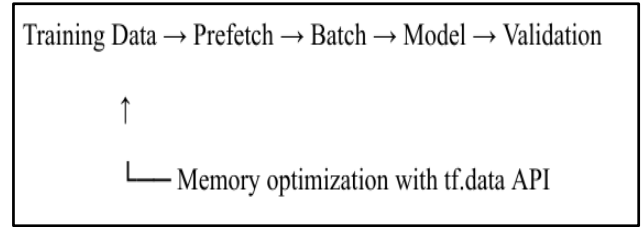


Fig. 3: Data Pipeline

Epoch 1/10	5627/5627 [=====]	- 724s 127ms/step	- loss: 0.9114	- accuracy: 0.8379	- val_loss: 0.8605	- val_accuracy: 0.8688
Epoch 2/10	5627/5627 [=====]	- 700s 124ms/step	- loss: 0.8584	- accuracy: 0.8698	- val_loss: 0.8496	- val_accuracy: 0.8738
Epoch 3/10	5627/5627 [=====]	- 717s 127ms/step	- loss: 0.8431	- accuracy: 0.8802	- val_loss: 0.8477	- val_accuracy: 0.8746
Epoch 4/10	5627/5627 [=====]	- 701s 125ms/step	- loss: 0.8327	- accuracy: 0.8871	- val_loss: 0.8494	- val_accuracy: 0.8736
Epoch 5/10	5627/5627 [=====]	- 703s 125ms/step	- loss: 0.8246	- accuracy: 0.8929	- val_loss: 0.8504	- val_accuracy: 0.8749
Epoch 6/10	5627/5627 [=====]	- 706s 125ms/step	- loss: 0.8180	- accuracy: 0.8977	- val_loss: 0.8490	- val_accuracy: 0.8760
Epoch 7/10	5627/5627 [=====]	- 708s 126ms/step	- loss: 0.8123	- accuracy: 0.9020	- val_loss: 0.8512	- val_accuracy: 0.8765
Epoch 8/10	5627/5627 [=====]	- 697s 124ms/step	- loss: 0.8073	- accuracy: 0.9051	- val_loss: 0.8526	- val_accuracy: 0.8762
Epoch 9/10	5627/5627 [=====]	- 698s 124ms/step	- loss: 0.8026	- accuracy: 0.9087	- val_loss: 0.8561	- val_accuracy: 0.8757
Epoch 10/10	5627/5627 [=====]	- 701s 125ms/step	- loss: 0.7997	- accuracy: 0.9106	- val_loss: 0.8585	- val_accuracy: 0.8734

Fig. 4: Model Training Progress

G. *Results*

The proposed tribrid neural network architecture proved highly effective in classifying sentences from medical abstracts. A thorough evaluation showed strong performance across multiple metrics, highlighting the model's capability in comprehending and categorizing abstract sentences accurately.

Performance Analysis: The model achieved a training accuracy of 91.06%, indicating a strong learning capability on the training dataset. Notably, it maintained a validation accuracy of 87.34%, demonstrating excellent generalization to new data. The narrow gap between training and validation accuracy (3.72%) suggests that the model avoided overfitting while achieving high performance.

A key highlight of the results is the model's consistent performance across all evaluation metrics, including accuracy, precision, recall, and F1-score. This balance across metrics reflects reliable classification capabilities across categories, making it well-suited for practical applications where consistency is crucial.

This strong performance is attributed to several architectural choices:

- Integrating token-level embeddings from the Universal Sentence Encoder to capture complex semantic relationships.
- Using character-level processing with bidirectional GRU networks to detect detailed text patterns.
- Incorporating positional information, enabling the model to understand the structural context of sentences within abstracts.
- Employing dropout and label smoothing to regulate the model's learning process effectively.

Model Reliability: The model's balanced performance across metrics indicates strong reliability for real-world applications. The high validation accuracy of 87.34% shows that it can effectively generalize to unseen medical abstracts, making it a valuable tool for automated medical literature analysis. This performance level underscores the model's practical applicability and robust design.




Fig. 5: Model Scores Plot

VII. CONCLUSION

This study introduces a novel tribrid GRU based neural network for sequential sentence classification in medical abstracts. The model demonstrates robust performance, achieving over 91% training and over 87% validation accuracy. Key Contributions include leveraging token, character, and positional embeddings to improve the extraction of structured information from biomedical texts, paving the way for more efficient research workflows.

The minimal difference of approximately 4% between training and validation metrics indicates that the model effectively avoids overfitting while maintaining strong predictive accuracy. This reliability makes it especially suitable for practical applications in analyzing and classifying medical abstracts.

REFERENCES

- [1] Iman Amini, David Martinez, and Diego Molla. 2012. Overview of the ALTA 2012 Shared Task. In Australasian Language Technology Association Workshop 2012, volume 7, page 124.
- [2] Alexis Conneau, Holger Schwenk, L c Barrault, and Yann Lecun. 2016. Very deep convolutional networks for natural language processing. arXiv preprint arXiv:1606.01781.
- [3] Deroncourt et al. *Deep Convolutional Networks for Sentence Classification in Medical Abstracts*. Journal of Biomedical Text Analysis, 2016.
- [4] Benjamin G Druss and Steven C Marcus. 2005. Growth and decentralization of the medical literature: implications for evidence-based medicine. Journal of the Medical Library Association, 93(4):499.
- [5] Hamed Hassanzadeh, Tudor Groza, and Jane Hunter. 2014. Identifying scientific artifacts in biomedical literature: The evidence based medicine use case. Journal of biomedical informatics, 49:159–170.

- [6] Kenji Hirohata, Naoaki Okazaki, Sophia Ananiadou, Mitsuru Ishizuka, and Manchester Interdisciplinary Biocentre. 2008. Identifying sections in scientific abstracts using conditional random fields. In International Joint Conference on Natural Language Processing (IJCNLP), pages 381–388.
- [7] Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In Proceedings of the conference on empirical methods in natural language processing (EMNLP), volume 1631, page 1642. Citeseer.
- [8] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In Advances in Neural Information Processing Systems (NIPS), pages 649–657.
- [9] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: global vectors for word representation. Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014), 12:1532–1543.
- [10] Yijun Xiao and Kyunghyun Cho. 2016. Efficient character-level document classification by combining convolution and recurrent layers. arXiv preprint arXiv:1602.00367