

# Make Satire Boring Again: Reducing Stylistic Bias of Satirical Corpus by Utilizing Generative LLMs

Asli Umay Ozturk<sup>1</sup>, Recep Firat Cekinel<sup>1</sup>, Pinar Karagoz<sup>1</sup>

<sup>1</sup> Department of Computer Engineering  
Middle East Technical University (METU)  
{auozturk,rfcekinel,karagoz}@ceng.metu.edu.tr

Correspondence: auozturk@ceng.metu.edu.tr

## Abstract

Satire detection is essential for accurately extracting opinions from textual data and combating misinformation online. However, the lack of diverse corpora for satire leads to the problem of stylistic bias which impacts the models' detection performances. This study proposes a debiasing approach for satire detection, focusing on reducing biases in training data by utilizing generative large language models. The approach is evaluated in both cross-domain (irony detection) and cross-lingual (English) settings. Results show that the debiasing method enhances the robustness and generalizability of the models for satire and irony detection tasks in Turkish and English. However, its impact on causal language models, such as Llama-3.1, is limited. Additionally, this work curates and presents the Turkish Satirical News Dataset with detailed human annotations, with case studies on classification, debiasing, and explainability.

## 1 Introduction

There are no universally agreed definitions for satire, sarcasm, and irony in NLP literature. However, Cambridge Dictionary<sup>1</sup> defines **satire** as a way of *criticizing* people or ideas in a *humorous* way, especially in order to make a political point. Whereas **sarcasm** is defined as the use of remarks that clearly *mean the opposite* of what they say, made in order to hurt someone's feelings or to *criticize* something in a *humorous* way. Moreover, **irony** is related to the use of words that are the *opposite of what you mean*, as a way of being *funny*. The overlaps in these definitions cause different studies to use these terms interchangeably, and borrow ideas from other studies (Barbieri et al., 2014; Van Hee et al., 2016a, 2018; Carvalho et al., 2020).

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

At first glance, it might not be very apparent why detecting satirical content is important. However, with the increased usage of social media, the primary source of news and information for many people has become the shared news articles in their social media feeds. Even though this makes the information more accessible, it can also cause misinformation to spread at fast rates (Allcott and Gentzkow, 2017; Aïmeur et al., 2023). It is not uncommon for regular social media users to take fake or satirical content as the truth (Wu et al., 2019), which is specifically problematic when it comes to news content. Hence, satire detection can offer a solution to this misinformation problem since automated detection of satirical content can be used to create automated warnings that inform social media users about the reliability of a piece of information.

For the last couple of years, with the rise of LLMs, an improvement in the performance of NLP tasks has been seen in the literature (Li et al., 2022). Even though recent studies have demonstrated high performances for satire, sarcasm, and irony classification tasks using multilingual LLMs, it remains unclear whether these concepts are represented similarly across different languages and domains (Ortega-Bueno et al., 2023; Maladry et al., 2023).

Moreover, focusing on a data-centric approach when training LLMs also raises problems since it is not easy to find a diverse set of resources for low-resourced languages (Doğru et al., 2018; Hangya et al., 2022; Acikgoz et al., 2024) or specific tasks. For example, annotating data that can be labeled as satirical, sarcastic, or ironic requires extensive human labor. Instead, automatic data collection processes may be employed, and the data would be collected from a limited number of sources that are already known to belong to the target label. As a result, this creates stylistically unbalanced corpora to be used to train and fine-tune LLMs. This may result in a bias or misalignment in the model (Xu et al., 2024). In other words, potential stylistic bias

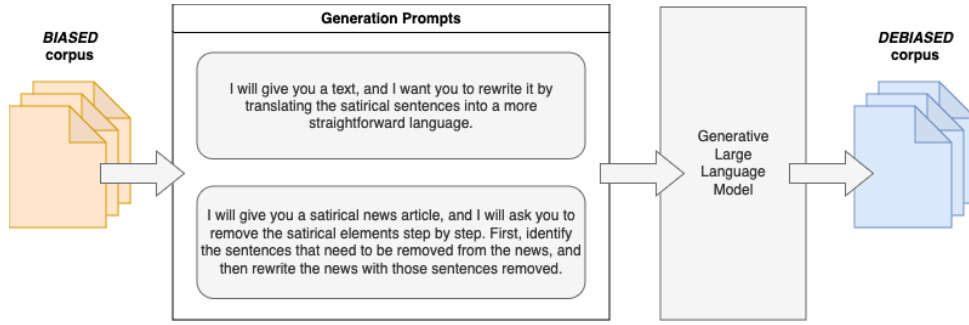


Figure 1: The proposed debiasing pipeline

in the curated dataset impacts the robustness of the models that use these datasets during training and fine-tuning.

This work aims to reduce the effect of stylistic bias stemming from a single-sourced satirical corpus, proposing a debiasing method that utilizes generative LLMs to reduce the stylistic bias of the instances in the biased corpus.

The proposed method works by generating satirical texts that are stylistically more natural, making the generated corpus more parallel to the non-satirical corpus. This method is demonstrated on a curated dataset for satirical news detection in Turkish. Our contributions can be summarized as follows:

- Curating the *Turkish Satirical News Dataset* with human annotations, and analyzing its stylistic bias and usability<sup>2</sup>,
- Proposing a *debiasing pipeline* (Figure 1) for combating stylistic bias and improving model generalizability,
- Analyzing the cross-lingual and cross-domain performance of Turkish satire detection model for irony and English.

## 2 Related Work

Starting in the early 2010s, the related literature began to focus on the problems of binary detection for satire, sarcasm, and irony. Earlier works generally utilize traditional supervised learning methods such as Support Vector Machine (SVM) or Naive Bayes (NB) based classifiers and propose different feature extraction methods for different languages and tasks (Buschmeier et al., 2014; Barbieri et al., 2014; Van Hee et al., 2016b; Pamungkas and Patti, 2018; Baloglu et al., 2019; Ozturk et al., 2021; Onan and Toçoğlu, 2020). Another approach

explored in the earlier studies is utilizing neural network based architectures such as LSTM (Long-Short Term Memory) networks (Wu et al., 2018; Zhang et al., 2019a). Later works started to utilize transformer architectures such as BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2019a) and reported improved results (Buyukbas et al., 2021).

In the last couple of years, newer studies have focused more and more on multimodal approaches and LLM-based models. In their study, Tomás et al. (2023) explore the irony detection performance of transformer-based models with both textual and visual inputs. On the other hand, Lin et al. (2024) combines transformer-based models with prompt engineering to improve the irony detection performance, specifically focusing on different features of the text.

A recurring problem in the literature for the aforementioned tasks is the lack of labeled data and openly available datasets. There are curated datasets for irony detection (Van Hee et al., 2016a) and sarcastic news detection (Barbieri et al., 2014) but they are mostly in English. Even though some other datasets curated for other languages do exist (Ortega-Bueno et al., 2019; Xiang et al., 2020; Ozturk et al., 2021; Ghanem et al., 2020; Joshi et al., 2017), they are much smaller than the available English corpus for irony, sarcasm, and satire.

Literature on irony, satire, and sarcasm detection also includes studies utilizing explainable artificial intelligence (AI) and interpretable machine learning (ML) methods (Buyukbas et al., 2021) such as LIME (Local Interpretable Model-Agnostic Explanations) (Ribeiro et al., 2016) and SHAP (SHapley Additive exPlanations) (Lundberg and Lee, 2017) or proposing task-based structures to understand model decisions (Shu et al., 2019). These studies aim to analyze the model decisions to improve the performance, fairness, and generalizability of the

<sup>2</sup><https://github.com/auotomaton/satiretr>

model, as well as reduce the bias of the model.

With the advancements in generative LLMs, there have been studies utilizing them to augment training data or synthetically generate data from scratch, to overcome the shortness of training data available for a diverse set of tasks (Hangya et al., 2022; Long et al., 2024).

A couple of works have tried to utilize this data generation approach to overcome biases in datasets. Qian et al. (2021) focus on dataset bias and propose a framework for debiasing using counterfactual inference. They show that their approach improves the effectiveness, generalisability, and fairness of the classifier. In another study, Schlicht et al. (2024) utilize conversational LLMs to reduce textual bias in news articles. Their findings show that even though they are compelling in some cases, they tend to leave out vital and contextual information during the debiasing process.

Differing from the existing works, this work explores the effects of synthetically generated data on debiasing binary classifiers trained on low-resource languages, focusing on satire detection in Turkish. We propose a debiasing pipeline that aims to improve the cross-lingual and cross-domain performance of a model trained on a stylistically biased dataset.

### 3 Dataset

One contribution of this study is to curate an open dataset for satire and satirical news detection in Turkish. Utilizing the satirical news publication Zaytung<sup>3</sup> and Turkish news agency Anadolu Agency (AA)<sup>4</sup>, *Turkish Satirical News Dataset* is curated.

#### 3.1 Curation of Turkish Satirical News Dataset

As a source of satirical news articles, the Turkish satirical news publication Zaytung is used. By crawling the Zaytung website archive, 2825 satirical articles are collected with *timestamp*, *title*, *body* and *header image* information. To improve the representative nature of the dataset, articles dated before 2014 are discarded.

As a source of non-satirical news articles, AA archives are crawled between the dates 2022-2023, and 4781 articles are collected with *id*, *title*, *sub-title*, *author*, *date*, *category*, *header image*, *city*, *body* and *tag* information.

<sup>3</sup><https://zaytung.com/>

<sup>4</sup><https://www.aa.com.tr/>

The final dataset includes 2202 SATIRICAL and 4781 NON-SATIRICAL articles. Additionally, 40 of the SATIRICAL instances have word-by-word human annotations aiming to capture the effect of each word on the satirical meaning of the article. The annotation process and a case study comparing the annotations with model explanations are presented in Appendix A and B.

The curated dataset is available publicly on GitHub<sup>5</sup> for other researchers’ use. The GitHub repository also includes the debiased articles generated from the Zaytung corpus, described in Section 4. All data is scraped from publically available sources and Zaytung has specifically been reached for consent to use their articles.

#### 3.2 Bias Analysis

Since both the SATIRICAL and the NON-SATIRICAL corpora are taken from a single publication each, it might be expected to see a stylistic and statistical bias that results in an easily separable dataset. This may result in classifiers that are trained using such datasets to become biased for the style elements of the corpus instead of identifying satire/non-satire indications of the text. One simple way to observe this is by conducting statistical analysis on both the SATIRICAL and NON-SATIRICAL corpora.

##### 3.2.1 Average Word and Sentence Count

A primary statistical analysis is performed and reported in Table 1 to better understand the data instances. It can be seen that SATIRICAL corpus has an average of 329 words per instance and 44 sentences per instance. On the other hand, NON-SATIRICAL corpus has an average of 313 words per instance and 43 sentences per instance. Even though the numbers are close, on average, we see that the SATIRICAL corpus has more words per sentence.

Statistic	SATIRICAL	NON-SATIRICAL
avg #of words	329	313
avg #of sentences	44	43

Table 1: Statistics of the *Turkish Satirical News Dataset* by corpus

##### 3.2.2 Top 10 Words

To have a general idea about the content of the news belonging to both the SATIRICAL and NON-SATIRICAL corpora, top-10 terms are extracted per label by TF-IDF (Term Frequency - In-

<sup>5</sup><https://github.com/automaton/satiretr>

Label	Top 10 Words
SATIRICAL	<i>almak (take), bir (one/a), demek (say), etmek (make), gelmek (come), iş (work/job), olarak (being), vermek (give), türkiye (Turkiye), yapmak (do)</i>
NON-SATIRICAL	<i>ülke (country), yıl (year), açıklama (explanation), ifade (expression), fotoğraf (photograph), spor (sport), bölge (region), başkan (president), konu (issue), çalışmak (work)</i>

Table 2: Statistics of the *Turkish Satirical News Dataset* by Top 10 Words

verse Document Frequency) scoring. These terms are shown in Table 2. It is visible that the top 10 words for the SATIRICAL and NON-SATIRICAL corpora do not have any words in common. This also follows the idea that the tones of the two corpora are different.

## 4 The Proposed Debiasing Method

Bias introduced from the style of the source of the corpus is a serious concern that is hard to eliminate without extensive human annotation. This is specifically prevalent in fairly low-resource languages (Shen et al., 2024) and text qualities such as satirical, sarcastic, or ironic meaning (Maladry et al., 2023; Ortega-Bueno et al., 2023). Using a single source as a corpus makes the models more prone to bias. Hence, this study proposes the *debiasing pipeline* utilizing LLMs for generating less biased counterparts of texts.

### 4.1 Pipeline Design

The *debiasing pipeline* proposed in this work utilizes prompt engineering and synthetic data generation to remove the effect of the bias coming from the heavily stylistic language of the SATIRICAL corpus. The proposed pipeline is created to improve the usability of the curated dataset, however, we believe the pipeline can be generalized for any biased dataset by adapting the generation prompts according to the task and the bias.

The *debiasing pipeline*, summarised in Figure 1, generates a new set of data to be used to replace the training instances from the biased corpus. In the scope of this work, the stylistically biased data in the train set is the SATIRICAL class coming from the scraped Zaytung corpus.

### 4.2 Generating "Debiased" Articles Using Prompt Engineering

This subsection goes over the prompt engineering and debiased instance generation using sample articles translated so that it will be easier for readers to follow. The original Turkish articles and Turkish prompts are available in Appendix C. In this study, the data generation is done via the GPT web interface, ChatGPT<sup>6</sup>. However, the generation pipeline can easily be automatized by using the OpenAI API to fully automate the debiasing process. The following two prompts are used to generate stylistically less biased articles:

**Prompt 1:** *"I will give you a satirical news article, and I will ask you to remove the satirical elements step by step. First, identify the sentences that need to be removed from the news, and then rewrite the news with those sentences removed.*

*Article text:"*

**Prompt 2:** *"I will give you a text, and I want you to rewrite it by translating the satirical sentences into a more straightforward language.*

*Article text:"*

The first prompt, Prompt 1, asks the generative LLM to identify the satirical elements in the text and rewrite the article by excluding them. On the other hand, Prompt 2 asks the generative LLM to directly rewrite the article by modifying the sentences with satirical meaning to have a more straightforward language. Both prompts put the generative LLM to test in terms of its language understanding.

First, consider Sample Article 1 and the article generated from it using Prompt 1, shown in Figure 2. Even though the prompt clearly asks for the removal of satirical elements, the satire created by the fake social media expectation narrative is still present in the generated article. However, we see that the style of the writing is less exaggerated. This makes Prompt 1 a good candidate for generation in some cases.

On the other hand, consider Sample Article 2 and the articles generated from it using both Prompt 1 and Prompt 2, shown in Figure 3. It can be seen that Prompt 1 removes most of the sentences contributing to the satire in the text which significantly reduces the satirical meaning of the overall article. Without background knowledge about the financial situation revolving around Turkish Lira (TL), the article becomes vaguely satirical or even non-satirical to the uninformed reader. However, the

<sup>6</sup><https://chatgpt.com/>



#### NASA Announces Termination of Mars Program Due to Photos Not Reaching Expected Like Counts...

The American Space Agency NASA is facing troubled days due to the photos sent by the Perseverance rover following its landing on Mars. NASA's one-ton Rover-type spacecraft, Perseverance, landed successfully on Mars' Jezero Crater at 15:55 Eastern Time on Thursday after a journey of approximately 7 months. However, 24 hours after the landing, when NASA officials shared the first photos from the red planet on their Instagram account, they expressed surprise and disappointment over the comments they received. The photos, which were expected to garner at least 10 million likes, only managed to accumulate 2 million, leading to significant disillusionment within the community. Comments such as "Did you send an iron to Mars to take these photos?", "You'd think they'd put a decent camera on it", and "There was no need to go all the way to Mars for this, we could have sorted it out in Yozgat city center" were particularly hurtful, according to James Watzin, Director of NASA's Mars Program. Watzin admitted, "What's truly upsetting is that our followers are absolutely right. Honestly, we weren't satisfied either. If it wasn't our own rover, we wouldn't have liked those photos. It's just not appealing. . ."

Watzin pointed out that the photos failed to receive as many likes as a makeup-free selfie by Cardi B, noting, "If we don't receive something like a group of Martian mating photos in the next few days, it seems impossible to reach the number of likes needed to cover the program's costs. The effort we put in seems to have missed its mark. Why did this happen? Did we choose the wrong cameras? Is Mars itself not photogenic? These are questions that need answering." He announced that the Mars program will be put on hold until the source of the issue is understood.

(a) The original article, translated to English

The American Space and Aeronautics Agency NASA is going through difficult days due to the photos sent after the Perseverance rover successfully landed on Mars. Contrary to expectations, NASA officials encountered low engagement on social media following the successful landing of the Mars program. Program director James Watzin admitted that the lack of expected interest has led to disappointment and that the program has been deemed a failure. The Mars program has been paused to understand the source of the problem.

(b) Generated article, translated to English

Figure 2: Debiasing of Sample Article (1) with Prompt 1

#### The Crisis in LUNA is Growing: It Has Been Revealed to Have Lost Value Even Against the Turkish Lira...

Following a 99.8% loss in value, where Terra Luna Coin dropped from \$64 to \$0.20 within two days, bad news continues to emerge. After this dramatic drop that shook the cryptocurrency market, Do Kwon, CEO of Terraform Labs, released a statement today. He noted that their investigation revealed Luna had lost value not only against the US dollar but also against the Turkish Lira. Kwon said, "Everything we thought couldn't happen in the last two days has happened. I am very sorry."

Kwon explained that the realization came during a meeting he held to boost the team's morale, saying, "I gathered the team to give a motivational speech along the lines of 'Look, it's not so bad. At least we are still gaining value in TL terms.' I hadn't felt the need to check if we were worse off compared to TL. That was the real mistake." He shared that the painful truth became apparent when they opened the graph on the screen.

"At that moment, half the team left the building. We still can't reach them. The remaining team members are just looking around the office to see if there's anything left to sell," Kwon described the current situation at Terra Labs. The experienced CEO also mentioned that they expect the depreciation of the TL to eventually surpass Luna, whether due to some absurd move by Turkey or the Central Bank of the Republic of Turkey depleting its remaining foreign currency reserves. He added that he has not neglected to give optimistic messages for the future.

(a) The original satirical article, translated to English

Terra Luna Coin has experienced a significant loss in value in recent days. Terraform Labs CEO Do Kwon stated that their investigation revealed that Luna has lost value not only against the US dollar but also against the Turkish Lira (TL). This situation has caused a morale decline among company employees. Kwon expressed that in the short term, they expect the depreciation of the TL to surpass that of Luna.

(b) Generated article with Prompt 1, translated to English

In the past two days, Terra Luna Coin has experienced a 99.8% loss in value, dropping from \$64 to \$0.2, and bad news continues to emerge. Following this major drop, Terraform Labs CEO Do Kwon stated in a recent announcement that they have found evidence that Luna has lost value not only against the US dollar but also against the Turkish Lira. Kwon said, "Everything we didn't expect to happen in these last two days happened. I'm very sorry."

Kwon mentioned that they realized the situation during a meeting he organized to boost the team's morale. "I wanted to tell the team, 'At least we're gaining value against the TL,' but I hadn't checked the situation. That was the real mistake," he said. He also mentioned that they faced the harsh reality when they opened the chart.

"During the meeting, half of the team left the building, and we still can't reach them. The rest are just looking for things to sell in the office," Kwon said. He also noted that in the short term, the loss of value of the TL might surpass that of Luna, and he did not forget to provide an optimistic message for the future.

(c) Generated article with Prompt 2, translated to English

Figure 3: Debiasing of Sample Article (2) with Prompt 1 and Prompt 2

generated article with Prompt 2 can still be classified as satirical easily, while stripped from some of the more dramatic and stylistically strong phrases (e.g. "*...absurd move by Turkey or the Central Bank of the Republic of Turkey...*" etc.). This makes Prompt 2 also a good candidate for generation.

Another successful example that shows an original-generated article pair generated using Prompt 2 can be seen in Appendix C. With these explorations done on the prompts and their performances, both prompts are used during the generation process of 200 debiased articles.

### 4.3 Comparing Generated and Original Articles

The proposed debiasing method is expected to keep the satirical value of the articles while making them

stylistically less biased. To check if this is the case, two analyses are conducted.

Firstly, the generated articles are checked manually (by one annotator) to see if they still can be classified as SATIRICAL. Additionally, they are analyzed to see if the original context of the article is still understandable or if some of the context (such as the events, relationships between people, and such) has been lost during the generation process. It is seen that:

- Out of the 200 articles, 29 of them can be labeled as NON-SATIRICAL by an unsuspecting reader with not enough knowledge of the Turkish political landscape.<sup>7</sup>

<sup>7</sup>Authors are aware that this definition of such a reader is highly subjective. This statistic is obtained to have a surface-level understanding of the quality of the generated content.

- Out of the aforementioned 29 articles, all 29 of them differ from their original counterparts with a loss of important contextual information, such as the total erasure of events or people occurring during the generation process.
- Out of the aforementioned 29 articles, 28 of them are generated with Prompt 1. Since Prompt 1 explicitly asks the model to remove satirical sentences, it is understandable to see a major loss in contextual information. This also points to Prompt 2 as being a better option for article generation.

As a second way of verifying whether the content of the articles in the corpus is maintained, the BERTScore (Zhang et al., 2019b) evaluation metric, which calculates the pairwise semantic similarity of tokens in the given pair of sentences using BERT’s contextual embeddings, is employed. We used the BERTurk model for extracting contextual embeddings and the cosine similarity of the original sarcastic news and the debiased counterparts is 0.6852 (in terms of F1-binary). This similarity score implies that the debiased text mostly retains the original content.

## 5 Experiments

The experiments were conducted on 4 Nvidia A6000 GPUs. We employed three training strategies to evaluate the proposed debiasing pipeline’s effectiveness.

### 5.1 Models and Parameters

We employed multilingual BERT (Devlin et al., 2019b), BERTurk (Schweter, 2020), XLM-Roberta-large (Conneau et al., 2020) and Llama-3.1-8B-Instruct (Dubey et al., 2024) models for conducting the experiments. These models were mainly selected because they were also pre-trained in Turkish. Llama-3.1-8B is a generative causal language model and we used sampling-based decoding to make predictions. As a result, for some test cases, the model’s inferences do not comply with the expected labels and we reported the rate of nonresponses in Table 3. The samples where the model failed to return any labels were excluded during the evaluation of the models on the test sets.

During the training, we randomly selected 10% of the training data for validation. A grid search was conducted to explore the following hyperparameters: *learning rate* 2e-5, 5e-5, 1e-4 and *batch*

Dataset	COMBINED	BIASED	DEBIASED
Zaytung	0.009	0.002	0.000
Onion	0.013	0.207	0.030
IronyTR	0.057	0.012	0.002

Table 3: Rate of nonresponses of Llama-3.1-8B-Instruct model

*size* 8, 16, 32. The best results, as determined through this process, are reported in Section 5.3. All models were trained on the training datasets for two epochs.

For the Llama-3.1 model, we set the sequence length to 2048, while for other models, the maximum sequence length was set to 512. Additionally, we employed QLoRA (Dettmers et al., 2024) for training the Llama model, with the LoRA rank and LoRA alpha set to 32 and the LoRA dropout to 0.05. We used the Adam optimizer with a cosine scheduler and trained the models with fp16 precision.

### 5.2 Training Setups and Test Datasets

The models were trained on three different setups as follows (illustrated in Figure 4):

**BIASED:** 200 instances from the SATIRICAL corpus and 200 instances from the NON-SATIRICAL corpus in *Turkish Satirical News Dataset* (detailed in Section 3) are selected.

**DEBIASED:** Same sets of instances as the *BIASED* setup are selected first, then the selected SATIRICAL instances are passed through the proposed *debiasing pipeline*. Final *DEBIASED* setup includes 200 instances from the SATIRICAL corpora and 200 instances from the NON-SATIRICAL corpora, where the SATIRICAL instances are debiased stylistically.

**HYBRID:** To have an intermediate setup between *BIASED* and *DEBIASED*, the same 200 SATIRICAL instances are selected, but only 100 of them are passed through the *debiasing pipeline*. Hence, the final *HYBRID* setup consists of 200 instances from NON-SATIRICAL corpora, 100 instances from original SATIRICAL corpora, and 100 instances from debiased SATIRICAL corpora.

After training, the models were evaluated in same-domain, cross-domain, and cross-lingual settings using the following test datasets:

**Zaytung + AA:** All instances from the *Turkish Satirical News Dataset* that are not used in

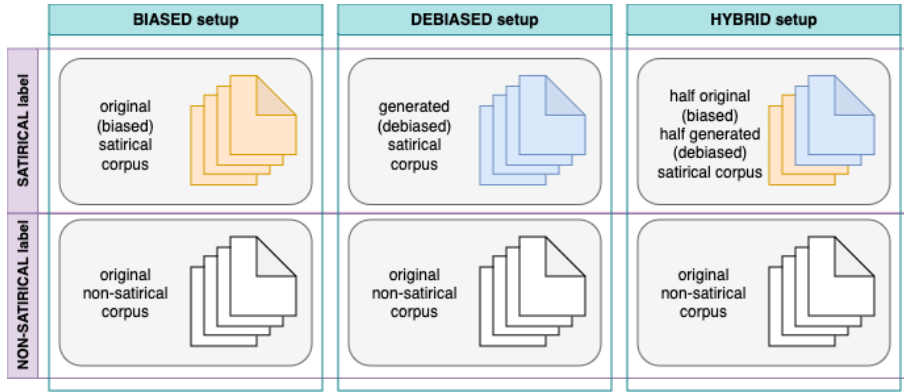


Figure 4: Training setups in the experiments

the training sets are included in this test set. Hence, this set consists of 4581 instances from the NON-SATIRICAL corpora and 2002 instances from the SATIRICAL corpora of the *Turkish Satirical News Dataset*.

**The Onion + HuffPost:** A fairly balanced set of 29000 English news article headlines from the American satirical news website the Onion<sup>8</sup> and HuffPost<sup>9</sup>, taken from the openly available News Headlines Dataset For Sarcasm Detection<sup>10</sup>.

**IronyTR:** IronyTR (Ozturk et al., 2021) is a Turkish social media irony detection dataset annotated by humans. It contains 300 Turkish ironic short texts and 300 non-ironic social media posts. This dataset is used to evaluate the cross-domain performance of the models trained under different settings.

### 5.3 Results

Table 4 presents the fine-tuning results of the selected language models on the *BIASED*, *DEBIASED*, and *HYBRID* setups which were evaluated on the Zaytung test set. According to the results, for each model except the Llama-3.1-8B, the *BIASED* setup achieved the highest F1-macro score. Since the training and test sets are from the same domain and the writing styles of the satirical and non-satirical news are significantly different, the outcome is expected. However, the proposed debiasing pipeline significantly reduced the F1-macro score across all language models. This reduction is attributed to the pipeline’s goal of changing the writing style and diminishing the sarcastic tone in

the texts, by making it more challenging for the models to differentiate between non-satirical and satirical news.

Secondly, Table 5 presents the cross-lingual evaluation results where the models trained on the Turkish dataset were tested using The Onion dataset. The results show that, except for Llama-3.1-8B, the proposed debiasing approach positively improved the F1-macro scores. In other words, the XLM-RoBERTa model achieved the highest score on the *HYBRID* dataset, which includes debiased instances, while the BERT models performed best on the *DEBIASED* setup.

Finally, Table 6 presents the models’ performance in a cross-domain setting using the IronyTR dataset which contains ironic and non-ironic social media posts. Since the social media posts are short texts, whereas the training instances are long-form news articles, the models’ performance was significantly lower than the results in Table 4. The proposed debiasing pipeline positively impacted the BERTurk and XLM-RoBERTa models; however, the highest F1-macro scores for multilingual BERT and Llama-3.1-8B were observed in the *BIASED* setup. While the Llama-3.1-8B model achieved its best F1-macro scores on both The Onion and IronyTR datasets using the *BIASED* setup, the scores were very close to those obtained in the other setups.

### 5.4 Discussion

Following the *DEBIASED* training, the masked language models demonstrated improved robustness in both cross-lingual (see Table 5) and cross-domain settings (see Table 6). However, for the Llama-3.1-8B model, *BIASED* training achieved the highest score, though the margin compared to other setups was minimal. More specifically in Table 5, *BIASED* setup outperformed *DEBIASED*

<sup>8</sup><https://theonion.com/>

<sup>9</sup><https://www.huffpost.com/>

<sup>10</sup><https://www.kaggle.com/datasets/rmisra/news-headlines-dataset-for-sarcasm-detection/data>

Models	BIASED				DEBIASED				HYBRID			
	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro
berturk	93.50%	93.95%	93.50%	<b>92.56%</b>	78.74%	77.90%	78.74%	72.26% (-20.30%)	90.52%	90.45%	90.52%	88.41% (-4.15%)
mbert-base	95.23%	95.59%	95.23%	<b>94.53%</b>	56.11%	56.21%	56.11%	56.01% (-38.55%)	94.68%	94.66%	94.68%	93.68% (-0.85%)
xlm-roberta large	97.83%	97.86%	97.83%	<b>97.39%</b>	93.61%	93.96%	93.61%	92.01% (-5.38%)	96.63%	96.73%	96.63%	95.90% (-1.49%)
llama-3.1-8B	67.26%	84.18%	67.26%	67.14%	65.28%	83.79%	65.28%	65.21% (-1.93%)	89.21%	91.40%	89.21%	<b>88.17%</b> (+21.03%)

Table 4: Evaluation on Zaytung + AA dataset

Models	BIASED				DEBIASED				HYBRID			
	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro
berturk	48.44%	48.09%	48.44%	47.78%	52.38%	52.68%	52.38%	<b>52.37%</b> (+4.59%)	50.41%	51.39%	50.41%	49.56% (+1.78%)
mbert-base	49.29%	48.44%	49.29%	47.34%	55.02%	56.26%	55.02%	<b>54.49%</b> (+7.15%)	56.08%	58.26%	56.08%	49.27% (+1.93%)
xlm-roberta large	52.45%	73.32%	52.45%	34.58%	56.53%	56.43%	56.53%	55.09% (+20.51%)	63.36%	63.72%	63.36%	<b>62.34%</b> (+27.76%)
llama-3.1-8B	70.95%	71.39%	70.95%	<b>70.70%</b>	69.54%	72.63%	69.54%	68.81% (-2.14%)	64.87%	66.36%	64.87%	63.23% (-7.72%)

Table 5: Evaluation on The Onion + HuffPost datasets

Models	BIASED				DEBIASED				HYBRID			
	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro	accuracy	precision	recall	f1-macro
berturk	53.10%	53.28%	53.10%	52.67%	59.46%	60.81%	59.46%	<b>58.01%</b> (+5.34%)	51.59%	51.58%	51.59%	51.56% (-1.11%)
mbert-base	62.48%	64.51%	62.48%	<b>61.23%</b>	56.11%	56.21%	56.11%	56.01% (-5.22%)	58.29%	58.41%	58.29%	58.08% (-3.15%)
xlm-roberta large	52.93%	55.27%	52.93%	46.43%	64.99%	70.61%	64.99%	62.56% (+16.13%)	68.01%	68.77%	68.01%	<b>67.72%</b> (+21.29%)
llama-3.1-8B	64.58%	64.59%	64.58%	<b>64.55%</b>	62.58%	70.90%	62.58%	58.69% (-5.86%)	64.83%	64.98%	64.83%	64.47% (-0.08%)

Table 6: Evaluation on the IronyTR dataset

setup by 2.14%, while in Table 6, it outperformed the *HYBRID* setup by 0.08%. This performance in *BIASED* setups for Llama can be attributed to the model’s pretraining knowledge. In other words, articles from The Onion, HuffPost, and instances of IronyTR might have been included in the pretraining data for the Llama model. Furthermore, Llama significantly outperformed the masked language models on The Onion + HuffPost dataset, further suggesting potential exposure during pretraining. Lastly, on *HYBRID* setup, XLM-RoBERTa outperformed the *BIASED* setup on both cross-lingual and cross-domain evaluations (see Table 5 and Table 6) with a significant margin. This result indicates that combining biased and debiased articles contributes to the model’s robustness.

## 6 Conclusions

The problem of satire detection demands a human in the loop by its nature since the labeling process cannot be automated. The only automatization possible is finding a satirical resource (such as Zaytung for Turkish) and assuming all scraped content is satirical by default. Unfortunately, this causes the data to be biased stylistically and trickles down this bias to the model where the model learns to identify the style of the corpus instead of the satire.

This work proposes a debiasing method utilizing LLM-based text generation within ethical limits. We show that generating data that is stylistically neutral to replace the biased data in the training set decreases the model performance significantly and improves the cross-lingual and cross-domain

robustness of the model for satire detection in Turkish. However, additional experimentation is needed to see if this method is generalizable as a debiasing method for different language tasks. Yet, the obtained results are promising to demonstrate the applicability of the proposed method.

## 7 Limitations

We tested a limited number of models which may not fully capture the variability across different models and configurations. Furthermore, there is a potential risk that some dataset instances may overlap with the training data of the LLMs (especially for the Llama-3.1-8B model) which could bias the evaluation results. Moreover, for the Zaytung dataset, the text field exceeded the sequence length of the masked language models (for BERT and RoBERTA). Therefore, we cropped the text fields for such instances.

It should also be noted that, using LLMs, specifically generative LLMs, ethical and environmental concerns should always be kept in mind. Generating textual data is an ethically convoluted topic, and should not be taken lightly. We believe that LLM-generated data should not be contextualized as if a real human has generated that content. These concerns may be limiting factors for the scalability of this study.

## Acknowledgments

The authors would like to thank Mustafa Ozan Alpay and Selin Özgürsoy for their hard work and contributions to data collection and annotation.



This work is supported by the EuroHPC Development Access Call with Project ID DD-24-99. Parts of this research received the support of the EXA4MIND project, funded by the European Union’s Horizon Europe Research and Innovation Programme, under Grant Agreement N° 101092944. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

## References

- Emre Can Acikgoz, Mete Erdogan, and Deniz Yuret. 2024. Bridging the bosphorus: Advancing turkish large language models through strategies for low-resource language adaptation and benchmarking. *arXiv preprint arXiv:2405.04685*.
- Esma Aïmeur, Sabine Amri, and Gilles Brassard. 2023. Fake news, disinformation and misinformation in social media: a review. *Social Network Analysis and Mining*, 13(1):30.
- Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–236.
- Ulas Baran Baloglu, Bilal Alatas, and Harun Bingol. 2019. Assessment of supervised learning algorithms for irony detection in online social media. *1st International Informatics and Software Engineering Conference (UBMYK)*, pages 1–5.
- Francesco Barbieri, Horacio Saggion, and Francesco Ronzano. 2014. Modelling sarcasm in twitter, a novel approach. *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 50–58.
- Konstantin Buschmeier, Philipp Cimiano, and Roman Klinger. 2014. An impact analysis of features in a classification approach to irony detection in product reviews. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 42–49.
- Ege Berk Buyukbas, Adnan Harun Dogan, Asli Umay Ozturk, and Pinar Karagoz. 2021. Explainability in irony detection. In *Big Data Analytics and Knowledge Discovery: 23rd International Conference, DaWaK 2021, Virtual Event, September 27–30, 2021, Proceedings 23*, pages 152–157. Springer.
- Paula Carvalho, Bruno Martins, Hugo Rosa, Silvio Amir, Jorge Baptista, and Mário J Silva. 2020. Situational irony in farcical news headlines. In *International Conference on Computational Processing of the Portuguese Language*, pages 65–75. Springer.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019a. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019b. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gökhan Doğru, Adrià Martín Mor, and Anna Aguilar-Amat. 2018. Parallel corpora preparation for machine translation of low-resource languages: Turkish to english cardiology corpora. In *Proceedings of the LREC 2018 Workshop ‘MultilingualBIO: Multilingual Biomedical Text Processing’*, pages 12–15.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Bilal Ghanem, Jihen Karoui, Farah Benamara, Paolo Rosso, and Véronique Moriceau. 2020. Irony detection in a multilingual context. In *European Conference on Information Retrieval*, pages 141–149. Springer.
- Viktor Hangya, Hossain Shaikh Saadi, and Alexander Fraser. 2022. Improving low-resource languages in pre-trained multilingual language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11993–12006.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark J Carman. 2017. Automatic sarcasm detection: A survey. *ACM Computing Surveys (CSUR)*, 50(5):1–22.

- Qian Li, Hao Peng, Jianxin Li, Congying Xia, Renyu Yang, Lichao Sun, Philip S Yu, and Lifang He. 2022. A survey on text classification: From traditional to deep learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(2):1–41.
- Yucheng Lin, Yuhan Xia, and Yunfei Long. 2024. Augmenting emotion features in irony detection with large language modeling. *arXiv preprint arXiv:2404.12291*.
- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. On llms-driven synthetic data generation, curation, and evaluation: A survey. *arXiv preprint arXiv:2406.15126*.
- Scott M Lundberg and Su-In Lee. 2017. [A unified approach to interpreting model predictions](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc.
- Aaron Maladry, Els Lefever, Cynthia Van Hee, and Veronique Hoste. 2023. A fine line between irony and sincerity: Identifying bias in transformer models for irony detection. In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 315–324.
- Aytuğ Onan and Mansur Alp Toçoğlu. 2020. Satire identification in turkish news articles based on ensemble of classifiers. *Turkish Journal of Electrical Engineering and Computer Sciences*, 28(2):1086–1106.
- Reynier Ortega-Bueno, Francisco Rangel, D Hernández Farias, Paolo Rosso, Manuel Montes-y-Gómez, and José E Medina Pagola. 2019. Overview of the task on irony detection in Spanish variants. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2019), Co-Located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2019)*. CEUR-WS. Org, volume 2421, pages 229–256.
- Reynier Ortega-Bueno, Paolo Rosso, and Elisabetta Fersini. 2023. Cross-domain and cross-language irony detection: The impact of bias on models’ generalization. In *International Conference on Applications of Natural Language to Information Systems*, pages 140–155. Springer.
- Asli Umay Ozturk, Yesim Cemek, and Pinar Karagoz. 2021. Ironytr: Irony detection in turkish informal texts. *International Journal of Intelligent Information Technologies (IJIT)*, 17(4):1–18.
- Endang Wahyu Pamungkas and Viviana Patti. 2018. # NonDicevoSulSerio at SemEval-2018 task 3: Exploiting emojis and affective content for irony detection in english tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 649–654.
- Chen Qian, Fuli Feng, Lijie Wen, Chunping Ma, and Pengjun Xie. 2021. Counterfactual inference for text classification debiasing. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5434–5445.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Ipek Baris Schlicht, Defne Altiok, Maryanne Taouk, and Lucie Flek. 2024. Pitfalls of conversational llms on news debiasing. *arXiv preprint arXiv:2404.06488*.
- Stefan Schweter. 2020. [Berturk - bert models for turkish](#).
- Lingfeng Shen, Weiting Tan, Sihao Chen, Yunmo Chen, Jingyu Zhang, Haoran Xu, Boyuan Zheng, Philipp Koehn, and Daniel Khashabi. 2024. The language barrier: Dissecting safety challenges of llms in multi-lingual contexts. *arXiv preprint arXiv:2401.13136*.
- Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. Defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 395–405.
- David Tomás, Reynier Ortega-Bueno, Guobiao Zhang, Paolo Rosso, and Rossano Schifanella. 2023. Transformer-based models for multimodal irony detection. *Journal of Ambient Intelligence and Humanized Computing*, 14(6):7399–7410.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2016a. Exploring the realization of irony in Twitter data. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1794–1799.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2016b. Monday mornings are my fave:)# not exploring the automatic recognition of irony in english tweets. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2730–2739.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 39–50.
- Chuhan Wu, Fangzhao Wu, Sixing Wu, Junxin Liu, Zhigang Yuan, and Yongfeng Huang. 2018. THU-NGN at SemEval-2018 task 3: Tweet irony detection with densely connected LSTM and multi-task learning. *Proceedings of the 12th International Workshop on Semantic Evaluation (SemEval-2018)*, pages 51–56.
- Liang Wu, Fred Morstatter, Kathleen M. Carley, and Huan Liu. 2019. [Misinformation in Social Media: Definition, Manipulation, and Detection](#). *ACM SIGKDD Explorations Newsletter*, 21(2):80–90.

Rong Xiang, Xuefeng Gao, Yunfei Long, Anran Li, Emmanuele Chersoni, Qin Lu, and Chu-Ren Huang. 2020. Ciron: A new benchmark dataset for Chinese irony detection. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5714–5720. European Language Resources Association.

Yuemei Xu, Ling Hu, Jiayi Zhao, Zihan Qiu, Yuqi Ye, and Hanwen Gu. 2024. A survey on multilingual large language models: Corpora, alignment, and bias. *arXiv preprint arXiv:2404.00929*.

Shiwei Zhang, Xiuzhen Zhang, Jeffrey Chan, and Paolo Rosso. 2019a. Irony detection via sentiment-based transfer learning. *Information Processing & Management*, 56(5):1633–1644.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019b. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

## A Curating Human Annotations

As is, the satirical class of the *Turkish Satirical News Dataset* is labelled as satirical since it is known to be collected from a satirical online newspaper. However, it is not analyzed to see what properties of the news articles make them satirical in the first place.

To extend the usability fields of the curated dataset and obtain more information about the satirical corpus in the dataset as a baseline for explainability tasks, a subset of the SATIRICAL instances are further annotated by a human annotator. The annotation process is as follows:

1. The main annotator goes through the whole article body and identifies the REAL and FAKE parts.
2. The REAL and FAKE markings are done according to the objective facts and events. The annotator is asked to fact-check and cite related information as needed.
3. Four volunteers from different age demographics cross-check the annotations to have a higher coverage of news landscape knowledge.
4. News articles with annotations that have a unified agreement are accepted, the rest is discarded.

Finally, the human-annotated corpus consists of 40 satirical articles. Three selected annotations are shown in Figures 6a, 7a, and 5a. The red text stands for the FAKE parts of the article, whereas the blue parts are marked as REAL.

## B Comparing Model Explanations with Human Annotations

The curated *Turkish Satirical News Dataset* includes human annotations for 40 SATIRICAL articles for utilizing explainable AI and interpretable ML methods on trained models. However, to draw a comparison between the human annotations and model explanations, it is needed to define a relation between satirical and fake content. Considering the nature of satirical news articles, it is assumed that the parts that are labeled as FAKE in the annotation are likely to contribute to the satirical meaning of the article. This can be in the form of a fake person, a fake quote, or a fake event.

Similarly, the parts that are annotated to be REAL are less likely to contribute to the overall satire in the text. For example, the event described in an article may be real, therefore it can be annotated as REAL, but there may be a fake quote in the rest of the article that contributes to the satirical meaning. Following these parallels, a binary classifier is trained on *Turkish Satirical News Dataset*, and the model decisions are explained using SHAP.

The SHAP (Lundberg and Lee, 2017) explainability method uses Shapley values to understand the relative importance of different features for a prediction instance of a model. In other words, it assigns importance values to the features relative to each other that show their weight in the final decision.

As a binary classifier model to be explained using SHAP, BERTurk (Schweter, 2020) is fine-tuned. Later, SHAP-based explanations are extracted from the model. Three selected articles that have been classified correctly and have also been annotated by the human annotator are compared. One of these comparisons is reported in this section, and the other two are discussed in detail in Appendix B, with English translations for all three articles.

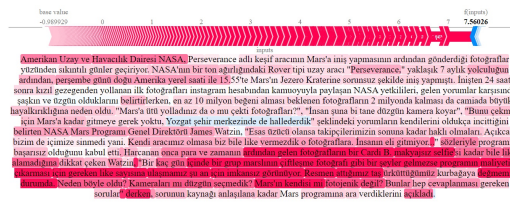
The selected explanation is shown in Figure 6b, and its human-annotated counterpart is shown in Figure 6a. The red highlights in the human annotation stand for the parts of the texts that are annotated as FAKE and the blue highlights specify the parts that are annotated as REAL. Similarly, for the SHAP output, red highlighted parts are explained as the *important* parts of the texts that the classifier focuses on when identifying a data instance as SATIRICAL. Blue highlights in the SHAP output indicate that those parts of the texts are pulling the label towards LEGITIMATE, and the parts that are



Amerikan Uzay ve Havacılık Dairesi NASA, Perseverance adlı keşif aracının Mars'a iniş yapmasını ardından gönderdiği fotoğraflar **yüzünden sıkıntılı günler geçiyor**. NASA'nın bir ton ağırlığındaki Rover tipi uzay aracı "Perseverance", yaklaşık 7 aylık yolculuğundan ardından, perşembe günü Jezero Kraterine sorunsuz şekilde iniş yaptı. İnişten 24 saat sonra kuzul gezegenden yollanan ilk fotoğrafları Instagram hesabından kamuoyuyla paylaştı NASA yetkilileri, gelen yorumlar karşısında şaşkın ve üzgün olduklarını belirtirken, en az 10 milyon beğeni olması beklenen fotoğrafların 2 milyonda kaldığına da camiaya büyük hayal kırıklığına neden oldu.

"Mars'a ültü yolladınız da o mu çekti fotoğraflar?", "İnsan şuna bi tane düzgün kamera koyar", "Bunu çekmek için Mars'a kadar gitmeye gerek yoktu. Yozgat şehri merkezinde de hallederdik" şeklindeki yorumların kendilerini oldukça incittiğini belirten NASA Mars Programı Genel Direktörü James Watzin, "Esas üzücü olansa takipçilerimizin sonuna kadar haklı olmaları. Açıkçası bizim de içimize sinmedi yani. Kendi aracımız olmasa biz bile like vermezdik o fotoğraflara. İnsanın eli gitmiyor..." sözleriyle programın başarısız olduğunu kabul etti. Harcanan onca para ve zamanın ardından gelen fotoğrafların bir Cardi B. makyajsız selfie'si kadar bile like alamadığına dikkat çeken Watzin, "Bir kaç gün içinde bir grup marslının oflenseme fotoğrafı gibi bir şeyler gelmezse programın maliyetini çıkarması için gereken like sayısına ulaşmazmış ya son için imkansız görünüyor. Resmen attığımız taş kırılmıştığımız kurbajaya değmişiz durumda. Neden böyle oldu? Kameraları ni düzgün seçmedik? Mars'ın kendisi mi fotojenik değil? Bunlar hep cevaplanmasa gereken sorular" derken, sorunun kaynağı anlaşılana kadar Mars programına ara verileceğini açıkladı.

(a) Human annotation for the article

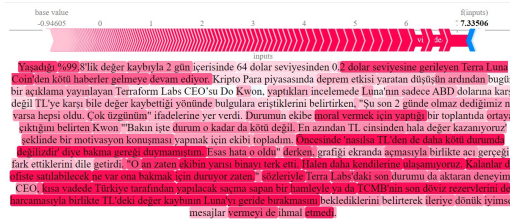


(b) SHAP output for the article

Figure 5: Human annotation and SHAP annotation for Sample Article (1)

Yaşadığı %99.8'lik değer kaybıyla 2 gün içerisinde 64 dolar seviyesinden 0.2 dolar seviyesine gerileyen Terra Luna Coin'den kötü haberler gelmeye devam ediyor. Kripto Para piyasasında deprem etkisi yaratan düşüşün ardından bugün bir açıklama yayınlayan Terraform Labs CEO'su Do Kwon, yaptıkları incelemede Luna'nın sadece ABD dolarına karşı değil TL'ye karşı bile değer kaybettiği yönünde bulgulara eriştiklerini belirttikten, "Şu son 2 günde olmaz dedüğümüz ne varsa hepsi oldu. Çok üzgünüz" ifadelerine yer verdi. Durumun ekibe moral vermek için yaptığı bir toplantıda ortaya çıktığını belirten Kwon "Bakın işte durum o kadar da kötü değil. En azından TL cinsinden hala değer kazanıyoruz" şeklinde bir motivasyon konuşması yapmak için ekibi topladım. Öncesinde "nasılsa TL'den de daha kötü durumda değilizdir" diye bakma gereği duymamıştım. Esas hata o oldu" derken, grafiği ekranda açmasıyla birlikte acı gereği fark ettiklerini dile getirdi. "O an zaten ekibin yarısı binayı terk etti. Hala daha kendilerine ulaşıyoruz. Kalemizi da ofiste satılabilecek ne var ona bakmak için duruyor zaten." sözleriyle Terra Labs'daki son durumu da aktaran deneyimli CEO, kısa vadede Türkiye tarafından yapılacak saçma sapan bir hamleyle ya da TCMB'nin son döviz rezervlerini de harcamasıyla birlikte TL'deki değer kaybını Luna'ya geride bırakmasını beklediklerini belirtirken ileride döviz kıymetli mesajlar veremeyi de ihmal etmedi.

(a) Human annotation for the article




(b) SHAP output for the article

Figure 6: Human annotation and SHAP annotation for Sample Article (2)

Will Smith'in dün gece düzenlenen 94. Oscar ödül töreninde sunucu Chris Rock'a attığı tokatın yankıları sürerken, gözler bu konuda henüz görüşünü bildirmemiş son sosyal medya kullanıcıları olan size çevirdi. Tokatın üzerinden 24 saat uzun bir süre geçmesine rağmen hala Will Smith haklı mı yoksa tam bir barzo mu, Chris Rock ayıp mı etti yoksa müthiş bir beyefendilik örneği mi sergiledi soruları hala net bir yanıt bulamazken, olayın açıklığı kavuşturulup kamuoyu vicdanının rahatlayabilmesi için bir noktada artık sizin de görüşünüzü bildirmeniz gerekiyor. Şu ana dek yaklaşık 2.4 milyar kişinin tarafını seçtiği olayla ilgili olarak son derece kritik fikrinizi açıklamadan önce bilmeniz gereken önemli bilgiler ise şöyle:

1. Will Smith'in karısı kanser değilmiş. Saçkıran mı ne öyle dandik bir hastalık yüzünden saçını kazmış
2. Evet hakkatın varmış. Ama yumruk değil tokat
3. "Toksik maskülinite" kalıbını cümle içinde kullanırken dikkat edin. Yanlış yazan çok var.
4. Will Smith tokadı attıktan sonra gidip ağlaya ağlaya Oscar aldı
5. Olay kurgu değil. Ama ola dabilir. Ya da yok ya değil...
6. T'vazetle kağıdında KDV %8'e indirildi (Bekli bir faydası olur)
7. Chris Rock haklı. Bunda düşünmek bir şey yok

(a) Human annotation for the article



(b) SHAP output for the article

Figure 7: Human annotation and SHAP annotation for Sample Article (3)

not highlighted are not important for the decision of the model.

According to Figures 6a, 6b, 7a, 7b, 5a, and 5b it can be seen that the SHAP output and the human annotation overlap for most of the red highlights, meaning that the expected match between SATIRICAL and FAKE annotations is observed. On the other hand, this seems not to be the case for blue highlights, i.e. for the parts that are annotated as REAL by the human annotators. It is observed that the model sometimes considers these parts as an indication of the SATIRICAL label or does not use those parts in the prediction at all.

A closer look at article A (Figure 6a and Figure 6b) shows that the red highlights for both the human annotation and model explanation generally match, but the blue highlights of the human annotator, i.e. the parts that are annotated as real correspond to the parts that are highlighted as slightly red or neutral by the model. Ideally, we would expect neutrality or blue highlights in the corre-

sponding parts of the SHAP output.

Additionally, a closer examination of article B (Figure 7a and Figure 7b) shows that the model and human annotator are in disagreement for most of the annotations. Even though the more saturated reds highlighted by the model match the red highlights of the annotator, that is not the case for all of the red highlights. The SHAP output even shows blue highlights where the corresponding human annotation is red, or vice versa.

Finally, a closer look at article C (Figure 5a and Figure 5b) shows that both the human annotations and the SHAP output are less continuous than the previous two comparisons. Here, the highlights mostly line up with the human annotations, but the model generally misses the sudden truth value changes in a sentence, which is expected.

Since the three examples all belong to the set of instances where the model predicts the class correctly, and the model already predicts almost every instance correctly, ideally it is expected to



see a consistent overlap between the annotations and the SHAP explanations, which is not the case. This supports the claim that a model trained with a stylistically biased dataset may not be learning the representation of satire, but the style of the corpus instead.

Additionally, translations of Sample Articles (1) and (2) can be found in Appendix C. Translation for Sample Article (3) can be found in D.

## C Prompts and Generated Articles in Turkish and English

Following two are the prompts in their original language, Turkish:

- **(Prompt 1)** *"Sana satirik bir haber vereceğim, adım adım bu haberdeki satirik unsurları kaldırmanı isteyeceğim. Önce bunun için haberdan çıkarılması gereken cümleleri tespit et, sonra da cümleler çıkarılmış haliyle haberi tekrar yaz.  
Haber metni:"*
- **(Prompt 2)** *"Sana bir metin vereceğim, içindeki satirik cümleleri daha düz bir dile çevirip tekrar yaz.  
Haber metni:"*

### C.1 Sample Article 1 generated with Prompt 1

In Turkish:

*"Amerikan Uzay ve Havacılık Dairesi NASA, Perseverance adlı keşif aracının Mars'a iniş yapmasının ardından gönderdiği fotoğraflar yüzünden sıkıntılı günler geçiriyor. Mars programının başarıyla sonuçlanan inişinden sonra, NASA yetkilileri beklenen ilginin aksine, sosyal medyada düşük bir etkileşimle karşılaştı. Program direktörü James Watzin, beklenen ilgiyi alamalarının hayal kırıklığına yol açtığını ve programın başarısız olduğunu kabul etti. Sorunun kaynağını anlamak için Mars programına ara verildi."*

In English:

*"The American Space and Aeronautics Agency NASA is going through difficult days due to the photos sent after the Perseverance rover successfully landed on Mars. Contrary to expectations, NASA officials encountered low engagement on social media following the successful landing of the Mars program. Program director James Watzin admitted that the lack of expected interest has led to disappointment and that the program has been deemed*

*a failure. The Mars program has been paused to understand the source of the problem."*

### C.2 Sample Article 2 generated with Prompt 1 and Prompt 2

Original article in Turkish:

Yaşadığı %99.8'lik değer kaybıyla 2 gün içerisinde 64 dolar seviyesinden 0.2 dolar seviyesine gerileyen Terra Luna Coin'den kötü haberler gelmeye devam ediyor. Kripto Para piyasasında deprem etkisi yaratan düşüşün ardından bugün bir açıklama yayınlayan Terraform Labs CEO'su Do Kwon, yaptıkları incelemede Luna'nın sadece ABD dolarına karşı değil TL'ye karşı bile değer kaybettiği yönünde bulgulara eriştiklerini belirtirken, "Şu son 2 günde olmaz dediğimiz ne varsa hepsi oldu. Çok üzgünüm" ifadelerine yer verdi.

Durumun ekibe moral vermek için yaptığı bir toplantıda ortaya çıktığını belirten Kwon "'Bakın işte durum o kadar da kötü değil. En azından TL cinsinden hala değer kazanıyoruz' şeklinde bir motivasyon konuşması yapmak için ekibi topladım. Öncesinde 'nasılsa TL'den de daha kötü durumda değildir' diye bakma gereği duymamıştım. Esas hata o oldu" derken, grafiği ekranda açmasıyla birlikte acı gerçeği fark ettiklerini dile getirdi.

"O an zaten ekibin yarısı binayı terk etti. Halen daha kendilerine ulaşamıyoruz. Kalanlar da ofiste satılabilecek ne var ona bakmak için duruyor zaten." sözleriyle Terra Labs'daki son durumu da aktaran deneyimli CEO, kısa vadede Türkiye tarafından yapılacak saçma sapan bir hamleyle ya da TCMB'nin son döviz rezervlerini de harcamasıyla birlikte TL'deki değer kaybının Luna'yı geride bırakmasını beklediklerini belirterek ileriye dönük iyimser mesajlar vermeyi de ihmal etmedi.

Original article in English:

Following a 99.8% loss in value, where Terra Luna Coin dropped from \$64 to \$0.20 within two days, bad news continues to emerge. After this dramatic drop that shook the cryptocurrency market, Do Kwon, CEO of Terraform Labs, released a statement today. He noted that their investigation revealed Luna had lost value not only against the US dollar but also against the Turkish Lira. Kwon said, "Everything we thought couldn't happen in the last two days has happened. I am very sorry."

Kwon explained that the realization came during a meeting he held to boost the team's morale, saying, "I gathered the team to give a motivational speech along the lines of 'Look, it's not so bad. At least we are still gaining value in TL terms.' I

hadn't felt the need to check if we were worse off compared to TL. That was the real mistake." He shared that the painful truth became apparent when they opened the graph on the screen.

"At that moment, half the team left the building. We still can't reach them. The remaining team members are just looking around the office to see if there's anything left to sell," Kwon described the current situation at Terra Labs. The experienced CEO also mentioned that they expect the depreciation of the TL to eventually surpass Luna, whether due to some absurd move by Turkey or the Central Bank of the Republic of Turkey depleting its remaining foreign currency reserves. He added that he has not neglected to give optimistic messages for the future.

### C.2.1 Generated with Prompt 1

In Turkish:

"Terra Luna Coin son günlerde büyük bir değer kaybı yaşadı. Terraform Labs CEO'su Do Kwon, yaptıkları incelemede Luna'nın sadece ABD dolarına karşı değil TL'ye karşı da değer kaybettiğini belirtti. Bu durum, şirket çalışanları arasında moral bozukluğuna neden oldu. Kwon, kısa vadede TL'deki değer kaybının Luna'yı geride bırakmasını beklediklerini ifade etti."

In English:

"Terra Luna Coin has experienced a significant loss in value in recent days. Terraform Labs CEO Do Kwon stated that their investigation revealed that Luna has lost value not only against the US dollar but also against the Turkish Lira (TL). This situation has caused a morale decline among company employees. Kwon expressed that in the short term, they expect the depreciation of the TL to surpass that of Luna."

### C.2.2 Generated with Prompt 2

In Turkish:

"Son iki günde %99.8 değer kaybı yaşayan Terra Luna Coin, 64 dolardan 0.2 dolara kadar geriledi ve kötü haberler gelmeye devam ediyor.

Bu büyük düşüşün ardından Terraform Labs CEO'su Do Kwon, bugün yaptığı açıklamada, Luna'nın sadece ABD dolarına değil, Türk Lirası'na karşı da değer kaybettiğini tespit ettiklerini söyledi. Kwon, "Bu son iki günde olmasını beklemediğimiz her şey oldu. Çok üzgünüm," dedi.

Kwon, durumu ekibe moral vermek için düzenlediği bir toplantıda fark ettiklerini belirtti. "Ekibe, 'En azından TL karşısında değer kazandık' demek iste-

dim ama durumu kontrol etmemiştim. Asıl hata bu oldu," dedi. Grafiği açtıklarında ise acı gerçekle yüzleştiklerini ifade etti.

"Toplantı sırasında ekibin yarısı binayı terk etti ve hala onlara ulaşamıyoruz. Kalanlar ise ofiste satılacak bir şeyler arıyor," diye konuşan Kwon, kısa vadede TL'nin değer kaybının Luna'ninkini geride bırakabileceğini belirterek, geleceğe yönelik iyimser bir mesaj vermeyi de ihmal etmedi."

In English:

"In the past two days, Terra Luna Coin has experienced a 99.8% loss in value, dropping from \$64 to \$0.2, and bad news continues to emerge.

Following this major drop, Terraform Labs CEO Do Kwon stated in a recent announcement that they have found evidence that Luna has lost value not only against the US dollar but also against the Turkish Lira. Kwon said, "Everything we didn't expect to happen in these last two days happened. I'm very sorry."

Kwon mentioned that they realized the situation during a meeting he organized to boost the team's morale. "I wanted to tell the team, 'At least we're gaining value against the TL,' but I hadn't checked the situation. That was the real mistake," he said. He also mentioned that they faced the harsh reality when they opened the chart.

"During the meeting, half of the team left the building, and we still can't reach them. The rest are just looking for things to sell in the office," Kwon said. He also noted that in the short term, the loss of value of the TL might surpass that of Luna, and he did not forget to provide an optimistic message for the future."

### C.3 Sample Article 4 generated with Prompt 2

Original article in Turkish:

Bu sabah Taraf gazetesi tarafından ortaya atılan "Tutsaklara karşılık Süleyman Şah Türbesi IŞİD'a verilecek" şeklindeki şok haber, Dışişleri Bakanlığı tarafından kesin bir dille yalanlandı. Önce internet sitesinden yapılan açıklama sonra da Basın Sözcüsü Kamuran Aydılen aracılığı ile kamuoyunu aydınlatan Dışişleri Bakanlığı, "Ortadaki yanlış anlaşılmaları gidermek için söylüyoruz, Süleyman Şah Türbesi'nin yıkılarak yerine AVM yapılması konusunda IŞİD'le görüştüğümüz doğru. Neticede türbe yıkımında kendilerinden daha tecrübeli bir ekip yok. Ancak bunun dışında herhangi bir pazarlık söz konusu değil" ifadeleri ile iddiaları reddetti.

Bakanlık binasında gazetecilerin sorularını yanıt-

layan Bakanlık Sözcüsü Aydilen, türbenin yıkım ihalesi için İŞİD ile pazarlık masasında oturulduğunu itiraf ederken, konunun rehinlerle doğrudan bir ilgisi bulunmadığını ise şu sözlerle savundu:

"Arkadaşlar 12 yıllık iktidarımızda artık bizi biraz tanımış olmanız lazım. Bütün dünya bilir ki biz, öyle 49 kişi için bir karış toprak vermeyiz. Hele de öyle bir toprağı, tam kupon arazi orası, deli misiniz ya? Mümkün mü böyle bir pazarlık? Türbeyi de geç, sırf arsası 4 milyar dolar eder. Orada nöbet tutan askerlerimize de sorduk, çevrede başka AVM de yokmuş. 'Çarşı izninde gidecek yer bulamıyoruz' diyorlar. Şu inşaat bir başlasın, Allah'ın izniyle para basacak orası..."

İŞİD'in özellikle türbe yıkım işinde uzmanlaşmış, işlerini severek yapan ve sahiplenen bir örgüt olduğunun altını çizen Basın Sözcüsü, "Şu an bizden haber bekliyorlar, tamam dediğimiz anda havanlarla falan girişecekler. Alimallah 1 saatte taş üstüne taş koymayız dediler. Rehinelere konusunu öyle özel olarak konuşmadık ama o konuda bir jest yaparlarsa biz bunu geri çevirmeyiz elbette. Neticede birlikte iş yapan insanlarız, yarın öbür gün başka yıkım ihaleleri de olur... Bunları da değerlendireceklerdir" ifadelerine yer verdi.

Mevcut anlaşmanın devletin kasasından bir kuruş çıkmadan halledileceğinin üzerinde duran Aydilen, yapılması planlanan AVM'nin detaylarını da basın mensuplarıyla paylaştı:

"Bakın buradan bööyle şimdiki türbenin kubbesi şeklinde bir tavan geliyor. Orası food court olacak... Alt katta SHAH'S SPORT adında bir fitness salonu ve atış poligonu var. Ta buraya kadar da meydan, forum mantığı gibi düşünün siz. Şimdi tabii aklınıza hemen ulaşım işi geliyor... Onu da düşündük. Hızlı treni 2017'de Marmaray'la Halkalı'ya bağladıktan sonra, Halkalı Ankara arası 4.5 saate inmiş olacak. Ankardan da ring seferiyle tak Halep'tesin. Son olarak Halep - Karakoza arası İDO'nun motorlarına binecek vatandaşlarımız anında AVM'de olacak. Bu kadar basit. Ayrıca oradaki askerlerimizi de özel güvenlik ve otopark görevlisi olarak AVM'de istihdam etmeyi düşünüyoruz. Gördüğünüz gibi bu projede kaybeden yok..."

Bir soru üzerine Suleymanium AVM'yi, yaşasaydı Suleyman Şah'ın da takdirle karşılayacağını sözlerine ekleyen Dışişleri Sözcüsü, son olarak şunları kaydetti:

"Yani düşünün tarihe geçmiş bir şahıssınız, arkanızda bir tanecik kullanılmayan türbe kalıyor. Ne sineması var, ne otoparkı... Böyle mi anmalıyız

eccadımızı? Ayrıca son dönemde biliyorsunuz TOKİ'nin mevcut tarihi yapılar etrafında çeşitli çalışmaları mevcut. Sosyal medyada tarihi kümbetle iç içe geçmiş yurtlarımız büyük ilgi gördü. Bu şekilde alışveriş keyfini manevi iklimle birleştiren bir çalışma halkımızın da ilgisini çekecektir..."

Original article in English:

This morning, the shocking news reported by Taraf newspaper, claiming "The Süleyman Shah Tomb will be handed over to ISIS in exchange for the hostages," was firmly denied by the Ministry of Foreign Affairs. The Ministry, which clarified the issue first through an announcement on its website and later through its Spokesperson Kamuran Aydilen, stated: "To clear up the misunderstandings, we are saying this: it is true that we have discussed with ISIS the demolition of the Süleyman Shah Tomb and replacing it with a shopping mall. After all, there is no team more experienced than them when it comes to demolitions. However, there is no bargaining or negotiation beyond this."

While answering journalists' questions at the Ministry building, Spokesperson Aydilen admitted that negotiations with ISIS had taken place regarding the demolition of the tomb, but he defended that the issue had nothing to do with the hostages, saying:

"Friends, you should have known us by now after 12 years in power. The whole world knows that we wouldn't give up an inch of land for 49 people. Especially not such a prime piece of land, are you crazy? How could there be such a deal? Forget the tomb, just the land itself is worth 4 billion dollars. We even asked the soldiers stationed there, and they said there aren't any other shopping malls around. They say, 'We can't find a place to go on our leave.' Once construction starts, God willing, that place will be printing money..."

Emphasizing that ISIS is a group specialized in demolitions, especially of tombs, and one that enjoys and takes ownership of their work, the Spokesperson added: "Right now, they are waiting for our confirmation. As soon as we say 'yes,' they'll come in with mortars and the like. I swear, they said they'll finish the job in an hour. We didn't specifically discuss the hostages, but if they make a gesture in that regard, we certainly won't turn it down. After all, we're people who work together, and who knows, there may be other demolition tenders in the future... They'll certainly evaluate those as well."

Aydilen stressed that the current agreement

would be handled without spending a penny from the state's coffers and shared details of the planned shopping mall with the press:

"Look, we're bringing in a ceiling shaped like the current dome of the tomb. That area will be a food court... On the lower floor, there will be a fitness center called SHAH'S SPORT and a shooting range. Think of it like a forum or square. Now, of course, you're probably thinking about transportation... We've thought of that too. After Marmaray connects to Halkalı in 2017, the travel time from Halkalı to Ankara will be reduced to 4.5 hours. From Ankara, a shuttle service will quickly take you to Aleppo. Lastly, our citizens traveling between Aleppo and Karakoza will take the İDO ferries and will be at the shopping mall in no time. It's that simple. We also plan to employ our soldiers stationed there as security and parking attendants in the shopping mall. As you can see, no one loses in this project..."

In response to a question, the Ministry Spokesperson added that the Süleymanium Shopping Mall would have been approved by Süleyman Shah himself, had he been alive:

"Imagine, you're a historical figure, and the only thing left behind is an unused tomb. It has no cinema, no parking lot... Is this how we should remember our ancestors? Also, as you know, TOKİ (Housing Development Administration of Turkey) has been working around existing historical structures. Our dormitories, integrated with the historical domes, have garnered a lot of attention on social media. This project, which combines the joy of shopping with a spiritual atmosphere, will also attract the public's interest..."

Generated article in Turkish:

"*Taraf gazetesinin "Tutsaklara karşılık Süleyman Şah Türbesi İŞİD'e verilecek" iddiası Dışişleri Bakanlığı tarafından kesin bir dille yalanlandı. Bakanlık, Süleyman Şah Türbesi'nin yıkılarak yerine AVM yapılması konusunda İŞİD'le görüşmelerini, ancak bu görüşmenin rehinelere ilgisi olmadığını belirtti.*

*Bakanlık Sözcüsü Kamuran Aydilen, "İŞİD yıkım işinde deneyimli bir ekip. Türbeyi yıkmak için onlarla görüşüyoruz ama toprak verme pazarlığı yok. Türbenin arsası çok değerli, rehinelere için böyle bir pazarlık olmaz. Orada AVM yaparak ekonomiye katkı sağlamayı planlıyoruz" dedi.*

*Aydilen, yapılacak AVM'nin türbenin kubbesi şeklinde bir tavanı olacağını, food court, fitness salonu ve atış poligonu gibi bölümler içereceğini, ulaşımın*

*hızlı tren ve İDO motorları ile sağlanacağını açıkladı. Askerlerin de AVM'de özel güvenlik ve otopark görevlisi olarak istihdam edileceğini ekledi. Ayrıca, bu projeye kimsenin kaybetmeyeceğini ve halkın ilgisini çekeceğini belirtti."*

Generated article in English:

*"The claim made by Taraf newspaper that the Süleyman Shah Tomb will be handed over to ISIS in exchange for hostages has been strongly denied by the Ministry of Foreign Affairs. The ministry clarified that they are in talks with ISIS regarding the demolition of the tomb to build a shopping mall in its place, but emphasized that these discussions are unrelated to the hostages.*

*Ministry spokesperson Kamuran Aydilen stated, "ISIS has an experienced demolition team. We're negotiating with them to demolish the tomb, but there is no land-for-hostages deal. The land where the tomb is located is very valuable, and we wouldn't negotiate it for hostages. We plan to contribute to the economy by building a shopping mall there."*

*Aydilen also mentioned that the mall will have a dome-shaped ceiling inspired by the tomb's dome and will include sections such as a food court, fitness center, and shooting range. Transportation to the mall will be provided by high-speed trains and İDO ferries. He added that the soldiers stationed there will be employed as security personnel and parking attendants at the mall. Furthermore, he emphasized that this project would not cause any losses and would attract public interest."*

## D Other articles

### D.1 Translation of Sample Article (3)

Warning: You're the Only One Who Hasn't Weighed In on Will Smith's Slap...

As the repercussions of Will Smith's slap on comedian Chris Rock at the 94th Oscars last night continue to ripple, all eyes have turned to you, the last social media user who has yet to express an opinion on the matter. Over 24 hours have passed since the slap, and the questions of whether Will Smith was justified or simply out of line, and whether Chris Rock was rude or demonstrated exemplary gentlemanliness, still lack clear answers. For the sake of clarifying the situation and easing the public conscience, it's time for you to share your view.

Before you disclose your crucial opinion on this event, which has involved approximately 2.4 billion



people choosing sides, here are some important details you need to know:

1. Will Smith's wife is not suffering from cancer. She shaved her head due to some trivial disease like alopecia. 2. Yes, he really did hit him. But it was a slap, not a punch. 3. Be careful when using the term "toxic masculinity" in a sentence. Many people spell it wrong. 4. After delivering the slap, Will Smith went on to cry and then won the Oscar. 5. The incident is not staged. But it could be. Or not... who knows. 6. The VAT rate on toilet paper has been reduced to 8% (Maybe this will help). 7. Chris Rock is right. There's nothing more to think about.