UmuTeam at TA1C 2025: Leveraging Large Language Models for Identifying and Spoiling Clickbait in Spanish Language

Tomás Bernal-Beltrán[1], Ronghao Pan[1], José Antonio García-Díaz[1] and Rafael Valencia-García[1]

[1]Facultad de Informática, Universidad de Murcia, Campus de Espinardo, 30100, Spain

Abstract

Clickbait headlines are designed to attract attention by creating an information gap, often prioritizing engagement over transparency. Their widespread use on social media platforms has raised concerns about their impact on the credibility of digital journalism and the spread of misinformation. This paper describes our participation in the TA1C 2025 shared task on clickbait detection and spoiling in Spanish. For the detection subtask, we developed an ensemble-based approach that combines three fine-tuned encoder-only Transformer models (MarIA, BERTIN and ALBETO) with the decoder-only Large Language Model Gemma-2-2B-it. We fine-tuned this last model using QLoRA for efficient adaptation. In this subtask, our system achieved the highest score of all the participants, demonstrating the effectiveness of combining multiple architectures under a unified classification framework. For the spoiling subtask, we proposed a two-step, zero-shot pipeline based entirely on in-context learning with Gemma-2-2B-it. In the first stage, the model generates a guiding question from the headline. In the second stage, it generates a concise spoiler by answering the question using the article body. While our system did not achieve the best performance in this subtask, it highlights the potential of prompt-only approaches for information gap resolution without any task-specific training. Our results demonstrate that, while zero-shot prompting can deliver competitive performance with minimal resources, combining it with supervised signals or hybrid techniques may be essential for more complex generation tasks, such as clickbait spoiling.

Keywords

Clickbait Identification, Clickbait spoiling, In-context Learning, Prompt-Tuning, Natural Language Processing

1. Introduction

Social media platforms; such as Facebook, X, TikTok, YouTube and Instagram, have become the main source of information for much of the population. They provide instant access to news, opinions and social interactions, and facilitate the proliferation of clickbait headlines. These are sensationalist headlines designed to exploit readers' curiosity and encourage them to click on a link. However, clickbait can also manifest in other ways, such as misleading thumbnails or vague summaries. This has significantly transformed the way users interact with news on digital platforms. Designed to capture attention and encourage clicks, these headlines often sacrifice accuracy and informational transparency, which can undermine public trust in the media [1]. While certain clickbait features can increase user engagement on platforms such as Facebook and X, studies have also found that this type of content decreases the perceived credibility of news websites [2]. Moreover, repeated exposure to misleading headlines can contribute to the spread of misinformation, as users tend to share content without verifying its accuracy.

Therefore, there is a critical need for effective mechanisms to identify and spoil this type of news, in order to mitigate its impact on the perceived credibility of news websites. However, due to the huge volume of content produced every second on social networks and news websites, performing this

IberLEF 2025, September 2025, Zaragoza, Spain
tomas.bernalb@um.es (T. Bernal-Beltrán); ronghao.pan@um.es (R. Pan); joseantonio.garcia8@um.es (J. A. García-Díaz); valencia@um.es (R. Valencia-García)
0009-0006-6971-1435 (T. Bernal-Beltrán); 0009-0008-7317-7145 (R. Pan); 0000-0002-3651-2660 (J. A. García-Díaz); 0000-0003-2457-1791 (R. Valencia-García)

identification manually becomes increasingly impractical. These processes are not only time-consuming but also resource-intensive, which underscores the need for automated and scalable solutions.

To address these challenges, the research community has been developing automated clickbait detection systems that combine advanced Natural Language Processing (NLP) with machine learning (ML) and deep learning (DL) techniques. Traditional clickbait detection methods analyze the linguistic features of headlines in an attempt to identify common clickbait patterns. For example, DL models with sentence embeddings have been employed to detect clickbait in various languages, including low-resource languages such as Urdu [3] and Indonesian [4]. Additionally, approaches combining semantic analysis and ML techniques have been explored to enhance the accuracy of identifying misleading headlines [5]. These advancements are essential for curbing the spread of clickbait, mitigating the potential misinformation it can cause, and restoring people trust in digital media.

The TA1C (Te Ahorré Un Click) Clickbait detection and spoiling in Spanish [6] shared task; which is part of IberLEF 2025 [7]), aims to efficiently identify and generate short texts (spoils) of Spanish language news articles that have previously been identified as clickbait. The overarching goal is to reduce redundancy in clickbait impact mitigation efforts by detecting which news articles need to be summarized/spoiled, and by satisfying users' curiosity and filling the information gap created by headlines. The task is divided into two "subtasks": (1) **Clickbait detection**. Given a tweet that links to a news article, the goal is to determine whether the article's content is clickbait; and (2) **Clickbait spoiling**. Given a clickbait teaser (tweet and title) and the corresponding news article, the goal is to generate or extract from the article, a short text that fills the information gap as concisely as possible (280 characters maximum) to satisfy the curiosity generated, or otherwise indicate that the article does not address it.

We participated in both subtasks of this shared task. For the Clickbait detection subtask, we propose an approach based on an ensemble of four fine-tuned models, most of which are encoder-only models. For the Clickbait spoiling subtask, we propose an approach based on the in-context learning capabilities of Large Language Models (LLMs). Specifically, we use the Gemma-2-2B-it [8] model with zero-shot learning techniques. We formulated the task as a Question Answering (QA) problem, for which the relevant questions had to be generated in advance. To this end, we used zero-shot prompting techniques with the Gemma-2-2B-it model, providing it with the news headline along with specific instructions to generate the corresponding question. Then, we addressed the generation of the spoiler by once again applying zero-shot prompting techniques using the same model. In this phase, the model was provided with the subtitle and body of the news article, along with detailed instructions, to generate the text revealing the information hidden behind the headline.

The rest of this paper is organized as follows: Section 2 reviews related work and recent advances relevant to the clickbait detection and spoiling task. Section 3 presents our system architecture and describes the approach used for both subtasks, as well as the prompting strategies employed. Section 4 outlines the dataset, preprocessing steps and evaluation metrics. Section 5 reports and discusses the results obtained. Section 6 summarizes our findings and outlines possible directions for future research.

2. Background Information

The increasing volume of clickbait content generated on social networks and news websites, coupled with the critical need for effective mechanisms to identify and spoil this type of news, has driven the development of the automation of the clickbait detection and spoiling process. Two key steps in this process are the detection of whether a news article is designed to exploit readers' curiosity and drive clicks; that is, if it's clickbait, (Clickbait detection) and the generation or extraction of a short text that fills the information gap, satisfying the curiosity generated by the misleading headline (Clickbait spoiling). These tasks have been extensively studied in recent literature in both in monolingual and multilingual contexts with the aim of creating robust systems that can operate efficiently across languages and domains, including informal domains such as social networks.

Early approaches tackled clickbait detection as a binary classification problem, whereby a headline, and sometimes its accompanying content, was analysed to determine whether or not it was clickbait. Initial efforts involved curating datasets by targeting publishers known for their clickbait content, such as BuzzFeed, Huffington Post and Upworthy, or by using crowdsourcing for data annotation. For example, in [9] the authors compiled 15.000 news headlines labeled via a publisher-based heuristic and in [10], the authors used article informality cues to detect clickbait. However, these early datasets were biased, for example, because they relied on the reputation of publishers, and many publishers with a low reputation still publish significant amounts of news without resorting to clickbait teaser messages.

In recent years, deep learning techniques have achieved state-of-the-art results in clickbait detection. Researchers have experimented with a variety of approaches. In [11], the authors use a hybrid categorization approach that integrates different features, sentence structure and clustering. In [12], the authors propose a detection model that integrates headline semantic and POS tag information.

An important and unique aspect of clickbait detection is headline–article consistency. Clickbait often occurs when a headline does not match the content, or when key information is withheld. To address this issue, summarization-assisted methods have been proposed. For example, in [13] the authors adopt this approach by employing prompt-based tuning. They generate a summary of the article using a pre-trained language model and then input the headline and the summary into a prompt that instructs the model to evaluate whether the headline is clickbait.

Another recent trend is the use of ensemble models. Given the variety of clickbait styles, ensembles of classifiers can improve robustness. For example, combining encoder-only transformers (for headline classification) with models that consider auxiliary information, such as the tweet text or metadata, can boost accuracy. In [14], the authors proposed a hybrid approach that integrates textual content with auxiliary features such as sentiment scores and engagement metrics. Their ensemble model, which combines transformer-based language models with stacking classifiers, demonstrated enhanced performance in identifying misinformation on Twitter.

The task of automatic clickbait spoiling was first introduced in [15], where the authors compiled the Webis Clickbait spoiling Corpus 2022. This corpus consists of 5.000 social media posts from Twitter, Reddit and Facebook, each of which is paired with a manually crafted spoiler text that reveals the information behind the tease. The proposed pipeline treats the spoiling task as a two-step process: (1) classifying the type of spoiler required. Either a short phrase (extracted directly from the text) or a longer passage (extracted from different parts of the text); and (2) generating the spoiler text itself. For the spoiler generation step, the authors employed a QA approach, reformulating the clickbait article headline as a question and using an extractive QA model to retrieve the answer from the article body.

Subsequently, in [16], the authors introduced the Jack-Ryder system in SemEval 2023 Task 5 [17], which embraced a zero-shot QA approach. Their method automatically rewrites clickbait headlines as questions and uses pre-trained QA models to extract spoilers directly from the text, without requiring any additional training. Furthermore, they strategically reorder the sentences within each article based on semantic similarity and select the most suitable model for each spoiler type, achieving competitive results purely through prompting and selection strategies.

An alternative strategy to solve this task is to treat it as a sequence-to-sequence learning problem, similar to summarization, in which LLMs are fine tuned to map the article body (or part of it) onto the spoiler text. For instance, [18] adopted this approach in the same SemEval Task, experimenting with various T5 models and finding that a T5-Large model fine-tuned for article-to-spoiler generation task could effectively generate spoilers when provided with the article content as input. Another example is [19], where the authors adopted a zero-shot prompt setup to evaluate a diverse set of LLMs in the clickbait summarization task. These studies have shown that abstractive approaches, such as QA or summarization, can be effective, particularly when an initial step is taken to determine the expected answer format.

The emergence of LLMs has recently opened up new possibilities for addressing the spoiling task through their in-context learning capabilities. LLMs can perform tasks without explicit fine-tuning, relying instead on carefully designed prompts with examples provided at inference time. This approach is particularly well-suited to generation tasks such as clickbait spoiling, where the model is asked to synthesize concise and informative answers from long article texts.

In fact, researchers have started using LLMs by simply instructing them to produce spoilers to clickbait articles. The QA system described earlier in [16] is an example of this approach: the authors used a text generation model as part of a zero-shot prompting spoiler extraction pipeline, crafting prompts such as "Provide the answer to the headline: [headline]?" to guide the model. Similarly, in [20], the authors used a prompt-based approach with a pre-trained language model, providing it with the article body alongside a directive to answer the headline's question; essentially building a prompt that instructs the model to perform the task like a user of "I saved you a click" website would do: reading the article and extracting the key detail. These studies demonstrated that LLMs can perform the spoiler extraction task effectively when given the right context and instructions, even without explicit training on a clickbait dataset or the clickbait spoiling task. This approach achieves performance that is competitive with fine-tuned systems while requiring significantly lower computational resources.

3. System Overview

Figure 1 illustrates the overall system architecture of the detection subtask. At the core of the system is an ensemble of fine-tuned models designed to perform binary text classification on the input text. Given a tweet that links to an news article, the system predicts whether the news article is clickbait ("Clickbait", label 1) or not ("No", label 0).

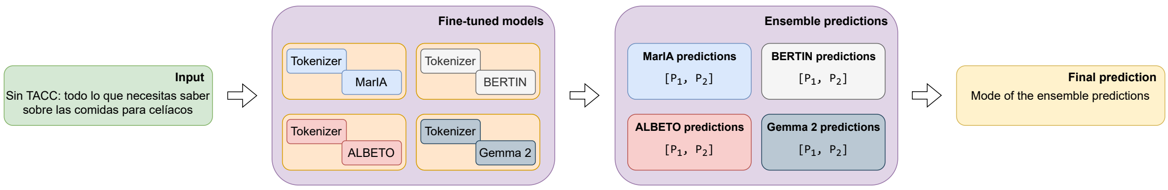

Figure 1: System architecture for the detection subtask.

The ensemble comprises four Transformer-based models: three encoder-only models: MarIA [21], BERTIN [22] and ALBETO [23], and one decoder-only large language model: Gemma-2-2B-it [8]. Each model was fine-tuned using specific hyperparameters, including the number of epochs, the batch size, the learning rate and the weight decay. These values were individually selected for each model via grid search within the following ranges: epochs $\in$ [1, 5, 10, 15, 20], batch size $\in$ [1, 2, 4, 8, 16, 32], learning rate $\in$ [1e-5, 1e-6, 3e-5, 5e-5] and weight decay $\in$ [0, 0.01, 0.05, 0.1].

We used the HuggingFace Transformers library to perform the fine-tuning for the detection subtask, formulating it as a binary text classification problem. Each piece of input text was assigned a single label: 1 for clickbait news articles, and 0 for non clickbait ones. In the case of encoder-only models, the classification head added to the encoder consists of a single linear layer on top of the `[CLS]` token, which produces logits for the two classes. For the decoder-only model, Gemma-2-2B-it, we adopted a prompt-based classification approach. During fine-tuning, the model was trained to generate a predefined label token in response to a natural language prompt containing the input headline. This token is then mapped to the corresponding binary label.

Fine-tuning was performed using the Trainer API, with cross-entropy loss as the optimization objective. Standard preprocessing steps included tokenizing the input texts, truncating or padding them to a fixed maximum length and batching them efficiently for GPU training. The tokenizer automatically handled special tokens and padding, and no additional label masking was required since classification was performed at the sequence level. For encoder-only models, fine-tuning was performed with full precision. For Gemma-2-2B-it, we applied parameter-efficient fine-tuning using QLoRA [24], which allowed us to adapt the model to the classification task while significantly reducing memory usage and training time.

Figure 2 illustrates the overall system architecture of the spoiling subtask. Our solution follows a two-step pipeline, entirely based on zero-shot prompting using the Gemma-2-2B-it model [8], without any additional fine-tuning.

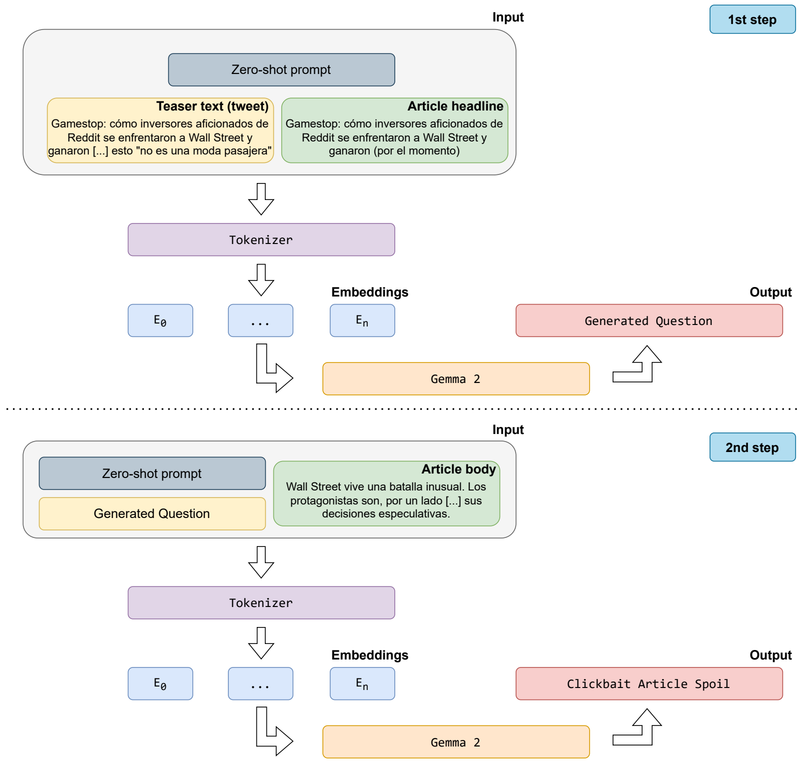

Figure 2: System architecture for the spoiling subtask.

In the first stage, since we formulated the spoiling subtask as a QA problem, our system addresses the generation of the questions that will be used to instruct the model to generate the spoiler. We then prompt the model with carefully crafted instructions designed to generate an explicit question that captures the information gap created by the clickbait article headline.

In the second stage, we use a different zero-shot prompt to generate the actual spoiler for the clickbait article. The model receives the previously generated question, as well as the subtitle and body of the news article, as input. The prompt instructs the model to generate a concise and informative response that addresses the question and satisfies the reader's curiosity. This two-step process enables the system to separate the reformulation of the headline from the generation of the answer, resulting in spoilers that are more focused and relevant.

We used the HuggingFace Transformers library to implement our solution for the spoiling subtask. Unlike the detection subtask, we did not perform any fine-tuning. Instead, we relied entirely on zero-shot prompting with the Gemma-2-2B-it model [8]. In this process, the input text was pre-processed by being lowercased and having extra whitespace trimmed, and then tokenized using the model's tokenizer.

4. Experimental Setup

For the detection subtask, our experiments were based exclusively on the training and development sets provided by the task organizers. Table 1 shows the number of examples and class distribution for each split.

To build the ensemble used to solve this subtask, we fine-tuned four different pre-trained language models, each with a different set of hyperparameters. For MarIA, fine-tuning was performed for 10 epochs with a learning rate of 1e-5, no weight decay and a batch size of 16. BERTIN was fine-tuned for 20 epochs with a learning rate of 1e-6, no weight decay and a batch size of 16. ALBETO was fine-tuned for 15 epochs with a learning rate of 3e-5, a weight decay of 0.01 and a batch size of 32.

Table 1
Detection subtask dataset distribution. Each cell reports the number of examples and the percentage they represent.

Dataset	Total examples	Clickbait	Non-Clickbait
Training	2800	798 (28.5%)	2002 (71.5%)
Development	700	203 (29%)	497 (71%)

For Gemma-2-2B-it, we applied parameter-efficient fine-tuning based on QLoRA, training the model for 10 epochs with a learning rate of 5e-5, a weight decay of 0.05 and a batch size of 4. LoRA-specific hyperparameters included a rank (r) of 8, an LoRA alpha of 16 and a dropout rate of 0.1.

For the spoiling subtask, our experiments were based exclusively on the training and development sets provided by the task organizers. Table 2 shows the number of examples for each split.

Table 2
Spoiling subtask dataset distribution. Each cell reports the number of examples.

Dataset	Number of examples
Training	300
Development	100

The first step in the pipeline designed to solve the spoiling subtask is to generate the questions that will guide the model during the spoiler generation stage. To achieve this, we used the prompt shown in Figure 1, which includes the article headline and a set of instructions for the model. Text generation was performed using the HuggingFace Transformers library's pipeline utility. We set do_sample=False, corresponding to greedy decoding. This means the model always selects the most probable token at each generation step, ensuring the generated questions are coherent and consistent.

Listing 1: Prompt used for question generation

```
Tarea: Dado un título de noticia sensacionalista en español, genera
una pregunta informativa y específica que un lector curioso podría
hacerse tras leerlo. La pregunta debe ser breve, directa y captar
la esencia de lo que se quiere saber.
Título: {headline}
```

The objective of the second step in this pipeline is to generate the actual spoiler for the clickbait article. To achieve this, we used the prompt shown in Figure 2. This prompt includes the previously generated question and the full article content (subtitle and body), as well as a set of instructions for the model. Text generation was performed using the generate() method with stochastic decoding enabled (*do_sample=True*). A combination of nucleus sampling (*top_p=0.95*), top-k sampling (*top_k=10*) and temperature-based scaling (*temperature=0.7*) was used to enable diverse yet controlled responses. The number of output tokens was limited to a maximum of 280 to match the task requirements. Decoding was performed by removing the prompt and any special tokens to extract the final spoiler.

Listing 2: Prompt used for spoiler generation

```
Eres un experto en detectar y revelar el contenido real detrás de
artículos periodísticos sensacionalistas en español (clickbait).
Tu tarea es leer la pregunta indicada y el texto del artículo
completo, para luego responder con una sola frase breve a dicha
pregunta basándote en el contenido del artículo.
**Instrucciones:**
- La respuesta debe ser clara, directa, neutral y concisa, sin
introducir información que no esté en el artículo.
```

```
- Usa solo una frase sin listas, subtítulos ni adornos.
- No repitas el contenido de la pregunta.
- La respuesta debe tener como máximo 280 caracteres
(incluyendo espacios).
- Si el artículo no contiene una respuesta clara, responde con:
"No hay respuesta".
Pregunta: {generated_question}
Texto: {article_body}
```

5. Results

In this section, we present and analyze the results obtained by our top-performing submission on the official test set. Throughout the development phase, we conducted multiple experiments, testing different fine-tuning configurations and alternative model combinations for the ensemble used in the detection subtask, and several prompt variations for the spoiling subtask. The results discussed here correspond to the best-performing setup submitted for each subtask.

Table 3 shows the official results of the detection subtask, ordered by overall F1-score. Our system (*tomasbernal01*) achieved the highest F1-score, with a value of 0.81564, outperforming the second and third-ranked submissions by a small margin. In terms of individual metrics, our system achieved a recall of 0.87425 and a precision of 0.76439. These results demonstrate the effectiveness of our ensemble-based approach, which combines multiple fine-tuned encoder-only models and a decoder-only LLM under a unified classification framework.

Table 3
Detection subtask results, the table shows the F1-score (evaluation metric), recall and precision obtained.

Name	F1-score	Recall	Precision
tomasbernal01	81.564	87.425	76.439
escom	81.524	83.233	79.885
viahes	81.481	85.628	77.717
dcere	80.480	80.239	80.722
julian_zsa	80.346	83.233	77.653
gsdeyson	80.115	83.233	77.222
Omar.Garcia	79.558	86.227	73.846
gaspai	77.747	86.826	70.388
danielrod99	77.562	83.832	72.164
John94	75.379	74.251	76.543
noetorres	75.379	74.251	76.543
ChristianRuizU	74.635	76.646	72.727
Carmen_Garcia	73.529	74.850	72.254

Table 4 shows the official results of the spoiling subtask, ordered by overall BLEU score. Our system (*tomasbernal01*) ranked third (last) with a BLEU score of 0.27980, behind the top-performing systems submitted by dcere and Omar.Garcia, which obtained scores of 0.43589 and 0.42814, respectively. Additionaly, our system achieved a ROUGE_L score of 0.64995 and a BERTscore of 0.65040. While our method did not achieve the best performance, it is notable that our approach was entirely zero-shot, relying solely on prompt-based generation with a single LLM with no additional fine-tuning or supervised training. These results highlight both the potential and current limitations of pure in-context learning strategies for this task.

Table 4
Spoiling subtask results, the table shows the BLEU (evaluation metric), ROUGE_L and BERTscore obtained.

Name	BLEU	ROUGE_L	BERTscore
dcere	43.589	74.970	73.744
Omar.Garcia	42.814	75.968	72.703
tomasbernal01	27.980	64.995	65.040

6. Conclusions and Further Work

In this work, we presented our system for the TA1C 2025 shared task, which addresses both the detection and spoiling subtasks. For the detection subtask, our ensemble-based approach combining three fine-tuned encoder-only models and a decoder-only LLM, achieved the highest score of all participants. Integrating Gemma-2-2B-it using parameter-efficient fine-tuning based on QLoRA allowed us to effectively incorporate a decoder-only architecture into a classification ensemble at low computational cost.

For the spoiling subtask, we proposed a fully zero-shot pipeline based on in-context learning with Gemma-2-2B-it. Our system framed the task as a two-step QA problem. First, it generated a guiding question from the news headline using prompting. Second, it generated the spoiler by conditioning on both the generated question and the article content. Despite lacking any task-specific fine-tuning or supervised examples, our approach achieved competitive performance, demonstrating the viability of prompt-only setups for resolving information gaps in Spanish-language news.

Although our system performed well in the detection subtask, the results of the spoiling subtask suggest that zero-shot prompting alone may not consistently generate high-quality spoils, particularly when the answer is implicit or scattered throughout the article.

Several directions remain open for future work, particularly for the spoiling subtask, where we plan to explore few-shot prompting. In particular, integrating example selection strategies, such as those proposed in recent work on Spanish hate speech detection using Gemma models [25], may improve performance compared to random few-shot or zero-shot setups. Another promising approach is to apply multi-task learning strategies, which have been successful in related fields such as hate speech detection [26]. This could unify subtasks such as question generation, spoiler generation and spoiler-type prediction under a single model. Furthermore, we recognize the potential of incorporating emotional and rhetorical analysis. Since many clickbait headlines rely on emotional manipulation to drive engagement, integrating emotion recognition techniques could enhance both detection and spoiler generation. In this regard, recent efforts such as the Spanish MEACorpus 2023 [27] open the door to enriching our models with emotional and prosodic features, which may help characterize the persuasive intent and psychological impact of clickbait more precisely.

Acknowledgments

This work is part of the research project LaTe4PoliticES (PID2022-138099OB-I00) funded by MCIN/AEI/10.13039/501100011033 and the European Fund for Regional Development (ERDF)-a way to make Europe. Mr. Tomás Bernal-Beltrán is supported by University of Murcia through the predoctoral programme.

Declaration on Generative AI

During the preparation of this work, the authors used DeepL for grammatical and spelling correction. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication’s content.

References

[1] N. Firdaus, J. Jumroni, A. Aziz, E. Sumartono, A. Purwanti, The influence of social media, misinformation, and digital communication strategies on public perception and trust, The Journal of Academic Science 1 (2024) 131–138.

[2] A.-K. Jung, S. Stieglitz, T. Kissmer, M. Mirbabaie, T. Kroll, Click me…! the influence of clickbait on user engagement in social media and the role of digital nudging, Plos one 17 (2022) e0266743.

[3] A. Muqadas, H. U. Khan, M. Ramzan, A. Naz, T. Alsahfi, A. Daud, Deep learning and sentence embeddings for detection of clickbait news from online content, Scientific Reports 15 (2025) 13251.

[4] N. P. I. Maharani, A. Purwarianti, A. F. Aji, Low-resource clickbait spoiling for indonesian via question answering, in: 2023 10th International Conference on Advanced Informatics: Concept, Theory and Application (ICAICTA), IEEE, 2023, pp. 1–6.

[5] M. Bronakowski, M. Al-Khassaweneh, A. Al Bataineh, Automatic detection of clickbait headlines using semantic analysis and machine learning techniques, Applied Sciences 13 (2023) 2456.

[6] G. Mordecki, L. Chiruzzo, R. Laguna, J. Prada, A. Rosá, I. Sastre, G. Moncecchi, Overview of TA1C at IberLEF 2025: Detecting and Spoiling Clickbait in Spanish-Language News, Procesamiento del Lenguaje Natural 75 (2025).

[7] J. Á. González-Barba, L. Chiruzzo, S. M. Jiménez-Zafra, Overview of IberLEF 2025: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2025), co-located with the 41st Conference of the Spanish Society for Natural Language Processing (SEPLN 2025), CEUR-WS. org, 2025, pp. –.

[8] M. Rivière, S. Pathak, P. G. Sessa, C. Hardin, S. Bhupatiraju, L. Hussenot, T. Mesnard, B. Shahriari, A. Ramé, J. Ferret, et al., Gemma 2: Improving open language models at a practical size, CoRR (2024).

[9] A. Chakraborty, B. Paranjape, S. Kakarla, N. Ganguly, Stop clickbait: Detecting and preventing clickbaits in online news media, in: 2016 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM), IEEE, 2016, pp. 9–16.

[10] P. Biyani, K. Tsioutsiouliklis, J. Blackmer, "8 amazing secrets for getting more clicks": detecting clickbaits in news streams using article informality, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 30, 2016, pp. –.

[11] A. Pujahari, D. S. Sisodia, Clickbait detection using multiple categorisation techniques, Journal of information science 47 (2021) 118–128.

[12] T. Liu, K. Yu, L. Wang, X. Zhang, X. Wu, Wcd: A new chinese online social media dataset for clickbait analysis and detection, in: 2021 7th IEEE International Conference on Network Intelligence and Digital Content (IC-NIDC), IEEE, 2021, pp. 368–372.

[13] H. Deng, Y. Zhu, Y. Wang, J. Qiang, Y. Yuan, Y. Li, R. Zhang, Prompt-tuning for clickbait detection via text summarization, CoRR (2024).

[14] R. A. Dar, R. Hashmy, A multimodal ensemble machine learning approach to covid-19 misinformation detection in twitter, in: ITM Web of Conferences, volume 54, EDP Sciences, 2023, p. 01015.

[15] M. Hagen, M. Fröbe, A. Jurk, M. Potthast, Clickbait spoiling via question answering and passage retrieval, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2022, pp. 7025–7036.

[16] D. Wangsadirdja, J. Pfister, K. Kobs, A. Hotho, Jack-ryder at semeval-2023 task 5: Zero-shot clickbait spoiling by rephrasing titles as questions, in: Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023), 2023, pp. 1090–1095.

[17] M. Fröbe, B. Stein, T. Gollub, M. Hagen, M. Potthast, SemEval-2023 task 5: Clickbait spoiling, in: A. K. Ojha, A. S. Doğruöz, G. Da San Martino, H. Tayyar Madabushi, R. Kumar, E. Sartori (Eds.), Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 2275–2286. URL: https://aclanthology.org/2023.semeval-1.312/. doi:10.18653/v1/2023.semeval-1.312.

[18] T. Bilgis, N. B. Bozdag, S. Bethard, Gallagher at semeval-2023 task 5: Tackling clickbait with

seq2seq models, in: Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023), 2023, pp. 1650–1655.
[19] I. García-Ferrero, B. Altuna, Noticia: A clickbait article summarization dataset in spanish, Procesamiento del Lenguaje Natural 73 (2024) 191–207.
[20] O. Johnson, B. Lou, J. Zhong, A. Kurenkov, Saved you a click: Automatically answering clickbait titles, arXiv preprint arXiv:2212.08196 (2022).
[21] A. G. Fandiño, J. A. Estapé, M. Pàmies, J. L. Palao, J. S. Ocampo, C. P. Carrino, C. A. Oller, C. R. Penagos, A. G. Agirre, M. Villegas, Maria: Spanish language models, Procesamiento del Lenguaje Natural 68 (2022). URL: https://upcommons.upc.edu/handle/2117/367156#.YyMTB4X9A-0.mendeley. doi:10.26342/2022-68-3.
[22] J. D. la Rosa y Eduardo G. Ponferrada y Manu Romero y Paulo Villegas y Pablo González de Prado Salas y María Grandury, Bertin: Efficient pre-training of a spanish language model using perplexity sampling, Procesamiento del Lenguaje Natural 68 (2022) 13–23. URL: http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/6403.
[23] J. Cañete, S. Donoso, F. Bravo-Marquez, A. Carvallo, V. Araujo, Albeto and distilbeto: Lightweight spanish language models. arxiv, arXiv preprint arXiv:2204.09145 (2022).
[24] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, Qlora: Efficient finetuning of quantized llms, Advances in neural information processing systems 36 (2023) 10088–10115.
[25] R. Pan, J. A. García-Díaz, R. Valencia-García, Optimizing few-shot learning through a consistent retrieval extraction system for hate speech detection, Procesamiento del Lenguaje Natural 74 (2025) 241–252.
[26] R. Pan, J. A. Garcia-Diaz, R. Valencia-Garcia, Spanish mtlhatecorpus 2023: Multi-task learning for hate speech detection to identify speech type, target, target group and intensity, Computer Standards & Interfaces 94 (2025) 103990.
[27] R. Pan, J. A. García-Díaz, M. Á. Rodríguez-García, R. Valencia-García, Spanish meacorpus 2023: A multimodal speech–text corpus for emotion analysis in spanish from natural environments, Computer Standards & Interfaces 90 (2024) 103856.