

# SaRoCo: Detecting Satire in a Novel Romanian Corpus of News Articles

Ana-Cristina Rogoz, Mihaela Găman, Radu Tudor Ionescu\*

University of Bucharest

14 Academiei Street, Bucharest, Romania

\*raducu.ionescu@gmail.com

## Abstract

In this work, we introduce a corpus for satire detection in Romanian news. We gathered 55,608 public news articles from multiple real and satirical news sources, composing one of the largest corpora for satire detection regardless of language and the only one for the Romanian language. We provide an official split of the text samples, such that training news articles belong to different sources than test news articles, thus ensuring that models do not achieve high performance simply due to overfitting. We conduct experiments with two state-of-the-art deep neural models, resulting in a set of strong baselines for our novel corpus. Our results show that the machine-level accuracy for satire detection in Romanian is quite low (under 73% on the test set) compared to the human-level accuracy (87%), leaving enough room for improvement in future research.

## 1 Introduction

According to its definition in the Cambridge Dictionary, satire is “a humorous way of criticizing people or ideas”<sup>1</sup>. News satire employs this mechanism in the form of seemingly legitimate journalistic reporting, with the intention of ridiculing public figures, politics or contemporary events (McClenen and Maisel, 2014; Peters and Broersma, 2013; Rubin et al., 2016). Although the articles pertaining to this genre contain fictionalized stories, the intent is not to mislead the public into thinking that the discussed subjects are real. On the contrary, satirical news articles are supposed to reveal their nature by the writing style and comedic devices employed, such as irony, parody or exaggeration. Thus, the intention behind the writing differentiates satirical news (Rubin et al., 2016) from fake

news (Meel and Vishwakarma, 2019; Pérez-Rosas et al., 2018; Sharma et al., 2019). However, in some rare cases, the real intent might be deeply buried in the complex irony and subtleties of news satire (Barbieri et al., 2015a), which has the effect of fiction being deemed as factual (Zhang et al., 2020). Even so, there is a clear distinction between satirical and fake news. In fake news, the intent is to deceive the readers in thinking that the news is real, while presenting fake facts to influence the readers’ opinion. Since our study is focused on satire detection, we consider discussing research on fake news detection as being out of our scope. At the same time, we acknowledge the growing importance of detecting fake news and the fact that an accurate differentiation of satirical from legitimate journalistic reports might be seen as a starting point in controlling the spread of deceptive news (De Sarkar et al., 2018).

Satire detection is an important task that could be addressed prior to the development of conversational systems and robots that interact with humans. Certainly, the importance of understanding satirical (funny, ridiculous or ironical) text becomes obvious when we consider a scenario in which a robot performs a dangerous action because it takes a satirical comment of the user too literally. Given the relevance of the task for the natural language processing community, satire detection has already been investigated in several well-studied languages such as Arabic (Saadany et al., 2020), English (Burfoot and Baldwin, 2009; De Sarkar et al., 2018; Goldwasser and Zhang, 2016; Yang et al., 2017), French (Ionescu and Chifu, 2021; Liu et al., 2019), German (McHardy et al., 2019), Spanish (Barbieri et al., 2015b) and Turkish (Toçoğlu and Onan, 2019). Through the definition of satire, the satire detection task is tightly connected to irony and sarcasm detection. These tasks strengthen or broaden the language variety with languages such as Ara-

<sup>1</sup><https://dictionary.cambridge.org/dictionary/english/satire>

Data Set	Language	#articles		
		Regular	Satirical	Total
(Burfoot and Baldwin, 2009)	English	4,000	233	4,233
(Frain and Wubben, 2016)	English	1,705	1,706	3,411
(Goldwasser and Zhang, 2016)	English	10,921	1,225	12,146
(Ionescu and Chifu, 2021)	French	5,648	5,922	11,570
(Li et al., 2020)	English	6,000	4,000	10,000
(Liu et al., 2019)	French	2,841	2,841	5,682
(McHardy et al., 2019)	German	320,219	9,643	329,862
(Ravi and Ravi, 2017)	English	1,272	393	1,665
(Saadany et al., 2020)	Arabic	3,185	3,710	6,895
(Toçoğlu and Onan, 2019)	Turkish	1,000	1,000	2,000
(Yang et al., 2017)	English	168,780	16,249	185,029
SaRoCo (ours)	Romanian	27,980	27,628	55,608

Table 1: Number of regular and satirical news articles in existing corpora versus SaRoCo.

Set	Regular		Satirical		Total	
	#articles	#tokens	#articles	#tokens	#articles	#tokens
Training	18,000	8,174,820	17,949	11,147,169	35,949	19,321,989
Validation	4,986	2,707,026	4,878	3,030,055	9,864	5,737,081
Test	4,994	2,124,346	4,801	1,468,199	9,795	3,592,545
Total	27,980	13,006,192	27,628	15,645,423	55,608	28,651,615

Table 2: Number of samples (#articles) and number of tokens (#tokens) for each subset in SaRoCo.

bic (Karoui et al., 2017), Chinese (Jia et al., 2019), Dutch (Liebrecht et al., 2013) and Italian (Giudice, 2018).

In this work, we introduce SaRoCo<sup>2</sup>, the **Satire detection Romanian Corpus**, which comprises 55,608 news articles collected from various sources. To the best of our knowledge, this is the first and only data set for the study of Romanian satirical news. Furthermore, SaRoCo is also one of the largest data sets for satirical news detection, being surpassed only two corpora, one for English (Yang et al., 2017) and one for German (McHardy et al., 2019). However, our corpus contains the largest collection of satirical news articles (over 27,000). These facts are confirmed by the comparative statistics presented in Table 1.

Along with the novel data set, we include two strong deep learning methods to be used as baselines in future works. The first method is based on low-level features learned by a character-level convolutional neural network (Zhang et al., 2015), while the second method employs high-level semantic features learned by the Romanian version of BERT (Dumitrescu et al., 2020). The gap between the human-level performance and that of the deep learning baselines indicates that there is enough room for improvement left for future studies. We make our corpus and baselines available online for

<sup>2</sup><https://github.com/MihaelaGaman/SaRoCo>

Sample Part	Average #tokens
Title	24.97
Full Articles	515.24

Table 3: Average number of tokens in full news articles and titles from SaRoCo.

nonprofit educational and research purposes, under an open-source noncommercial license agreement.

## 2 Corpus

SaRoCo gathers both satirical and non-satirical news from some of the most popular Romanian news websites. The collected news samples were found in the public web domain, i.e. access is provided for free without requiring any subscription to the publication sources. The entire corpus consists of 55,608 samples (27,628 satirical samples and 27,980 non-satirical samples), having more than 28 million tokens in total, as illustrated in Table 2. Each sample is composed of a title (headline), a body and a corresponding label (satirical or non-satirical). As shown in Table 3, an article has around 515.24 tokens on average, with an average of 24.97 tokens for the headline. We underline that the labels are automatically determined, based on the fact that a publication source publishes either regular or satirical news, but not both.

We provide an official split for our corpus, such that all future studies will use the same training, val-

Category	Example	Translation
Regular	“Tragedie în zi de sărbătoare”	“Tragedy during celebration day”
	“Demisia lui \$NE\$ \$NE\$ se amână”	“\$NE\$ \$NE\$’s resignation is post-poned”
	“Premierul bulgar \$NE\$ \$NE\$ are \$NE\$”	“Bulgarian prime-minister \$NE\$ \$NE\$ has \$NE\$”
	“A murit actorul \$NE\$ \$NE\$”	“The actor \$NE\$ \$NE\$ died”
	“Metroul din \$NE\$ \$NE\$ se deschide azi”	“Subway to \$NE\$ \$NE\$ opens up today”
Satirical	“Comedia cu pălărioară de staniol”	“Comedy with little tin-foil hat”
	“10 restricții dure pe care \$NE\$ le pregătește pe ascuns”	“10 harsh restrictions that \$NE\$ is planning in secrecy”
	“Câți pokemoni ai prins azi?”	“How many pokemons did you catch today?”
	“Biserica \$NE\$ lansează apa sfințită cu aromă”	“The \$NE\$ Church launches flavored holy water”
	“Dragostea în vremea sclerozei”	“Love in the time of sclerosis”

Table 4: Examples of news headlines from SaRoCo.

idation and test sets, easing the direct comparison with prior results. Following [McHardy et al. \(2019\)](#), we use disjoint sources for training, validation and test, ensuring that models do not achieve high performance by learning author styles or topic biases particular to certain news websites. While crawling the public news articles, we selected the same topics (culture, economy, politics, social, sports, tech) and the same time frame (between 2011 and 2020) for all news sources to control for potential biases induced by uneven topic or time distributions across the satirical and non-satirical genres.

After crawling satirical and non-satirical news samples, our first aim was to prevent discrimination based on named entities. The satirical character of an article should be inferred from the language use rather than specific clues, such as named entities. For example, certain sources of news satire show preference towards mocking politicians from a specific political party, and an automated system might erroneously label a news article about a member of the respective party as satirical simply based on the presence of the named entity. Furthermore, we even noticed that some Romanian politicians have certain mocking nicknames assigned in satirical news. In order to eliminate named entities, we followed a similar approach as the one used for the

MOROCCO ([Butnaru and Ionescu, 2019](#)) data set. Thus, all the identified named entities are replaced with the special token \$NE\$. Besides eliminating named entities, we also substituted all whitespace characters with space and replaced multiple consecutive spaces with a single space. A set of processed satirical and regular headlines are shown in Table 4.

### 3 Baselines

**Fine-tuned Ro-BERT.** Our first baseline consists of a fine-tuned Romanian BERT ([Dumitrescu et al., 2020](#)), which follows the same transformer-based model architecture as the original BERT ([Devlin et al., 2019](#)). According to [Dumitrescu et al. \(2020\)](#), the Romanian BERT (Ro-BERT) attains better results than the multilingual BERT on a range of tasks. We therefore assume that the Romanian BERT should represent a stronger baseline for our Romanian corpus.

We use the Ro-BERT encoder to encode each text sequence into a list of token IDs. The tokens are further processed by the model, obtaining the corresponding 768-dimensional embeddings. At this point, we add a global average pooling layer to obtain a Continuous Bag-of-Words (CBOW) representation for each sequence of text, followed by a Softmax output layer with two neural units, each predicting the probability for one category, either non-satirical or satirical. To obtain the final class label for a text sample, we apply *argmax* on the two probabilities. We fine-tune the whole model for 10 epochs on mini-batches of 32 samples, using the Adam with decoupled weight decay (AdamW) optimizer ([Loshchilov and Hutter, 2019](#)), with a learning rate of  $10^{-7}$  and the default value for  $\epsilon$ .

**Character-level CNN.** The second baseline model considered in the experiments is a Convolutional Neural Network (CNN) that operates at the character level ([Zhang et al., 2015](#)). We set the input size to 1,000 characters. After the input layer, we add an embedding layer to encode each character into a vector of 128 components. The optimal architecture for the task at hand proved to be composed of three convolutional (conv) blocks, each having a conv layer with 64 filters applied at stride 1, followed by Scaled Exponential Linear Unit (SELU) activation. From the first block to the third block, the convolutional kernel sizes are 5, 3 and 1, respectively. Max-pooling with a filter size of 3 is applied after each conv layer. After each conv block, we insert a Squeeze-and-Excitation block

Method	Validation						Test					
	Acc.	Macro $F_1$	Satirical		Regular		Acc.	Macro $F_1$	Satirical		Regular	
			Prec.	Rec.	Prec.	Rec.			Prec.	Rec.	Prec.	Rec.
Ro-BERT	0.8241	0.8160	0.9260	0.6991	0.7633	0.9462	0.7300	0.7150	0.8750	0.5250	0.6700	0.9250
Char-CNN	0.7342	0.7475	0.8023	0.6138	0.6928	0.8520	0.6966	0.7109	0.7612	0.5551	0.6606	0.8326

Table 5: Validation and test results of the character-level CNN and the fine-tuned Ro-BERT applied on SaRoCo.

with the reduction ratio set to  $r = 64$ , following Butnaru and Ionescu (2019). To prevent overfitting, we use batch normalization and Alpha Dropout (Klambauer et al., 2017) with a dropout rate of 0.5. The final prediction layer is composed of two neural units, one for each class (i.e. legitimate and satirical), with Softmax activation. We use the Nesterov-accelerated Adaptive Moment Estimation (Nadam) optimizer (Dozat, 2016) with a learning rate of  $2 \cdot 10^{-4}$ , training the network for 50 epochs on mini-batches of 128 samples.

## 4 Experiments

**Evaluation.** We conducted binary classification experiments on SaRoCo, predicting if a given piece of text is either satirical or non-satirical. As evaluation metrics, we employ the precision and recall for each of the two classes. We also combine these scores through the macro  $F_1$  and micro  $F_1$  (accuracy) measures.

**Results.** In Table 5, we present the results of the two baselines on the SaRoCo validation and test sets. We observe that both models tend to have higher precision scores in detecting satire than in detecting regular news. The trade-off between precision and recall is skewed towards higher recall for the non-satirical news class. Since both models share the same behavior, we conjecture that the behavior is rather caused by the particularities of the satire detection task.

**Discriminative feature analysis.** We analyze the discriminative features learned by the character-level CNN, which is one of the proposed baseline systems for satire detection. We opted for the character-level CNN in favor of the fine-tuned BERT, as the former method allows us to visualize discriminative features using Grad-CAM (Selvaraju et al., 2017), a technique that was initially used to explain decisions of CNNs applied on images. We adapted this technique for the character-level CNN, then extracted and analyzed the most predictive patterns in SaRoCo. The motivation behind this was to validate that the network’s decisions are not based on some biases that escaped

Category	Example	Translation
Slang	“ <i>cel mai marfă serial din lume</i> ” “ <i>cocalar</i> ”	“ <i>the dopest TV show in the world</i> ” “ <i>douche</i> ”
Insult	“ <i>odiosul primar</i> ” “ <i>bunicuț retardat</i> ” “ <i>dugongul ăla slinos de la sectorul 4</i> ”	“ <i>the odious mayor</i> ” “ <i>retarded grandpa</i> ” “ <i>that slender dugong in the 4th sector</i> ”
Repetition	“ <i>Mii de gunoaie care lasă gunoaie au remarcat că [...] plajele [...] s-au umplut de gunoaie, lăsate [...] de gunoaiele care au venit înaintea lor</i> ”	“ <i>Thousands of scums who leave garbage noticed that [...] beaches [...] got full of garbage, left behind [...] by the scums who were there before them</i> ”
Exaggeration Exclamation Irony	“ <i>Ne-am săturat!</i> ” “ <i>Ruşine să le fie!</i> ” “ <i>Chiar nu suntem o naţie de hoji!</i> ”	“ <i>We’re sick of it!</i> ” “ <i>Shame on them!</i> ” “ <i>We’re totally not a nation of thieves!</i> ”
Popular Saying	“ <i>a sărit calul</i> ” “ <i>a făcut-o de oaie</i> ” “ <i>minte de găină</i> ”	“ <i>went overboard</i> ” “ <i>messed up</i> ” “ <i>bird brain</i> ”

Table 6: Examples of predictive patterns of satire learned by the character-level CNN.

Category	Example	Translation
Stats	“ <i>Importurile au scăzut cu 2.1% [...] pentru o creștere de 0.1% și prelungirea scăderii de 1.4% din iulie.</i> ”	“ <i>Imports decreased by 2.1% [...] for an increase of 0.1% and the prolongation of the decrease of 1.4% since July.</i> ”
Legal terms	“ <i>asasinat</i> ” “ <i>l-au denunțat pe autorul atacului</i> ”	“ <i>assassinated</i> ” “ <i>denounced the perpetrator</i> ”
Weather	“ <i>temperatura în timpul nopții a scăzut</i> ”	“ <i>the temperature has dropped during the night</i> ”
Political terms	“ <i>scrutinul prezidențial</i> ” “ <i>prefectura informează că</i> ”	“ <i>presidential election</i> ” “ <i>prefecture informs that</i> ”

Table 7: Examples of predictive patterns of legitimate news learned by the character-level CNN.

our data collection and cleaning process.

In Tables 6 and 7, we present a few examples of interesting patterns considered relevant for predicting satire versus regular news, respectively. A broad range of constructions covering a great variety of styles and significant words are underlined

Method	Acc.	Macro $F_1$	Satirical		Regular	
			Prec.	Rec.	Prec.	Rec.
Ro-BERT	0.6800	0.6750	0.7800	0.5100	0.6350	0.8550
Char-CNN	0.6500	0.6510	0.6389	0.6900	0.6630	0.6100
Humans	0.8735	0.8711	0.9416	0.7970	0.8332	0.9500

Table 8: Averaged performance of ten human annotators versus deep learning baselines on 200 news headlines from SaRoCo.

via Grad-CAM in the satirical news samples. The network seems to pick up obvious clues such as slang, insults and popular sayings rather than more subtle indicatives of satire, including irony or exaggeration. At the same time, for the real news in SaRoCo, there are fewer categories of predictive patterns. In general, the CNN deems formal, standard news expressions as relevant for regular news. These patterns vary across topics and domains. The CNN also finds that the presence of numbers and statistical clues is indicative for non-satirical content, which is consistent with the observations of Yang et al. (2017). Our analysis reveals that the discriminative features are appropriate for satire detection, showing that our corpus is indeed suitable for the considered task.

**Deep models versus humans.** Given 100 satirical and 100 non-satirical news headlines (titles) randomly sampled from the SaRoCo test set, we asked ten Romanian human annotators to label each sample as satirical or non-satirical. We evaluated the deep learning methods on the same subset of 200 samples, reporting the results in Table 8. First, we observe that humans have a similar bias as the deep learning models. Indeed, for both humans and models, the trade-off between precision and recall is skewed towards higher precision for the satirical class and higher recall for the non-satirical class. We believe this is linked to the way people and machines make a decision. Humans look for patterns of satire in order to label a sample as satire. If a satire-specific pattern is not identified, the respective sample is labeled as regular, increasing the recall for the non-satirical class. Although humans and machine seem to share the same way of thinking, there is a considerable performance gap in satire detection between humans and machines. Indeed, the average accuracy of our ten human annotators is around 87%, while the state-of-the-art deep learning models do not surpass 68% on the same news headlines. Even on full news articles (see Table 5), the models barely reach an accuracy of 73% on the test set. Hence, we conclude there is

a significant performance gap between humans and machines, leaving enough room for exploration in future work on Romanian satire detection.

We would like to emphasize that our human evaluation was performed by casual news readers, and the samples were shown after named entity removal, thus having a fair comparison with the AI models. We underline that named entity removal makes the task more challenging, even for humans.

## 5 Conclusion

In this work, we presented SaRoCo, a novel data set containing satirical and non-satirical news samples. To the best of our knowledge, SaRoCo is the only corpus for Romanian satire detection and one of the largest corpora regardless of language. We trained two state-of-the-art neural models as baselines for future research on our novel corpus. We also compared the performance of the neural models with the averaged performance of ten human annotators, showing that the neural models lag far behind the human-level performance. Our discriminative feature analysis confirms the limitations of state-of-the-art neural models in detecting satire. Although we selected a set of strong models from the recent literature as baselines for SaRoCo, significant future research is necessary to close the gap with respect to the human-level satire detection performance. Designing models to pick up irony or exaggerations could pave the way towards closing this gap in future work.

## Acknowledgments

The authors thank reviewers for their useful remarks. This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI – UEFISCDI, project number PN-III-P1-1.1-TE-2019-0235, within PNCDI III. This article has also benefited from the support of the Romanian Young Academy, which is funded by Stiftung Mercator and the Alexander von Humboldt Foundation for the period 2020-2022.

## References

- Francesco Barbieri, Francesco Ronzano, and Horacio Saggion. 2015a. Do we criticise (and laugh) in the same way? Automatic detection of multi-lingual satirical news in Twitter. In *Proceedings of IJCAI*, pages 1215–1221.
- Francesco Barbieri, Francesco Ronzano, and Horacio Saggion. 2015b. Is this tweet satirical? A computational approach for satire detection in Spanish. *Procesamiento de Lenguaje Natural*, 55:135–142.
- Clint Burfoot and Timothy Baldwin. 2009. Automatic Satire Detection: Are You Having a Laugh? In *Proceedings of ACL-IJCNLP*, pages 161–164.
- Andrei Butnaru and Radu Tudor Ionescu. 2019. MO-ROCO: The Moldavian and Romanian Dialectal Corpus. In *Proceedings of ACL*, pages 688–698.
- Sohan De Sarkar, Fan Yang, and Arjun Mukherjee. 2018. Attending Sentences to detect Satirical Fake News. In *Proceedings of COLING*, pages 3371–3380.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*, pages 4171–4186.
- Timothy Dozat. 2016. Incorporating Nesterov Momentum into Adam. In *Proceedings of ICLR Workshops*.
- Ștefan Daniel Dumitrescu, Andrei-Marius Avram, and Sampo Pyysalo. 2020. The birth of Romanian BERT. In *Findings of EMNLP*, pages 4324–4328.
- Alice Frain and Sander Wubben. 2016. SatiricLR: a Language Resource of Satirical News Articles. In *Proceedings of LREC*, pages 4137–4140.
- Valentino Giudice. 2018. Aspie96 at IronITA (EVALITA 2018): Irony Detection in Italian Tweets with Character-Level Convolutional RNN. In *Proceedings of EVALITA*, pages 160–165.
- Dan Goldwasser and Xiao Zhang. 2016. Understanding Satirical Articles Using Common-Sense. *Transactions of the Association for Computational Linguistics*, 4:537–549.
- Radu Tudor Ionescu and Adrian Gabriel Chifu. 2021. Fresada: A french satire data set for cross-domain satire detection. *arXiv preprint arXiv:2104.04828*.
- Xiuyi Jia, Zhao Deng, Fan Min, and Dun Liu. 2019. Three-way decisions based feature fusion for Chinese irony detection. *International Journal of Approximate Reasoning*, 113:324–335.
- Jihen Karoui, Farah Zitoune, and Véronique Moriceau. 2017. SOUKHRIA: Towards an Irony Detection System for Arabic in Social Media. In *Proceedings of ACLing*, volume 117, pages 161–168.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-Normalizing Neural Networks. In *Proceedings of NIPS*, pages 972–981.
- Lily Li, Or Levi, Pedram Hosseini, and David Broniatowski. 2020. A Multi-Modal Method for Satire Detection using Textual and Visual Cues. In *Proceedings of NLP4IF*, pages 33–38.
- Christine Liebrecht, Florian Kunneman, and Antal van den Bosch. 2013. The perfect solution for detecting sarcasm in tweets #not. In *Proceedings of WASSA*, pages 29–37.
- Zhan Liu, Shaban Shabani, Nicole Glassey Balet, and Maria Sokhn. 2019. Detection of Satiric News on Social Media: Analysis of the Phenomenon with a French Dataset. In *Proceedings of ICCCN*, pages 1–6.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Proceedings of ICLR*.
- Sophia A. McClennen and Remy M. Maisel. 2014. I’m Not Laughing at You, I’m Laughing With You: How to Stop Worrying and Love the Laughter. In *Is Satire Saving Our Nation? Mockery and American Politics*, pages 189–201. Springer.
- Robert McHardy, Heike Adel, and Roman Klinger. 2019. Adversarial Training for Satire Detection: Controlling for Confounding Variables. In *Proceedings of NAACL*, pages 660–665.
- Priyanka Meel and Dinesh Kumar Vishwakarma. 2019. Fake news, rumor, information pollution in social media and web: A contemporary survey of state-of-the-arts, challenges and opportunities. *Expert Systems with Applications*, 153:112986.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. Automatic Detection of Fake News. In *Proceedings of COLING*, pages 3391–3401.
- Chris Peters and Marcel Jeroen Broersma. 2013. *Rethinking Journalism: Trust and Participation in a Transformed News Landscape*. Routledge.
- Kumar Ravi and Vadlamani Ravi. 2017. A novel automatic satire and irony detection using ensembled feature selection and data mining. *Knowledge-Based Systems*, 120:15–33.
- Victoria Rubin, Niall Conroy, Yimin Chen, and Sarah Cornwell. 2016. Fake News or Truth? Using Satirical Cues to Detect Potentially Misleading News. In *Proceedings of CADD*, pages 7–17.
- Hadeel Saadany, Constantin Orasan, and Emad Mohamed. 2020. Fake or Real? A Study of Arabic Satirical Fake News. In *Proceedings of RDSM*, pages 70–80.

- Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. [Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization](#). In *Proceedings of ICCV*, pages 618–626.
- Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. [Combating fake news: A survey on identification and mitigation techniques](#). *ACM Transactions on Intelligent Systems and Technology*, 10(3):1–42.
- Mansur Alp Toçoğlu and Aytuğ Onan. 2019. [Satire Detection in Turkish News Articles: A Machine Learning Approach](#). In *Proceedings of Innovate-Data*, pages 107–117.
- Fan Yang, Arjun Mukherjee, and Eduard Dragut. 2017. [Satirical News Detection and Analysis using Attention Mechanism and Linguistic Features](#). In *Proceedings EMNLP*, pages 1979–1989.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level Convolutional Networks for Text Classification](#). In *Proceedings of NIPS*, pages 649–657.
- Yigeng Zhang, Fan Yang, Eduard Constantin Dragut, and Arjun Mukherjee. 2020. [Birds of a Feather Flock Together: Satirical News Detection via Language Model Differentiation](#). *arXiv preprint arXiv:2007.02164*.