FakeRecogna Anomaly: Fake News Detection in a New Brazilian Corpus

Gabriel Lino Garcia[1][a], Luis C. S. Afonso[1][b], Leandro A. Passos[2][c], Danilo S. Jodas[1][d], Kelton A. P. da Costa[1][e] and João P. Papa[1][f]

[1]*School of Sciences, São Paulo State University, Bauru, Brazil*

[2]*CMI Lab, School of Engineering and Informatics, University of Wolverhampton, Wolverhampton, England, U.K.*

Keywords: Fake News, Corpus, Portuguese.

Abstract: The advances in technology have allowed digital content to be shared in a very short time and reach thousands of people. Fake news is one of the content shared among people and it has a negative impact on our society. Therefore, its detection has become a research topic of great importance in the natural language processing and machine learning communities. Besides the techniques employed for detection, it is also important a good corpus so that machine learning techniques can learn to differentiate between real and fake news. One can find corpora in Brazilian Portuguese; however, they are either outdated or balanced, which does not reflect a real-life situation. This work presents a new updated and imbalanced corpus for the detection of fake news where the detection can be treated as an anomaly detection problem. This work also evaluates the proposed corpus by using classifiers designed for anomaly detection purposes.

1 INTRODUCTION

The advent of social media has made possible to easily create and share information and reach millions of people around the world. As information can be easily shared, it can also be manipulated and distorted to shape opinions about products and services as well. Sensational headlines, satirical texts with misleading content, and usually with no theoretical basis on the subject addressed are a few examples. This bad practice is often referred to as "fake news" and defined by Allcott and Gentzkow (Allcott and Gentzkow, 2017) as news articles that are intentionally and demonstrably false and may mislead readers. Lazer et al. (Lazer et al., 2018) define "fake news" as fabricated information that mimics news media content in form but not in organizational process or intent.

The first significant worldwide impact caused by fake news in this century occurred in 2016 during the U.S. elections which influenced the choice of candidates for the presidency of that country. More recently, the current pandemic situation has also generated a considerable number of false news about treatments and vaccines. It is worth mentioning that fake news became popular in this century, but false reports are circulating in the community since the 6th century (Allcott and Gentzkow, 2017).

In recent years, fake news detection has caught the attention of a large number of academics to help online users identify what is real or fake information. Many of the works take advantage of the advances in artificial intelligence (AI) and machine learning (ML) along with natural language processing (NLP) techniques to address this problem. Singhal (Singhal et al., 2019) proposes a framework to identify fake news by exploring textual and visual resources using language models and image resources. Oshikawa (Oshikawa et al., 2018) outlines the challenges involved in detecting fake news by reviewing and comparing systems, as well as highlighting the importance of NLP solutions for detecting fake news. Kesarwani (Kesarwani et al., 2020) presents an approach to detecting fake news on social media with the help of the K-Nearest Neighbors (KNN) classifier. Ruchansky (Ruchansky et al., 2017) works with a model called CSI, which stands for Capture, Score, and Integrate, that comprises a module that extracts the temporal representation of news and a module for representing and scoring user behavior. Zhou (Zhou

[a] https://orcid.org/0000-0003-1236-7929
[b] https://orcid.org/0000-0002-5543-3896
[c] https://orcid.org/0000-0003-3529-3109
[d] https://orcid.org/0000-0002-0370-1211
[e] https://orcid.org/0000-0001-5458-3908
[f] https://orcid.org/0000-0002-6494-7514

Garcia, G., Afonso, L., Passos, L., Jodas, D., P. da Costa, K. and Papa, J.
FakeRecogna Anomaly: Fake News Detection in a New Brazilian Corpus.
DOI: 10.5220/0011660700003417
In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2023) - Volume 4: VISAPP*, pages 830-837
ISBN: 978-989-758-634-7; ISSN: 2184-4321

and Zafarani, 2020) reviews and evaluates methods that detect fake news from four perspectives: the false knowledge it carries, writing style, propagation patterns, and the credibility of the source.

Many works deal with news published in English. However, we focus on the detection of fake news among texts published in Portuguese. One of the main disadvantages of works related to the detection of fake news in Portuguese is the lack of corpus and mainly the amount of data. Garcia et al. (Garcia et al., 2022) proposes creating a new fake news dataset named FakeRecogna that contains a greater number of samples, more up-to-date news and covers a few of the most important categories. Monteiro et al (Monteiro et al., 2018) developed a corpus called Fake.Br, that has $7,200$ that has news extracted between 2015-2018, while The FakePedia corpus (Charles et al., 2022), on the other hand, gathers news collected between 2013 and 2021 with a total of $12,398$ samples.

However, all mentioned corpora are balanced, that is, they have the same amount of real and fake news, which does not reflect the real world. Taking notice of this point, this work proposes a new large corpus named FakeRecogna Anomaly, which is highly imbalanced and interesting for studies in the context of anomaly detection. The corpus is comprised of $101,000$ samples where $100,000$ are real news and $1,000$ are fake ones. This work also presents an experiment using six classifiers designed for the context of anomaly (outlier) detection evaluated over the proposed corpus.

The main contributions of our work are three-fold: (i) to create a dataset in Portuguese for anomaly detection called FakeRecogna Anomaly, (ii) to apply and quantitatively compare the results of the anomaly detection techniques in our dataset, and (iii) to foster the research on fake news in Brazilian Portuguese.

The remainder of this paper is organized as follows: Section 2 provides a review of related works. Section 3 introduces the proposed dataset. Sections 4 and 5 present the methodology and results of a toy-evaluation performed over the proposed dataset, respectively. Finally, Section 6 states conclusions.

2 RELATED WORKS

Concerning fake news detection, Zhang and Ghorbani (Zhang and Ghorbani, 2020) present a comprehensive overview of the online fake news detection ecosystem. Mishra et al. (Mishra et al., 2022) present a research on fake news using the probabilistic latent semantic analysis approach; besides this, a comparison of different machine learning and deep learning techniques assesses the performance of fake news detection. Oshikawa et al. (Oshikawa et al., 2018) describe the challenges involved in fake news detection and also compare the task formulations, datasets, and NLP solutions. In Brazil, Faustini and Covões (Faustini and Covões, 2019) proposed to detect fake news by training a model with only fake samples in the training dataset through One-Class Classification. Endo et al. (Endo et al., 2022) explore the use of machine learning and deep learning techniques to identify fake news in online communications in the Brazilian Portuguese language relating to the COVID-19 pandemic. Differently from many works, Li et al. (Li et al., 2021) proposed an unsupervised fake news detector using an autoencoder-based method.

In the field of anomaly detection, Wang et al. (Wang et al., 2019) presented a review of outlier detection methods from 2000 to 2019. Chalapathy and Chawla (Chalapathy and Chawla, 2019) presented a comprehensive overview of deep learning-based methods for anomaly detection and review the adoption of such methods in various domains and assess their effectiveness. Later, Chalapathy et al. (Chalapathy et al., 2018) proposed a one-class neural network to detect anomalies in complex datasets. Mohaghegh and Abdurakhmanov (Mohaghegh and Abdurakhmanov, 2021) proposed a character-level representation of unsupervised text datasets for anomaly detection problems. Pang et al. (Pang et al., 2021) reviewed twelve diverse modeling perspectives on harnessing deep learning techniques for anomaly detection. Kannan et al. (Kannan et al., 2017) presented a matrix factorization method, which is naturally able to distinguish the anomalies with the use of low-rank approximations of the underlying data. Ruff et al. (Ruff et al., 2019) introduced a new anomaly detection method—Context Vector Data Description which builds upon word embedding models to learn multiple sentence representations that capture multiple semantic contexts via the self-attention mechanism. In Brazil, Carvalho et al. (de Carvalho et al., 2018) evaluate how machine learning techniques can be employed in the task of identifying anomalies in bee behavior. Kintopp (Kintopp, 2017) demonstrated the application of anomaly detection to find anomalies in the expenditure information disclosed by Brazilian municipalities, thus being able to indicate suspected cases of administrative improbity. However, the use of anomaly detection techniques for textual data is not common in the Brazilian context.

3 FakeRecogna ANOMALY

The development of the FakeRecogna Anomaly corpus as an imbalanced dataset of real and fake news is based on the FakeRecogna Corpus (Garcia et al., 2022). The main idea of the FakeRecogna Anomaly corpus is to explore a real-life situation where the number of real news is greater than the number of fake ones. The proposed imbalance between real and fake news in the contextualizes the Brazilian scenario, in which real news are produced in a greater amount compared to fake ones.

The FakeRecogna Anomaly corpus has all the attributes contained in the FakeRecogna corpus, but with a ratio of 100 real news to 1 fake news, that is, our corpus has 100,000 true news and 1,000 false news. The collection of news was carried out by crawlers developed to mine news pages from well-known agencies of great national relevance.

Following the pattern established in the FakeRecogna corpus, the main news agencies used as a basis for extracting news were G1[1], UOL[2] and Extra[3], which are publicly recognized as reliable news outlets, in addition to the Ministry of Health do Brasil[4]. For the selection of fake news, the main means of fact-checking already extracted in the development of FakeRecogna were analyzed, and we chose to randomly extract data based on the proportion of news from each site. Table 1 presents the news agencies as well as the number of fake news collected from each source.

Table 1: The fact-checking agencies used in the FakeRecogna Anomaly.

agency	web address	# news
AFP Checamos	https://checamos.afp.com/afp-brasil	11
Boatos.org	https://boatos.org	504
E-farsas	https://www.e-farsas.com	167
Fato ou Fake (“Fact or Fake”)	https://oglobo.globo.com/fato-ou-fake	110
Projeto Comprova	https://projetocomprova.com.br	91
UOL Confere	https://noticias.uol.com.br/confere	117
total		1,000

Differently from FakeRecogna, each sample of FakeRecogna Anomaly has seven metadata fields instead of eight, the ”Subtitle” column was not placed in the dataset because the vast majority of real news does not have a subtitle, which are described in Table 2.

Another difference between the corpora is that texts in FakeRecogna Anomaly are distributed in five categories. Table 3 presents the distribution of news by category and its respective quantity.

[1]https://g1.globo.com/

[2]https://www.uol.com.br/

[3]https://extra.globo.com/

[4]https://www.gov.br/saude/pt-br

Table 2: Metadata used to describe each sample.

columns	description
Title	Title of article
News	Information about the article
Category	News grouped according to your information
Author	Publication author
Date	Publication date
URL	Article web address
Class	0 for fake news and 1 for real news

Table 3: Distribution of news per category in the FakeRecogna Anomaly.

category	# news
Brazil	2,912
Entertainment	20,081
Health	37,599
Politics	33,311
Science	7,097
total	101,000

4 METHODOLOGY

This section presents the steps employed for the evaluation of the proposed corpus.

4.1 Pre-Processing

The preprocessing is responsible for preparing the text for the next steps. This step comprises four operations:

- Truncation: It intends to avoid bias that the length of news may cause in the training and classification steps since real news are usually longer than fake ones.
- Removal of specific terms: It removes any words that may characterize a piece of news as a fake one, such as, “enganoso”, “boato”, “#fake” and, so on. This operation also removes punctuation, special characters, and URLs.
- Lemmatization: The lemmatization comprises the morphological analysis of the words by taking into account the context, that is, differentiating the meaning of identical words depending on the context.
- Removal of stop words: It removes any words considered irrelevant for the understanding of the news, such as articles and prepositions.

4.2 Text Representation

This step is responsible for computing a numerical representation for each text over the output of the pre-

processing step. The evaluation takes into account the following three techniques:

- Bag-of-Words: The Bag-of-Words model is one of the most popular representation methods for object categorization. The key idea is to quantize each extracted key point into one of the visual words, and then represent each image by a histogram of the visual words (Zhang et al., 2010).
- Term Frequency-Inverse Document Frequency (TF-IDF): It is a statistical analysis method for keywords, used to evaluate the importance of a word to a document or a corpus (Li, 2021). TF-IDF for word representation is computed using the 1,000 most frequent words.
- FastText: It extracts morphological features by processing subwords of each word, where "subword" is a character-level n-gram of the word (Choi and Lee, 2020). The representation of a word is given by the sum of the numerical features of all it is n-grams. This approach allows extracting the meaning of shorter words and allows embeddings to understand suffixes and prefixes. The parameter values for FastText were: embedding size equals 200 to dimensions, the maximum number of unique words as 1,000, and the maximum amount of tokens for each sentence equals to 1,000.

Note that this step outputs three representations for each text and they are not combined. The main idea is to perform the detection using different representations.

4.3 Classification

The experiment employs seven classifiers designed for the problem of anomaly detection, which are categorized as follows:

1. Probabilistic Models:
   (a) Empirical-Cumulative-distribution-based Outlier Detection (ECOD) (Li et al., 2022): This algorithm is inspired by the fact that outliers are often considered events that appear in the tails of a distribution. Briefly, ECOD first assesses the underlying distribution of the input data non-parametrically by calculating the empirical cumulative distribution for each of to estimate the tail probabilities for each data point. Finally, it calculates an outlier score for each data point by aggregating the estimated tail probabilities across dimensions.
2. Proximity-Based Models:
   (a) Local Outlier Factor (LOF) (Cheng et al., 2019): This classifier is a density-based discrepancy detection algorithm that finds discrepancies by calculating the local deviation of a given data point. The outlier determination is judged based on the density between each data point and its neighboring points; the lower the density of the point, the more likely it is to be identified as the outlier.
   (b) Clustering-Based Local Outlier Factor (CBLOF) (He et al., 2003): It is a measure to identify the physical significance of an outlier and gives importance to the behavior of local data.
3. Linear Models:
   (a) One-class SVM (Géron, 2019): The algorithm seeks to separate the instances in the high-dimensional space of the origin. In a high-dimensional space, the algorithm tries to find a hyperplane that separates the training set samples from the origin, this will correspond to finding a small region that encompasses all instances, so if a new instance does not fit in this region, it is classified as an anomaly.
   (b) One-Class SVM using Stochastic Gradient Descent (SGDOneClassSVM) (Pedregosa et al., 2011): The algorithm implements an online linear version of the One-Class SVM using stochastic gradient descent. Combined with kernel approximation techniques, this algorithm can be used to approximate the solution of a kernelized One-Class SVM with a linear complexity in the number of samples.
4. Outlier Ensembles:
   (a) Isolation Forest (Géron, 2019): The algorithm builds a Random Forest in which each Decision Tree grows randomly: at each node, it picks a feature randomly, then it picks a random threshold value (between a minimum and a maximum value) to split the dataset into two. The dataset gradually gets chopped into pieces until all instances end up isolated from the other instances. An anomaly is usually far from other instances and it tends to get isolated in fewer steps than normal instances.
5. Graph-Based Model:
   (a) Optimum-Path Forest Algorithm: The Optimum-Path Forest for anomaly detection, namely OPF-AD, was proposed by Passos et al. (Passos et al., 2016) and concerned with an adaption of the unsupervised OPF for the task. In this context, the unsupervised

OPF is employed to cluster a training dataset composed of positive samples only. Regarding the testing step, each new instance is associated with the closest cluster using a *k*-NN-based approach, as performed by the standard unsupervised OPF. Further, the density of this new sample is computed and compared against a threshold, which is pre-defined as a hyperparameter. Finally, the instance is considered regular if the density is higher than the threshold. Otherwise, it receives an anomaly label.

4.4 Experimental Setup and Evaluation Metrics

Anomaly detection methods are unsupervised learning models that learn what is normal from training data. Thus, any other data that does not follow the pattern determined by each method is identified as an anomaly. To carry out the experiments, we use the real news samples for training, making the classifier understand this pattern. Thus, when the testing samples containing fake news appear, the classifier can flag the information focuses that deviate from the standard properties. The classifiers are trained using 80% of the real news from the corpus, whereas the testing set comprises the remaining 20% of real news plus all fake news, Table 4. The idea of partitioning idea of using usch a partitioning configuration is due to the fact that the amount of fake news is much lower than the real news. Thus, this division is fairer and generates a greater impact when verified at a training percentage greater than 20%.

The performance is assessed through three metrics: (i) f1-score micro, (ii) f1-score macro, and (iii) precision. The f1-score micro calculates the overall mean f1-score by counting the sums of true positives (TP), false negatives (FN) and false positives (FP). Macro f1-score treats all class equally, regardless their supporting values because is calculated using the arithmetic mean (also known as unweighted average) of all f1-scores per class. The application of such metrics is motivated by the type of database (i.e., unbalanced) used in the experiments.

Table 4: Details of each experimental set.

set	types of news	# samples
train	real	80,000
test	real and fake	21,000

5 EXPERIMENTAL RESULTS

This section discusses the experimental results using FakeRecogna Anomaly, as presented Table 5. As aforementioned, each classifier is evaluated using three text representations as input to analyze their impact on the detection task.

The codes of each classifier can be found either in the scikit-learn library[5] or in the PyOD library[6], which is one of the most comprehensive and scalable Python library for detecting peripheral objects in multivariate data (Zhao et al., 2019).

Table 5: Experimental results over FakeRecogna Anomaly.

model	text representation	f1-score		precision	
		micro	macro	real	fake
ECOD	BoW	0.866	0.522	96.0%	8.0%
	TF-IDF	0.869	0.519	96.0%	8.0%
	FastText	0.867	0.519	96.0%	8.0%
LOF	BoW	0.904	0.703	99.0%	31.0%
	TF-IDF	0.928	0.517	95.0%	9.0%
	FastText	0.866	0.500	95.0%	5.0%
CBLOF	BoW	0.863	0.526	96.0%	9.0%
	TF-IDF	0.902	0.717	100.0%	32.0%
	FastText	0.866	0.501	95.0%	5.0%
OneClassSVM	BoW	0.488	0.358	94.0%	4.0%
	TF-IDF	0.519	0.413	100.0%	9.0%
	FastText	0.494	0.354	94.0%	3.0%
SGDOneClassSVM	BoW	0.591	0.457	100.0%	10.0%
	TF-IDF	0.518	0.413	100.0%	9.0%
	FastText	0.951	0.491	95.0%	9.0%
IsolationForest	BoW	0.912	0.548	96.0%	13.0%
	TF-IDF	0.912	0.515	95.0%	8.0%
	FastText	0.909	0.522	95.0%	9.0%
OPF-AD	BoW	0.056	0.054	99.0%	5.0%
	TF-IDF	0.069	0.068	100.0%	5.0%
	FastText	0.049	0.047	97.0%	5.0%

Concerning text representation, TF-IDF achieved the best precision results. The fact that we use stemming in the pre-processing phase may have hindered FastText, since it works with n-grams, prefixes, and suffixes, which are information that can be lost when words are stemmed. Additionally to precision, when analyzing an unbalanced base, one of the main metrics is the macro f1-score, in which CBLOF classifier with TF-IDF representation achieved the highest average among the other experiments, followed by LOF with BoW representation and, much further down, IsolationForest also with BoW.

Regarding the models, the proximity-based outlier detection methods obtained the best results to identify real and fake news. Two other points worth mentioning are: (i) the proximity-based models for outlier detection obtained better results than the other techniques, and (ii) the ECOD classifier did not present any great performance variation for any of the textual representations.

[5]https://scikit-learn.org/stable/

[6]https://pyod.readthedocs.io/en/latest/

5.1 Comparison Experiments: Standard vs Anomaly Detection classifier

This section presents the results that evaluate the performance of both standard and anomaly detection based technique classifiers over the FakeRecogna Anomaly dataset. However, the data was normalized, that is, we used a pre-processing technique that places the data on the same scale, thus transforming the samples and distributing the data in an range.

To carry out this comparison, we created different folds from the tests presented above, but applying the same methodology. This set of experiments focus on the micro and macro f1-score metrics. Table 6 shows the results obtained with the previously mentioned classifiers, the table being divided first with the supervised algorithms and then below, the unsupervised algorithms.

Table 6: Algorithm results using FakeRecogna Anomaly.

classifier	f1-score		precision	
	micro	macro	real	fake
LinearSVC	0.993	0.498	99.0%	0.0%
Naive Bayes	0.990	0.499	99.0%	0.0%
SVM	0.990	0.498	99.0%	25.0%
OPF	0.983	0.575	99.0%	15.0%
RandomForest	0.990	0.418	96.0%	0.52%
ECOD [NB]	0.331	0.285	97.0%	5.0%
IsolationForest [RF]	0.931	0.497	95.0%	2.5%
LOF [NB]	0.619	0.413	94.0%	3.0%
OneClassSVM [SVC]	0.931	0.502	100.0%	32.0%
OPF-AD [OPF]	0.070	0.069	100.0%	0.05%
SGDOneClassSVM [SVM]	0.952	0.502	95.0%	0.0%

The achieved results show that even for an unbalanced base, standard classifiers achieve a good performance. However, the use of unsupervised algorithms had a better performance analyzing the applied metrics, since the f1-score micro gives equal importance to each observation, that is, when the classes are unbalanced, those classes with more observations will have a greater impact. in the final score, resulting in a score that can hide the performance of minority classes and amplifies the majority.

On the other hand, the F1-macro score gives equal importance to each class, so a majority class will contribute equally to the minority, allowing the f1-macro score to still return objective results in unbalanced datasets. In short, if we analyze the results, we can see that the unsupervised algorithms obtained a better result involving the FakeRecogna Anomaly dataset than the supervised algorithms, thus demonstrating that the detection of fake news can be treated as an anomaly detection problem.

6 CONCLUSIONS AND FUTURE WORKS

This work proposed a new corpus for the detection of fake news in Portuguese called FakeRecogna Anomaly. We carried out experiments with classical textual representations and different classifiers covering different working methods, in which we managed to produce interesting results considering the difficulty of the problem. A far we know, no specific databases were found that present the proportion developed in this work, i.e, according to their characteristics of corpus formation with one false news for every 1,000 real news, coming from internationally recognized agencies and that have their news updated in comparison to other databases. By building the dataset, we hope to encourage research on anomaly detection involving fake news in Portuguese.

We also performed experiments over the proposed corpus using classifiers designed for outlier detection, where the models achieved interesting results, in addition to the performance comparison involving standard algorithms.

The main scope of the work is to foster research on the detection of fake news in Portuguese, working with new approaches and providing a corpus to be further developed by the scientific community. As future work, we intend to explore our dataset with new approaches, use other word embeddings, and use more utilities provided by the PyOD library. Regarding deep learning, we intend to consider new architectures such as unsupervised neural attention models, unsupervised Bidirectional Encoder Representations from Transformers (BERT) and other neural network models.

ACKNOWLEDGEMENTS

The authors are grateful to FAPESP grants #2013/07375-0, #2014/12236-1, #2019/07665-4, #2019/18287-0, and #2021/05516-1, and CNPq grant 308529/2021-9.

REFERENCES

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31:211–236.

Chalapathy, R. and Chawla, S. (2019). Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*.

Chalapathy, R., Menon, A. K., and Chawla, S. (2018). Anomaly detection using one-class neural networks. *arXiv preprint arXiv:1802.06360*.

Charles, A. C., Ruback, L., and Oliveira, J. (2022). Fakepedia corpus: A flexible fake news corpus in portuguese. In Pinheiro, V., Gamallo, P., Amaro, R., Scarton, C., Batista, F., Silva, D., Magro, C., and Pinto, H., editors, *Computational Processing of the Portuguese Language*, pages 37–45, Cham. Springer International Publishing.

Cheng, Z., Zou, C., and Dong, J. (2019). Outlier detection using isolation forest and local outlier factor. In *Proceedings of the conference on research in adaptive and convergent systems*, pages 161–168.

Choi, J. and Lee, S.-W. (2020). Improving fasttext with inverse document frequency of subwords. *Pattern Recognition Letters*, 133:165–172.

de Carvalho, H. V. F., Carvalho, E. C., Arruda, H., Imperatriz-Fonseca, V., de Souza, P., and Pessin, G. (2018). Detecção de anomalias em comportamento de abelhas utilizando redes neurais recorrentes. In *Anais do IX Workshop de Computação Aplicada a Gestão do Meio Ambiente e Recursos Naturais*, Porto Alegre, RS, Brasil. SBC.

Endo, P. T., Santos, G. L., de Lima Xavier, M. E., Nascimento Campos, G. R., de Lima, L. C., Silva, I., Egli, A., and Lynn, T. (2022). Illusion of truth: Analysing and classifying covid-19 fake news in brazilian portuguese language. *Big Data and Cognitive Computing*, 6(2).

Faustini, P. and Covões, T. (2019). Fake news detection using one-class classification. In *2019 8th Brazilian Conference on Intelligent Systems (BRACIS)*, pages 592–597.

Garcia, G. L., Afonso, L., and Papa, J. P. (2022). Fakerecogna: A new brazilian corpus for fake news detection. In *International Conference on Computational Processing of the Portuguese Language*, pages 57–67. Springer.

Géron, A. (2019). *Maos a Obra: Aprendizado de Maquina com Scikit-Learn & TensorFlow*. O'Reilly Media.

He, Z., Xu, X., and Deng, S. (2003). Discovering cluster-based local outliers. *Pattern recognition letters*, 24(9-10):1641–1650.

Kannan, R., Woo, H., Aggarwal, C. C., and Park, H. (2017). Outlier detection for text data. In *Proceedings of the 2017 siam international conference on data mining*, pages 489–497. SIAM.

Kesarwani, A., Chauhan, S. S., and Nair, A. R. (2020). Fake news detection on social media using k-nearest neighbor classifier. In *2020 International Conference on Advances in Computing and Communication Engineering (ICACCE)*, pages 1–4.

Kintopp, P. M. (2017). Aplicação de técnicas de aprendizado de máquina em dados públicos para detecção de anomalias. B.S. thesis, Universidade Tecnológica Federal do Paraná.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., et al. (2018). The science of fake news. *Science*, 359(6380):1094–1096.

Li, D., Guo, H., Wang, Z., and Zheng, Z. (2021). Unsupervised fake news detection based on autoencoder. *IEEE Access*, 9:29356–29365.

Li, K. (2021). Haha at fakedes 2021: A fake news detection method based on tf-idf and ensemble machine learning. In *IberLEF@ SEPLN*, pages 630–638.

Li, Z., Zhao, Y., Hu, X., Botta, N., Ionescu, C., and Chen, G. (2022). Ecod: Unsupervised outlier detection using empirical cumulative distribution functions. *IEEE Transactions on Knowledge and Data Engineering*.

Mishra, S., Shukla, P., and Agarwal, R. (2022). Analyzing machine learning enabled fake news detection techniques for diversified datasets. *Wireless Communications and Mobile Computing*, 2022.

Mohaghegh, M. and Abdurakhmanov, A. (2021). Anomaly detection in text data sets using character-level representation. In *Journal of Physics: Conference Series*, volume 1880, page 012028. IOP Publishing.

Monteiro, R. A., Santos, R. L., Pardo, T. A., Almeida, T. A. d., Ruiz, E. E., and Vale, O. A. (2018). Contributions to the study of fake news in portuguese: New corpus and automatic detection results. In *International Conference on Computational Processing of the Portuguese Language*, pages 324–334. Springer.

Oshikawa, R., Qian, J., and Wang, W. Y. (2018). A survey on natural language processing for fake news detection. *arXiv preprint arXiv:1811.00770*.

Pang, G., Shen, C., Cao, L., and Hengel, A. V. D. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys (CSUR)*, 54(2):1–38.

Passos, L. A., Ramos, C. C. O., Rodrigues, D., Pereira, D. R., Souza, A. N., Costa, K. A. P., and Papa, J. (2016). Unsupervised non-technical losses identification through optimum-path forest. *Electric Power Systems Research*, 140:413–423.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Ruchansky, N., Seo, S., and Liu, Y. (2017). Csi: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 797–806.

Ruff, L., Zemlyanskiy, Y., Vandermeulen, R., Schnake, T., and Kloft, M. (2019). Self-attentive, multi-context one-class classification for unsupervised anomaly detection on text. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4061–4071, Florence, Italy. Association for Computational Linguistics.

Singhal, S., Shah, R. R., Chakraborty, T., Kumaraguru, P., and Satoh, S. (2019). Spotfake: A multi-modal framework for fake news detection. In *2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM)*, pages 39–47.

Wang, H., Bah, M. J., and Hammad, M. (2019). Progress in outlier detection techniques: A survey. *Ieee Access*, 7:107964–108000.

Zhang, X. and Ghorbani, A. A. (2020). An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 57(2):102025.

Zhang, Y., Jin, R., and Zhou, Z.-H. (2010). Understanding bag-of-words model: a statistical framework. *International journal of machine learning and cybernetics*, 1(1):43–52.

Zhao, Y., Nasrullah, Z., and Li, Z. (2019). Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20(96):1–7.

Zhou, X. and Zafarani, R. (2020). A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Computing Surveys (CSUR)*, 53(5):1–40.