

This document is published at:

L. Ibañez-Lissen, L. González-Manzano, J. M. de Fuentes and M. Goyanes, "On the Feasibility of Predicting Volumes of Fake News—The Spanish Case," in IEEE Transactions on Computational Social Systems

DOI: [10.1109/TCSS.2023.3297093](https://doi.org/10.1109/TCSS.2023.3297093)

© The Authors, 2023



This work is licensed under a [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/)

On the Feasibility of Predicting Volumes of Fake News—The Spanish Case

Luis Ibañez-Lissen^{ID}, Lorena González-Manzano^{ID}, José M. de Fuentes^{ID}, and Manuel Goyanes^{ID}

Abstract—The growing amount of news shared on the Internet makes it hard to verify them in real-time. Malicious actors take advantage of this situation by spreading fake news to impact society through misinformation. An estimation of future fake news would help to focus the detection and verification efforts. Unfortunately, no previous work has addressed this issue yet. Therefore, this work measures the feasibility of predicting the volume of future fake news in a particular context—Spanish contents related to Spain. The approach involves different artificial intelligence (AI) mechanisms on a dataset of 298k real news and 8.9k fake news in the period 2019–2022. Results show that very accurate predictions can be reached. In general words, the use of long short-term memory (LSTM) with attention mechanisms offers the best performance, being headlines useful when a small amount of days is taken as input. In the best cases, when predictions are made for periods, an error of 10.3% is made considering the mean of fake news. This error raises to 28.7% when predicting a single day in the future.

Index Terms—Fake news, machine learning, prediction.

I. INTRODUCTION

FAKE news are articles that are intentionally and verifiably false [1]. They have been shared and spread as long as the very first newspapers were released. They have been used to destabilize countries, states and manipulate public opinion [2]. Indeed, 86% of online users have been exposed to these contents at some point in time [3], [4].

Cases such as the 2016 U.S. elections or the recent Ukrainian war reveals how social media platforms can be used to instantly spread this kind of misinformation [5], [6]. Fact-checkers, public institutions, and social media platforms are being held responsible for verifying information [7]. For example, the European Union is countering this threat [8].

To address this issue, a vast array of research efforts have been focused on automatic detection of fake news. Thus, they train machine learning techniques to identify

abnormal features in a piece of news [9], [10], [11], [12]. Other works have characterized different aspects of fake news, such as the profile of the victims [13] or their spread pattern [14], [15], [16].

Despite prior efforts, no comprehensive research currently focuses on predicting the volume of fake news. This can be attributed to the scarcity of publicly available datasets that encompass temporal information regarding verified instances of fake news. However, the emergence of private fact-checking organizations opens up an unprecedented opportunity for collaboration by providing new data sources. By learning from these resources, valuable trends in the domain of fake news can be uncovered, enabling a deeper understanding and more effective countermeasures.

Thus, this work differs from other efforts (e.g., fake news detectors), but it is complementary to them. Mainly, it addresses this matter by applying different artificial intelligence (AI) techniques for forecasting the amount of fake news concerning the events happening in a particular case study—news written in Spanish related to Spain. Being the fourth worldwide language as of 2022 [17], it is an interesting choice for attackers when creating fake news. Interestingly, our models are trained only with real news from reputed media, as qualifying news as fake is typically not achievable in real-time. Focusing on a particular context is necessary as fake news is a local effect with intrinsic cultural factors [18]. Therefore, our approach is intended to be illustrative enough to inspire future efforts related to other countries or languages.

Several studies have focused on preventing the spread of misinformation by analyzing user behavior to identify potential topics [13] or sources of fake news [15]. However, these studies do not attempt to understand the trends or quantify the expected amount of fake news given the events occurring within a particular country. Other works, such as [14], [16], have investigated how fake news evolve over time. In contrast, our study aims to identify potential hidden trends in the quantification of fake news and, consequently, the evolution of the overall volume of fake news.

The research question at stake and related contributions of this work is as follows.

Research question—Is it feasible to predict the future amount of fake news?

- 1) Different AI models are configured and trained, exhibiting relevant differences in their reliability.
- 2) Input data required to perform the prediction are also characterized.

Manuscript received 20 February 2023; revised 21 April 2023, 7 June 2023, and 7 July 2023; accepted 13 July 2023. This work was supported in part by the Universidad Carlos III de Madrid (UC3M) and the Government of Madrid [Community of Madrid (CAM)] under Grant DEPROFAKE-CM-UC3M; in part by the CAM through the Project CYNAMON, co-funded by the European Research Development Fund (ERDF), under Grant P2018/TCS-4566-CM; and in part by the Spanish Ministry of Science and Innovation (MICINN) of Spain under Grant PID2019-111429RB-C21. (Corresponding author: José M. de Fuentes.)

Luis Ibañez-Lissen, Lorena González-Manzano, and José M. de Fuentes are with the Computer Science and Engineering Department, Universidad Carlos III de Madrid, Leganes, ES28911 Madrid, Spain (e-mail: jfuentes@inf.uc3m.es).

Manuel Goyanes is with the Communication Studies Department, Universidad Carlos III de Madrid, Leganes, ES28911 Madrid, Spain.

Digital Object Identifier 10.1109/TCSS.2023.3297093

- 3) The reliability of short, medium, and long-term predictions is assessed.
- 4) The trained models are publicly released to foster further research.

Answering this question is useful to understand the evolving nature of fake news. The obtained knowledge can benefit journalists, intelligence agencies, and trusted sources such as fact-checkers, allowing them to gain an advantage over misinformation campaigns. In this way, reactive countermeasures could be established such as the search of a particular type of news in the case of fact-checkers or the need to highlight the reality of a fact for journalists. Furthermore, other disciplines within social sciences such as sociologists may also benefit from anticipating fake news.

This research can augment existing data-driven fake news detection models by providing an additional feature. The ability to uncover hidden trends in fake news dissemination can enhance the allocation of prediction efforts. By identifying emerging patterns employed by purveyors of fake news, detection models can adapt and improve their accuracy. By leveraging predictive analysis on the volume of fake news, the capabilities of fact-checkers and journalists could be enhanced by staying ahead of the curve, ensuring they are adequately prepared to tackle the surge in fake news during predicted periods of increased activity.

Article Organization: Related works are analyzed in Section II. Section III gives the background to understand the proposal. Afterward, Section IV describes the proposal. The preparation of experiments is addressed in Section V, whereas the assessment is shown in Section VI. Section VII identifies the limitations of this work. Finally, Section VIII concludes this article.

II. RELATED WORK

Fake news and fake content have always been a primary concern for states, publishers, and social media platforms. Discerning when a given text, image, or video is fake or not has been the main concern of academia in the field of fake news. This is why most of the data-driven efforts have been focused on detection and classification tasks, especially analyzing texts using machine learning and natural language processing techniques. There are many examples [19], [20], [21], [22] of how the latest machine-learning approaches have been successfully deployed for these purposes. While Hakak et al. [19] and Jiang et al. [20] try to mix different machine-learning techniques stacking the results of different simpler models in order to obtain more accurate predictions, Raza and Ding [21] and Ajao et al. [22] rely on the sole application of more advanced models like transformers or hybrid convolutional recurrent neural networks (RNNs).

However, as this proposal focuses on fake news predictions, some works have been done in this regard, though not directly related. Del Vicario et al. [13] proposed a mechanism for the early detection of the most probable topic on which fake news may appear by designing a framework to extract and analyze posts and topics in Italy from Facebook. Murayama et al. [14] investigated how fake news spread on Twitter, trying to approximate when it is more probable that a new fake news

is posted on Twitter. Rath et al. [15] used an attention-based graph neural network to predict whether an actor is more likely to spread fake news.

In this vein, Guarino et al. [16] presented a framework for tracing the origins of fake news in social media. Their research places particular emphasis on comprehending the dynamics that underlie the evolution of fake news.

Table I shows a summary of different goals and techniques previously covered by academia. In sum, some studies focus on identifying the truth of news articles using data-driven techniques on top of one or several models. Others try to detect the most likely target or topic that may be a victim of fake news and predict possible future fake news spreaders. The last line of the research is devoted to understanding the development of fake news within social media. In contrast, our research aims to predict the amount of fake news by discovering trends in misinformation, rather than determining the legitimacy of individual news. Moreover, we apply different datasets and techniques in comparison to those used until now.

III. BACKGROUND

This section provides the main notions related to the proposal, namely the data sources and AI techniques at stake.

A. GDELT

The Global Data on Events, Location, and Tone (GDELT) project [24] is an online dataset that monitors a wide range of world's broadcast, print, and web news in a variety of languages and countries to identify all the events happening on each of them. GDELT categorizes all events into 300 types, from protests to accusations or riots; every 15 min, the public GDELT event data table is updated to keep track of them. Each entry is confirmed by a set of 58 fields (61 in the 2.0 version) following the Conflict and Mediation Event Observations (CAMEO) format proposed by Gerner et al. [25]. The CAMEO framework is a widely used system for coding event data, with a primary emphasis on political events. This ontology encompasses four primary categories: verbal cooperation, verbal conflict, material cooperation, and material conflict between two or more actors.

Beyond the URLs of the news pointing to the source media, some fields like the GoldsteinScale or the NumSources measure the importance and the possibility of impacting society on different scales. Moreover, other fields like Actor1Geo_CountryCode or Actor1Geo_Fullname identify the involved actors in a categorical way. GoldsteinScale scale measures the potential impact of a given type of event on the stability of a country on a scale of -10 to $+10$. AvgTone measures the tone of all documents mentioning an event. Scales range from highly negative (-100) to positive ($+100$). Table II shows a simplified example of the type of information available in the records.

This dataset has been previously used to measure the feasibility of detecting upcoming events [23], [26], [27], [28] like riots or social unrest or detecting levels of violence [29].

TABLE I
RELATED WORK ANALYSIS

	Del vicario <i>et al.</i> [13]	Guarino <i>et al.</i> [16]	Murayama T <i>et al.</i> [14]	Bhavtosh <i>et al.</i> [15]	S Deng <i>et al.</i> [23]	Ours
Goal	Early warning	Understand the dynamics of (fake) news in social networks.	Modeling of spread of fake news	Predict fake news sources	Predict next events	Trends and fake news prediction through media.
Dataset	Facebook pages from official news papers and fake news website spreaders	Data extracted from different fact-checkers and online news papers about Italian elections in 2016.	Recent fake news dataset, Tohoku earthquake and tsunami fake news dataset	Twitter, www.altnews.in	GDELT	GDELT, www.maldita.es
Technique	Linear Regression, Logistic Regression, Support Vector Machine, K-Nearest, Neural networks	Custom module-based framework	Time-Dependent Hawkes process	Graph neural network with attention	Contextualized Multilevel Feature learning framework	LSTM, LSTM with attention, Encoder-decoders, multi-modal models.
Year	2019	2019	2021	2021	2021	2023
Description	Identify the topics that are most likely to become a target for future fake news	Create a framework to understand the behavior of (fake) news in social networks	Understanding the dynamics of the spread of fake news on social media	Identify nodes or actor which are more likely to produce fake news	Predict next type of events	Predict volumes of fake news leveraging real news

TABLE II
SIMPLIFIED EXAMPLE OF A RECORD IN GDELT DATASET

Date	NumSources	Actor[Geo_CountryCode]	Actor[Geo_Fullname]	AvgToneAvgTone	GoldsteinScale	NumArticles	Source URL
02-04-2022	12	SP	SPAIN	100	10.0	20	https://www.lavozdegalicia.es/noticia/galicia/2022/04/02/xunta-alega-contra-ley-movilidad-gobierno-de-espana-espana-rural-solo-atende-urbana
13-04-2016	0	JP	JAPAN	16	1	1	https://www.eldia.es/2016-04-13/palma/5-Espana-Japon-firman-instalacion-cuatro-museos-telecomunicaciones-Palma.htm

B. Maldita

Maldita.es¹ is a Spanish fact-checker. They collect intelligence to verify news by leveraging experts and official sources. They maintain a private dataset of proven and disproved fake news provided to us for this investigation. All fake news are collected in a community-based approach. Maldita.es counts with different channels such as a WhatsApp channel, Twitter, their application, and their own sources. As well as in GDELT, Maldita.es records follow the CAMEO format, which has been enriched with multiple extra features such as the recurrence and sources of the fake news.

C. AI Algorithms

There are a plethora of AI techniques that can be used to carry out predictions [30]. In the following, we provide an overview of the main concepts related to our proposal. For an insightful definition of these notions, the reader may refer to [31].

In recent years, there has been a surge of interest in neural networks owing to their demonstrated capacity for self-learning, adaptivity, fault tolerance, nonlinearity, and superior efficiency in input–output mapping [32]. This has led to their outperformance of human capabilities on specific tasks, such as object recognition [33].

Among neural networks, one of the most commonly used types for spatial–temporal prediction tasks is the RNN. First proposed by Rumerhalt et al. [34], RNNs are a particular set of neural networks with cyclic connections which allows the input sequences of input vectors [35]. In mathematical terms, RNNs can be described by the following equation, where h_t

is a hidden state at any given time step t , x_t is the input at any time step t , W is a weight matrix, and U is a hidden-state-to-hidden state matrix:

$$h_t = \sigma(W_{x_t} + U h_{t-1}). \quad (1)$$

Due to their higher computational complexity, training these RNNs is more expensive than traditional artificial neural networks (ANNs) due to their recurrent nature and the need to process sequential data. Moreover, new challenges known as gradient disappearance and gradient explosion have emerged during the training procedures, as described in [36]. These issues could impact the stability of the training process and make it harder to achieve good results. To solve those problems, long short-term memory (LSTM) cells were introduced in [37]. LSTM cells present a set of three gates—forget, input, and output gates which allow the network selectively remember or forget information.

In the input gate, the network decides which information from the current input vector should be passed on to the next time step. The network decides the information that should be discarded or forgotten in the forget gate. Finally, in the output gate, the network decides how much of the remembered information should be used to produce the output for the current input vector.

These kinds of RNN LSTM networks have been successfully deployed in a wide variety of tasks, including classification, prediction, sequence generation, or recognition [38]. However, researchers noticed that long-term dependencies could not be captured in long texts when using LSTM cells [39]. To address this issue, the so-called attention mechanisms were introduced. These layers can be used as an information allocation reservoir, allowing the models to access all the information produced during the input processing. Instead of having a single context vector as before, the formula is modified as shown in the following equation, where α_{ij} denotes the attention weight or the importance of each hidden state h_i to the next hidden state:

$$c_t = \sum_{j=1}^T \alpha_{tj} h_j. \quad (2)$$

¹<https://maldita.es/nosotros-maldita/>, last access February 20, 2023.

These neural networks may take the form of an encoded–decoder architecture, as presented in [40]. The main feature of this architecture is that the encoder part compresses the input, reducing its dimensionality to form the so-called latent vector. This latent representation of the original data is subsequently passed to the decoder, which aims to recover the original information. By leveraging this approach, the architecture can acquire crucial information, thereby improving the neural network’s comprehension of the data.

This architecture has demonstrated promising results in various tasks, including machine translation, anomaly detection [41], [42], [43], and image generation [44], [45], [46].

IV. PROPOSAL

This section focuses on the proposed approach to predict the count of future fake news. In particular, Section IV-A provides an overview. Afterward, the proposed AI models are introduced in Section IV-B. Finally, the pursued goals are presented in Section IV-C.

A. Approach Overview

This proposal focuses on studying the feasibility of predicting the amount of expected fake news based on media data. Different models are used for this purpose and, in particular, the following steps are initially considered in the process (see Fig. 1).

The process consists of four phases as follows.

- 1) *Data Extraction*: Events from GDELT are extracted, filtering by the target country, Spain. A subset of all sources contained therein has been selected based on their reputation. Moreover, fake news are obtained from Maldita database.
- 2) *Data Preprocessing and Feature Extraction/Generation*: In this step, the features at stake are selected and normalized. It must be noted that some additional features may be computed based on the original dataset information. The sentiment extraction task from headlines is an example of this.
- 3) *Dataset Preparation*: Once preprocessed, input data are gathered in groups of different sizes, to understand the most suitable input size. Moreover, different features of real news are considered in each experiment. The order of input data is also considered to measure its relevance. With all these requirements, datasets are prepared for each experiment, splitting them into training and testing subsets.
- 4) *Training and Assessment*: The model is trained and performance results are computed.

Remarkably, it must be noted that the models must make predictions based on data from real news only. Thus, the amount of fake news is actually the predicted value. This increases the realism of the proposal, as typically fake news are tagged as such after some time. While timing has been shown as a dramatic factor to counter fake news, it is unlikely that they will be immediately discovered [47].

B. AI Models

In order to address this issue, a set of three models (plus one variant) have been trained. First, a basic RNN LSTM model is designed [see Fig. 2(a)]. It consists of a model with an LSTM layer in charge of encoding all the information from the tabular data of GDELT. This model contains a single LSTM layer and sends its output to a multilayer perceptron.

A variant of this model is considered by adding an attention layer after the LSTM layer. The idea is to understand whether adding an attention layer to the basic LSTM improves the quality of the predictions.

Third, an encoder–decoder is designed [see Fig. 2(b)], where this architecture’s latent spaces may help extract features from the inputs. The architecture forces the model to compress the input to, later on, reconstruct the initial input from the compressed representation. This may allow the model to extract essential features that help reach better predictions.

Moreover, a multimodal encoder–decoder architecture is designed [see Fig. 2(c)] in order to be able to encode headlines information from the processed URLs. We included an LSTM layer in charge of encoding the headlines after an embedding layer. All this information is concatenated together and sent through a multilayer perceptron. The idea behind this model is to encode textual features which may be relevant to the model.

C. Goals

The general goal of this proposal is the prediction of the amount of fake news by observing the ongoing events in the society covered by the media. Beyond the mere accuracy of predictions, the pursued goals are the following.

- O1 *Diversity of Prediction Horizons*: The mechanism should produce predictions for short to long-term periods.
- O2 *Data Affordability*: The amount and type of input data must be affordable.

V. EXPERIMENT PREPARATION

This section describes the preparation for the experiments. In particular, Section V-A describes the data preparation. Afterward, the experimental settings are introduced in Section V-B.

A. Data Preparation

1) *GDELT Processing*: First, all events available between 2019 and 2022 in GDELT filtering by Spain are collected. Afterward, only those produced by reputed news sources in Spain (see our online repository, as explained in Section V-B) are considered. Therefore, a dataset of 298 242 events is produced. Each event refers to one news, so the title for each news is retrieved from its URL. The top ten headlines of each day are extracted taking into account the number of articles related to that event, the number of mentions of that specific event, and the number of sources where that event has been extracted. The mean number of headlines linked to an important event is 15 per day. The importance of GDELT can be measured by the AvgTone feature (recall Section III) of the

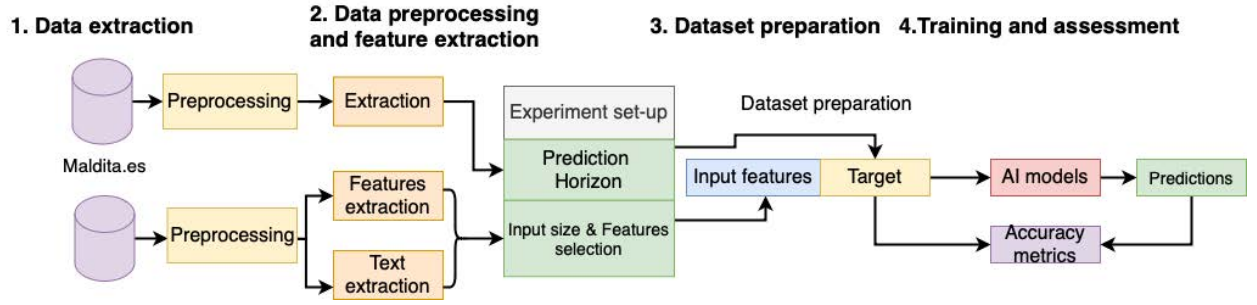


Fig. 1. Approach overview.

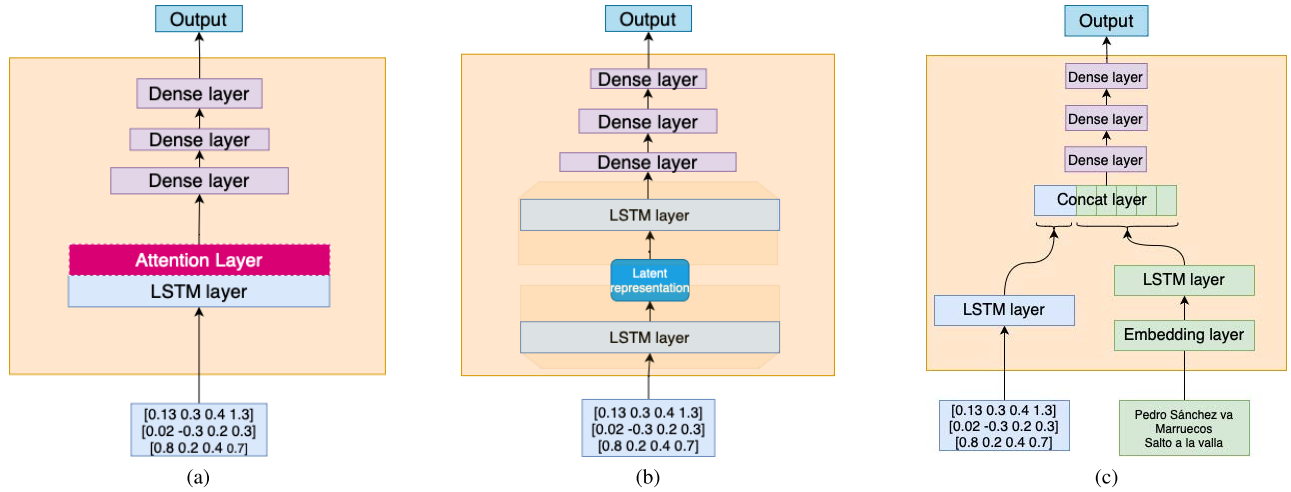


Fig. 2. Trained AI models. (a) LSTM and LSTM+attention. (b) Encoder-decoder. (c) LSTM + headlines.

articles of a given headline. This amount of headlines ensures that the most important news are considered each day while involving an affordable amount of computational and memory resources.

2) *Maldita.Es Processing*: The dataset contains 8990 fake news collected over four years, from 2019 until mid-2022. Precisely, they comprise a range of 1018 days (September 30, 2019, to July 14, 2022). For each fake news, timestamps of the first appearances are recorded. In 76.22% of the days at least one fake news appeared, and values are distributed with high variability—a mean of 7.2 fake news per day and a standard deviation of 8.16. Therefore, predicting fake news is a meaningful task. All the fake news per day are grouped, computing the total per day.

To ensure the representativeness of these fake news, each fake news is assigned a topic. The predominant ones are politics (32.9%), social (12.4%), health (11.5%), racism (8.6%), and science (7.8%).

3) *Inferring Sentiments*: To enrich the set of input variables, the sentiment of each real or fake news is extracted. For this purpose, a pretrained BERT model fine-tuned using transfer-learning techniques over a dataset of tweets is used [48]. This model was fine-tuned to understand different languages, so it can be used to classify headlines in Spanish following the standard parameters proposed in [49].

Table III shows a nearly similar distribution of sentiments between both real and fake news. Most news in GDELT and Maldita.es are positive or neutral, so fake news tend to mimic

TABLE III
DISTRIBUTION OF SENTIMENTS OF THE TEXTS IN GDELT AND MALDITA

	Positive	Neutral	Negative
Maldita.es	61.2%	35.2%	3.3%
GDELT	61.9%	29.2%	8.9%

the nature of the real news to go unnoticed. This enables them to produce misinformation while leveraging stealthiness.

4) *Selection of Variables*: The set of potential candidates comes either from GDELT or from the sentiment analysis. Particularly, they are either the values per day or their statistical descriptors, namely maximum and minimum per day and sum or average depending on the nature of the data.

Based on the domain knowledge, four GDELT variables were selected: *isRootEvent* and *numArticles*, which measure the relevance of each news piece, and *AvgTone* and *GoldsteinScale*, which describe its forcefulness. These variables were complemented with the extracted sentiments of positive, negative, and neutral. Statistical descriptors were then applied to these seven variables, resulting in 21 values. Additionally, the variables such as *NumberEvents*, *Year*, and *Month* were included to provide a temporal context for the news. Furthermore, the *Headline* was considered to preserve the semantics of the news.

Out of this set of 25 variables, a couple of measurements were carried out to guide the variable selection. Both of them are related to the intercorrelation of variables shown in Fig. 3. For this analysis, variables such as *Month*, *Year*,

TABLE IV
DESCRIPTION OF THE SELECTED VARIABLES

Feature	Description	Origin
	Input	
count_neutral (sentiments)	Number of neutral news in a day	Generated
count_negative (sentiments)	Number of negative news in a day	Generated
count_positive (sentiments)	Number of positive news in a day	Generated
IsRootEvent (Mean, Count)	Mean of root events per day and Number of root events in a day	GDELT
AvgTone (Mean, Max, Min)	Mean of the tone of the documents in a day	GDELT
GoldsteinScale (Mean, Max)	Maximum and minimum tone of the documents in a day	GDELT
	Mean of the GoldsteinScale of the events in a day	GDELT
	Maximum GoldsteinScale of the day	GDELT
	Mean of the number of articles per event in day	GDELT
NumArticles (Mean, Max, Min)	Maximum and minimum number of articles about an event in that day	GDELT
NumberEvents	Number of total events in a day	GDELT
year	Year of the event	GDELT
month	Month of the event	GDELT
(Headlines model only)	Headlines	Media URLs from GDELT
	Target values	
Global_number_fake_news	Actual number of fake news in the target period	Maldita.es

and Headline are excluded—although they are essential for the approach, they do not exhibit any valuable correlation. On the one hand, the Pearson Correlation [50] with the predicted value (i.e., the amount of fake news) was computed (first row/column of Fig. 3). Not only the correlation is quite diverse among all of them, but the maximum absolute value is 0.45. Then, a discarding threshold was set to 0.2, as the significance of lower values of correlation is very limited. This led to a set of 17 variables listed in Table IV. On the other hand, to further ensure the appropriateness of the established threshold, experiments were carried out with all 25 variables as well. The obtained results (omitted for space constraints) confirm that the inclusion of these additional variables does not lead to any valuable gain. The second measurement was the cross correlation among variables (see Fig. 3). Results show that there were no spurious variables as the degree of correlation is not consistently high with all the remaining values.

B. Experimental Settings

This section provides training settings of the described models. The training was conducted using a Google Colab pro² subscription, and the models have been implemented with Keras and Tensorflow. For training, the dataset is divided into training and testing subsets with a 90%–10% random distribution. This choice was made after a trial-and-error process—lower training rates led to slightly higher errors. As an optimization, the classical ADAM optimizer was used [51] with a learning rate of 10^{-3} and making use of early stopping as recommended in other works [52]. In what comes to the batch size, our preliminary tests show that smaller sizes do not affect the accuracy while they harm performance. Thus, the batch size has been set to 512.

The training procedure consists of creating sets of a given amount of input days. These sets gather from 2 to 20 days for the sake of completeness. Each set is associated with the number of fake news observed in the target period. This period may comprise one or several days. We have considered both periods of 1, 2, 4, and 7 days as they illustrate short, medium, and long-term predictions. Similarly, these values have also been used to refer to single future days as the target of predictions.

During training, we shuffle the sets, so the model does not learn with consecutive data. In this way, we prevent stationary effects. All experiments have been repeated five times after random shuffling with a limit of 3000 epochs. Results reported herein refer to the mean of all executions. The performance is assessed every 500 epochs, stopping the process if no improvements are achieved.

To foster further research, all models are publicly released.³

VI. ASSESSMENT

This section measures the achievement of the established goals (recall Section IV-C). Before addressing them, Section VI-A introduces the metrics and Section VI-B focuses on a preliminary test to ensure the proper operation of the models at stake. In Section VI-C, an initial analysis is carried out to set the grounds for the assessment. In particular, the proposed approach relies on predicting fake news by considering real news only. Therefore, it is necessary to analyze how effective are fake news themselves to predict future appearances, to clarify the relevance of the proposed approach. Afterward, the results obtained for different prediction horizons (goal O1) is addressed in Section VI-D. In what comes to data affordability (goal O2), it is divided into two parts—characterizing the amount of input data and measuring the impact of semantic data from headlines. For the sake of clarity, the former is addressed together with prediction horizons, whereas Section VI-E focuses on the latter. Finally, results are discussed in Section VI-F.

A. Metrics

Three error functions are considered herein, namely the mean absolute error (MAE), the root mean squared error (RMSE), and the coefficient of determination (R^2). They are introduced in the following equations, where y_i and \hat{y}_i refer to the real and predicted values, respectively,

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i| \quad (3)$$

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{n}} \quad (4)$$

$$R^2 = 1 - \frac{\sum (y_i - \hat{y})^2}{\sum (y_i - \bar{y})^2}. \quad (5)$$

Concerning MAE, it is computed as the average of the absolute difference between the actual and predicted values. RMSE is the square root of the mean of the square of all of the errors. Both MAE and RMSE serve to characterize the mean and standard deviation of the prediction errors and are measured in the predicted unit (i.e., number of fake news) [53]. Thus, their optimal value is zero. Finally, the coefficient of determination (R^2) represents the strength of the relationship or the portion of common variation in two time-series or variables [54]. The closer to 1, the better, whereas negative values indicate that a horizontal line would be a better fit than the model at stake.

²<https://colab.research.google.com>, last access February 20, 2023.

³https://github.com/Luisibear98/prediction_volume_fakenews

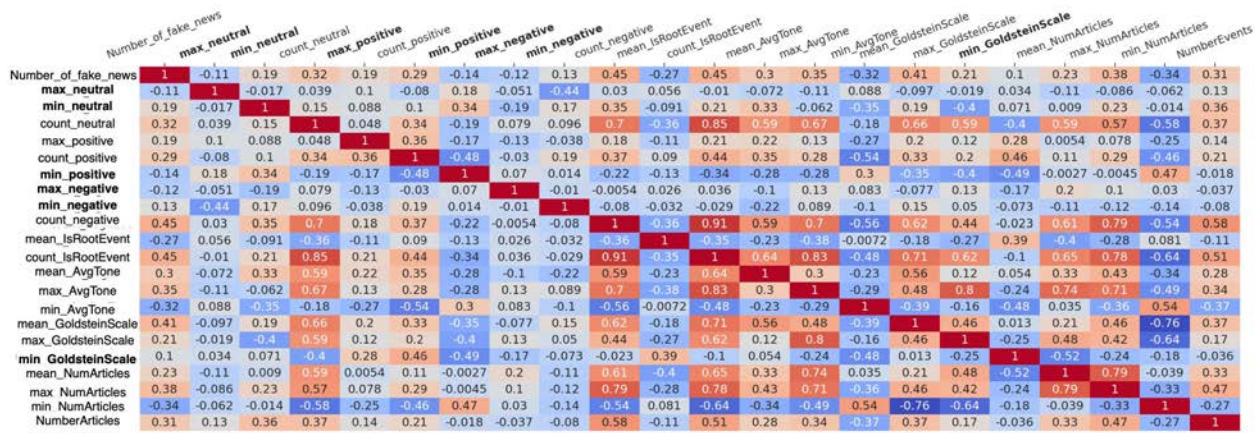


Fig. 3. Correlation matrix of all the variables. Discarded variables are marked in bold.

B. Preliminary Test

This test aims to assess whether the trained models are able to perform predictions that are not only valid for isolated points in time, but also for a number of consecutive days. It must be noted that the core of the experiments introduced later measures accuracy by testing the models in a big amount of days—more than 100. However, as these days are not meant to be consecutive, to avoid biases, this preliminary test complements the analysis by measuring the accuracy in a more minor (but representative) amount of successive days.

More precisely, March 2022 has been selected as the target of predictions. It has been chosen as it is the month with more fake news after the COVID 2019 period within the dataset. Thus, models were trained with the rest of the dataset and asked to predict the whole month.

Although different input sizes and prediction horizons were applied, for the sake of brevity only one setting is reported herein. Thus, Fig. 4 shows the predictions made for the next day by the different models, when seven days are used as input. In general, models are able to resemble the actual trend of fake news. Interestingly, the encoder–decoder model is the most similar, just missing some information in the highest peaks.

C. Baseline (Fake News Self-Predictions)

In this section, the ability of fake news themselves to predict further appearance is modeled as a time-series prediction task. Such a task is based on predicting the subsequent more likely output given a set of consecutive input sequences.

Before applying any particular technique, it is necessary to characterize the distribution. In this regard, a set of statistical tests has been carried out to check the stationarity of the data regarding the number of fake news. The augmented Dickey–Fuller test (ADF) and the Kwiatkowski–Phillips–Schmidt–Shin (KPSS) were both employed [55]. The results presented in Table V suggest that the number of fake news is nonstationary—the mean, variance, or autocorrelation of this series changes over time.

Among the classical methodologies employed for time series forecasting, the autoregressive moving average (ARMA) [56] models have found widespread application in addressing a diverse array of forecasting challenges [57], [58].

TABLE V
STATIONARITY STATISTICAL TESTS RESULTS

	ADF	KPSS
Statistic	-3.06	4.37
p-value	0.029	0.01
Critical Values 1%	-3.43	0.73
Critical Values 5%	-2.86	0.46
Critical Values 10%	-2.56	0.34

TABLE VI
FAKE NEWS SELF-PREDICTION FOR PERIODS

	Next day	2 days	4 days	7 days
RSME	3.96	8.88	16.71	18.40
MAE	3.39	7.38	15.12	15.61
R^2	-0.003	-0.934	-3.5	-3.95

The extensive adoption of this methodological paradigm has led to the development of refined variants such as time-varying (TV) ARMA [59], autoregressive fractionally integrated moving average (ARFIMA) [60], and autoregressive integrated moving average (ARIMA) [61], [62]. All these models have been computed, but we focus on ARIMA for providing better results. A statistical model is autoregressive if it predicts future values based on past values. Although stationarity is desirable for ARIMA, it can be applied over nonstationary series as well [63].

The ARIMA model was computed for different time horizons using the auto ARIMA methodology.⁴ It is useful to automatically find the best parameters and transformations to assist in achieving the most accurate prediction. Table VI shows a summary of the results. The model resulted in significant errors across all predictions. Notably, the R^2 value was negative, indicating that the model's adaptation was unreliable.

D. Impact of Time Horizons and Input Sizes

Results are depicted in Table VII. It shows the prediction errors per model, time horizon, and window size. Moreover, the mean and the standard deviation of fake

⁴<http://alkaline-ml.com/pmdarima/>, last access April 14, 2023.




Fig. 4. Preliminary test results.

news (i.e., Maldita dataset) are included for comparison purposes.

This analysis is carried out in two steps of increasing difficulty. Thus, the first part is devoted to predicting the amount of fake news in a given period, whereas the second one targets a particular day in the future.

1) *Predictions for Periods*: For this first set of experiments, two to ten input days were considered for the sake of representativeness. All models (excluding the one with headlines, which is discussed later) report consistent values—when increasing the amount of input days, models learn better the trends. The longer the window sizes, the better the results. This improvement in error is especially interesting when predicting more extended periods.

The target period size (up to seven days, recall Section V-B) affects the model performance. When predicting the number of fake news in the next four and seven days, models do not have to face the complexity of understanding the possible abrupt changes that may happen from one day to another. The variability of the data is reduced and hence, the model does not have to face that much daily life uncertainty. When predicting the next day, all models show a relatively worse accuracy in direct comparison with the prediction of the amount of fake news in the next seven days. Indeed, in terms of MAE, for horizon one day, MAE is around 3 and given a mean of 7.72 fake news in the dataset, the error is close to 38%. The situation improves for larger horizons and large windows size, for instance, the horizon of four days and windows of eight days, the mean MAE is 6.1 which leads to 21% of error.

Regarding the applied models, it has been observed that the LSTM model incorporating attention outperforms its simple LSTM counterpart. Remarkably, the R^2 outcomes corroborate the efficacy of these models in capturing the underlying pattern of fake news. This improvement can be attributed to the attention mechanism, which facilitates the model's comprehension of the interdependencies among inputs via attention alignments. In this regard, the optimal performance is attained by employing the LSTM model with attention to a temporal horizon of seven days with an input span of ten days, leading to a 10.3% error.

2) *Predictions for Particular Dates*: Based on previous results, models are now fed with longer spans of data as they lead to a better output. Thus, input sizes between 5 and 20, in steps of 5, are considered herein.

Table VIII presents the results. As expected, the higher the distance of the prediction, the more error is produced. This prediction seems to be more challenging for the encoder-decoder, which despite showing an inline MAE and RMSE, the metric R^2 falls apart the rest of the models. Indeed, its value is negative so the model is not reliably fitting the data. The LSTM shows better behavior, given the relative architectural simplicity compared to the rest. As in the previous case, the use of attention mechanisms is beneficial.

Results confirm that increasing the input size has a slightly positive impact on most of the models, especially on the simple LSTM. In the best case, the error produced as compared to the mean value is $1.91/6.65 = 28.7\%$.

E. Impact of the Use of Headlines

Headlines models show slightly better R^2 in some of the input sizes as compared to the rest of the models when predicting periods (see Table VII, rightmost column). This model is better at capturing the data variability, especially in shorter window sizes. Headlines seem to supply the lack of statistical information on shorter window sizes adding in-context features that help the model better understand short-term trends.

When increasing the prediction horizon, the headline positively impacts the results until four days. The decrease in performance may explain that having too much textual information from the past adds extra noise. For example, when using window sizes of ten days, the model consumes 100 headlines, which talk about many different events and topics leading to an information overload. However, though errors are quite similar in most settings, the mean and standard deviation in Maldita show that though error is bigger in larger horizons, such configurations are preferable.

When predictions for specific dates are at stake (see Table VIII, rightmost column), most of these findings are consistent. Again, longer input values are not beneficial, especially for medium to long-term predictions (i.e., four to seven days). Some settings lead to negative R^2 , which shows its unreliability. However, for short-term predictions, the accuracy is slightly better than that of LSTM with attention. This might be explained as the news at stake may have a closer semantic connection with the predicted fake news.

TABLE VII
PREDICTIONS FOR PERIODS

Maldita dataset				LSTM			Encoder-decoder			LSTM + attention			LSTM + Headlines		
		Horizon	Input Size	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2
Mean	7.72	Next day	2 Days	4.08	5.78	0.385	4.03	5.61	0.417	3.83	5.45	0.440	4.18	6.17	0.451
			4 Days	3.88	5.84	0.435	3.93	5.83	0.438	3.76	5.47	0.505	4.04	5.95	0.489
			6 Days	3.71	5.58	0.473	3.69	5.58	0.494	3.58	5.43	0.497	4.11	6.03	0.475
			8 Days	3.94	5.87	0.492	3.83	5.71	0.513	3.61	5.45	0.565	4.07	5.85	0.506
			10 Days	3.68	5.59	0.437	3.59	5.29	0.496	3.35	5.06	0.544	3.96	5.76	0.522
Standard deviation	8.23	Next 2 days	2 Days	6.78	9.43	0.492	6.50	9.25	0.511	6.55	9.20	0.520	4.59	6.62	0.747
			4 Days	6.24	8.51	0.496	6.20	8.42	0.501	5.95	8.33	0.504	4.95	7.00	0.716
			6 Days	5.83	8.00	0.557	5.96	8.17	0.532	5.51	7.68	0.578	5.01	6.84	0.731
			8 Days	5.23	7.34	0.710	5.36	7.27	0.717	4.61	6.53	0.770	4.98	6.87	0.725
			10 Days	5.41	7.56	0.627	5.26	7.70	0.617	4.88	7.00	0.681	4.82	6.76	0.737
Mean	14.40	Next 4 Days	2 Days	8.95	12.30	0.645	8.64	12.10	0.657	8.91	12.30	0.642	5.55	8.02	0.868
			4 Days	8.27	11.30	0.679	8.22	11.50	0.667	8.08	11.20	0.681	5.58	7.88	0.872
			6 Days	7.44	9.96	0.788	6.97	9.51	0.806	6.42	9.29	0.810	5.91	8.25	0.860
			8 Days	6.86	9.32	0.805	6.28	8.77	0.826	6.09	8.58	0.843	5.79	8.46	0.853
			10 Days	6.08	8.49	0.821	5.84	8.29	0.828	5.71	7.68	0.847	5.61	7.90	0.871
Standard deviation	21.90	Next 7 Days	2 Days	11.70	15.50	0.772	11.40	15.40	0.775	11.30	14.90	0.791	7.07	10.00	0.947
			4 Days	10.50	14.60	0.788	10.20	14.80	0.783	9.27	13.20	0.829	7.76	10.50	0.943
			6 Days	8.50	11.90	0.862	7.92	11.00	0.879	7.31	11.10	0.878	7.76	10.70	0.939
			8 Days	7.08	10.10	0.904	7.57	10.80	0.890	6.03	9.19	0.920	7.91	11.00	0.936
			10 Days	5.95	8.13	0.934	6.35	8.62	0.926	5.18	7.33	0.947	8.01	11.20	0.934

TABLE VIII
PREDICTION FOR SPECIFIC DATES

Maldita dataset			LSTM			Encoder-decoder			LSTM + attention			LSTM + headlines				
			Horizon	Input Size	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2
Mean	7.2	In 2 days	5 Days	2.53	3.53	0.115	2.55	3.60	0.082	2.29	3.37	0.194	2.13	3.23	0.173	
			10 Days	2.28	3.29	0.238	2.40	3.51	0.133	2.20	3.22	0.271	2.24	3.45	0.099	
			15 Days	2.35	3.35	0.197	2.28	3.37	0.187	2.10	3.16	0.286	2.24	3.45	0.116	
			20 Days	2.20	3.27	0.246	2.38	3.49	0.136	2.13	3.21	0.274	1.96	2.90	0.128	
		In 4 Days	5 Days	3.63	5.30	0.602	2.77	3.75	0.005	2.33	0.77	0.175	1.91	2.91	0.079	
			10 Days	4.00	5.65	0.502	2.79	3.69	0.045	2.19	0.71	0.288	2.10	3.80	-0.043	
			15 Days	2.58	3.85	0.690	2.60	3.68	0.035	2.12	0.83	0.256	2.33	3.42	0.021	
			20 Days	3.47	5.15	0.605	2.62	3.73	0.015	2.16	0.78	0.264	2.40	3.51	0.046	
		In 7 Days	5 Days	2.61	3.56	0.103	2.89	3.83	-0.035	2.36	3.40	0.182	3.63	5.30	0.122	
			10 Days	2.49	3.47	0.145	2.85	3.84	-0.042	2.23	3.29	0.235	4.00	5.65	-0.067	
			15 Days	2.44	3.52	0.117	2.84	3.78	-0.015	2.15	3.25	0.252	2.58	3.85	-0.070	
			20 Days	2.62	3.61	0.084	2.82	3.79	-0.009	2.20	3.35	0.212	3.47	5.15	-0.032	
Standard deviation		8.16														

F. Discussion

Fake news volume prediction is a complex task as many factors may happen in the real world, such as unexpected riots. However, despite the complexity, the models have been shown to learn trends and produce meaningful predictions.

The horizon and input size are particularly relevant when predicting for periods, where the higher they are, the better the predictions and especially using LSTM with attention. However, including headlines, though the horizon is a key element to maximize, the input size is not that important as MAE and R^2 differ just around 0.5 and 0.05 between the best and the worst input size.

In the prediction for specific dates, neither the horizon nor the input size seems to be remarkable features because error values are quite close to each other. Nonetheless, considering R^2 , the better choice is the horizon four days and the LSTM model. In contrast to prediction for periods, headlines do not lead to improvements as R^2 is quite small or negative.

Looking at Table VIII, it can be observed that some configurations like the encoder-decoder and LSTM + headlines are not able to reproduce the variability of the data in windows of four days or more. A negative R^2 suggests that the best fit the models were able to find in those cases is a simple straight line which tries to minimize the prediction error.

Finally, experiments are compared with ARIMA (recall Section VI-C, Table VI) for being a common statistical technique. Our proposal demonstrates superior performance

in all the metrics versus ARIMA in the majority of cases, especially when considering predictions beyond the next day.

ARIMA's MAE and RMSE (refer to Table VI) show better numerical results when predicting the volume of fake news for the next day in some cases. Particularly, when comparing ARIMA against LSTM, encoder-decoder, and LSTM + Headlines models, it performs slightly better. However, in terms of MAE the LSTM + attention model outperforms ARIMA (MAE = 3.39) by achieving a lower value of 3.35 with input data of ten days. When predictions are made for a horizon of two days and beyond, our results outperform ARIMA in the vast majority of cases. Thus, excluding LSTM, encoder-decoder and LSTM + attention, RMSE with two days input size on a horizon of next two days, any other numerical comparison with ARIMA demonstrates better performance across the three aforementioned metrics.

When comparing predictions for seven-day periods, our models lead to a $18.4/5.18 \cdot 100 = 355\%$ of improvement. Moreover, ARIMA uses historical fake news' data, while proposed models apply real news, being an advantage from a real setting in which fake news are not easily or immediately recognized. Additionally, ARIMA uses one variable (i.e., amounts of daily fake news) as input and for prediction and results show the complexity of the problem and the need for considering different variables. Indeed, highlighting the data variability metric (R^2), all the models and configurations

presented in Table VII outperform ARIMA's R^2 , which fails to capture the data variability of the problem effectively.

Two issues endorse the significance of these results. On the one hand, a substantial amount of news is at stake, totaling around 307k. On the other hand, the long time span—around four years. Nonetheless, these findings are not easy to extrapolate to different cultures, regions, or languages, as media content and fake news may heavily differ.

As compared to existing works (recall Table I), our results show that our approach is the only one addressing the goals stated in this article, involving different prediction horizons while keeping data affordability.

VII. LIMITATIONS

The high variability of the data given the uncertainty inherited in real-world data models a really complex problem that impacts the capacities of the models in certain set-ups. Some results rendered in Table VIII suggest that in some cases, the models are not able to find a solution that represents this variability—their conservative nature makes it difficult to handle steep changes during the forecasting process.

The scope of this work is restricted to the case of Spain, since the data have been provided by a fact-checker that is focused on this country. Different countries can have unique social, cultural, and political contexts that influence the nature of information and fact-checking dynamics. Futures collaboration with more fact-checkers companies is needed in order to contrast results and improve the generalization of the solution and the performance in order to better understand the misinformation dynamics.

In what comes to the data at stake, the LSTM + headlines model focuses solely on the headlines of the pieces of news. This limitation arises due to the nature of GDELT dataset, which does not provide the actual body of the news in their records. GDELT dataset information is collected from a wide range of diverse sources. Given the vast number of sources and the lack of automatic access to the full text of the news articles, it becomes challenging to automatically extract and analyze all the information contained within the articles. Not accessing the full body of the news limits a deeper analysis and potentially restricts the model's ability to capture subtle nuances that may be present in the mere wording of the news. The headlines provide some insights, but they may not capture the entire breadth of the news content.

VIII. CONCLUSION AND FUTURE WORK

Predicting the amount of expected fake news may help one to optimize the use of detection techniques. However, this daunting task has received little attention from the research community. In this regard, this work has focused on how techniques may leverage data from real news to perform predictions on fake news in Spain. Our results support the feasibility of this approach. Although dramatic differences exist between models, input data, and prediction horizons, the applied techniques have shown their ability to reproduce the fake news trend along a period of four years.

Regarding future work, applying these techniques and models in other contexts and countries is relevant to

compare results, considering increasing the amount of textual information sources, i.e., news corpus and more news sources. However, it would require collecting comprehensive datasets suitable for this specific purpose, which is challenging. Deploying more advanced encoders, for example, transformers, to extract semantics information from the headlines or for sentiment feature extraction could be a good exercise for comparison. The length of the provided dataset with real fake news could be larger, facilitating the training of the RNN.

ACKNOWLEDGMENT

The authors thank Fundación Maldita.es for sharing their fake news database and Prof. Angel Sanchez for his recommendations on error measurements. They thank the anonymous reviewers for their insightful comments.

CONFLICT OF INTEREST

The authors declare they have no conflict of interest.

REFERENCES

- [1] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *SIGKDD Explor. Newslett.*, vol. 19, no. 1, pp. 22–36, Jun. 2017, doi: [10.1145/3137597.3137600](https://doi.org/10.1145/3137597.3137600).
- [2] B. Yeoman, "That's fake news!" *Saturday Evening Post*, 2022. [Online]. Available: <https://www.saturdayeveningpost.com/2022/07/thats-fake-news/>
- [3] S. Simpson, *Fake News: A Global Epidemic Vast Majority (86%) of Online—Ipsos*. Accessed: Nov. 24, 2022. [Online]. Available: <https://www.ipsos.com/en-us/news-polls/cigi-fake-news-global-epidemic>
- [4] C. Reuter, K. Hartwig, J. Kirchner, and N. Schlegel, "Fake news perception in Germany: A representative study of people's attitudes and approaches to counteract disinformation," in *Human Practice. Digital Ecologies. Our Future: 14. Internationale Tagung Wirtschaftsinformatik (WI 2019)*. Darmstadt, Germany: Association for Information Systems AIS, 2022, pp. 1069–1083, doi: [10.26083/tuprints-00020795](https://doi.org/10.26083/tuprints-00020795).
- [5] Á. Figueira and L. Oliveira, "The current state of fake news: Challenges and opportunities," *Proc. Comput. Sci.*, vol. 121, pp. 817–825, Jan. 2017.
- [6] J. Allen, B. Howland, M. Mobius, D. Rothschild, and D. J. Watts, "Evaluating the fake news problem at the scale of the information ecosystem," *Sci. Adv.*, vol. 6, no. 14, Apr. 2020, Art. no. eaay3539, doi: [10.1126/sciadv.aay3539](https://doi.org/10.1126/sciadv.aay3539).
- [7] L. Graves and F. Cherubini, "The rise of fact-checking sites in Europe," *Digit. News Project Rep.*, 2016.
- [8] *Action Plan Against Disinformation*, Eur. Union, Brussels, Belgium, 2018.
- [9] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *ACM SIGKDD Explor. Newslett.*, vol. 19, no. 1, pp. 22–36, 2017.
- [10] R. R. Rajalaxmi, L. V. N. Prasad, B. Janakiramaiah, C. S. Pavankumar, N. Neelima, and V. E. Sathishkumar, "Optimizing hyperparameters and performance analysis of LSTM model in detecting fake news on social media," *Trans. Asian Low-Resource Lang. Inf. Process.*, 2022.
- [11] S. I. Manzoor and J. Singla, "Fake news detection using machine learning approaches: A systematic review," in *Proc. 3rd Int. Conf. Trends Electron. Inform. (ICOEI)*, Tirunelveli, India, 2019, pp. 230–234, doi: [10.1109/ICOEI.2019.8862770](https://doi.org/10.1109/ICOEI.2019.8862770).
- [12] X. Zhou, A. Jain, V. V. Phoha, and R. Zafarani, "Fake news early detection: A theory-driven model," *Digit. Threats, Res. Pract.*, vol. 1, no. 2, pp. 1–25, Jun. 2020.
- [13] M. Del Vicario, W. Quattrociocchi, A. Scala, and F. Zollo, "Polarization and fake news: Early warning of potential misinformation targets," 2018, *arXiv:1802.01400*.
- [14] T. Murayama, S. Wakamiya, E. Aramaki, and R. Kobayashi, "Modeling the spread of fake news on Twitter," *PLoS ONE*, vol. 16, no. 4, pp. 1–16, Apr. 2021, doi: [10.1371/journal.pone.0250419](https://doi.org/10.1371/journal.pone.0250419).
- [15] B. Rath, X. Morales, and J. Srivastava, "SCARLET: Explainable attention based graph neural network for fake news spreader prediction," in *Advances in Knowledge Discovery and Data Mining (Lecture Notes in Computer Science)*, vol. 12712, K. Karlapalem et al., Eds. Cham, Switzerland: Springer, 2021, doi: [10.1007/978-3-030-75762-5_56](https://doi.org/10.1007/978-3-030-75762-5_56).

- [16] S. Guarino, N. Trino, A. Chessa, and G. Riotta, "Beyond fact-checking: Network analysis tools for monitoring disinformation in social media," in *Complex Networks and Their Applications VIII* (Studies in Computational Intelligence), vol. 881, H. Cherifi, S. Gaito, J. Mendes, E. Moro, and L. Rocha, Eds. Cham, Switzerland: Springer, 2020, doi: [10.1007/978-3-030-36687-2_36](https://doi.org/10.1007/978-3-030-36687-2_36).
- [17] Statista. (2022). *World Languages by Number of Speakers*. [Online]. Available: <https://www.statista.com/statistics/266808/the-most-spoken-languages-worldwide/>
- [18] G. Rampersad and T. Althiyabi, "Fake news: Acceptance by demographics and culture on social media," *J. Inf. Technol. Politics*, vol. 17, no. 1, pp. 1–11, Jan. 2020.
- [19] S. Hakak, M. Alazab, S. Khan, T. R. Gadekallu, P. K. R. Maddikunta, and W. Z. Khan, "An ensemble machine learning approach through effective feature extraction to classify fake news," *Future Gener. Comput. Syst.*, vol. 117, pp. 47–58, Apr. 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167739X20330466>
- [20] T. Jiang, J. P. Li, A. U. Haq, A. Saboor, and A. Ali, "A novel stacking approach for accurate detection of fake news," *IEEE Access*, vol. 9, pp. 22626–22639, 2021.
- [21] S. Raza and C. Ding, "Fake news detection based on news content and social contexts: A transformer-based approach," *Int. J. Data Sci. Anal.*, vol. 13, no. 4, pp. 335–362, May 2022, doi: [10.1007/s41060-021-00302-z](https://doi.org/10.1007/s41060-021-00302-z).
- [22] O. Ajao, D. Bhowmik, and S. Zargari, "Fake news identification on Twitter with hybrid CNN and RNN models," in *Proc. SMSociety*. New York, NY, USA: Association for Computing Machinery, Jul. 2018, pp. 226–230.
- [23] S. Deng, H. Rangwala, and Y. Ning, "Understanding event predictions via contextualized multilevel feature learning," in *Proc. 30th ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2021, pp. 342–351.
- [24] K. Leetaru and P. A. Schrodt, "GDELT: Global data on events, location, and tone," in *Proc. ISA Annu. Conv.*, vol. 2, no. 4. Princeton, NJ, USA: Citeseer, 2013, pp. 1–49.
- [25] D. J. Gerner, P. A. Schrodt, O. Yilmaz, and R. Abu-Jabr, "Conflict and mediation event observations (CAMEO): A new event data framework for the analysis of foreign policy interactions," *Int. Stud. Assoc.*, New Orleans, LA, USA, 2002.
- [26] P. Candito et al., "Predicting social unrest events with hidden Markov models using GDELT," *Discrete Dyn. Nature Soc.*, vol. 2017, May 2017, Art. no. 8180272.
- [27] D. Galla and J. Burke, "Predicting social unrest using GDELT," in *Machine Learning and Data Mining in Pattern Recognition* (Lecture Notes in Computer Science), vol. 10935, P. Perner, Ed. Cham, Switzerland: Springer, 2018, doi: [10.1007/978-3-319-96133-0_8](https://doi.org/10.1007/978-3-319-96133-0_8).
- [28] Q. Fengcai, D. Jinsheng, and W. Li, "An online framework for temporal social unrest event prediction using news stream," in *Proc. Int. Conf. Cyber-Enabled Distrib. Comput. Knowl. Discovery (CyberC)*, Oct. 2020, pp. 176–182.
- [29] J. E. Yonamine, "Predicting future levels of violence in Afghanistan districts using GDELT," *Unpublished Manuscript*, 2013. [Online]. Available: <http://data.gdelproject.org/documentation/Predicting-Future-Levels-of-Violence-in-Afghanistan-Districts-using-GDELT.pdf>
- [30] O. I. Abiodun, A. Jantan, A. E. Omolara, K. V. Dada, N. A. Mohamed, and H. Arshad, "State-of-the-art in artificial neural network applications: A survey," *Heliyon*, vol. 4, no. 11, Nov. 2018, Art. no. e00938.
- [31] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [32] D. Wang, H. He, and D. Liu, "Intelligent optimal control with critic learning for a nonlinear overhead crane system," *IEEE Trans. Ind. Informat.*, vol. 14, no. 7, pp. 2932–2940, Jul. 2018.
- [33] R. Geirhos, D. H. J. Janssen, H. H. Schütt, J. Rauber, M. Bethge, and F. A. Wichmann, "Comparing deep neural networks against humans: Object recognition when the signal gets weaker," 2017, *arXiv:1706.06969*.
- [34] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986.
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Sequence Modeling: Recurrent and Recursive Net*. Cambridge, MA, USA: MIT Press, 2016, pp. 367–407.
- [36] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Trans. Neural Netw.*, vol. 5, no. 2, pp. 157–166, Mar. 1994.
- [37] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- [38] K. Smagulova and A. P. James, "A survey on LSTM memristive neural network architectures and applications," *Eur. Phys. J. Special Topics*, vol. 228, no. 10, pp. 2313–2324, Oct. 2019.
- [39] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, *arXiv:1409.0473*.
- [40] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 1–9.
- [41] Z. Chen, C. K. Yeo, B. S. Lee, and C. T. Lau, "Autoencoder-based network anomaly detection," in *Proc. Wireless Telecommun. Symp. (WTS)*, Apr. 2018, pp. 1–5.
- [42] H. Xu et al., "Unsupervised anomaly detection via variational auto-encoder for seasonal KPIs in web applications," in *Proc. World Wide Web Conf. World Wide Web (WWW)*. New York, NY, USA: ACM Press, 2018, pp. 187–196, doi: [10.1145/3178876.3185996](https://doi.org/10.1145/3178876.3185996).
- [43] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lect. IE*, vol. 2, no. 1, pp. 1–18, 2015.
- [44] I. Gulrajani et al., "PixelVAE: A latent variable model for natural images," 2016, *arXiv:1611.05013*.
- [45] W. Xu, S. Keshmiri, and G. Wang, "Adversarially approximated adversarial AutoEncoder for image generation and manipulation," *IEEE Trans. Multimedia*, vol. 21, no. 9, pp. 2387–2396, Sep. 2019, doi: [10.1109/TMM.2019.2898777](https://doi.org/10.1109/TMM.2019.2898777).
- [46] Y. Zhao, B. Deng, J. Huang, H. Lu, and X.-S. Hua, "Stylized adversarial AutoEncoder for image generation," in *Proc. 25th ACM Int. Conf. Multimedia*. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 244–251.
- [47] N. M. Brashier, G. Pennycook, A. J. Berinsky, and D. G. Rand, "Timing matters when correcting fake news," *Proc. Nat. Acad. Sci. USA*, vol. 118, no. 5, Feb. 2021, Art. no. e2020043118.
- [48] J. M. Pérez, J. C. Giudici, and F. Luque, "Pysentimiento: A Python toolkit for sentiment analysis and SocialNLP tasks," 2021, *arXiv:2106.09462*.
- [49] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics (Long Papers)*, vol. 1. Melbourne, VIC, Australia: Association for Computational Linguistics, 2018, pp. 328–339.
- [50] P. Sedgwick, "Pearson's correlation coefficient," *BMJ*, vol. 345, Jul. 2012, Art. no. e4483, doi: [10.1136/bmj.e4483](https://doi.org/10.1136/bmj.e4483).
- [51] D. P. Kingma and J. B. Adam, "A method for stochastic optimization," in *Proc. ICLR*, 2015, p. 5.
- [52] S. Ruder, "An overview of gradient descent optimization algorithms," 2016, *arXiv:1609.04747*.
- [53] A. Botchkarev, "A new typology design of performance metrics to measure errors in machine learning regression algorithms," *Interdiscipl. J. Inf., Knowl., Manage.*, vol. 14, pp. 45–76, Jan. 2019.
- [54] D. Chicco, M. J. Warrens, and G. Jurman, "The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation," *PeerJ Comput. Sci.*, vol. 7, p. e623, Jul. 2021.
- [55] D. Kwiatkowski, P. C. B. Phillips, P. Schmidt, and Y. Shin, "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?" *J. Econometrics*, vol. 54, nos. 1–3, pp. 159–178, Dec. 1992.
- [56] J. Gurland, "Review of Hypothesis Testing in Time Series Analysis," by P. Whittle," *J. Amer. Stat. Assoc.*, vol. 49, no. 265, pp. 197–200, 1954, doi: [10.2307/2281054](https://doi.org/10.2307/2281054).
- [57] P. Cortez, M. Rocha, and J. Neves, "Evolving time series forecasting ARMA models," *J. Heuristics*, vol. 10, no. 4, pp. 415–429, Jul. 2004.
- [58] W. Polasek, "Review of time series analysis and its applications: With R examples, third edition, by R. H. Shumway & D. S. Stoffer," *Int. Stat. Rev.*, vol. 81, no. 2, pp. 323–325, 2013. [Online]. Available: <http://www.jstor.org/stable/43298885>
- [59] Y. Grenier, "Time-dependent ARMA modeling of nonstationary signals," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-31, no. 4, pp. 899–911, Aug. 1983.
- [60] C. W. J. Granger, "Long memory relationships and the aggregation of dynamic models," *J. Econometrics*, vol. 14, no. 2, pp. 227–238, Oct. 1980.
- [61] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*. Hoboken, NJ, USA: Wiley, 2015.
- [62] M. H. Pesaran, *Time Series and Panel Data Econometrics*. London, U.K.: Oxford Univ. Press, Oct. 2015, doi: [10.1093/acprof:oso/9780198736912.001.0001](https://doi.org/10.1093/acprof:oso/9780198736912.001.0001).
- [63] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting* (Springer Texts in Statistics). Springer, 2016, doi: [10.1007/978-3-319-29854-2](https://doi.org/10.1007/978-3-319-29854-2).