DOI: 10.36558/rsc.v14i1.8621

Fact or Fake? Comparing News Classifications Using Different Artificial Neural Network Architectures

Gabriel Baldo
Fatec Adamantina
gabriel.baldo21@gmail.com

Heitor Henrique F. Cardoso
Fatec Adamantina
heitor.faccas@gmail.com

Paulo Roberto S. Ruiz
Fatec Adamantina
paulo.ruiz2@fatec.sp.gov.br

ABSTRACT

This article addresses the truthfulness news through artificial intelligence, applying different architectures of Artificial Neural Networks to classify news as false or true. It seeks to provide insight into machine learning in fake news detection, contributing to the understanding of the role of fact-checking. The article discusses the popularization of technology in the dissemination of digital information and emphasizes the importance of digital news in the construction of collective knowledge. It highlights concerns about the spread of fake news and showcases fact-checking initiatives on various platforms. The theoretical foundation explores concepts of Machine Learning, Natural Language Processing, and Artificial Neural Networks. The methodology details the use of the "Fake.Br Corpus" database and describes data preprocessing, its division, and the construction of classification models. The results indicate training, testing, and validation accuracies, as well as the comparison of ROC curves between models. The conclusion emphasizes the feasibility of news classification by artificial intelligence, with accuracies exceeding 90%. The study suggests continuing tests, focusing on the exploration of deep learning architectures with appropriate vectorization in the preprocessing stage.

Keywords

Artificial Neural Network; Fake News; Machine Learning; Natural Language Processing.

1. INTRODUCTION

The popularization of technology in recent decades has given rise to a culture of digital news consumption, making it common to share summaries and even complete texts via social networks [28]. The utilization of electronic message chains (email), digital forums, and news aggregator portals across diverse social media platforms has become pervasive. In this context, digital news plays a crucial role in shaping the knowledge foundation of modern society [24]. It offers nearly instantaneous information dissemination through social networks, contributing to the development of collective knowledge by assembling a substantial body of information that is readily accessible and analyzable.

Despite enabling various advancements in information dissemination and digital interaction strategies [3], this phenomenon has also created an opportunity for malicious use of culture and technology. This trend facilitates the spread of fake news and hinders the validation of their accuracy, especially when the user lacks familiarity with technology or fails to verify information from other sources.

In recent years, fact-checking initiatives have been implemented on various platforms and by different companies. These services can be manual or collaborative, such as Twitter's misinformation alert program [38], the 'Fact or Fake' service by the Globo Group [8], Lupa [18], FactCheck.org [6], among others. There are also automated services, such as Facebook's portal, which uses Machine Learning (ML) to assist teams in detecting fraud and applying anti-spam policies that block millions of fake accounts daily [20].

On the flip side, some of these initiatives might cast doubt on their reliability within the population. According to the Edelman Trust Barometer 2022 [5], trust in the media was viewed negatively, with 47% of Brazilians stating they trust it, revealing that less than half of the respondents trust media outlets that provide fact-checking services. In contrast, a study involving Brazilians conducted in 2023 by the Australian institution KPMG, titled 'Trust in Artificial Intelligence: A global study' [9], states that 84% of respondents trust Artificial Intelligence (AI), potentially indicating that the intersection of this technology and fact-checking initiatives could generate more trust among the population.

In this context, the objective of this work is to use different Artificial Neural Network (ANN) architectures to classify news as true or false. The article will also provide a detailed explanation of the production steps, evaluation metrics, and a comparative analysis of results using different processing architectures. Additionally, this work aims to stimulate a discussion about the relevance and appropriateness of using AI in contemporary society, particularly regarding questioning the trustworthiness of fake news classification results. Consequently, the specific objectives are defined as follows: implement three different ANN architectures, comparatively evaluate the results with performance metrics, and discuss the intersection between society and technology.

Finally, this work is justified in light of the growing use of technology in news dissemination, coupled with the increasing use of technology in fake news detection. Therefore, this work posits the following hypothesis: Is it possible to build reliable solutions for detecting news with false content?

2. BACKGROUND

To elucidate the objective and definition of Machine Learning (ML), it is relevant to mention Arthur Samuel, considered one of the pioneers in this field. In his 1959 work, titled 'Some Studies in Machine Learning Using the Game of Checkers,' Samuel highlighted: 'Enough work has been done to verify the fact that a computer can be programmed so that it learns to play checkers better than the person who wrote the program' [33]. Samuel asserted that, by employing ML, it is possible to reach a state where the trained program surpasses or equals human capability in performing specific tasks. According to [22], the goal of ML is to construct programs that enhance their performance through examples. It involves programming computers to use previous data and records to achieve novel results, such as predictions, pattern identification, value classification, among other tasks.

ML programs generally fall into the categories of supervised, unsupervised, or reinforcement learning. In supervised learning, the aim is to fit the data to a known class or value range provided by the training agent [29]. Unsupervised learning seeks to understand relationships between data and group them based on identified patterns, without the explicit presentation of classes to the model. Meanwhile, reinforcement learning involves the collection and penalization of results, characterized by trial and error [22]. In contrast to supervised methods, in reinforcement learning, the agent does not receive explicitly labeled examples. Instead, it explores the environment, takes actions, and collects the results of those actions.

2.1 Natural Language Processing

Text processing poses a challenging task for computers, as they inherently work with numerical inputs and lack native capabilities for handling textual data. [1]. However, techniques have been developed to overcome this limitation, transforming texts into numbers, commonly referred to as vectorization techniques.

However, addressing all linguistic features in Natural Language Processing (NLP) is extremely challenging, especially semantic challenges such as detecting figurative language, irony, sarcasm, symbols, among others. According to [4], no model truly understands texts or human language; instead, it performs a statistical analysis of language structure, associating patterns with numerical values.

One of the main vectorization techniques is known as one-hot encoding. It involves associating a single integer index with each word and then transforming this index into a binary vector of the vocabulary size. The vector is mainly composed of zeros, except for the i-th entry, which is filled with the value 1 [4]. This is one of the most basic forms of vectorization.

However, when it comes to ANN, the word embedding technique can be applied. This involves generating values by learning from data, resulting in denser vectors compared to those generated by one-hot encoding. Thus, the dense vectors produced by word embedding store more information in smaller dimensions [30].

2.2 Artificial Neural Networks

ANNs are algorithms designed to emulate the human neural system, aiming at acquiring knowledge through computational parallelism [16]. The learning process of a neural network unfolds in three stages: neural network stimulation, involving the input of data; subsequent adjustment and updating of weights; and finally, the network responds in a novel manner due to the alteration in its structure [14].

In 1943, Warren McCulloch and Walter Pitts introduced a simplified neural model [19], which significantly influenced the creation of the Perceptron by Rosenblatt in 1957 [31]. During the 1980s, the Backpropagation algorithm emerged, facilitating the adjustment of weights based on the discrepancy between the obtained and desired errors. Geoffrey Hinton, David Rumelhart, and Ronald Williams were instrumental in applying this algorithm, enabling the efficient training of multi-layered ANNs, commonly known as Multi-Layer Perceptrons (MLP) [32].

Lecun et al. [17], in their article titled 'Gradient-based learning applied to document recognition,' showcased the use of a deep learning architecture for pattern recognition, specifically handwritten letters presented to the algorithm as images. The proposed architecture, LeNet-5, not only contributed to the field of image recognition but also provided insights into how data features are extracted within an ANN context.

Subsequent advancements are evident in works such as 'A Fast Learning Algorithm for Deep Belief Nets' [15], which introduced the Deep Belief Nets (DBN) model, applied to an unsupervised problem. This model combines visible and hidden layers to represent complex probabilistic distributions, pushing the boundaries of hierarchical and abstract data representation learning while aiding in pattern recognition.

Deep Learning can be defined as a subset of ML, emphasizing an approach focused on learning successive layers of increasingly meaningful representations. The term 'deep' does not allude to a deeper understanding but rather signifies the incorporation of successive layers [4]. The depth of a model, indicating the number of layers it comprises, serves as an indicator of its complexity and learning capacity. Presently, deep models can encompass dozens or even hundreds of successive layers.

2.3 Fake news detection

Due to the inherent distrust towards traditional media, rumors spread rapidly in the online environment, fueling individuals' desire to stay informed without being manipulated [26]. Traditional methods based on human verification do not scale to the same volume of fake news generated on social media [37], while stated that the accessibility provided by social networks allows news to be disseminated independently of the author's reputation [7]. Considering this, there is a noticeable concern and focus on the detection of fake news, which is a current societal necessity.

In this context, various studies are conducted on combating misinformation, including the use of AI. Works such as [21] are already exploring the optimization of fake news detection through the application of the Ensemble method, which involves a set of AI models used to reduce variance and increase prediction accuracy.

Moreover, there are studies examining more common AI methods which compares the use of methods like Decision Trees, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN) [2]. Classification algorithms, notably Decision Trees and SVM, demonstrated superior performance, while KNN revealed less effective performance. The corresponding accuracy rates were 99.6% and 99.5% for Decision Trees and SVM, respectively, in contrast to KNN, which had an accuracy of 60.84%.

3. MATERIALS AND METHODS

3.1 Materials

This study utilized the "Fake.Br Corpus" dataset [23], a compilation of news in Portuguese created by students from the University of São Paulo, São Carlos campus. The dataset is accessible on the GitHub platform [10]. It consists of 7200 texts processed by its creators, with 3600 representing fake news and the remaining half comprising genuine news. Figure 1 displays an excerpt from the file.

```
index,label,preprocessed_news
0,fake,katia abreu diz vai colocar expulsao moldura nao reclamar senadora katia abreu disse expulsao pmdb resultado acao cupula atual legenda segundo oportunista amanha vou botar moldura dourada expulsao po
1,fake,ray peita bolsonaro conservador fake entrevista danilo gentili divide direita site vem avisando jair bolsonaro deveria abandonar pauta estatista vez fazer discurso convincente boa parte liberais conservadores
2,fake,reinaldo azevedo desmascarado policia federal ferrenho critico presidenciavel jair bolsonaro desmascarado grampo telefonico conversas irma aecio neves conversa debochavam juiz sergio moro operacao lav
3,fake,relatorio assustador bndes mostra dinheiro publico brasil jorrando paises comunistas relatorio intitulado bndes transformado robin hood avessas mostra bndes influenciado governo petista conseguiu benefi
4,fake,radialista americano fala sobre pt vendem ilusao povo brasileiro so pensam divertir sitio radialista americano alex jones chorou falar criancas saude brasil jones citou microcefalia espalha mundo partir descas
5,fake,willian waack continuara internado sirio libanes apresentador jornal globo ficara afastado funcoes tempo indeterminado waack continuara internado hospital sao paulo onde submetido ultima procedimento
6,fake,problema cardiaco tira william waack jornal globo apresentador sera operado amanha waack ficara afastado temporariamente jornal globo apos submeter exames rotina medicos determinaram jornalista pass
7,fake,edir macedo podera disputar eleicoes presidenciais caso lula absolvido instancia dificil acontecer eleicoes segundo pesquisas preliminares apontam segundo turno disputadissimo petista deputado jair bolso
8,fake,boechat sobre lula alma pode ser honesta governo diabo lula todo direito dizer alma pode ser honesta porem nao significa nada disse governo diabo partido diabo tesoureiro pt vaccari tentou isentar cana ze c
9,fake,pt diz lula nao dinheiro sobreviver chama moro mentiroso mesquinho conspirador partido trabalhadores soltou nota repudio decisao juiz sergio moro sobre bloqueio patrimonios r mil partido alega lula pobre
10,fake,numero fuzilados durante regime fidel castro pode ter chegado casa mil acordo documentos projeto cuba archive organizado residentes nova jersey chamado paredon cubano vitimou menos acordo fonte po
...
7183,true,onu temer fala reducao desmatamento amazonia reformas tornarao brasil aberto presidente brasileiro fez discurso abertura assembleia geral onu novo brasil surgindo reformas pais aberto mundo preside
7184,true,fase operacao ouro verde cumpre mandados prisao apreende veiculos luxo campinas rio preto campinas prefeitura presos medico preso sao jose rio preto advogado alvo mandado prisao ainda nao localiza
7185,true,fome mudou visao propaganda coreia norte hoje satiriza pais song byeok fazia ilustracoes glorificando lideres incitando confronto imperialismo ate desiludir regime tentar fugir preso traicao libertado pes
7186,true,corte verbas pf ja afeta lava jato afirma coordenador nesta semana ministro justica afirmou pf vai ter remanejar verbas manter operacoes prioritarias governo faz contingenciamento tentar conter rombo co
7187,true,retrospectiva relembre fatos marcaram acre ano guerra faccoes menino acre extremos clima relembre noticias marcaram acre historias emocionaram surpreenderam chocaram ano tanto acre ano guerra fa
7188,true,abril aprendi desde cedo ainda infancia sertao rio peixe anos ser contra defeito combatido sal frutas eno receitava anuncio sucesso publicado revistas pai trazia rio janeiro sao paulo viagens caminhao hoje
7189,true,caminhada paz rr promove debates sobre redes sociais reformas migracao concentracao passeata ocorre cruzamento avenidas bairro santa luzia fieis catolicos participam nesta tradicional caminhada paz an
7190,true,dancarina morta apos show atingida veiculo parado localizou motorista tentou ajudar alessandra apos atropelamento carro parado apos colisao anchieta jogado cima motorista parou tentar socorrer profes
7191,true,achamos mundo pior realmente pesquisa paises mostra percepcoes sobre questoes peso midia caracteristicas populacao tendem ser sombrias brasil taxa homicidios hoje bem alta ano quase metade meni
```

Figure 1. Excerpt from the database.

To carry out this research, the Python programming language was employed for tasks such as data processing, model training, classification, and result evaluation. Python was chosen for its simplicity in implementing machine learning algorithms, facilitated by libraries like Pandas [27], Scikit-learn [34], and Keras [36]. The Google Collaboratory platform, commonly known as Google Colab [11], served as the development environment, primarily due to its ability to leverage cloud computing resources.

3.2 Methods

Figure 2 illustrates the methodological flowchart adopted in this study. The following steps were encompassed within its scope: pre-processing, dataset splitting, training, model construction, classification, and validation of the classification, all of which will be addressed in the following sections.

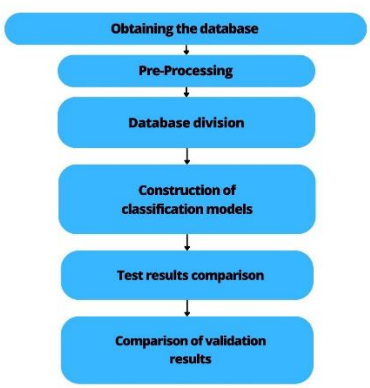

Figure 2. Methodological flowchart.

3.2.1 Preprocessing

The news has two classes, either false or true. Since they are in textual format, it was necessary to apply the LabelEncoder technique [34] to designate the classes as 0 for false and 1 for true. To vectorize the news the CountVectorizer technique [34] was applied, creating a binary matrix based on the token count, i.e., words.

For the Deep Learning model, two types of preprocessing were applied with the aim of comparing their results. The first method used was the CountVectorizer [34], as explained earlier. The second form of vectorization employed the Tokenizer method [36], which creates vectors of tokens without transforming the values into numeric. This transformation is done through a layer in the deep neural network called Embedding or Word Embedding, which performs dense vectorization [4].

3.2.2 Database division

The database was partitioned into three distinct sets to facilitate the construction and evaluation of models. The chosen strategy involved allocating 70% for training, 15% for testing the training, and the remaining 15% for validating the classification models. This approach ensures a well-balanced distribution of data throughout the various phases of model development.

3.2.3 Construction of classification models

For the construction of Model 1, the Sequential() function from the Keras library was employed. This model comprises a Dense layer, fully connected, with 10 neurons and Rectified Linear Unit (ReLU) activation [12] for processing input data. Subsequently, a Dropout layer with a rate of 0.1 is added, randomly deactivating neurons to prevent overfitting [35]. Following this, another Dense layer with 8 neurons and ReLU activation is included. A second Dropout layer with a rate of 0.1 is incorporated. The final layer is a Dense layer with 1 neuron and sigmoid activation, suitable for binary classification problems.

The model is compiled using the following parameters: loss - "mean_squared_error," seeking to minimize the mean of the squared differences between predictions and actual values; optimizer - "adam," adjusting the neural network weights to minimize the loss function; evaluation metric - "accuracy," used to assess prediction accuracy.

To execute model training, the data, divided into text files (X_train) and their respective classes (y_train), are provided. Additionally, the following parameters are configured: epochs, defining the number of times the algorithm passes through the entire training set; batch_size, specifying the number of samples processed per iteration; verbose, indicating the training visualization mode; and validation data, where 15% of the data set aside for training testing is passed. Figure 3 illustrates the construction of Model 1, compilation parameters, and training execution.

```
model = Sequential()
model.add(Dense(units=10, activation="relu",
                input_dim=X_train.shape[1]))
model.add(Dropout(0.1))
model.add(Dense(units=8,activation="relu"))
model.add(Dropout(0.1))
model.add(Dense(units=1, activation= 'sigmoid'))
model.compile(loss="mean_squared_error",
              optimizer="adam", metrics=["accuracy"])
model.fit(X_train, y_train, epochs=20, batch_size=10,
          verbose=True, validation_data=(X_test, y_test))
```

Figure 3. Construction of Model 1.

Model 2 uses pre-vectorized data with the CountVectorizer() function. It begins with a dense layer of 10 neurons and ReLU activation. Subsequently, a Flatten layer is added to reduce the data's dimensionality. Following that, another dense layer with 10 neurons and ReLU activation is introduced. Then, a Dropout layer is applied, and finally, a dense output layer with sigmoid activation is employed for classification. The compilation parameters remain consistent with those of Model 1, namely the "mean_squared_error" loss function, "adam" optimizer, and "accuracy" evaluation metric (Figure 4).

```
model = Sequential()
model.add(Dense(units=10, activation="relu",
                input_dim=X_train.shape[1]))
model.add(Flatten())
model.add(Dense(units = 10, activation  = 'relu'))
model.add(Dropout(0.1))
model.add(Dense(units = 1, activation = 'sigmoid'))
model.compile(loss = 'mean_squared_error',
              optimizer = 'adam', metrics = ['accuracy'])
```

Figure 4. Construction of Model 2.

On the other hand, Model 3 features an Embedding type input layer. This layer receives parameters such as the vocabulary size, defined by text tokenization, and its input and output dimensions. The subsequent layers and compilation are identical to the previous configuration (Figure 5).

The training execution parameters for the three deep learning models were set similarly, except for using a batch_size of 128 in Models 2 and 3 to enhance the efficiency of gradient updates during the training process.

```
model = Sequential()
model.add(Embedding(input_dim=len(tokenizer.word_index) + 1,
                    output_dim=50, input_length=500))
model.add(Flatten())
model.add(Dense(units=10, activation='relu'))
model.add(Dropout(0.1))
model.add(Dense(units=1, activation='sigmoid'))
model.compile(loss='mean_squared_error', optimizer='adam',
              metrics=['accuracy'])
```

Figure 5. Construction of Model 3.

3.2.4 Validation of Classification Models

Model validation was conducted using 1,080 news articles, corresponding to 15% of the dataset. This step allowed for assessing the accuracy of the classification models on novel data that were not used during the training phase.

The results of the classifications obtained from each algorithm were organized into confusion matrices. A confusion matrix presents true classes in rows while the classifier's output is in columns, placing correct predictions on the diagonal and errors elsewhere [29]. Evaluation was measured through two key metrics: accuracy and ROC curve (Receiver Operator Characteristic Curve).

Accuracy is a metric that measures the proportion of correct predictions in relation to the total predictions made by the model. In other words, it is the overall correctness rate of the model. High accuracy indicates that the model is making accurate predictions in most cases [4]. In contrast, the ROC curve is a metric that assesses the performance of binary classification models at different probability cutoff points. It represents the relationship between the True Positive Rate (TPR), i.e., the proportion of correctly classified examples, and the False Positive Rate (FPR), i.e., the proportion of incorrectly classified examples. There is also the area under the ROC curve, referred to as AUC-ROC (Area Under the Curve), a metric quantifying the overall performance of the model. The higher the AUC-ROC, the better the model distinguishes classes; thus, closer to 1 signifies better classification, while closer to 0.5 indicates higher randomness [13].

4. RESULTS

Figure 6 shows the accuracy and training loss of the models. It can be observed that Models 1 and 2 exhibit more abrupt and swift convergence, occurring around epoch 3. On the other hand, Model 3 demonstrates a gradual and somewhat delayed convergence, around the seventh epoch.

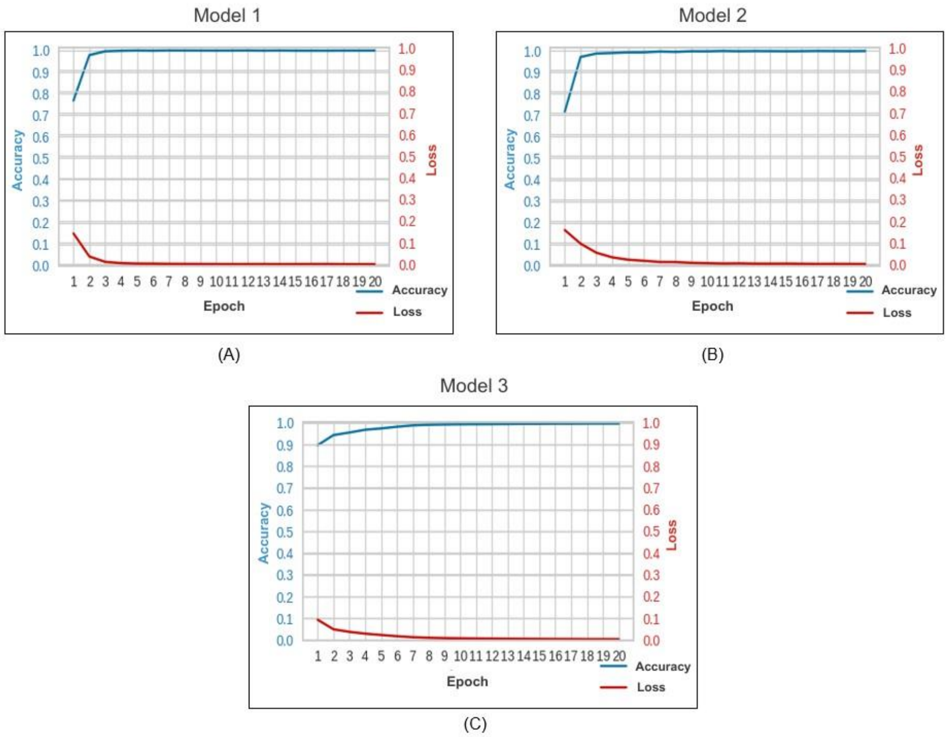

Figure 6. Training accuracy and loss: (A) – Model 1; (B) – Model 2; (C) – Model 3.

Figure 7 presents the accuracy values observed during the training and validation stages of the model. In all three cases, training accuracies of 99% were achieved. However, during validation, slight differences were observed. Model 1 achieved an accuracy of 95.37%, while Models 2 and 3 showed very close indices, with 96.38% and 96.11%, respectively.

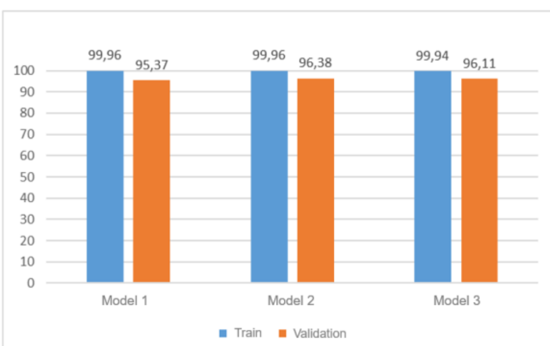

Figure 7. Accuracy values of training and validation.

Figure 8 presents the confusion matrices for the validation of the models. Model 1 correctly classified 1030 records, showing a True Positive Rate (TPR) of 93.88% and a False Positive Rate (FPR) of 2.76%. Model 2 correctly classified 1041 records, with a TPR of 96.88% and an FPR of 4.11%. In contrast, Model 3 correctly classified 1038 records, with a TPR of 95.36% and an FPR of 3.08%.

Model 2 achieved the highest TPR, with an advantage of 3% over Model 1 and 1.52% over Model 3. Model 1 achieved the lowest FPR, being 1.35% lower than Model 2 and 0.32% lower than Model 3. Therefore, it is noted that Models 1 and 2 are less balanced than Model 3, despite having higher values in certain rates. Considering the validation results, Model 2 achieved the highest accuracy. However, the gains compared to Model 1 were 1.01%, while the gains compared to Model 3 were only 0.27%.

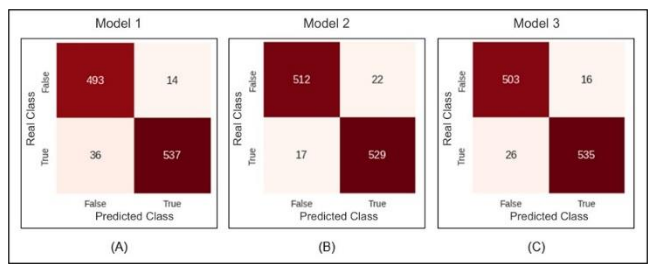

Figure 8. Confusion matrices for the validation: (A) – Model 1; (B) – Model 2; (C) – Model 3.

Regarding the ROC curve results, all models achieved a value of 99% in testing, showing no apparent superiority. However, Model 1 has a 4% lower rate in validation results when compared to the other two models, as observed in Table 1.

Table 1. ROC curve results.

Stage	Model 1	Model 2	Model 3
Testing	99%	99%	99%
Validation	95%	99%	99%

The consistency in the test results suggests that the models are effectively learning patterns in the data, while the differences between the models in the validation results suggest that Models 2 and 3 are indeed providing an advantage in terms of generalization compared to Model 1. These results indicate the effectiveness of the additional layers applied in Models 2 and 3.

To validate the results of this study, one can compare the accuracy rates obtained with previous works, such as [21] and [2], as well as other studies like [7], which applied their models to

COVID-19 databases, and [25], who unified various databases and methodologies, presenting the highest accuracies in their works as 91.83%, 99%, 94.02%, and 91%, respectively. In other words, the present study yields result consistent with recent literature.

5. CONCLUSIONS AND FUTURE WORK

This study conducted three classifications of news using different neural network architectures. The aim was to compare the efficiency of architectures in identifying the truthfulness of news. The results demonstrated accuracies exceeding 90%, consistent with recent literature.

Applications involving neural networks in the task of text/news classification can bring numerous benefits to society. The development and application of verification tools have the potential to contribute to combating the spread of false information and strengthening trust in online news and information sources.

With the completion of this study, its hypothesis can be answered: yes, it is possible to build reliable solutions for detecting news with false content. Classification models based on ANNs can make this detection, validating the hypothesis. However, it is important to note that there is still an error rate. In this sense, it is feasible to expand the tests to a larger number of news to produce more results with different samples and continue the work of training for fake news detection. Similarly, it is also necessary to emphasize the verification of news through reliable sources.

6. REFERENCES

[1] Araujo, D. C. 2019. *Avaliação de comitês com classificadores tradicionais e profundos para análise de sentimentos*. Master Dissertation, Federal University of Pernambuco, Recife, BR.

[2] Bastos, L. M. and Lima, R. M. 2023. Detecção de Fake News usando os algoritmos Decision Tree, Support Vector Machine e K-Nearest Neighbors. *Revista de Estudos Multidisciplinares UNDB*. 3, 1 (fev. 2023).

[3] Carvalho, R. B., Reis, A. M. P., Larieira, C. L. C. and Pinochet, L. H. C. 2021.Transformação digital: desafios na formação de um constructo e cenários para uma agenda de pesquisa. *Revista de Administração Mackenzie*. 22, 6, 1-15. DOI= https://doi.org/10.1590/1678-6971/eRAMD210400.

[4] Chollet, F. 2018. Deep Learning with Python. Manning Publications Co.

[5] Edelman. 2022. Edelman Trust Barometer. Available at: https://www.edelman.com.br/edelman-trust-barometer-2022. Accessed on: September 6, 2023.

[6] FactCheck. 2023. Our Mission. Available at: https://www.factcheck.org/about/our-mission/. Accessed on: September 6, 2023.

[7] Freire, P. and Goldschmidt, R. 2019. Combatendo Fake News nas Redes Sociais via Crowd Signals Implícitos. In: *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, 16., 2019, Salvador. Sociedade Brasileira de Computação, Porto Alegre, 424-435. DOI: https://doi.org/10.5753/eniac.2019.9303.

[8] G1. 2023. G1 Fato ou Fake - O serviço de checagem de fatos do Grupo Globo. Available at: https://g1.globo.com/fato-ou-fake/. Accessed on: September 6, 2023.

[9] Gillespie, N., Lockey, S., Curtis, C., Pool, J., & Akbari, A. 2023. *Trust in Artificial Intelligence*: A Global Study. The University of Queensland and KPMG Australia. DOI= https://doi.org/10.14264/00d3c94.

[10] Github. 2023. Build software better, together. Available at: https://github.com. Accessed on: November 15, 2023.

[11] Google Colaboratory. 2023. Colab. Available at: https://colab.research.google.com. Accessed on: November 15, 2023.

[12] Glorot, X., Bordes, A. and Bengio, Y. 2011. Deep sparse rectifier neural networks. In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, p. 315-323.

[13] Hand, D. J. and Till, R. J. 2001. A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Machine Learning*, 45, 171–186. DOI= https://doi.org/10.1023/A:1010920819831.

[14] Haykin, S. 1994. *Neural Networks: A Comprehensive Foundation*. India: Prentice Hall.

[15] Hinton, G. E., Osindero, S., and Teh, Y. W. 2006. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554. DOI= https://doi.org/10.1162/neco.2006.18.7.1527.

[16] Kovács, Z. L. 2006. *Redes neurais artificiais*. Editora Livraria da Física.

[17] Lecun, Y., Bottou, L., Bengio, Y. and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324. DOI= https://doi.org/10.1109/5.726791.

[18] Lupa. 2015. Como fazemos nossas checagens. Available at: https://lupa.uol.com.br/institucional/2015/10/15/como-fazemos-nossas-checagens. Accessed on: Novembre 10, 2023.

[19] McCulloch, W. S. and Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115–133. DOI= https://doi.org/10.1007/BF02478259.

[20] Meta. 2018. Questões complexas: Qual é a estratégia do Facebook para combater notícias falsas? Available at: https://about.fb.com/br/news/2018/05/questoes-complexas-qual-e-a-estrategia-do-facebook-para-combater-noticias-falsas/. Accessed on: November 10, 2023.

[21] Merola, U., Freire, P. M. S., Goldschmidt, R. R. and Soares, J. 2023. Métodos de Detecção de Fake News: Uma Comparação entre as Abordagens de Crowd Signals e Ensembles. In: *Anais do Simpósio Brasileiro de Banco de Dados (SBBD)*, 38, 2023, Belo Horizonte/MG. Porto Alegre: Sociedade Brasileira de Computação, 372-377. DOI= https://doi.org/10.5753/sbbd.2023.233398.

[22] Mitchell, T. M. 1997. *Machine learning*. New York: Mcgraw Hill.

[23] Monteiro, R. A., Santos, R. L. S., Pardo, T. A. S., De Almeida T. A., Ruiz E. E. S. and Vale, O. A. 2018. Contributions to the Study of Fake News in Portuguese: New Corpus and Automatic Detection Results. In: *Computational Processing of the Portuguese Language. PROPOR 2018*.

Lecture Notes in Computer Science, vol 11122. Springer, Cham. DOI= https://doi.org/10.1007/978-3-319-99722-3_33.

[24] Nielsen, R. K. 2017. Digital News as Forms of Knowledge: A New Chapter in the Sociology of Knowledge. In: *Remaking the News: Essays on the Future of Journalism Scholarship in the Digital Age*. Cambridge, Mass: MIT Press.

[25] Oshita, I. T. 2021. *Classificação de Fake News por Mineração de Texto.* Undergraduate Thesis (Bachelor's in Production Engineering) - Federal Technological University of Paraná, Londrina, BR.

[26] Paula, L. T., Silva, T. D. R. S. and Blanco, Y. A. 2018. Pós-verdade e Fontes de Informação: um estudo sobre fake news. Revista Conhecimento em Ação, 3(1), 93–110.

[27] Pandas. 2023. Python Data Analysis Library - pandas: Python Data Analysis Library. Available at: https://pandas.pydata.org/. Accessed on: November 10, 2023.

[28] Prado, J. A. B. 2022. Abordagem formal e multimodal da notícia como gênero, estrutura e metalinguagem em contextos de mídia social e digital. O exemplo do Twitter. *Bakhtiniana, Rev. Estud. Discurso*. 17, 4 (out-dez 2022), e16635. DOI= https://doi.org/10.1590/2176-45734316635

[29] Priddy, K. L. and Keller, P. E. 2005. *Artificial neural networks: an introduction*. SPIE press.

[30] Rodrigues da Silva, J. and De Medeiros Caseli, H. 2020. Generating Sense Embeddings for Syntactic and Semantic Analogy for Portuguese. *arXiv e-prints*, arXiv:2001.07574.

[31] Rosenblatt, F. 1957. *The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain*. Technical Report. Buffalo, N.Y.

[32] Rumelhart, D. E., Hinton, G. E. and Williams, R. J. 1986. Learning representations by back-propagating errors. *Nature* 323, 6088 (1986), 533–536. DOI= https://doi.org/10.1038/323533a0.

[33] Samuel, A. L. 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*. 3, 3 (jul 1959), 210–229. DOI= https://doi.org/10.1147/rd.33.0210.

[34] Scikit-Learn. 2023. Scikit-Learn: machine learning in Python. Available at: https://scikit-learn.org/stable/. Accessed on: Mar 10, 2023.

[35] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*. 15, 56 (2014), 1929–1958. DOI= https://doi.org/10.5555/12345678.

[36] Tensorflow core. 2023. Tensorflow. Available at: https://www.tensorflow.org/guide/keras?hl=pt-br. Accessed on: Mar 10, 2023.

[37] Tschiatschek, S., Singla, A., Rodriguez, M. G., Merchant, A., and Krause, A. 2018. Fake news detection in social networks via crowd signals. In: *Proceedings of the 2018 World Wide Web Conference (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 499–509. DOI= https://doi.org/10.1145/3184558.3188722.

[38] Twitter. 2023. How we address misinformation on Twitter. Available at: https://help.twitter.com/en/resources/addressing-misleading-info. Accessed on: September 06, 2023.