Overview of FakeDeS at IberLEF 2021: Fake News Detection in Spanish Shared Task

Resumen de FakeDeS en IberLEF 2021: Tarea compartida para la detección de noticias falsas en español

Helena Gómez-Adorno[1], Juan Pablo Posadas-Durán[2],
Gemma Bel Enguix[3], Claudia Porto Capetillo[1]
[1]Instituto de Investigaciones en Matemáticas Aplicadas y en Sistemas,
Universidad Nacional Autónoma de México, Mexico
[2]Instituto de Ingeniería, Universidad Nacional Autónoma de México, Mexico
[3]Escuela Superior de Ingeniería Mecánica y Eléctrica, Unidad Zacatenco,
Instituto Politécnico Nacional, Mexico
[4]Posgrado en Ciencia e Ingeniería de la Computación,
Universidad Nacional Autónoma de México, Mexico
helena.gomez@iimas.unam.mx, jposadasd@ipn.mx,
gbele@iingen.unam.mx, clauporto@comunidad.unam.mx

Abstract: This paper presents the overview of FakeDeS 2021, the second edition of this lab under the IberLEF conference. The FakeDeS shared task aims to explore different methodologies and strategies related to fake news detection in Spanish. This year edition brings two main challenges: thematic and language variation. For this purpose, we introduce a new testing corpus containing news related to COVID-19 and news from other Ibero-American countries.
Keywords: Fake news detection, FakeDeS, Iberlef.

Resumen: Este artículo hace un presentación general de la tarea compartida Fake-DeS 2021, cuya segunda edición ha tenido lugar en 2021 bajo el congreso IberLEF, aunque se trata de la primera vez con esta denominación. La tarea FakeDeS tiene por objetivo explorar diferentes métodos y estrategias relacionados con la detección de noticias falsas en español, principalmente en su variante de México. La edición de esta año propone dos desafíos principales: variación temática y variación lingüística. Para ello, se introduce un nuevo corpus de prueba que contiene noticias relacionadas con COVID 19 y noticias de otros países de Iber-América.
Palabras clave: Fake news detection, FakeDeS, Iberlef.

1 Introduction

In recent years, internet and social networks have increased the fast spreading of news and the power of individuals to create and share their own content, that is usually partial and un-verified. Traditional media have been forced to adapt to this new online scenario that favours "spectacle over restraint and veriication" (Chen, Conroy, and Rubin, 2015).

This situation has an impact in the arising of Fake News, which either have the objective to manipulate people or have not been conveniently fact-checked. Misinformation is present in every area: politics, science, even in sports. It spreads in seconds among thousands of people. Therefore, it is necessary to build computational systems in order to identify the false information in the web and social media.

As every task in NLP, Fake News detection has been developed mainly for English. However, and given the number of speakers of Spanish around the world, it seems urgent to design methods that can identify fake content in this language.

A fake news detection system aims to help users detect and filter out potentially deceptive news. An approach for the prediction of intentionally misleading news is based on the analysis of truthful and fraudulent previously reviewed news, i.e., annotated corpora.

The FakeDeS task is the second edition of the fake new detection task organized in

Helena Gómez-Adorno, Juan Pablo Posadas-Durán,
Gemma Bel Enguix, Claudia Porto Capetillo

the IberLeF conference. The first edition was integrated in the MEX-A3T (Aragón et al., 2020), a forum for the analysis of social media content in Mexican Spanish, at IberLeF.

The 2021 FakeDeS task at Iberlef 2021 (Montes et al., 2021) has taken into account the global pandemic situation. A new corpus containing news associated with COVID-19 will be used as a testing corpus, while the corpus used in the 2020 edition has been provided as the training set. Our aim is to explore the robustness of methods when trained on generic news and then evaluated in news associated with a very specific theme.

The rest of this paper is organized as follows. Section 2 presents the evaluation framework used at FakeDeS 2021. Section 3 shows an overview of the participating approaches. Section 4 reports and analyzes the results obtained by the participating teams. Finally, Section 5 presents our conclusions from this shared task.

2 FakeDeS 2021 Corpus and evaluation framework

For the development of the solution proposals for the FakeDeS task, the complete corpus, i.e. both training and testing partitions of the corpus, used in the previous edition of the task (Fake News Detection track at MEX-A3T) was made available to the participants. The corpus used for the training is a compilation of news collected between January and July 2018, using the following resources on the web: websites of established newspapers and media companies, special websites dedicated to the validation of fake news and websites designated by different journalists as sites that publish fake news regularly. The news were tagged as true and fake manually following the procedure described in (Posadas-Durán et al., 2019). The training corpus includes true-fake news pairs of diverse events to have a corpus as balanced as possible. The data contains 971 news divided into 491 real news and 480 fake news. The news compiled is related to 9 different topics (Science, Sports, Economy, Education, Entertainment, Politics, Health, Safety and Society) established by the source of the news. Table 2 shows the topic distribution in the train corpus.

Participants can evaluate their approach to achieve satisfying performance by applying the corpus split used in the previous editions

Topic	True	Fake
Science	46	43
Sport	66	58
Economy	24	19
Education	10	12
Entertainment	70	78
Politics	175	148
Health	23	23
Security	17	25
Society	60	74
Σ	**491**	**480**

Tabla 1: Topic distribution in the train corpus.

of the task (Aragón et al., 2020).

In order to evaluate of the submitted approaches, a new set of news were collected to conform the evaluation corpus. The test corpus contains pairs of fake and true publications about different events, that were collected from November 2020 to March 2021. Different sources from the web were used to gather the information, but mainly of two types: 1) newspapers and media companies websites, and 2) fact-cheking websites.

The following fact-cheking websites were used in the process of the corpus compilation:

1. *Maldito Bulo*
   (`https://maldita.es/malditobulo`)

2. *Colombian Check*
   (`https://colombiacheck.com/`)

3. *AFP Factual*
   (`https://factual.afp.com/`)

4. *EcuadorChequea*
   (`http://www.ecuadorchequea.com/`)

5. *Chequeado*
   (`https://chequeado.com/`)

6. *Verificado*
   (`https://verificado.com.mx/`)

7. *Animal Político*
   (`https://www.animalpolitico.com/`)

Most of the revised fact-checking sites used follow the recommendations of the International Fact-Checking Network[1](IFCN) at Poynter that seeks to promote good practice in fact-checking. Only two of the sites do not show the badge issued by IFCN, however

---

[1]`https://ifcncodeofprinciples.poynter.org/`

on their website they describe the methodology they follow for the fact-checking.

For the evaluation, the assembled test corpus has 572 instances. The instances were labeled using the same two classes considered in train set, i.e., true or fake. The test corpus is balanced with respect to these two classes. To compile the true-fake news pair of the test corpus, the following guidelines were followed:

- A fake news is added to the corpus if any of the selected fact-checking sites determines it.

- Given a fake news, its true news counterpart is added if there is evidence that it has been published in a reliable site (established newspaper site or media site).

The topics covered in the test corpus that match the training corpus are Science, Sport, Politics, and Society. Three new topics, COVID-19, Environment, and International, were added to the test corpus. Table 2 shows the topic distribution in the test corpus. The test corpus includes mostly news articles, however, on this occasion social media posts were also included in the category of fake news. Exactly 90 posts were included as fake news (15.73 % of the total). This posts were recovered mainly from Facebook and WhatsApp.

Topic	True	Fake	Posts
Environment	2	2	1
Science	6	7	3
COVID-19	118	119	59
Sport	1	1	0
International	7	7	0
Politics	53	54	17
Society	99	96	10
Σ	**286**	**286**	**90**

Tabla 2: Topic distribution in the test corpus.

The use of the various fact-checking sites involved consulting pages from different countries that offer content in Spanish in addition to Mexico, so different variants of Spanish are included in the test corpus. These sites included countries like Argentina, Bolivia, Chile, Colombia, Costa Rica, Ecuador, Spain, United States, France, Peru, Uruguay, England and Venezuela.

3 Overview of the Submitted Approaches

At this edition, twenty one teams submitted one or more solutions to the task through the codalab platform[2]. CodaLab Competitions is a powerful open source framework for running competitions that involve result or code submission. The evaluation methodology of a competition in this platform consists of receiving as input the predictive outputs of systems and returns a performance evaluation based on the metrics defined for each task.

This section presents a summary of their approaches regarding preprocessing steps, features, and classification algorithms. In Table 3 we indicate the general approach used for each team. It can be appreciated that participants used three general approaches: transformers, deep neural networks, and traditional representations like BoW and n-grams feeding a SVM classifier. Following, we briefly describe each of the participating methods

- *GDUF_DM at FakeDeS 2021: Spanish Fake News Detection with BERT and Sample Memory* (Huang, Xiong, and Jiang, 2021)

  - **Team name: GDUFS_DM**
  - **Summary:** The authors detect Spanish fake news using BERT and Sample Memory with an attention mechanism. The method consists of taking the first and last segments of the texts and feeding them into a BERT system, obtaining two embeddings (head and tail). Additionally, a so-called "sample memory" matrix is obtained by taking a random sample of the head and tail embeddings; this matrix is used in an attention mechanism with the rest of the input (the previously explained head and tail embeddings of each entry). The final embedding is formed from the concatenation of the head, tail, and memory embeddings.

- *HAHA at FakeDeS 2021: A Fake News Detection Method Based on TF-IDF and Ensemble Machine Learning* (Li, 2021)

---

[2]https://competitions.codalab.org/competitions/29545

Approach	GDUFS_DM	Haha	Lcad_UFES	CiTIUS-NLP	zk15120170770	ForceNLP	TSIA	Yeti	Bribones
Transformers	X			X	X		X	X	
Traditional Deep Neural Networks							X		
BoW, n-grams, Stylometrics		X	X	X		X			X

Tabla 3: General approach of each participating team.

– **Team name: Haha**

– **Summary:** The authors presented a classic bag-of-words feature extraction with TF-IDF weighting scheme and a stacking ensemble learning method based on weak classifiers. Their model not only analyzes the content of the news, but also combines information such as publishers and topics using label encoding to improve the performance of the model.

- *LCAD - UFES at FakeDeS 2021: Fake News Detection Using Named Entity Recognition and Part-of-Speech Sequences* (Spalenza et al., 2021)

  – **Team name: LCAD_UFES**

  – **Summary:** The authors implemented a method using named entity recognition and Part-of-Speech sequences with a machine learning approach. The authors analyze the effectiveness of different classifiers, support vector machines, random forest, gradient boosting and Wi-SARD; achieving the best performance with gradient boosting.

- *CiTIUS at FakeDeS 2021: A Hybrid Strategy for Fake News Detection* (Gamallo, 2021)

  – **Team name:CiTIUS-NLP**

  – **Summary:** The authors explored three different BERT-based supervised classification strategies: 1) a fine tuned model approach that adds a dense layer on top of the last layer of the pre-trained BERT model and then trains the whole model on the task specific dataset, 2)

a sentence similarity approach using sentence embeddings obtained with BERT model, and 3) an hybrid approach that combines sentence similarity with some linguistic heuristics (size of the news, the presence of sentences written in capital letters, or specific fake statements).

- *zk15120170770 at FakeDeS 2021: Fake news detection based on Pre-training Model* (Zhao, Zhou, and Li, 2021)

  – **Team name: zk15120170770**

  – **Summary:** The authors presented a method based on fine-tuning the XML-Roberta pretrained language model. They trained the classifier using Adam optimizer with a learning rate 2e-5 and cross entropy loss. The authors used 10 epochs for the training process and the maximum length of the news is set to 512.

- *ForceNLP at FakeDeS 2021: Analysis of Text Features Applied to Fake News Detection in Spanish* (Reyes-Magaña and Argota Vega, 2021)

  – **Team name: ForceNLP**

  – **Summary:** The authors presented an approach based on traditional supervised learning using of different types and of n-grams (of different size) to represent the news. The authors use classical classifiers: Logistic Regression, SVM and Multinomial Naive-Bayes to distinguish between Fake and True news. Their objective is to analyze the impact of text features in the task of fake news detection.

- *TSIA team at FakeDeS 2021: Fake News Detection in Spanish Using Multi-Model Ensemble Learning* (Guan, 2021)

  – **Team name: Tsia**

  – **Summary:** The author presented three model architectures based on the pre-trained model BETO and XLM-RoBERTa-Large. The author first fine-tuned the Spanish pre-trained model BETO and then used the multi-language model XLM-RoBERTa-Large to replace BETO, the third model includes the addition of CNN for feature extraction.

- *YETI at FakeDeS 2021: Fake News Detection in Spanish with ALBERT* (Luo, 2021)

  – **Team name: Yeti**

  – **Summary:** The authors presented a modification of ALBERT base model adding richer semantic content. ALBERT, unlike BERT, implements three improvements: Embedding layer, parameter sharing, and SOP task. The authors modified ALBERT to directly receive the news and concatenate the hidden state of the first token sequence of the last three hidden layers into the classifier. The results of the modification on 5-fold cross-validation show an improvement over the AL-BERT base model.

- *Bribones tras la esmeralda perdida@FakeDeS 2021: Fake news detection based on random forests, k-nearest neighbors, and n-grams for a Spanish corpora* (Lomas-Barrie et al., 2021)

  – **Team name: Bribones tras la esmeralda perdida**

  – **Summary:** The authors used a simple BOW-based classification system to detect fake news. The methods implemented are K-nearest neighbors and random forests. The paper starts from an overly naive hypothesis: true and fake news can be easily detected with the vocabulary they use.

This hypothesis is rejected in the conclusions of the paper.

4 Experimental evaluation and analysis of results

This section summarizes the results obtained by the the participants of FakeDeS at Iber-LEF 2021: Fake News Detection in Spanish Shared Task. We compare and analyze in detail the performance of the submitted solutions. For the final phase of the challenge, participants sent their predictions for the test partition, the performance on this data was used to rank them. We used the F_1 score over the fake class as the main evaluation measure.

For computing the evaluation scores we relied on the Codalab Competitions platform. In this section, we report the results obtained by participants as evaluated by Codalab Competitions and an analysis of their results.

As baseline methods, we implemented three popular approaches: i) a classification model trained on the bag of words (BoW) representation, ii) a classification model trained on the n-grams representation, and iii) a classification model based on transformers using BERT.

For the three baseline, the corpus was pre-processed using techniques described in the NLP literature, such as elimination of stop words, punctuation marks, and tokens related to numbers. Other techniques used are lowercasing, tokenization, and lemmatization.

The first baseline is a classification model trained on the BoW representation, where all vocabulary was used for the vectorization. Particular preprocessing techniques for this baseline includes tokenization (keeping the stopwords) and lowercasing. For classification we used an SVM classifier with linear kernel.

The second baseline is a classification model trained on the n-grams representation. The preprocessing techniques are the same as in the first baseline. The text is represented by character 3-grams and all vocabulary was used. For classification we used an SVM classifier with linear kernel.

The last baseline follows a deep learning approach, it is based on the Bidirectional Encoder Representation from Transformers (BERT) architecture(Devlin et al., 2019). We use the pre-trained model BETO: Spanish

BERT [3]. We trained the classifier using an Adam optimizer with a learning rate of $2e-5$. The epoch and maximum sentence length is 10 and 512, respectively. For the process of fine tuning and sequence classification, we use *Bert For Sequence Classification* as implemented in HuggingFace's Transformers (Wolf et al., 2020) which has a linear layer added at the top for classification to be used as a sentence classifier.

Table 4 shows a summary of the results obtained by each team in the FakeDeS shared task. We report the F_1 score in both fake and true classes, the macro F_1 score, and the accuracy. We used the F_1 over the fake class to rank participants. In this edition of the FakeDeS shared task, the approach submitted by the GDUFS_DM team outperformed all the other approaches and the baselines. GDUFS_DM used an approach based on a transformer architecture; however the second best approach (HAHA) used a classic approach with an ensemble classifier. These results show that for this task, classic approaches are still competitive with respect to deep learning. it is important to mention that the third and fourth best approaches did not send their system description papers.

All participated teams that send their system description papers used a machine-learning-based approach relying on style-based features. Neither team used fact-checking approach to verify the authenticity of the news.

For the analysis of the complementariness and the diversity of the predictions of the different approaches we use the Maximum Possible Accuracy (MPA) and Coincident Failure Diversity (CFD) metrics(Tang, Suganthan, and Yao, 2006). The MPA is defined as the ratio between the correctly classified instances over the total number of test instances. An instance is considered as correctly classified if at least one of the participating teams classified it correctly. On the other hand, the CFD metric give us the error diversity among the participants predictions. This measure gives a minimum value 0 when all teams simultaneously predict an instance correctly or wrongly, while a maximum value of 1, when all teams classify incorrectly different instances, i.e. misclassifications are all unique.

Table 5 shows the results of these me-

trics according to the approaches used by the participants. Transformers approach includes the following teams: GDUFS_DM, CiTIUS-NLP, zk15120170770, TSIA, and Yeti. Traditional Deep Neural Networks includes only the TSIA team. BoW, n-grams, and stylometrics approaches is used by the following teams: Haha, Lcas_UFES, CiTIUS-NLP, ForceNLP, and Bribones. The CFD for the traditional deep neural networks methodologies could not be calculated because only one participant team used this approach. The MPA for the row of all teams has the highest value, which means that the team's approaches complement each other. In terms of the different approaches, the Transformes approach obtained similar MPA over the *BoW, n-grams, Stylometrics* approach. However, the MPA of all approaches with paper submissions shows a 3% of increase over these individual approaches, suggesting that the transformers approaches and traditional approaches are complementary to each other. The values for the CFD score are comparable among all approaches, which means that their predictions are complementary to me extend, this lead us to conclude that there is still different information learned by traditional and transformer based approaches.

The Table 6 shows the results of the F_1 score for the fake class in the different topics of the test corpus. It can be appreciated that the Sports category achieved the overall lower results for all the evaluated approaches, even though there were six systems that correctly classified all the instances in this topic. This is related to the fact that in this category there are only two news items, one fake and one true. The performance of the systems does not seem related to the availability of the topics in the training set. While there were no news related to Covid-19 in the training set, the results achieved in this topic are comparable to those on topics available in the training set such as Politics and Society. The news belonging to the International topic were the easiest to classify in average, it can be observed that systems achieved F_1 scores above 15% and 20% in this topic.

We identified the common prediction errors across all the systems and find that there were only 2 news, both in the true class that none of the approaches classify correctly. All fake news were identified by at least one system. Table 7 shows the classified instan-

Team	Fake	True	F_{macro}	Accuracy
GDUFS_DM	**0.7666**	**0.7649**	**0.7666**	**0.7657**
Haha	0.7548	0.7522	0.7548	0.7535
Chats_	0.7514	0.7690	0.7514	0.7605
SINAI	0.7385	0.7821	0.7385	0.7622
baseline-BERT	0.7321	0.7432	0.7321	0.7378
baseline-BOW-SVM	0.7217	0.7359	0.7217	0.729
Lcad_UFES	0.7102	0.6837	0.7102	0.6976
CiTIUS-NLP	0.7098	0.4940	0.7098	0.6311
baseline-CHAR-3-GRAMS-SVM	0.7063	0.6883	0.7063	0.6976
zk15120170770	0.7053	0.6053	0.7053	0.6626
ForceNLP	0.6925	0.4739	0.6925	0.6119
GRX	0.6915	0.5624	0.6915	0.6381
TSIA	0.6860	0.5263	0.686	0.6224
FREE	0.6855	0.6519	0.6855	0.6696
LIMCA	0.6812	0.7027	0.6812	0.6923
ZZWEI	0.6737	0.6794	0.6737	0.6766
Premjithb	0.6576	0.7177	0.6576	0.6906
Sdamian	0.6542	0.75	0.6542	0.7098
Yeti	0.6316	0.609	0.6316	0.6206
Gulu	0.6226	0.476	0.6226	0.5612
Nicksss	0.6119	0.7592	0.6119	0.7028
Bribones tras la esmeralda perdida	0.5835	0.5878	0.5835	0.5857
WSSC	0.5118	0.6657	0.5118	0.6031
Skblaz	0.4838	0.649	0.4838	0.5822

Tabla 4: Results for the fake news detection task on the test set.

Approach	Best Accuracy	MPA	CFD	Number of Systems
Transformers	0.7657	0.9528	0.3834	5
Traditional Deep Neural Networks	0.6224	0.6224	-	1
BoW, n-grams, Stylometrics	0.7535	0.9580	0.3941	5
All teams (with submission)	0.7657	0.9895	0.3732	9
All teams	0.7657	0.9965	0.3382	21

Tabla 5: Comparison of MPA and CFD results between the different general approaches.

ces, it can be observed that one of the news belong to the Politics category and the other one to the Society category. It is interesting that all news (fake and true) in the Covid-19 topic were classified correctly by at least one team.

5 Conclusions

This paper described the design and results of the FakeDeS shared task collocated with IberLef 2021. FakeDeS stands for *Fake news Detection in Spanish*. This has been the second edition of the task, the first with this name, since the last year was presented as the fake news detection track in the wider shared task MEX-A3T (Aragón et al., 2020).

Although the best results in this shared task were reached by a team that proposed to

approach the problem with a method based in transformers, traditional machine learning strategies obtained very similar results. This seems to indicate the complexity of the task, that needs to be tackled by systems that consider multiple variables.

The results illustrate that fake news detection is a hard problem in the area of Natural Language Processing. The development of techniques especially designed to generate disinformation is a challenge for the area that has as a goal to identify such disinformation.

Summing up, the FakeDeS evaluation task promotes the work in Spanish in this crucial area of NLP, encourages the scientific exchange between researchers, and provides an annotated corpus in Spanish that is openly available for the scientific community. All this,

Helena Gómez-Adorno, Juan Pablo Posadas-Durán,
Gemma Bel Enguix, Claudia Porto Capetillo

Team	Covid-19	Politics	Society	International	Science	Sports	Environmental
GDUFS_DM	**0.78**	0.74	**0.76**	**0.93**	0.71	**1.00**	0.67
Haha	0.76	0.73	0.74	0.92	**0.88**	0.00	0.80
Chats_	0.77	0.76	0.71	0.92	0.62	**1.00**	**1.00**
SINAI	0.74	0.74	0.74	0.92	0.50	**1.00**	0.67
baseline-BERT	0.72	0.75	0.72	0.88	0.77	**1.00**	0.67
baseline-BOW-SVM	0.72	**0.77**	0.72	0.92	0.50	0.0	0.50
Lcad_UFES	0.71	0.74	0.70	0.77	0.71	0.67	0.50
CiTIUS-NLP	0.71	0.73	0.70	0.70	0.70	0.67	0.80
baseline-CHAR-3-GRAMS-SVM	0.68	0.73	0.71	0.92	0.75	0.0	0.80
zk15120170770	0.71	0.72	0.69	0.88	0.63	0.00	0.67
ForceNLP	0.69	0.67	0.71	0.67	0.67	**1.00**	0.80
GRX	0.70	0.72	0.66	0.71	0.78	0.00	0.50
TSIA	0.68	0.67	0.69	0.86	0.80	**1.00**	0.50
FREE	0.66	0.74	0.68	0.71	0.71	0.00	0.67
Limca	0.66	0.73	0.67	0.92	0.71	0.00	0.67
ZZWEI	0.66	0.70	0.68	0.80	0.67	0.00	0.67
Premjithb	0.66	0.63	0.64	0.83	0.80	0.67	0.67
Sdamian	0.64	0.68	0.68	0.83	0.62	0.00	0.50
Yeti	0.58	0.67	0.65	0.75	0.71	**1.00**	0.67
Gulu	0.61	0.59	0.64	0.74	0.67	0.67	0.80
Niksss	0.62	0.62	0.58	0.92	0.40	0.00	0.67
Bribones tras la esmeralda perdida	0.58	0.62	0.56	0.67	0.67	0.67	0.40
WSSC	0.35	0.49	0.66	0.92	0.55	0.00	0.00
Skblaz	0.40	0.34	0.61	0.86	0.46	0.00	0.50
Average	0.66	0.68	0.68	**0.83**	0.67	0.43	0.63

Tabla 6: Results for the fake news detection task by topics.

Topic	Source	Title
Politics	El Mundo	No, no ha habido un misterioso apagón en el Vaticano...
Society	El Comercio	La falsa noticia de la empleada que defeca en la mesa de su jefe ...

Tabla 7: True Instances Missclassified as Fake by all Systems.

and the cutting-edge methods implemented by the participants, help to place Spanish among the languages with an increasing number of resources and experiments oriented to fake news detection.

The use of transformers again achieved the best results this year showing the appropriateness of this method in a cross-topic scenario. Although this year's results are lower than last year's, they are really promising considering the added difficulty. The complimentariness analysis showed that hybrid approaches, using both transformers and traditional approaches (including traditional Deep Neural Networks) can help to increase the performance of automatic fake news detection methods.

Acknowledgments

References

Aragón, M., H. Jarquín-Vásquez, M. M. y Gómez, H. J. Escalante, L. V. Pineda, H. Gómez-Adorno, J. Posadas-Durán, and G. Bel-Enguix. 2020. Overview of MEX-A3T at IberLEF 2020: Fake news and aggressiveness analysis in mexican spanish. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2020) co-located with 36th Conference of the Spanish Society for Natural Language Processing (SEPLN 2020)*.

Chen, Y., N. Conroy, and V. Rubin. 2015. News in an online world: The need for an "automatic crap detector". *Proceedings of the Association for Information Science and Technology*, 52(1):1–4.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding.

Gamallo, P. 2021. CiTIUS at FakeDeS 2021: A hybrid strategy for fake news detection. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Guan, Z. 2021. TSIA team at FakeDeS 2021: Fake news detection in spanish using multi-model ensemble learning. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Huang, X., J. Xiong, and S. Jiang. 2021. GDUF_DM at FakeDeS 2021: Spanish fake news detection with BERT and sample memory. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Li, K. 2021. HAHA at FakeDeS 2021: A fake news detection method based on TF-IDF and ensemble machine learning. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Lomas-Barrie, V., N. Perez, V. M. Lara, and A. Neme. 2021. Bribones tras la esmeralda perdida@FakeDeS 2021: Fake news detection based on random forests, k-nearest neighbors, and n-grams for a spanish corpora. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Luo, H. 2021. YETI at FakeDeS 2021: Fake news detection in spanish with ALBERT. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Montes, M., P. Rosso, J. Gonzalo, E. Aragón, R. Agerri, M. Á. Álvarez-Carmona, E. Á. Mellado, J. C. de Albornoz, L. F. Luis Chiruzzo, H. Gómez-Adorno, Y. Gutiérrez, S. M. J. Zafra, S. Lima, F. M. P. de Arco, and M. Taulé, editors. 2021. *Proceedings of the Iberian Languages Evaluation Forum*, Iberlef 2021. CEUR Workshop Proceedings, september.

Posadas-Durán, J.-P., H. Gómez-Adorno, G. Sidorov, and J. J. M. Escobar. 2019. Detection of fake news in a new corpus for the spanish language. *Journal of Intelligent & Fuzzy Systems*, 36(5):4869–4876.

Reyes-Magaña, J. and L. E. Argota Vega. 2021. ForceNLP at FakeDeS 2021: Analysis of text features applied to fake news detection in spanish. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Spalenza, M. A., L. Lusquino-Filho, F. M. G. França, P. M. V. Lima, and E. de Oliveira. 2021. LCAD - UFES at FakeDeS 2021: Fake news detection using named entity recognition and part-of-speech sequences. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.

Tang, E. K., P. N. Suganthan, and X. Yao. 2006. An analysis of diversity measures. *Mach. Learn.*, 65(1):247–271.

Wolf, T., L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October. Association for Computational Linguistics.

Zhao, K., S. Zhou, and W. Li. 2021. zk15120170770 at FakeDeS 2021: Fake news detection based on pre-training model. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*.