Automatic detection of suicidal ideation on social media using a hybrid CNN–BiLSTM architecture

Sadala Nagib Salame Neto, Kaique da Silva Pinto, Silvio Tadeu Teles da Silva, and Miércio Cardoso de Alcântara Neto

***Abstract*—Suicide is a serious public health problem, responsible for more than 700,000 deaths annually worldwide, according to the World Health Organization. Early identification of signs of suicidal ideation can support prevention initiatives and reduce fatal outcomes. This study investigated the application of Natural Language Processing (NLP) and deep learning techniques for the automatic detection of suicidal ideation in social media posts. A hybrid architecture based on Bidirectional Long Short-Term Memory (BiLSTM), combined with convolutional layers and regularization strategies, was proposed. The evaluation was conducted on a Reddit dataset processed in two versions: with and without stopword removal. The results showed that the proposed architecture achieved high performance in both versions, reaching an accuracy of 95.77% and an F1-score of 95.79% on the test set without stopword removal. Despite the higher computational cost, this approach outperformed previous studies reported in the literature. In summary, the proposed model demonstrates potential to support early monitoring strategies for suicidal ideation in digital environments, offering a relevant contribution to public health and suicide prevention.**

Link to graphical and video abstracts, and to code: https://latamt.ieeer9.org/index.php/transactions/article/view/10144

***Index Terms*—Natural Language Processing (NLP), Suicidal Ideation, BiLSTM, Social Media, Data Mining.**

I. Introduction

SUICIDE represents one of the greatest challenges to global public health. According to the World Health Organization (WHO) [1], more than 700,000 people take their own lives every year, making it the third leading cause of death among young people aged 15 to 29. In addition, millions of suicide attempts are recorded annually, highlighting the severity and complexity of the problem.

In Brazil and other Latin American countries, suicide rates remain high even after decades of recognizing the phenomenon as a public health issue. As argued by Stavizki Junior [2], territorial understanding is central to addressing this challenge, as it reveals contradictions and potentialities in mental health policies within a regional context. This perspective reinforces the need for more effective and socially grounded strategies capable of identifying early warning signs and guiding appropriate interventions.

At the same time, the growth of social networks has created an environment where millions of users share personal experiences, feelings, and emotional states on a daily basis. Among these posts, many may reveal signs of psychological distress or suicidal ideation. However, the large volume of data, the presence of linguistic ambiguities, and the use of slang make manual analysis unfeasible, requiring advanced computational solutions [3].

In this context, Natural Language Processing (NLP) and deep learning techniques have proven to be promising for the automatic detection of linguistic patterns associated with suicidal ideation [7], [12]. Recent approaches based on recurrent neural networks, particularly Bidirectional Long Short-Term Memory (BiLSTM) architectures, stand out for their ability to capture contextual and temporal dependencies in textual sequences [23].

Given this perspective, the present study aims to propose and evaluate a hybrid architecture based on BiLSTM, combined with convolutional layers and regularization strategies, applied to the detection of suicidal ideation in social media posts. The investigation seeks to contribute to the development of computational tools that support early monitoring and assist suicide prevention initiatives in the field of public health.

This article is organized as follows: Section 2 presents the theoretical framework supporting the study; Section 3 discusses related work on the topic; Section 4 describes the methodology in detail; Section 5 presents and analyzes the results obtained; and finally, Section 6 provides the conclusions and final considerations.

II. Theoretical Foundation

The advancement of digital technologies and the popularization of social networks have transformed the way people express their feelings and thoughts. Platforms such as X (formerly Twitter), Facebook, Instagram, Reddit, and online forums have become spaces where users share their experiences, including accounts that may contain emotional distress and suicidal ideation [3]. The automated analysis of these texts offers a unique opportunity to support suicide prevention by enabling early identification of linguistic patterns associated with risk.

Data mining is a form of automated data analysis that involves collecting and analyzing large volumes of information with the aid of computer programs that process the collected

The associate editor coordinating the review of this manuscript and approving it for publication was Saul Zapotecas-Martinez (*Corresponding author: Sadala Nagib Salame Neto*).

This research was supported in part by Programa de Pós-Graduação em Engenharia Elétrica of Universidade Federal do Pará (PPGEE-UFPA).

Sadala Nagib Salame Neto, K. D. S. Pinto, S. T. T. D. Silva, M. C. D. A. Neto are with Institute of Technology, Universidade Federal do Pará, Rua Augusto Corrêa, 01, Guamá, Belém, PA, 66075-110, Brazil (e-mails: sadalaneto@ufpa.br, kaique.pinto@itec.ufpa.br, silvio.teles@itec.ufpa.br, and miercio@ufpa.br).

data to extract relevant information [4]. The data mining process seeks to discover new information from a specific set of data [5]. Furthermore, data mining is one of the stages within Knowledge Discovery in Databases, which also comprises the following steps: selection, preprocessing, transformation, and data interpretation [6]. In the mental health context, data mining aims to identify indicators of psychological disorders, depression, and suicidal behavior in posts and messages [7].

Affective computing is a relatively recent research field that can be defined as computing that arises from, relates to, or is influenced by emotions or other affective phenomena [8]. According to Rosalind Picard, affective computing is defined as the area that brings together research on the use of human emotions in different aspects of computational systems [9]. The intersection between data mining and affective computing enables the development of models that not only identify emotions but also enhance accuracy by considering contextual nuances. For this reason, affective computing plays a central role in the analysis of texts related to mental health.

Among the various branches of data mining, there is one linked to affective computing called opinion mining. Opinion mining can be defined as the task of detecting, extracting, and classifying opinions, feelings, and attitudes about different topics expressed in [10]. Opinion mining generally involves the use of text analysis techniques, Natural Language Processing (NLP), and computational linguistics to identify, extract, and understand subjective content [11].

In the mental health context, opinion mining can be applied to the analysis of texts with a focus on detecting suicidal ideation. Unlike traditional sentiment analysis, which is limited to identifying positive and negative polarities, this approach seeks to detect the presence or absence of suicidal ideation signs. Studies such as Gaur et al. (2021) demonstrate that certain writing patterns, such as references to death, feelings of hopelessness, self-deprecation, and farewells, can be important risk indicators [12].

From this perspective, artificial neural networks play a crucial role, offering architectures capable of learning complex representations of emotions expressed in text. Artificial Neural Networks (ANNs) are computational algorithms that use mathematical models inspired by the neural structure of intelligent organisms, attempting to simulate the functioning of the human brain on computers [13]. One of the main characteristics of an ANN is its ability to learn through examples and generalize the information learned, producing a non-linear model [14].

For tasks that require the use of neural networks on databases containing only text, recurrent neural networks (RNNs) are typically employed. RNNs are neural networks designed for processing sequential data. They differ from traditional ANNs because they include a feedback loop, where the information processed at one step is reused as input in the next step [15]. These deep learning algorithms are frequently applied to tasks involving temporal or ordinal problems, such as language translation and NLP [16].

To overcome the traditional limitations of RNNs, such as the vanishing and exploding gradient problem, specialized architectures were developed. Long Short-Term Memory (LSTM) is a recurrent neural network architecture proposed by Hochreiter and Schmidhuber (1997) [17], specifically designed to address the vanishing and exploding gradient issues that affect the training process of conventional RNNs [18], [19].

LSTM introduces an internal memory cell capable of storing information for long periods [20]. In addition, it employs control gates, which are mechanisms that regulate the flow of information selectively [19]. These gates include the input gate, which determines which input information should be stored; the forget gate, responsible for controlling what should be discarded from memory; and the output gate, which determines what will be emitted as output from the cell. This structure allows the network to decide when to retain and when to discard information, keeping gradients stable during training, even with very long sequences [21], [22].

BiLSTMs represent an extension of traditional LSTMs, as they process data sequences in both directions — from left to right and from right to left. This feature enhances the model's ability to capture contextual dependencies, enabling a more comprehensive understanding of the text, which makes them particularly effective for Natural Language Processing (NLP) tasks [23].

Therefore, the integration of data mining, affective computing, and neural network architectures enables the development of more robust and accurate systems for analyzing texts related to mental health.

III. Related Works

The application of NLP and deep learning to the detection of suicidal ideation has received increasing attention in recent years, driven by the need to identify early signs of suicidal ideation in individuals. Research in this field has explored different approaches, ranging from traditional data analysis techniques to advanced computational methods such as deep learning, aiming to improve accuracy for early detection of suicidal ideation. This section presents some studies addressing this topic.

Birjali et al. [24] proposed a method for predicting suicidal ideation on social networks using Twitter as the data source. Their methodology involved building a manually curated vocabulary related to the topic of suicide, collecting tweets via the Twitter4J API, and subsequently performing automated classification of these data. For analysis, the authors used the Weka data mining software, employing machine learning algorithms and a semantic sentiment analysis approach based on WordNet. Experimental results showed that 67% of the suspicious tweets presented suicide risk, with the highest accuracy achieved by the SMO algorithm at 89.5% for tweets classified as suicide risk.

Lim and Loo [25] analyzed linguistic features associated with different levels of suicide risk in Twitter (now X) posts. They proposed a machine learning approach to classify suicide risk into three levels (low, medium, and high). The study used a dataset containing 690 manually annotated tweets, evenly distributed among the three classes. Additionally, three feature extraction techniques were applied—TF-IDF, Part-of-Speech tagging, and sentiment analysis with VADER (Valence-Aware

Dictionary for Sentiment Reasoning)—to train a Random Forest model with an 80:20 train-test split. The proposed model achieved 86.23% accuracy, 86.71% precision, and 86.23% recall on the test set, and 89.33% accuracy, 92.98% precision, and 83.08% recall on real-time data. Analysis indicated that terms such as "die" and "kill" were more common in high-risk cases, while expressions like "suicidal thoughts" predominated in medium-risk cases. The authors concluded that the overall context of the text is the most determining factor for prediction, outweighing the importance of isolated grammatical markers.

Cao et al. [26] investigated new ways to detect suicide risk in social networks by analyzing users' internal thoughts and emotional changes. The authors proposed a methodology based on three subtasks: (1) inferring internal thoughts from open posts using GRU, ResNet, attention mechanisms, and suicide-specific embeddings; (2) capturing implicit emotional changes; and (3) integrating these results for the final suicide risk prediction. The study used datasets from Weibo (3,652 at-risk users and 3,652 normal users) and Reddit (392 at-risk and 108 normal users). Results showed that the proposed approach outperformed previous methods, achieving 95.02% accuracy and 94.45% F1-score on Weibo, and 90.47% accuracy and 90.54% F1-score on Reddit.

Nikhileswar et al. [27] explored the detection of suicidal ideation in social media forums. The study used 232,074 posts collected from the "SuicideWatch" subreddit (potentially suicidal content) and "teenagers" (non-suicidal content), balanced through under-sampling. The proposed methodology employed the Universal Sentence Encoder (USE), a Transformer-based approach, to convert texts into 512-dimensional vectors, feeding a Fully Connected Neural Network (FCNN). The study compared different deep learning and machine learning models using various text representation techniques. The proposed FCNN model with USE achieved superior performance, reaching an F1-score of 94.5%, precision of 94.5%, recall of 94.5%, and accuracy of 94.16%, outperforming other methods.

Li et al. [28] conducted an investigation focused on assessing suicide risk from Reddit data, specifically from the "Suicide Watch" subreddit. The work aimed to identify and extract information to evaluate suicide risk. The researchers proposed a multi-feature fusion recurrent attention network. This approach used a BiLSTM for text representation, a self-attention mechanism to highlight relevant information, and the fusion of external linguistic features. The model evaluation was conducted using the 2019 CLPsych dataset. Results showed that the proposed model improved risk-F1 by 3.3%, existence-F1 by 0.9%, and urgency-F1 by 3.7%.

Ghanadian et al. [29] analyzed social media data related to suicide to identify gaps between risk factors discussed online and those established in psychological literature. The applied methodology involved unsupervised topic modeling on Reddit data, a scoping review of psychological literature for extracting risk factors, guided topic modeling to assess the presence of these factors, and synthetic data generation using GPT-3.5 Turbo to augment the dataset. Results indicated that including synthetic data in classifiers improved suicidal ideation detection capability, increasing the F1-score from 0.87 to 0.91 in one test subset and from 0.70 to 0.90 in another. Finally, the study concluded that the use of synthetic data can enhance understanding of online discussions about suicide and contribute to more accurate machine learning models for its detection.

Naseem et al. [30] examined suicide risk identification in social media posts by proposing an explainable model based on a hybrid representation. The authors proposed a methodology combining Longformer to capture low-level word features and TensorGCN for high-level document-level features. These representations were fused and processed in a Transformer encoder with ordinal classification. Experiments were conducted on a dataset of posts from 500 Reddit users, labeled into five risk levels according to the C-SSRS scale. Results showed that the proposed approach achieved superior performance compared to baseline methods, with an F-score of 0.79 (an absolute gain of 15% over the best baseline, SISMO).

Haque et al. [31] conducted a comparative analysis between machine learning and deep learning models for detecting suicidal ideation in tweets. The methodology involved creating a dataset with 49,178 tweets collected via the Python Tweepy API using 18 keywords associated with suicidal thoughts. For processing, feature extraction techniques such as Count Vectorizer were used for machine learning models and word embeddings for deep learning models. Results showed that the best-performing machine learning model was Random Forest (RF), achieving 93% accuracy and an F1-score of 0.92. However, the deep learning model BiLSTM achieved the best overall performance, with 93.6% accuracy and an F1-score of 0.93.

Ryu et al. [32] developed machine learning-based predictive models to identify suicide attempts among individuals with suicidal ideation in South Korea. The study used data from 35,116 participants aged over 19, collected from the Korean National Health and Nutrition Survey. To address class imbalance, the Synthetic Minority Over-sampling Technique (SMOTE) was used, resulting in 1,324 cases of individuals who attempted suicide and 1,330 cases who did not. The machine learning model used was Random Forest, trained on a set of 1,858 samples and tested on 796 samples. Results indicated that the model achieved high performance, with an AUC of 0.947 and an accuracy of 88.9% in predicting individuals who attempted suicide.

Tadesse et al. [33] analyzed suicidal ideation in Reddit posts. The methodology included preprocessing with NLTK, feature extraction via TF-IDF, Bag of Words, and Word2Vec embeddings. The authors compared the performance of machine learning classifiers and deep learning models and proposed a hybrid LSTM-CNN architecture. This model combined LSTM's ability to capture long-term dependencies with CNN's capability to extract local patterns. Results showed that the LSTM-CNN architecture achieved the best performance, reaching 93.8% accuracy and an F1-score of 93.4%, outperforming LSTM (91.7% accuracy) and CNN (90.6% accuracy), as well as traditional methods.

The studies presented in this section highlight the diversity of approaches and the growing interest in applying NLP, machine learning, and deep learning for the automatic detection of suicidal ideation. A significant performance evolution can be observed, from classical machine learning techniques such as Random Forest to more advanced architectures like CNNs, LSTMs, and BiLSTMs, as well as hybrid models combining multiple mechanisms. In particular, BiLSTM neural networks have stood out due to their ability to capture contextual relationships in both directions of the textual sequence, which helps enhance the detection of subtle signals of suicidal ideation. Despite these advances, challenges remain related to the scarcity of large, balanced datasets, linguistic variability, and the interpretation of implicit expressions. In this scenario, the present work proposes the use of a hybrid architecture based on BiLSTM, combined with convolutional layers and regularization strategies, to classify social media posts according to the presence or absence of indicators of suicidal ideation.

IV. Methodology

In this study, a case study was conducted to evaluate the effectiveness of a hybrid CNN–BiLSTM architecture in detecting indicators of suicidal ideation in social media posts. The objective of this work is to classify social media posts into two classes: presence of suicidal ideation and absence of suicidal ideation. The adopted methodology was divided into the following stages:

1) Selection of the dataset;
2) Computational environment used;
3) Data preprocessing;
4) Configuration of the proposed architecture;
5) Evaluated metrics;
6) Evaluation of the obtained results (which will be presented in the results section);

The flowchart in Fig. 1 graphically illustrates the stages of this article's methodology.

A. Database Selection

For this study, the same dataset used in the work by Nikhileswar et al. (2021) was employed, as mentioned in the previous section, and it is available on the Kaggle platform. The choice of this dataset is due to its large number of records and balanced classes. The most recent version of this dataset contains 232,074 samples equally distributed into two classes: suicide and non-suicide.

Before splitting the data into training, testing, and validation sets, an automated check was performed on the dataset to identify possible duplicates. After this verification, each record was confirmed to be unique, and the distribution of the two classes was found to be equal.

In this work, the entire dataset was used, encompassing all 232,074 samples contained in the dataset. These samples were divided into three subsets: training, with 139,244 samples; testing, with 46,415 samples; and validation, with 46,415 samples.

B. Computational Environment used for Work

This study was developed using the Python programming language on the Google Colab platform, in addition to libraries such as TensorFlow/Keras, NLTK, and scikit-learn. Google Colab is a free Jupyter notebook platform that requires no configuration and runs code directly in the cloud, eliminating the need for local installation [34]. Furthermore, it is a tool that uses Python and can be accessed for free via a web browser without requiring installation.

C. Data Processing

In this step of the work, the dataset was processed in two distinct versions: one version with the removal of stopwords and another version without the removal of stopwords, meaning that all original terms were kept. This allows for an evaluation of the impact that the presence of stopwords can have on the performance of the neural network for the problem at hand.

Stopwords are textual elements or words considered to have little relevance for text analysis [35]. A stopword list includes grammatical words such as articles, prepositions, and conjunctions that, although essential for grammar, are generally discarded in text analysis because they do not add significant semantic value [39].

Additionally, in both versions of the dataset, HTML tags, unnecessary spaces between words, and special characters were removed, and all data in the dataset were converted to lowercase.

After text preprocessing, tokenization and padding techniques were applied to each dataset version. Tokenization is a text segmentation process that breaks down a text into smaller lexical units called tokens, delimiting word boundaries and preparing the dataset for subsequent computational analyses [36]. Tokenization allows the text to be converted into a structured sequence of more meaningful units, facilitating further processing such as vectorization, word counting, and syntactic analysis. In this work, tokenization was performed using the Tokenizer method from the tensorflow.keras.preprocessing.text library. For both dataset versions (with and without stopword removal), the parameter num_words was set to 50,000, meaning that only the 50,000 most frequent words were retained for analysis, while the remaining words were discarded.

Padding is a technique used to ensure that all input data have the same size [37]. To achieve this, a maximum desired sequence length is defined. If a sequence is shorter than the defined maximum length, zeros are added at the end until the desired length is reached. If a sequence exceeds the maximum length, it is truncated to the defined value, meaning any content beyond that value is discarded. For the version with stopword removal, the maximum padding length was set to 160 words, while for the version without stopword removal, the maximum padding length was set to 320, preserving more content per sequence.

The parameters adopted in the tokenization and padding stages were defined based on an analysis conducted after the preprocessing phase. The selection of these values was empirical, grounded in preliminary experiments and direct observation of the dataset's behavior, while adhering to established practices for handling long texts.

The num_words parameter was set to 50,000 in order to preserve a sufficiently large vocabulary capable of capturing

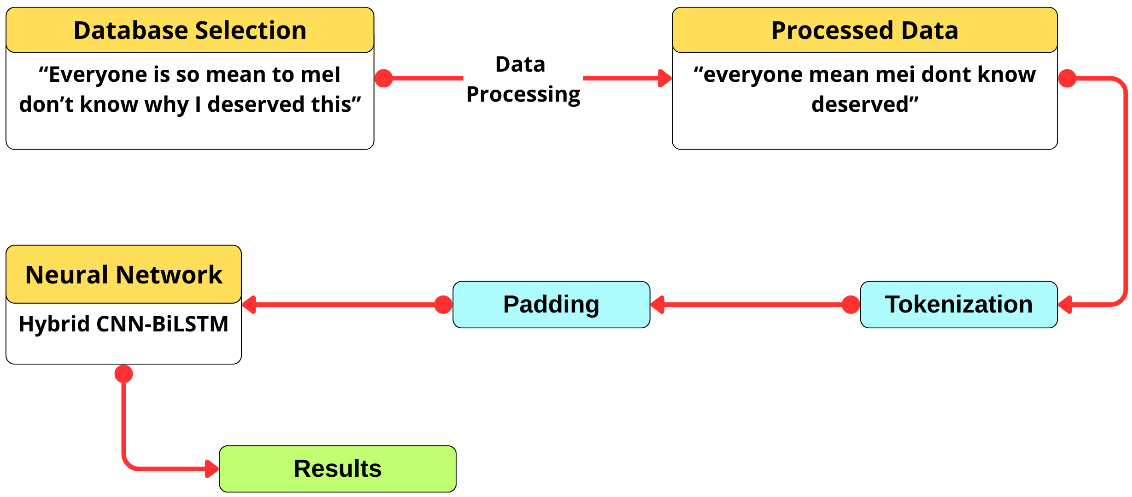

Fig. 1. Stages of the methodology adopted in this study, including database selection, text preprocessing, tokenization, padding, and classification using the proposed CNN–BiLSTM neural network architecture.

the lexical diversity present in the dataset, while avoiding the inclusion of extremely rare terms, which tend to introduce noise and unnecessarily increase computational cost.

Regarding padding, the maximum lengths of 160 and 320 tokens were defined based on the distribution of post lengths after preprocessing. It was observed that, in the version with stopword removal, more than 90% of the samples had lengths shorter than 160 tokens, whereas in the version without stopword removal this distribution shifted toward longer sequences, justifying the use of 320 tokens to minimize information loss due to truncation.

This choice is directly related to the nature of the analyzed dataset, as Reddit posts tend to be longer and more discursive when compared to short texts such as tweets. Therefore, both the adopted preprocessing strategy and the proposed CNN–BiLSTM architecture were designed to handle long texts, aiming to preserve broad contextual dependencies that may be essential for the accurate identification of indicators of suicidal ideation.

D. Neural Network Architecture used

The architecture implemented in this study was built using a sequential model, combining convolutional, recurrent, and dense layers to capture both local patterns and long-term dependencies in the analyzed texts. Additionally, regularization techniques (such as Dropout, Batch Normalization, and L2 regularization) were applied to reduce overfitting and improve the model’s generalization capability. Fig. 2 provides a simplified representation of the architecture developed in this work.

The model consists of the following layers:

1) Input and Embedding Layer: Tokenized texts were converted into dense vectors through an Embedding layer with an output dimension of 100. This dimension was chosen as it represents a good trade-off between semantic representation capacity and computational cost.
2) Dropout: Applied with a rate of 0.3 in order to reduce undesirable correlations between neurons.
3) 1D Convolutional Layer: With 32 filters, a kernel size of 5, and ReLU activation function. This layer is responsible for extracting local patterns from the textual sequences, such as frequent word combinations and semantically relevant expressions for the detection of suicidal ideation.
4) MaxPooling Layer: A Pooling layer with a window size of 2, used to reduce the dimensionality of the representations extracted by the convolutional layer while preserving the most informative features.
5) Batch Normalization Layer: Normalization of the previous layer’s activations, accelerating training and contributing to greater model stability.
6) BiLSTM Layer: With 32 neurons, dropout of 0.3, and recurrent dropout of 0.3. L2 regularization with a coefficient of 0.001 was also applied, followed by a Dropout layer with a rate of 0.5. The BiLSTM enables the capture of long-range contextual dependencies by simultaneously considering both the preceding and the subsequent context of each term in the textual sequence.
7) Dense Layer: With 16 neurons, ReLU activation function, and L2 regularization of 0.001. This layer acts as a refiner of the representations extracted by the previous layers, preparing them for the final classification decision. This layer is followed by a Batch Normalization layer and a Dropout layer with a rate of 0.4.
8) Output Layer: With 1 neuron and a sigmoid activation function, suitable for binary classification between posts

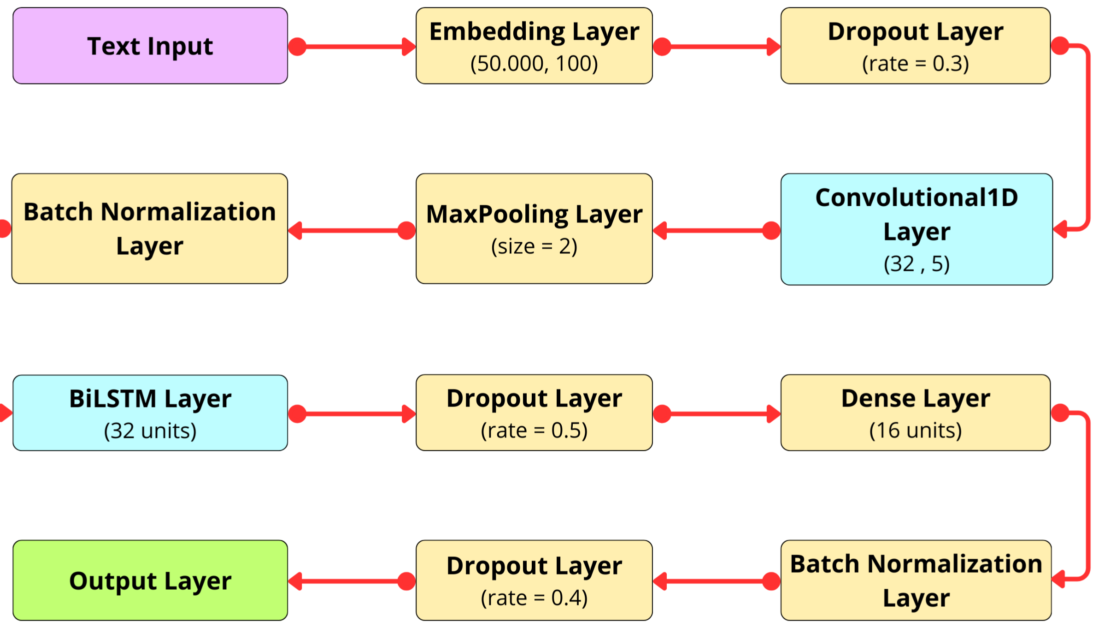

Fig. 2. Proposed hybrid CNN–BiLSTM architecture used for automatic detection of suicidal ideation in social media texts.

with the presence of suicidal ideation and posts with the absence of suicidal ideation.

The hyperparameters of the architecture were defined based on preliminary experiments and well-established practices in the literature on deep neural networks applied to text classification. The choice of 32 filters in the convolutional layer and 32 neurons in the BiLSTM layer aimed to balance model expressiveness with computational efficiency, avoiding excessively deep architectures that could increase the risk of overfitting.

Overfitting occurs when a predictive model fits the training data so closely that it not only captures important patterns but also memorizes noise and sample-specific peculiarities [40]. As a result, the model merely memorizes the observed patterns and does not truly "learn" the relevant underlying structures [40]. This means that, although it may perform very well on known data, its ability to generalize and make accurate predictions on new datasets is compromised. In the context of artificial neural networks, overfitting is evident when a given model achieves excellent performance on the training set but performs poorly when evaluated on test samples [41]. Therefore, overfitting is a problem directly related to the model's generalization capability.

The adopted dropout rates, ranging from 0.3 to 0.5, follow widely used recommendations for deep networks and contribute to reducing co-adaptation among neurons. L2 regularization was applied to penalize excessively large weights, helping to control model complexity.

The integration between convolutional layers and the BiLSTM plays a fundamental role in the proposed architecture. The convolutional layers act as local feature extractors, capable of identifying relevant patterns in short word subsequences, such as recurring expressions, specific linguistic constructions, and textual cues associated with suicidal ideation. The use of MaxPooling helps highlight the most relevant features and reduce the dimensionality of intermediate representations.

Subsequently, the BiLSTM layer processes these representations by considering long-range contextual dependencies in both directions of the textual sequence. This characteristic is particularly important in social media texts, in which the meaning of a word or expression may depend on both preceding and subsequent context. In this way, the hybrid CNN–BiLSTM architecture enables the simultaneous capture of local patterns and global contextual relationships, resulting in a more comprehensive and suitable representation for the task of automatic suicidal ideation detection.

In addition to the regularization techniques mentioned, an EarlyStopping criterion was also applied, configured to monitor the loss function on the validation set (val_loss). This mechanism automatically halts training if no improvements are observed after a specified number of consecutive epochs; in this work, the number of epochs used was set to five, restoring the model's best weights.

Complementarily, the ReduceLROnPlateau strategy was employed, which halves the learning rate whenever there is no improvement in val_loss for three consecutive epochs, down to a minimum of 1e-6. The combination of these mechanisms helped mitigate the risk of overfitting and ensured greater stability during the training process.

The model was trained using the Adam optimizer with a learning rate of 0.001 and the binary cross-entropy loss function.

E. Evaluated metrics

In the literature on applying machine learning models to sensitive domains such as financial fraud detection, widely used performance metrics include sensitivity, specificity, precision, F1-score, and area under the ROC curve (AUC), as well as unsupervised metrics such as Silhouette Coefficient, Rand Index, and Davies–Bouldin Index [38]. The same set of metrics was adopted here due to its ability to reflect both the accuracy and robustness of the model across different classification scenarios, guiding the selection of indicators considered in this work.

Thus, the evaluation metrics monitored during training were: accuracy, precision, recall, AUC, and F1-score.

V. Results Obtained

In this section, the results obtained after the data preprocessing steps and the training of the proposed architecture are presented. Table I shows the results achieved by the model for the metrics accuracy, precision, recall, AUC, F1-score, and loss, considering the training, validation, and test sets of the dataset version with stopword removal. Table II, on the other hand, presents the results for the same metrics using the dataset version without stopword removal, allowing a comparison of the impact of these terms on the neural network's performance. Table III presents the execution time (in seconds) for both versions of the dataset during training, the approximate RAM memory consumption, as well as the number of epochs reached before interruption by the EarlyStopping mechanism. This comparison allows the evaluation of the impact of preprocessing on training efficiency, in addition to enabling a comparison of the computational cost between the two approaches.

The evaluation of the results presented in Tables I and II shows that the proposed architecture achieved high performance in both versions of the dataset, reaching consistent values of accuracy, precision, recall, F1-score, and AUC across the training, validation, and test sets. A notable highlight is the AUC values, close to 0.99 in all cases, which reinforces the excellent capability of the model for binary classification tasks between posts with the presence of suicidal ideation and posts with the absence of suicidal ideation. This behavior is illustrated in the graphs shown in Figs. 3 and 4, which present the ROC curves for the test set for the models with and without stopword removal. Both ROC curves demonstrate the high discriminative capacity of the models, corroborating the results observed in the aggregated metrics.

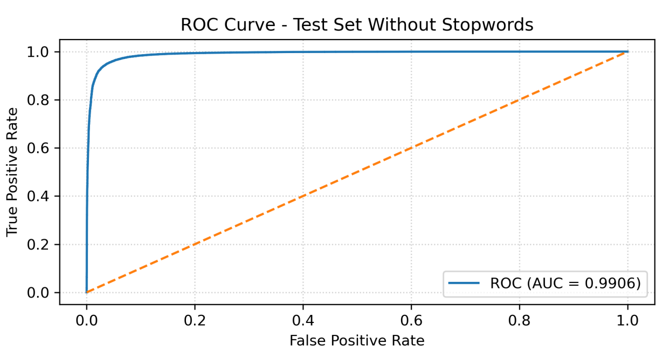

Fig. 3. ROC curve for the proposed model using the test dataset without stopwords.

When comparing the two approaches, it can be observed that both versions achieved very similar metrics, with a slight

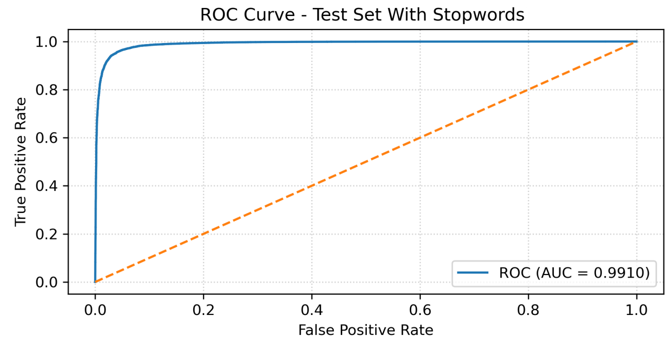

Fig. 4. ROC curve for the proposed model using the test dataset with stopwords.

advantage for the version without stopword removal, which reached an accuracy of 0.9577 and an F1-score of 0.9579 on the test set, compared to 0.9560 and 0.9560, respectively, for the version with stopword removal. Although small, this difference suggests that retaining stopwords may, in some cases, help the neural network capture contextual nuances, which can support the identification of relevant patterns for the detection of suicidal ideation. However, it is important to highlight that the computational cost during training of the model without stopword removal was higher compared to the model with removal, as shown in Table III. This increase is related to the larger vocabulary and the higher padding value used in this dataset, which makes the version with stopword removal more computationally efficient, given that its execution time and memory consumption were significantly lower than those of the model without stopword removal, even though the latter achieved slightly superior performance. This trade-off reinforces the importance of evaluating computational cost in practical large-scale monitoring applications.

In order to further deepen the performance analysis of the proposed model, the confusion matrices obtained on the test set for the dataset versions with and without stopword removal were analyzed, as shown in Figs. 5 and 6.

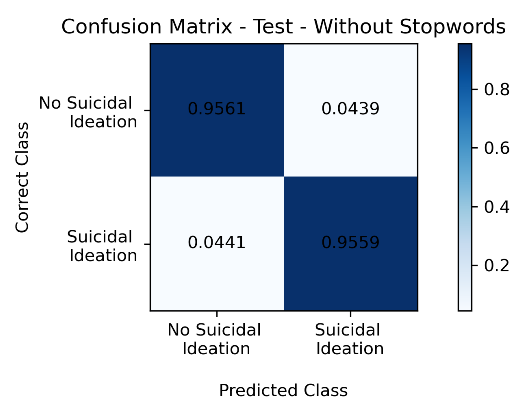

Fig. 5. Confusion matrix obtained from the test set without stopwords.

When comparing the results obtained by the two approaches, it is possible to observe that the model maintained a balance between false positives and false negatives, which is consistent with the similar values of precision, recall, and F1-

TABLE I
RESULTS OF THE PROPOSED MODEL USING THE DATASET WITHOUT STOPWORDS, CONSIDERING THE TRAINING, VALIDATION, AND TEST SETS

Dataset	Accuracy	Precision	Recall	AUC	F1-Score	Loss
Training Set	0.9923	0.9939	0.9907	0.9993	0.9923	0.0396
Validation Set	0.9554	0.9556	0.9553	0.9887	0.9554	0.1505
Test Set	0.9560	0.9561	0.9559	0.9906	0.9560	0.1453

TABLE II
RESULTS OF THE PROPOSED MODEL USING THE DATASET WITH STOPWORDS, CONSIDERING THE TRAINING, VALIDATION, AND TEST SETS

Dataset	Accuracy	Precision	Recall	AUC	F1-Score	Loss
Training Set	0.9896	0.9893	0.9898	0.9988	0.9896	0.0549
Validation Set	0.9546	0.9505	0.9592	0.9898	0.9548	0.1498
Test Set	0.9577	0.9536	0.9623	0.9910	0.9579	0.1442

TABLE III
EXECUTION TIME IN SECONDS, APPROXIMATE RAM MEMORY CONSUMPTION, AND NUMBER OF EPOCHS UNTIL EARLYSTOPPING FOR BOTH VERSIONS OF THE DATASET

Dataset versions	Execution time in seconds	Approximate RAM consumption (MB)	Number of epochs
Dataset without stopwords	715.2753	808.60	12
Dataset with stopwords	1954.9124	1061.72	12

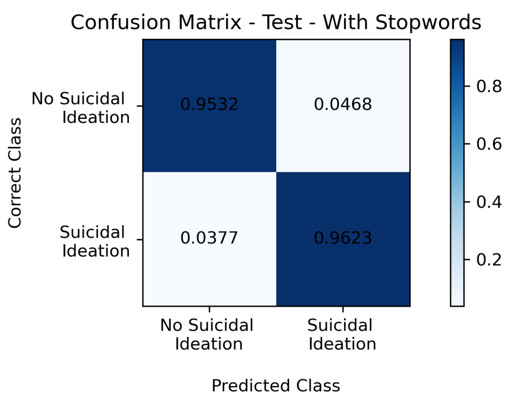

Fig. 6. Confusion matrix obtained from the test set with stopwords.

score, thus avoiding bias. The version with stopword removal presented a slight reduction in false positives, whereas the version without stopword removal showed a lower rate of false negatives. Although the numerical differences are small, the critical nature of the topic—mental health—makes the version without stopword removal preferable, as its greater effectiveness in detecting true cases—lower false negative rate—minimizes the risk of omission in situations of potential threat to life.

From a practical perspective, false negatives represent the most critical type of error, as they correspond to posts containing suicidal ideation that are not identified by the model. In this sense, the low false negative rates observed in both versions indicate that the proposed architecture has good capability for identifying potentially sensitive cases. On the other hand, false positives, although less critical from a clinical standpoint, may generate undue alerts in automatic monitoring systems, which reinforces the importance of maintaining a balance between sensitivity and specificity.

Thus, the analysis of the confusion matrices confirms that the model not only achieves high aggregated metrics but also exhibits stable behavior in the distribution of errors between classes.

In mental health–related applications, performance analysis must consider the asymmetry between different types of error. In particular, false negatives represent the most critical error, as they correspond to posts containing indications of suicidal ideation that are not identified by the model. From this risk-oriented perspective, the slight reduction in the false negative rate observed in the version without stopword removal assumes practical relevance, even in the presence of higher computational cost. This risk-oriented interpretation reinforces the suitability of the proposed model for screening and early monitoring scenarios.

Another relevant point is the small difference observed between the results of the training, validation, and test sets, which indicates that the model was able to maintain good generalization capability. This is corroborated by the low loss values obtained, as shown in Figs. 7 and 8, which present the loss curves for the training and validation sets for both versions of the dataset. This performance can be attributed to the use of regularization techniques, which proved effective in mitigating the risk of overfitting.

When comparing the results of this work with those reported by Nikhileswar et al. (2021), who used the same dataset and achieved an accuracy of 94.16% and an F1-score of 94.5% with a model based on USE and an FCNN, it can be observed that the architecture proposed in this study outperformed those results, achieving an accuracy of 95.77% and an F1-

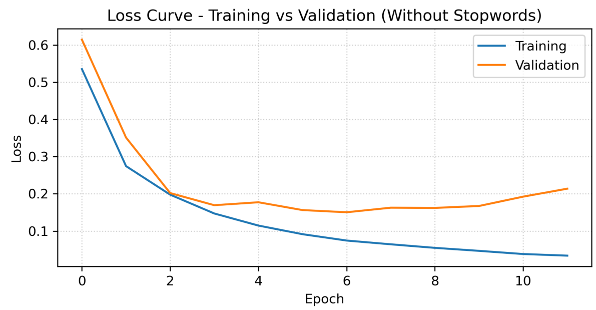

Fig. 7. Loss curve plot of the proposed model for the dataset without stopwords.

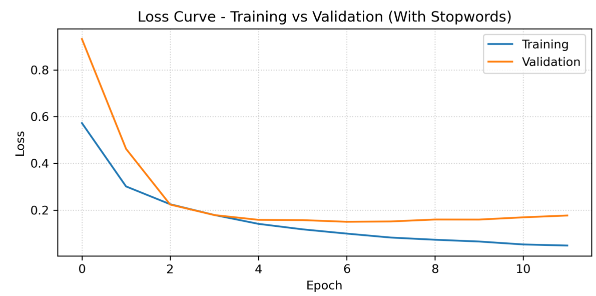

Fig. 8. Loss curve plot of the proposed model for the dataset with stopwords.

score of 95.79% in the version without stopword removal, as summarized in Table IV. This gain demonstrates that the proposed architecture is more efficient than the approach used by the aforementioned authors, resulting in better performance in detecting suicidal ideation in social media posts for the dataset under analysis.

In addition to the quantitative improvement observed, the present work differs from previous studies by combining a hybrid CNN-BiLSTM architecture with explicit regularization strategies, a risk-oriented analysis of classification errors, and an ethical discussion regarding the use of automatic models in mental health. These aspects reinforce that the contribution of the study is not limited to numerical performance, but also encompasses the methodological and conceptual suitability of the approach for sensitive contexts.

Therefore, it can be concluded that although both versions of the dataset produced consistent results, the approach with stopword removal proved to be more efficient, as it achieves metrics very close to those of the version without removal, but with significantly lower computational cost. Moreover, the comparison with previous studies further reinforces the relevance of the proposed architecture as a promising alternative for the analysis of sensitive texts related to suicidal ideation.

VI. Conclusion

This study aimed to investigate the application of Natural Language Processing (NLP) and deep learning techniques for the detection of suicidal ideation in social media posts. To this end, a hybrid architecture based on BiLSTM, combined with convolutional layers and regularization techniques, was proposed, using a dataset containing data collected from the Reddit platform.

The experiments conducted confirmed the effectiveness of the proposed architecture in the classification task, presenting high and consistent metrics in both versions of the dataset, with and without stopword removal. It was observed, however, that the version without stopword removal achieved slightly superior performance, reaching an accuracy of 95.77% and an F1-score of 95.79% on the test set, compared to the accuracy and F1-score of 95.60% obtained by the version with stopword removal. Despite achieving better results, the computational cost of the model without stopword removal was higher in terms of training time and RAM usage, indicating that the version with stopword removal provides a better balance between performance and computational cost.

In addition, the comparison with previous studies, such as that of Nikhileswar et al. (2021), which used the same dataset with an approach based on USE and FCNN, demonstrated that the hybrid CNN–BiLSTM architecture proposed in this work is more efficient in exploiting the temporal and contextual dependencies of the texts, outperforming the results previously reported in the literature.

In summary, the results obtained in this study reinforce the relevance of using BiLSTM-based architectures combined with convolutional layers as a promising approach for the automatic analysis of texts related to suicidal ideation on social media.

In this context, the proposed model should be understood as a computational screening and early-warning system, whose objective is to support large-scale monitoring processes rather than to replace clinical assessments performed by specialized professionals. The automatic identification of potentially sensitive posts can support subsequent stages of human analysis, integrating ethically and responsibly into mental health care workflows. This delimitation is essential for the appropriate use of deep learning techniques in sensitive domains such as suicide prevention.

Beyond its methodological and experimental contributions, this work reinforces the importance of using deep learning techniques as support tools for early monitoring strategies, potentially assisting suicide prevention initiatives and providing additional support for professionals and policymakers in the public health domain.

However, it is important to note that despite the promising results, this study has limitations, as the model was evaluated exclusively on data extracted from Reddit, which may limit the assessment of its generalization capability to data from other platforms, sociocultural contexts, or linguistic styles. Therefore, future work should consider validating the proposed architecture on independent datasets in order to assess the robustness and generalization capacity of the approach, as well as evaluating its performance in real-world application scenarios.

Finally, the use of automatic models for the detection of suicidal ideation on social media involves significant ethical challenges, particularly regarding user privacy, the risk of false positives, and the responsible interpretation of results. Thus, the proposed model should be understood as a support tool for monitoring and initial screening, not as a substitute for clinical evaluation, and should always be used in conjunction

TABLE IV
PERFORMANCE COMPARISON BETWEEN THE PROPOSED BILSTM MODEL AND THE MODEL BY NIKHILESWAR ET AL. (2021)

Model or Study	Dataset Version	Accuracy (%)	F1-Score (%)	Observations
Nikhileswar et al. (2021)	-	94.16	94.50	USE + FCNN Model
CNN-BiLSTM (from this work)	Without stopwords	95.60	95.60	-
CNN-BiLSTM (from this work)	With stopwords	95.77	95.79	Best result obtained

with appropriate ethical, legal, and professional policies.

REFERENCES

[1] World Health Organization, *Suicide worldwide in 2021: global health estimates*. World Health Organization, 2025, p. iv, 56. [Online]. Available: https://www.who.int/publications/i/item/9789240110069

[2] C. Stavizki Junior, "Sofrimento Social e Racionalidade Neoliberal: Contextos, Instituições e Atores das Políticas de Prevenção do Suicídio no Estado do Rio Grande do Sul – Brasil", Ph.D. dissertation, 2025. doi: 10.13140/RG.2.2.31160.46087.

[3] H. Allam et al., "AI-Driven Mental Health Surveillance: Identifying Suicidal Ideation Through Machine Learning Techniques", *Big Data and Cognitive Computing*, vol. 9, no. 1, p. 16, 2025. doi: 10.3390/bdcc9010016.

[4] L. Silva, S. Peres, and C. Boscarioli, *Introdução a Mineração de Dados com aplicações em R*. Rio de Janeiro: Elsevier, 2016. 296 p. ISBN: 978-85-352-8446-1.

[5] F. T. Fernandes and A. D. P. Chiavegatto Filho, "Perspectivas do uso de mineração de dados e aprendizado de máquina em saúde e segurança no trabalho", *Revista Brasileira de Saúde Ocupacional*, vol. 44, p. e13, 2019. doi: 10.1590/2317-6369000019418.

[6] M. A. Gualhano and A. P. V. Vasconcelos, "Análise dos cursos de licenciatura da Rede Federal utilizando mineração de dados", *Educação e Pesquisa*, vol. 46, p. e219576, 2020. doi: 10.1590/S1678-4634202046219576.

[7] M. Conway and D. O'Connor, "Social media, big data, and mental health: current advances and ethical implications", *Current Opinion in Psychology*, vol. 9, pp. 77–82, 2016. doi: 10.1016/j.copsyc.2016.01.004.

[8] L. Oliveira, C. Oliveira, L. Oliveira, and A. Pimentel. "Desafios no Ensino a Distância: Soluções Computacionais para a Aprendizagem Colaborativa com Computação Afetiva", in *Anais do VIII Workshop de Desafios da Computação aplicada à Educação*, Brasília, 2019, pp. 22-24, doi: https://doi.org/10.5753/desafie.2019.12185.

[9] M. Bercht, "Computação Afetiva: vínculos com a psicologia e aplicações na educação", *Produções do III PSICOINFO e II JORNADA do NPPI*, p. 106, 2006.

[10] J. G. Brandão, "Mineração de opinião em mídias sociais com aprendizado de máquina", M.S. thesis, Universidade Federal de Goiás, Goiânia, 2020. [Online]. Available: http://repositorio.bc.ufg.br/tede/handle/tede/11185

[11] C. V. Sundermann, "Extração de contexto de reviews para sistemas de recomendação utilizando mineração de textos e opiniões", Ph.D. dissertation, Universidade de São Paulo, São Carlos, 2019. doi: 10.11606/T.55.2020.tde-18032020-094344.

[12] M. Gaur et al., "Characterization of time-variant and time-invariant assessment of suicidality on Reddit using C-SSRS", *PLOS ONE*, vol. 16, no. 5, pp. 1-21, 2021. doi: 10.1371/journal.pone.0250448.

[13] C. Spörl, E. Castro, and A. Luchiari, "APLICAÇÃO DE REDES NEURAIS ARTIFICIAIS NA CONSTRUÇÃO DE MODELOS DE FRAGILIDADE AMBIENTAL", *Revista do Departamento de Geografia*, vol. 21, pp. 113–135, 2011. doi: 10.7154/RDG.2011.0021.0006.

[14] L. Fleck et al., "Redes neurais artificiais: princípios básicos", *Revista Eletrônica Científica Inovação e Tecnologia*, vol. 7, p. 47, 2016. doi: 10.3895/recit.v7.n15.4330.

[15] D. K. F. Matsumoto et al., "Estudo em séries temporais financeiras utilizando redes neurais recorrentes", 2019. [Online]. Available: http://www.repositorio.ufal.br/jspui/handle/riufal/6813

[16] E. Figueiredo, "Classificação de sentimentos em textos de e-commerce utilizando redes neurais artificiais," M.S. thesis, Universidade Estadual do Oeste do Paraná, Cascavel, PR, 2022. [Online]. Available: https://tede.unioeste.br/handle/tede/5924

[17] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory", *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997. doi: 10.1162/neco.1997.9.8.1735.

[18] K. Smagulova and A. P. James, "A survey on LSTM memristive neural network architectures and applications", *Eur. Phys. J. Spec. Top.*, vol. 228, pp. 2313–2324, 2019. doi: 10.1140/epjst/e2019-900046-x

[19] J. Du, Q. Liu, K. Chen, and J. Wang, "Forecasting stock prices in two ways based on LSTM neural network", in *2019 IEEE 3rd Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, 2019, pp. 1083–1086. doi: 10.1109/ITNEC.2019.8729026.

[20] G. Santos, "Uma aplicação de redes neurais recorrentes do tipo LSTM à previsão dos preços de curto prazo do mercado de energia elétrica brasileiro", M.S. thesis, Escola de Economia de São Paulo, 2019. [Online]. Available: https://hdl.handle.net/10438/28069

[21] F. A. Gers, J. Schmidhuber, and F. Cummins, "Learning to Forget: Continual Prediction with LSTM", *Neural Computation*, vol. 12, no. 10, pp. 2451–2471, 2000. doi: 10.1162/089976600300015015.

[22] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A Search Space Odyssey", *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 10, pp. 2222–2232, 2017. doi: 10.1109/TNNLS.2016.2582924.

[23] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM and other neural network architectures", *Neural Networks*, vol. 18, no. 5, pp. 602–610, 2005. doi: 10.1016/j.neunet.2005.06.042.

[24] M. Birjali, A. Beni-Hssane, and M. Erritali, "Machine Learning and Semantic Sentiment Analysis based Algorithms for Suicide Sentiment Prediction in Social Networks", *Procedia Computer Science*, vol. 113, pp. 65–72, 2017. doi: 10.1016/j.procs.2017.08.290.

[25] Y. Lim and Y. Loo, "Characteristics of Multi-Class Suicide Risks Tweets Through Feature Extraction and Machine Learning Techniques", *JOIV : International Journal on Informatics Visualization*, vol. 7, pp. 2297, 2023. doi: 10.62527/joiv.7.4.2284.

[26] L. Cao, H. Zhang, X. Wang, and L. Feng, "Learning Users Inner Thoughts and Emotion Changes for Social Media Based Suicide Risk Detection", *IEEE Transactions on Affective Computing*, vol. 14, no. 2, pp. 1280–1296, 2023. doi: 10.1109/TAFFC.2021.3116026.

[27] K. Nikhileswar, D. Vishal, L. Sphoorthi, and S. Fathimabi, "Suicide Ideation Detection in Social Media Forums", in *2021 2nd International Conference on Smart Electronics and Communication (ICOSEC)*, 2021, pp. 1741–1747. doi: 10.1109/ICOSEC51865.2021.9591887.

[28] J. Li, S. Zhang, Y. Zhang, H. Lin, and J. Wang, "Multifeature Fusion Attention Network for Suicide Risk Assessment Based on Social Media: Algorithm Development and Validation", *JMIR Med Inform*, vol. 9, no. 7, p. e28227, Jul. 2021. doi: 10.2196/28227.

[29] H. Ghanadian, I. Nejadgholi, and H. Al Osman, "Improving Suicidal Ideation Detection in Social Media Posts: Topic Modeling and Synthetic Data Augmentation Approach", *JMIR Form Res*, vol. 9, p. e63272, Jun. 2025. doi: 10.2196/63272.

[30] U. Naseem, M. Khushi, J. Kim, and A. G. Dunn, "Hybrid Text Representation for Explainable Suicide Risk Identification on Social Media", *IEEE Transactions on Computational Social Systems*, vol. 11, no. 4, pp. 4663–4672, 2024. doi: 10.1109/TCSS.2022.3184984.

[31] R. Haque, N. Islam, M. Islam, and M. M. Ahsan, "A Comparative Analysis on Suicidal Ideation Detection Using NLP, Machine, and Deep Learning", *Technologies*, vol. 10, no. 3, p. 57, 2022. doi: 10.3390/technologies10030057.

[32] S. Ryu et al., "Detection of Suicide Attempters among Suicide Ideators Using Machine Learning", *Psychiatry Investig.*, vol. 16, no. 8, pp. 588–593, Aug. 2019. doi: 10.30773/pi.2019.06.19

[33] M. M. Tadesse, H. Lin, B. Xu, and L. Yang, "Detection of Suicide Ideation in Social Media Forums Using Deep Learning", *Algorithms*, vol. 13, no. 1, p. 7, 2020. doi: 10.3390/a13010007.

[34] T. S. Gunawan et al., "Development of video-based emotion recognition using deep learning with google colab", *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, vol. 18, no. 5, pp. 2463–2471, 2020. doi: 10.12928/telkomnika.v18i5.16717.

[35] L. Cardozo and L. Freitas, “Análise de Sentimentos: Avaliando o Desempenho de Pré-Processamento e de Algoritmos de Aprendizagem de Máquina sobre o Dataset TweetSentBR”, in *Anais do X Brazilian Workshop on Social Network Analysis and Mining*, Evento Online, 2021, pp. 169–174. doi: 10.5753/brasnam.2021.16135.

[36] W. S. Oliveira Júnior et al., “O uso de data augmentation como técnica para o aprimoramento de redes neurais a fim de detectar notícias falsas sobre a covid-19”, 2021. [Online]. Available: https://ri.ucsal.br/handle/prefix/4508.

[37] V. R. Ferraz, “Análise de sentimentos e classificação multiclasse de textos aplicadas ao customer success”, 2020. [Online]. Available: https://dco-unesp-bauru.github.io/tcc-bcc-2020-1/ViniciusRF/thesis-ViniciusRF.pdf.

[38] L. H. Aros et al., “Financial fraud detection through the application of machine learning techniques: a literature review”, *Humanit Soc Sci Commun*, vol. 11, p. 1130, 2024. doi: 10.1057/s41599-024-03606-0.

[39] M. Rodríguez e B. Dantas Bezerra, “Processamento de Linguagem Natural para Reconhecimento de Entidades Nomeadas em Textos Jurídicos de Atos Administrativos (Portarias)”, *REPA*, vol. 5, no. 1, pp. 67-77, abr. 2020. doi: https://doi.org/10.25286/repa.v5i1.1204.

[40] Y. Peng and M. H. Nagata, “An empirical overview of nonlinearity and overfitting in machine learning using COVID-19 data”, *Chaos Solitons Fractals*, vol. 139, p. 110055, Oct. 2020. doi: 10.1016/j.chaos.2020.110055.

[41] J. M. Kernbach e V. E. Staartjes, “Foundations of Machine Learning-Based Clinical Prediction Modeling: Part II—Generalization and Overfitting”, *Acta Neurochir Suppl.*, vol. 134, pp. 15–21, 2022. doi: 10.1007/978-3-030-85292-4_3.

Sadala Nagib Salame Neto received a B.Sc. degree in Computer Engineering from the Federal University of Pará (UFPA), Brazil, in 2023. He is currently pursuing an M.Sc. degree at the same institution. He also works at UFPA as an Information Technology technician. His research interests include Natural Language Processing, Neural Networks, and Sentiment Analysis.

Kaique da Silva Pinto received the B.Sc. degree in Computer Engineering from the Federal University of Pará (UFPA), Brazil. He is currently a Master's student in Applied Computing at the Federal University of Pará, Brazil, with an expected completion in 2025. His research interests include software engineering, applied computing, and information systems development. He is currently working as a Systems Development Analyst at the Court of Accounts of the Municipalities of the State of Pará (TCMPA), Brazil, where he contributes to the design and implementation of information systems to support institutional processes and public sector accountability.

Silvio Tadeu Teles da Silva received a B.Sc. degree in Mathematics from the State University of Pará (UEPA, Brazil) in 2011 and an M.Sc. degree in Education from the same institution in 2013. He is currently a Ph.D. candidate in Electrical Engineering at the Federal University of Pará (UFPA, Brazil), working on the optimal dispatch of natural gas-based power generation under security constraints of energy systems. He is a faculty member at Estácio College of Belém, where he has served as Chair of the Institutional Self-Assessment Committee (2023–2024) and Coordinator of the PROJETE Extension Program (2021–2025). His previous affiliations include the Federal Institute of Pará (IFPA), where he contributed to the CNPq Research Group on Mathematical Modeling and Computational Simulation (GM2SC), and the Laboratory of Computing and Telecommunications (LCT/ITEC/UFPA), associated with the National Institute of Science and Technology on Wireless Communications (INCT-CSF). His research interests include applied mathematics, multi-criteria decision-making, fuzzy and hesitant fuzzy methods, optimization, and machine learning applications in energy systems and resource management. He has authored and co-authored books, journal articles, and book chapters on mathematics education, applied mathematics, and computational models. His work has been recognized with academic awards, such as the 1st place in the Student Satisfaction Index (ISA) at Estácio Belém. Dr. Silva has participated in interdisciplinary projects, including “Mathematical Intelligence Interpreting the Propagation of the Memory Wave” (IFPA, 2025), and serves as reviewer for journals and conferences in the fields of mathematics education, applied computational mathematics, and engineering.

Miércio Cardoso de Alcântara Neto received a B.Sc. degree in Telecommunications Engineering from the Instituto de Estudos Superiores da Amazônia (IESAM, Brazil), in 2010, and an M.Sc. degree (2013), Ph.D. degree (2015), and Post-Doctorate (2018) in Electrical Engineering with an emphasis on Telecommunications from the Federal University of Pará (UFPA, Brazil). Currently, he is an Assistant Professor at UFPA, where he also serves as Deputy Director of the Institute of Technology (ITEC) for the term 2022–2026. His previous affiliations include the School of Electrical and Biomedical Engineering at UFPA, where he worked as Director for two consecutive terms (2018–2022) and as Coordinator of the Biomedical Engineering undergraduate program (2017–2018). His research interests include Machine Learning, Electromagnetic Theory, and the Design of Resonant Devices and Circuits. He is a researcher affiliated with the National Institute of Science and Technology in Wireless Communication (INCT-CSF) and the Applied Electromagnetism Group at UFPA. He also serves as Communications Director of the Brazilian Society of Microwaves and Optoelectronics (SBMO), where he has been a member of the Deliberative Council.