

# A Browser-based Open Source Assistant for Multimodal Content Verification

Rosanna Milner<sup>1</sup>, Michael Foster<sup>1</sup>, Olesya Razuvayevskaya<sup>1</sup>, Ian Roberts<sup>1</sup>,  
Valentin Porcellini<sup>2</sup>, Denis Teyssou<sup>2</sup>, Kalina Bontcheva<sup>1</sup>,

<sup>1</sup>University of Sheffield, <sup>2</sup>AFP Medialab,

Correspondence: [rosanna.milner@sheffield.ac.uk](mailto:rosanna.milner@sheffield.ac.uk)

## Abstract

Disinformation and false content produced by generative AI pose a significant challenge for journalists and fact-checkers who must rapidly verify digital media information. While there is an abundance of NLP models for detecting credibility signals such as persuasion techniques, subjectivity, or machine-generated text, such methods often remain inaccessible to non-expert users and are not integrated into their daily workflows as a unified framework. This paper demonstrates the VERIFICATION ASSISTANT, a browser-based tool designed to bridge this gap. The VERIFICATION ASSISTANT, a core component of the widely adopted VERIFICATION PLUGIN (140,000+ users), allows users to submit URLs or media files to a unified interface. It automatically extracts content and routes it to a suite of backend NLP classifiers, delivering actionable credibility signals, estimating AI-generated content, and providing other verification guidance in a clear, easy-to-digest format. This paper showcases the tool’s architecture, its integration of multiple NLP services, and its real-world application to detecting disinformation.

## Acknowledgments

This work has been co-funded by the UK’s innovation agency (Innovate UK) grant 10039055 (approved under the Horizon Europe Programme as vera.ai, EU grant agreement 101070093) under action number 2020-EU-IA-0282.

## 1 Introduction and Related Work

Digital disinformation poses a significant threat to democratic societies. The rapid advancement of generative AI, which can produce plausible text, images, and videos in seconds (Zhou and Zafarani, 2021), exacerbates this problem. This technological shift has dramatically increased the volume and sophistication of “fake news”, making the manual

verification of online content a near-impossible task for journalists and fact-checkers (Guo et al., 2022).

In response, the NLP community has developed a broad range of automated methods for information verification (Sharma et al., 2019; Srba et al., 2026). Beyond high-level fake news detection, recent works have shifted toward more fine-grained, explainable assessments of content credibility levels that mirror the aspects of professional fact-checking (Shu et al., 2017; Srba et al., 2026). Some examples of such credibility signals are the presence of propaganda and persuasion techniques (Piskorski et al., 2023), bias and subjectivity (Maab et al., 2024; Piskorski et al., 2023), and AI-generated text (Gehrmann et al., 2019).

However, a significant gap persists between this state-of-the-art research and its practical application by journalists. Most verification tools require technical expertise, and are published as stand-alone models (Srba et al., 2026). This leaves non-technical users without efficient access to the very tools designed to help them. Among the existing systems to support journalistic work and assist general users, the majority focus on a single functionality, such as news bias detection by AllSides<sup>1</sup> and GroundNews<sup>2</sup>, claim verification (Hassan et al., 2017), or propaganda detection (Da San Martino et al., 2020).

To bridge this “research-to-practice” gap, we introduce the VERIFICATION ASSISTANT, a tool designed to integrate multiple backend NLP and media analysis microservices, making state-of-the-art research accessible within the user’s browser. It represents an NLP-focused component within the VERIFICATION PLUGIN, a Chrome extension with over 140,000 active users<sup>3</sup>. The VERIFICATION ASSISTANT provides a single, unified interface where a user can submit a URL (from a news

<sup>1</sup><https://www.allsides.com>

<sup>2</sup><https://ground.news/extension>

<sup>3</sup><http://u.afp.com/plugin>

article or social media post) or upload their own media. The system first extracts relevant content (text, metadata, images) and then dispatches this content to a configurable set of backend NLP microservices (e.g., machine generated text detection, topic and genre classification, URL domain analysis). The VERIFICATION ASSISTANT’s feature set is the result of continuous, participatory design sessions with a panel of journalists and fact-checkers, ensuring its real-world utility. A key focus of our design lies in supporting the multilingual nature of journalism. Finally, the tool aggregates and displays the results in an informative, user-friendly dashboard, moving beyond simple binary “fake/real” labels and providing nuanced, explainable insights.

## 2 Architecture

The VERIFICATION ASSISTANT is a verification “Swiss army knife” that provides a single entry point to various state-of-the-art verification models. There are three main components: (a) the frontend interface, (b) the assistant backend, and (c) the verification services. The frontend provides the main user interface and is written in JavaScript using React Redux. It is distributed as a Chrome extension with language support for English, French, Spanish, Greek, Italian, Arabic, German, Japanese, Portuguese and Hungarian. The code is open source under an MIT licence and is available at <https://github.com/AFP-Medialab/verification-plugin>. The assistant backend is a Quart server application written in Python, with the source code stored in a private repository hosted by the GATE<sup>4</sup> team at the University of Sheffield. It mediates interaction between the frontend and the verification services. Finally, the verification services host individual NLP models trained to assess specific credibility indicators, such as presence of framing in the input text or the likelihood that the content is machine-generated.

The main workflow is illustrated in Figure 1. To verify an article, post, or item of media, the user can submit a URL through the frontend interface. Alternatively, there is an option to upload an image or video file stored locally. The frontend then passes this through to the assistant backend, which scrapes the contents and extracts the text, images, videos, and links contained in it and returns these to the frontend. When the frontend receives the scraped results, it then sends several requests in

<sup>4</sup><https://gate.ac.uk/>

parallel to the various verification services. The remainder of this section introduces these services and provides details of their results in the order they appear in the frontend interface.

### 2.1 Database of Known Fakes

The Database of Known Fakes (DBKF)<sup>5</sup> stores the claims that have been previously debunked by trusted organisations, such as Snopes<sup>6</sup>. Its multilingual text service passes the extracted text through this database to find a potential match. It considers the first 100 characters from the text to imitate a title and prevent distant matches from occurring. If any matches are found, these are returned to the frontend along with a score for each match which are used for ranking purposes. For example, if one result has a higher score than a second result, then that first result is more relevant for the search. Matches with a score ranking over 40 are presented to the user along with a link to the original debunk in the *Detection of previously fact checked claims* section, as shown in Figure 2.

The Fact Check Semantic (FCSS)<sup>7</sup> is the second part of the component that matches a larger proportion of the text to other pieces of text that already exist in a collection of fact-checking databases, such as Snopes.

### 2.2 URL Domain Analysis

The URL Domain Analysis service<sup>8</sup> collects information about a domain from multiple sources to inform the user about its credibility. For example, the Duke Reporter’s Lab<sup>9</sup> maintains a database of known fact-checking sites. Analysis of social media links is performed on individual accounts rather than site domains. For example, if <https://x.com/BBCNews> was referenced in an article, the service would consider the BBCNews account rather than the entire “x.com” domain. The URL for a site that has been explicitly listed as unreliable is flagged to the user as a *Warning*, as shown in Figure 3a. Where the site has been mentioned as part of a debunk, but is not listed as unreliable, this is flagged to the user as a *Mention* together with the details. Known fact-checking sites are flagged to the user as a *Fact Checker*.

<sup>5</sup><https://www.ontotext.com/knowledgehub/current/weverify-project/>

<sup>6</sup><https://www.snopes.com/>

<sup>7</sup><https://kinit.sk/>

<sup>8</sup><https://cloud.gate.ac.uk/shopfront/displayItem/url-domain-analysis>

<sup>9</sup><https://reporterslab.org/>

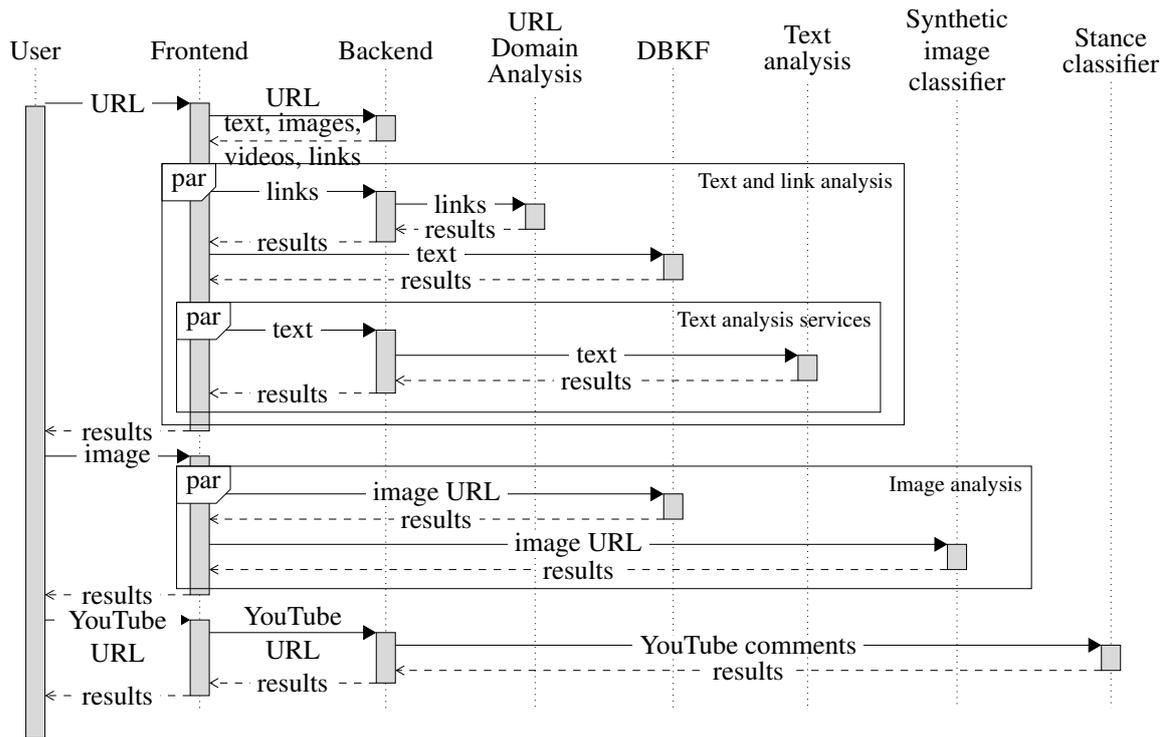
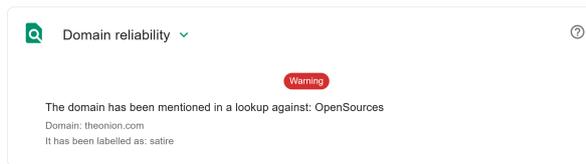


Figure 1: Requests sent by the assistant when checking a typical webpage.



Figure 2: DBKF text service and Fact Check Semantic Search for simulated data.



(a) URL domain analysis results for submitted URL.

ID	Status	Domain or Account	URL	Options
1	Fact checked	Domain: 3 URLs	https://www.welt.de/politik/deutschland/article250007304...	
2	Mentions	Social media account	x.com/BreitbartNews	
3	Warning	Domain: 1 URL	mxxl.media	
4	Warning, Mentions	Social media account	facebook.com/Breitbart	
5	Unlabelled		https://www.urhuz.org/user/emergencies/ukraine-emergency	

(b) List of extracted URLs with URL domain analysis results using mock data.

Figure 3: URL domain analysis results.

Extracted links are additionally passed to the source credibility service, with warnings, mentions,

and known fact-checking sites being flagged to the user. The results are then grouped into domains and social media accounts, and are presented in a sortable grid format in the *Extracted URLs with URL Domain Analysis* section, as illustrated in Figure 3b.

## 2.3 Media Analysis

The VERIFICATION ASSISTANT displays extracted image and video thumbnails in the *Extracted media files* section in the order in which they appear on the webpage, to facilitate an easy localisation of media in its original context. illustrated in Figure 4, users may click on an extracted image or video to view more details about it, including a list of tools within the VERIFICATION PLUGIN that are recommended for analysing the media. The possible recommended image analysis tools include image magnifier, metadata retrieval, forensic analysis, OCR, synthetic image detection, geolocalizer and provenance (C2PA). The video analysis tools include video analysis, keyframes, thumbnails, metadata and deepfake. The corresponding image and video analysis tools have been developed by AFP

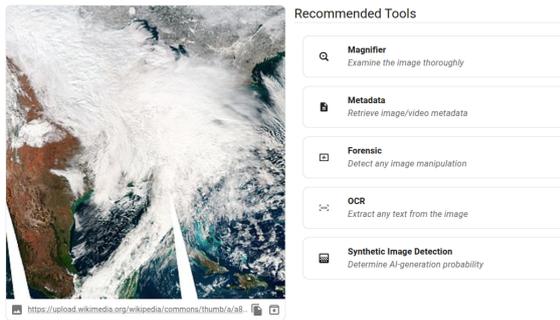


Figure 4: Recommended tools for an image.

Medialab<sup>10</sup>, Borelli Center<sup>11</sup>, ITI CERTH<sup>12</sup>, GRIP UNINA<sup>13</sup> and the University of Sheffield<sup>14</sup>.

Images are additionally sent to the DBKF image search service. Using the image similarity techniques, the service searches for matches in a database of already debunked fakes. If a match is found, this is highlighted to the user as a warning and displayed in the *Detection of previously fact checked claims* section, as described in Section 2.1.

## 2.4 Credibility Signals

The term *credibility signals* refers to a set of context- and content-based indicators that cumulatively contribute to the overall assessment of the credibility of textual information (Srba et al., 2026). The VERIFICATION ASSISTANT incorporates 5 such content-based signals proven to be vital for credibility assessment by prior research (Srba et al., 2026). More specifically, it integrates state-of-the-art classifiers for Framing (Razuvayevskaya et al., 2024; Wu et al., 2023), Genre (Razuvayevskaya et al., 2024; Wu et al., 2023), Persuasion Techniques (Razuvayevskaya et al., 2024; Wu et al., 2023), Subjectivity (Schlicht et al., 2023) and Machine Generated text (Macko et al., 2023) detection. Each signal detector is implemented through a classifier that outputs the associated signal label along with the location in text where the signal is most likely to be present. Together, these credibility signals can be interpreted by the user as an “information nutrition label” (Fuhr et al., 2018). The user can find the results for the credibility signals in the *Extracted text* section separated into different tabs.

<sup>10</sup><https://www.afp.com/en/fact-checking/fact-checking-afp/medialab>

<sup>11</sup><https://ens-paris-saclay.fr/en/research/research-laboratories/centre-borelli>

<sup>12</sup><https://caa.iti.gr/>

<sup>13</sup><https://www.grip.unina.it/>

<sup>14</sup><https://cloud.gate.ac.uk/>

España, aunque no es uno de los principales productores mundiales, ha experimentado un aumento del 23% en el consumo de leche vegetal en los últimos tres años, siendo la de almendras la preferida por el 37% de los consumidores de entre 30 y 40 años, según datos de Nielsen. Este cambio responde a un interés creciente por alternativas más saludables y sostenibles frente a la leche de vaca. Pero las autoridades ahora cuestionan si la leche de almendras es realmente la opción “verde” que muchos creían. Además de la huella hídrica, un nuevo estudio del Instituto Europeo de Nutrición advierte que la leche de almendras industrial contiene apenas un 2% de almendra en promedio, lo que ha suscitado dudas sobre su valor nutricional. “Estamos ante un producto que vende salud, pero ofrece agua con aditivos”, asegura la doctora Ana Moreno, especialista en salud pública. Aunque no se trata aún de una “prohibición formal”, los minoristas ya están siendo advertidos de una posible retirada escalonada de los productos a base de almendras, especialmente aquellos etiquetados como “leche”. Se espera que, bajo la nueva regulación, el término “leche” solo pueda emplearse para productos de origen animal, en línea con recientes directivas de la Unión Europea. La medida ha dividido a consumidores y expertos. Para algunos, se trata de una acción necesaria para frenar prácticas agrícolas insostenibles. Para otros, es una interferencia innecesaria en las

(a) News topic service.

España, aunque no es uno de los principales productores mundiales, ha experimentado un aumento del 23% en el consumo de leche vegetal en los últimos tres años, siendo la de almendras la preferida por el 37% de los consumidores de entre 30 y 40 años, según datos de Nielsen. Este cambio responde a un interés creciente por alternativas más saludables y sostenibles frente a la leche de vaca. Pero las autoridades ahora cuestionan si la leche de almendras es realmente la opción “verde” que muchos creían. Además de la huella hídrica, un nuevo estudio del Instituto Europeo de Nutrición advierte que la leche de almendras industrial contiene apenas un 2% de almendra en promedio, lo que ha suscitado dudas sobre su valor nutricional. “Estamos ante un producto que vende salud, pero ofrece agua con aditivos”, asegura la doctora Ana Moreno, especialista en salud pública. Aunque no se trata aún de una “prohibición formal”, los minoristas ya están siendo advertidos de una posible retirada escalonada de los productos a base de almendras, especialmente aquellos etiquetados como “leche”. Se espera que, bajo la nueva regulación, el término “leche” solo pueda emplearse para productos de origen animal, en línea con recientes directivas de la Unión Europea. La medida ha dividido a consumidores y expertos. Para algunos, se trata de una acción necesaria para frenar prácticas agrícolas insostenibles. Para otros, es una interferencia innecesaria en las

(b) Persuasion techniques service.

España, aunque no es uno de los principales productores mundiales, ha experimentado un aumento del 23% en el consumo de leche vegetal en los últimos tres años, siendo la de almendras la preferida por el 37% de los consumidores de entre 30 y 40 años, según datos de Nielsen. Este cambio responde a un interés creciente por alternativas más saludables y sostenibles frente a la leche de vaca. Pero las autoridades ahora cuestionan si la leche de almendras es realmente la opción “verde” que muchos creían. Además de la huella hídrica, un nuevo estudio del Instituto Europeo de Nutrición advierte que la leche de almendras industrial contiene apenas un 2% de almendra en promedio, lo que ha suscitado dudas sobre su valor nutricional. “Estamos ante un producto que vende salud, pero ofrece agua con aditivos”, asegura la doctora Ana Moreno, especialista en salud pública. Aunque no se trata aún de una “prohibición formal”, los minoristas ya están siendo advertidos de una posible retirada escalonada de los productos a base de almendras, especialmente aquellos etiquetados como “leche”. Se espera que, bajo la nueva regulación, el término “leche” solo pueda emplearse para productos de origen animal, en línea con recientes directivas de la Unión Europea. La medida ha dividido a consumidores y expertos. Para algunos, se trata de una acción necesaria para frenar prácticas agrícolas insostenibles. Para otros, es una interferencia innecesaria en las decisiones de los consumidores adultos. Sin duda, esta decisión parece demasiado severa para un país que valora la libertad de elección.

(c) Subjectivity service.

España, aunque no es uno de los principales productores mundiales, ha experimentado un aumento del 23% en el consumo de leche vegetal en los últimos tres años, siendo la de almendras la preferida por el 37% de los consumidores de entre 30 y 40 años, según datos de Nielsen. Este cambio responde a un interés creciente por alternativas más saludables y sostenibles frente a la leche de vaca. Pero las autoridades ahora cuestionan si la leche de almendras es realmente la opción “verde” que muchos creían. Además de la huella hídrica, un nuevo estudio del Instituto Europeo de Nutrición advierte que la leche de almendras industrial contiene apenas un 2% de almendra en promedio, lo que ha suscitado dudas sobre su valor nutricional. “Estamos ante un producto que vende salud, pero ofrece agua con aditivos”, asegura la doctora Ana Moreno, especialista en salud pública. Aunque no se trata aún de una “prohibición formal”, los minoristas ya están siendo advertidos de una posible retirada escalonada de los productos a base de almendras, especialmente aquellos etiquetados como “leche”. Se espera que, bajo la nueva regulación, el término “leche” solo pueda emplearse para productos de origen animal, en línea con recientes directivas de la Unión Europea. La medida ha dividido a consumidores y expertos. Para algunos, se trata de una acción necesaria para frenar prácticas agrícolas insostenibles. Para otros, es una interferencia innecesaria en las decisiones de los consumidores adultos. Sin duda, esta decisión parece demasiado severa para un país que valora la libertad de elección. Sea como fuere, todo indica que la leche de almendras, tal como la conocemos, tiene los días contados en España.

(d) Machine generated text service.

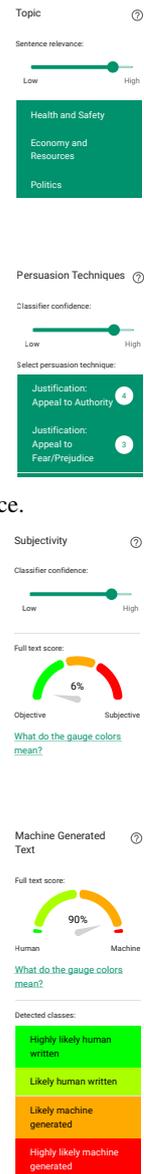


Figure 5: Credibility signals on an AI-generated article.

### 2.4.1 Framing

Framing, referred in VERIFICATION ASSISTANT as *Topic*, is a signal that represents the perspective from which information is presented (Srba et al., 2026). Its purpose is to “frame” information and guide readers toward a particular meaning. The framing classifier (Razuvayevskaya et al., 2024; Wu et al., 2023) has been trained to detect nine main frames: *Economy and Resources, Religious, Ethical and Cultural, Fairness, Equality and Rights, Law and Justice System, Crime and Punishment, Security, Defense and Well-being, Health and Safety, Politics and International Relations*. A single news article frequently incorporates several

overlapping frames simultaneously, and the tool returns all the frames above a certain threshold. For example, Figure 5a shows the top three that have a confidence score  $> 0.8$ . The classifier has been tested on six languages used during fine-tuning and three languages “unseen” during training (Razuvayevskaya et al., 2024). The model demonstrated average performance of  $F_{1_{macro}} = 59.9 \pm 3.1$  and  $F_{1_{micro}} = 61.7 \pm 7.5$  across all 9 languages. The interface highlights sentences that are deemed important by the underlying model in making a decision based on the normalised per-sentence attention scores. The interface allows the user to change the sentence importance threshold by moving a slider.

#### 2.4.2 Genre

Information *genre* is often defined as a way of distinguishing texts based on their writing style (Srba et al., 2026). For the purpose of information credibility detection, the distinction most relevant to the task concerns whether a text maintains an objective tone or incorporates manipulative language. To operationalize this, the incorporated classifier (Razuvayevskaya et al., 2024; Wu et al., 2023) distinguishes between three types of genre: *objective* reporting, *opinionated* pieces, and *satirical* content. Within this taxonomy, objective reporting is characterized by comprehensive coverage of pertinent facts and perspectives. Opinionated news, on the other hand, tends to rely on persuasive or propagandistic techniques, while satirical articles differ from both categories, as they intentionally employ fictionalized material for comedic or critical purposes. Similarly to the framing classifier, the models were tested on both “seen” and “unseen” languages (Razuvayevskaya et al., 2024), with the overall performance of  $F_{1_{macro}} = 49.2 \pm 7.4$  and  $F_{1_{micro}} = 56.7 \pm 6.1$ , averaged across genres and languages. The UI is identical to the framing classifier described in Section 2.4.1.

#### 2.4.3 Persuasion Techniques

*Persuasion techniques*, sometimes referred to as *propaganda techniques* (Piskorski et al., 2023), refer to communication strategies aimed at influencing or manipulating the reader’s opinions. The classifier (Razuvayevskaya et al., 2024; Wu et al., 2023) is trained to identify 23 different techniques which can be organised into six groups: *justification*, *simplification*, *distraction*, *call*, *manipulative wording* and *attack on reputation*. The model achieved an average performance of  $F_{1_{macro}} = 23.7 \pm 5.0$  and

$F_{1_{micro}} = 41.8 \pm 8.6$  across all persuasion techniques and 9 languages, 6 seen and 3 unseen.

Figure 5b presents the user interface. Sentences with a detected persuasion technique(s) are shown as highlighted. The interface also allows the user to hover over a certain highlighted sentence and see the corresponding detected technique(s). A list of the persuasion techniques detected across the complete text is shown on the right hand side. The user can select any technique to only highlight the sentences in which it appears. For this credibility signal, the algorithm provides a confidence score for each sentence-persuasion technique pair. Similarly to the framing classifier, a slider can be moved by the user to change the threshold of these confidence scores for highlighting the associated sentences.

#### 2.4.4 Subjectivity

*Subjectivity* signal refers to the degree to which a news article reflects personal opinions, biases, or emotions rather than strictly objective facts. The integrated classifier (Schlicht et al., 2023), trained based on the CLEF-2023 CheckThat! data challenge (Barrón-Cedeño et al., 2023), returns the score indicating the degree of subjectivity per sentence. The classifier was trained in a multilingual manner, with the test and training sets in English, Turkish and German. The model demonstrated relative robustness across the languages, with the performance of  $F_1 = 0.87$ ,  $F_1 = 0.78$  and  $F_1 = 0.74$  for Turkish, English and German respectively.

As shown in Figure 5c, the interface highlights the subjective sentences identified in text. The overall subjectivity score of the whole text is calculated as a percentage and displayed in a gauge format. The percentage falls into one of three levels: *objective*, *somewhat subjective* and *highly subjective*. Similar to persuasion techniques, a confidence score is provided for sentence subjectivity level and a slider can be moved by the user to change its threshold.

#### 2.4.5 Machine Generated Text

*Machine-generated text* detection as a credibility signal should be distinguished from legitimate automatically generated content, such as machine translation or grammar correction. The classifier (Macko et al., 2023) integrated into the tool, therefore, focuses solely on machine-generated text intended for malicious purposes, such as spreading misinformation. The VERIFICATION ASSISTANT

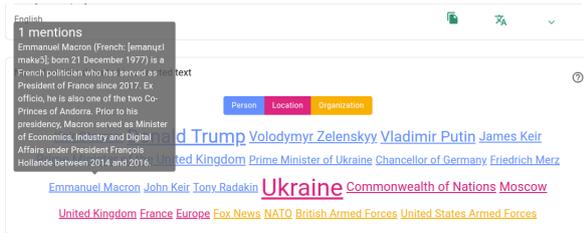


Figure 6: Named entities extracted from a news article on the war in Ukraine.

automatically identifies how likely a piece of text is to be machine generated (e.g. by a large language model), and presents this to the user as a percentage of the full text in a gauge format, as seen in Figure 5d. The text is split into sentences each of which belongs to one of four categories: *highly likely human*, *likely human*, *likely machine generated* and *highly likely machine generated*. The MGT service represents a call to a set of detectors, each trained to identify the generation by a various combinations of language models. Each model was fine-tuned on the data in English, Spanish and Russian languages, and subsequently tested on related languages for each of the training ones: Dutch and German for English, Czech and Ukrainian for Russian, Portuguese and Catalan for Spanish. The evaluation results (Macko et al., 2023) demonstrated high accuracy, with the best-performing model demonstrating  $F_{1_{macro}} = 0.8480$  and weighted  $F_1 = 0.9400$ , averaged across the classes, generator models and languages.

## 2.5 Named Entities

The named entity detector provided by the University of Sheffield identifies the names of people, locations, and organisations in text documents. To extract the named entities, the tool first identifies WikiData concepts in the text. These are then linked to their corresponding DBpedia articles, from which the `rdf:type` field is used for classification. Entities that are not classified as a Person, Location, or Organization are filtered out. The extracted entities are presented to the user as a word cloud, as illustrated in Figure 6. Here, users can select which classes of entities are displayed. The size of each entity in the cloud is proportional to the number of times that entity is mentioned in the text. When the user hovers over an entity, they are shown an abstract from DBpedia along with how many times the entity was mentioned in the text and the link to the corresponding DBpedia article.

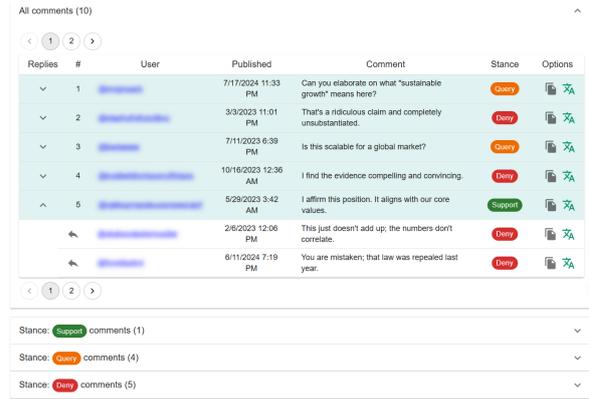


Figure 7: Synthesized comments with classifications. Usernames have been blurred to avoid accidental collisions with real (and potentially future) YouTube users.

## 2.6 YouTube Comments with Stance Classifier

When the user submits the URL of a YouTube video, the backend calls the YouTube API to extract the top 10 video comments with their replies. These comments and replies are sent to the stance classifier<sup>15</sup> which identifies comments that *support*, *deny*, or *question* the video, with respect to its title. For comment replies, the classifier identifies whether comments *support*, *deny*, or *query* the original comment. If either are found to be none of these, then the comment is simply labelled as *comment*. This is illustrated in Figure 7 which, to protect the privacy of real users, shows a set of synthesized comments from synthesized users, see Appendix B.

## 3 Evaluation

We collected quantitative feedback from 72 target users—participants of IFCN Global Fact 12<sup>16</sup> conference—in the form of the questionnaires. The survey consisted of 16 questions about the participants, their main occupation, their frequent use of the tool, the features’ usefulness in their workflow, and more qualitative questions about the user interface, their trust in the results or new features they would like to see in updates. Respondents (mainly from AFP, Al Jazeera, dpa, LeadStories, MythDetector, Deutsche Welle and France24) could declare several occupations. Their frequency of use of the tool measured 3.75 out of 5, with 19.44% of respondents using it very often and 41.67% often. Users were asked to rate each new (or enhanced) feature

<sup>15</sup><https://cloud.gate.ac.uk/shopfront/displayItem/stance-classification-multilingual>

<sup>16</sup>Global Fact is the main yearly worldwide gathering of fact-checkers

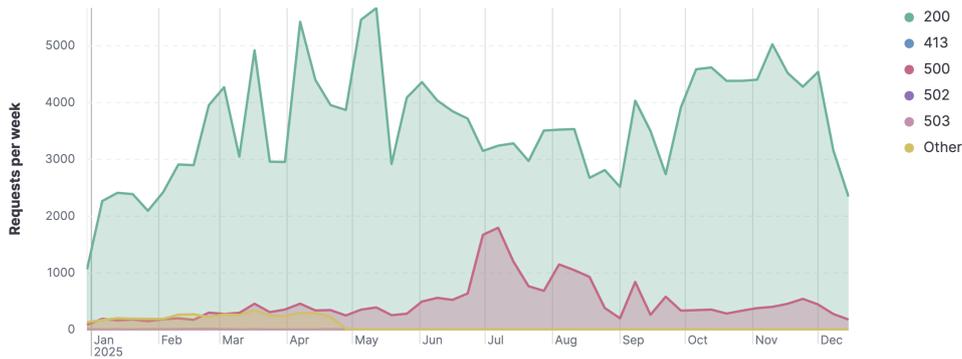


Figure 8: HTTP response code distribution through 2025.

in a scale between Excellent-Good-Average-Poor-Very poor, and this was later transformed into a Likert scale. The overall satisfaction with the VERIFICATION ASSISTANT scored 3.56 out of 5.

Additionally, regular user-centred design (UCD) evaluation sessions were conducted with our target users: professional fact-checkers, journalists, and disinformation researchers. During each session, participants first received a brief overview of the tool and its capabilities. They were then asked to verify a provided piece of information. In addition, participants were encouraged to submit their own content relevant to their daily work—for example, a news article in Greek—and discuss whether they considered the information trustworthy and described how the tool’s outputs influenced their judgment. The primary goal of these sessions was to assess the tool’s practical utility and usability, rather than just raw classifier accuracy. Qualitative feedback was gathered on result agreement, clarity of the output, and the tool’s overall value in their workflow. This feedback was central to our iterative development cycle, leading directly to practical UI refinements. Participants also suggested features or interface improvements that could enhance usability. For example, user feedback prompted the addition of explanatory pop-ups that clarify what each classifier measures and how to interpret its results.

Alongside continuously gathering user feedback, we regularly collect logs of service calls and returned errors to evaluate the robustness of the VERIFICATION ASSISTANT. Figure 8 shows the distribution of HTTP response status codes from the start until the end of 2025. As can be seen, only 14.50% of the requests received a server error response, with the majority of these errors being due to the scraping failures. Overall, 30.82% of user

requests resulted in scraper errors, which could be a result of the users submitting URLs that are either incorrect or cannot be accessed by the tool. Overall, Figure 9 (Appendix A) demonstrates the clear growing trend from the beginning until the end of 2025 in terms of the number of queries submitted by the end-users, with an overall of 18,106 requests submitted from the tool page.

#### 4 Conclusions and Future Work

This paper has presented the VERIFICATION ASSISTANT, a Chrome extension that directly bridges the gap between advanced NLP research and the daily workflow of journalists and fact checkers. By providing a common interface to a plethora of text classifiers, the VERIFICATION ASSISTANT empowers users to analyse content in terms of AI-generation, subjectivity, and other credibility signals. The VERIFICATION ASSISTANT’s integration within a plugin used by over 140,000 users demonstrates its real-world value. Its design is continually refined based on feedback from a panel of fact-checkers, journalists and researchers, ensuring it remains relevant to user needs.

Future work is divided into two main streams. From an engineering perspective, the primary challenge is to develop robust, long-term solutions for content scraping from social media platforms and news sites, which frequently alter their structure and limit access. From a research perspective, we aim to enhance the VERIFICATION ASSISTANT’s utility by integrating more complex AI models. Additionally, increasing the transparency and explainability of their outputs is paramount for building and maintaining user trust. Finally, the recent migration from Webpack to WXT has enabled browser-agnostic functionality, with Firefox support currently under development.

## References

- Alberto Barrón-Cedeño, Firoj Alam, Tommaso Caselli, Giovanni Da San Martino, Tamer Elsayed, Andrea Galassi, Fatima Haouari, Federico Ruggeri, Julia Maria Struß, Rabindra Nath Nandi, Gullal S. Cheema, Dilshod Azizov, and Preslav Nakov. 2023. [The clef-2023 checkthat! lab: Checkworthiness, subjectivity, political bias, factuality, and authority](#). In *Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III*, page 506–517, Berlin, Heidelberg. Springer-Verlag.
- Giovanni Da San Martino, Shaden Shaar, Yifan Zhang, Seunghak Yu, Alberto Barrón-Cedeño, and Preslav Nakov. 2020. [Prta: A system to support the analysis of propaganda techniques in the news](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 287–293, Online. Association for Computational Linguistics.
- Norbert Fuhr, Anastasia Giachanou, Gregory Grefenstette, Iryna Gurevych, Andreas Hanselowski, Kalervo Jarvelin, Rosie Jones, YiquN Liu, Josiane Mothe, Wolfgang Nejdl, Isabella Peters, and Benno Stein. 2018. [An information nutritional label for online documents](#). *SIGIR Forum*, 51(3):46–66.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. 2019. [GLTR: Statistical detection and visualization of generated text](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.
- Zhijiang Guo, Michael Sejr Schlichtkrull, and Andreas Vlachos. 2022. [A survey on automated fact-checking](#). *Trans. Assoc. Comput. Linguistics*, 10:178–206.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. [Toward automated fact-checking: Detecting check-worthy factual claims by claim-buster](#). In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17*, page 1803–1812, New York, NY, USA. Association for Computing Machinery.
- Jonas Maab, Eimara Marrese-Taylor, and Sarah Taylor. 2024. [Media bias detection across families of language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4083–4098, Mexico City, Mexico. Association for Computational Linguistics.
- Dominik Macko, Robert Moro, Adaku Uchendu, Jason Lucas, Michiharu Yamashita, Matúš Pikuliak, Ivan Srba, Thai Le, Dongwon Lee, Jakub Simko, and Maria Bielikova. 2023. [MULTITuDE: Large-scale multilingual machine-generated text detection benchmark](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9960–9987, Singapore. Association for Computational Linguistics.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. [SemEval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multilingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Olesya Razuvayevskaya, Ben Wu, João A. Leite, Freddy Heppell, Ivan Srba, Carolina Scarton, Kalina Bontcheva, and Xingyi Song. 2024. [Comparison between parameter-efficient techniques and full fine-tuning: A case study on multilingual news article classification](#). *PLOS ONE*, 19(5):e0301738.
- Ipek Baris Schlicht, Lynn Khellaf, and Defne Altiok. 2023. [Dwreco at checkthat!-2023: Enhancing subjectivity detection through style-based data sampling](#). 3497:306–317.
- Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. [Combating fake news: A survey on identification and mitigation techniques](#). *ACM Trans. Intell. Syst. Technol.*, 10(3).
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. [Fake news detection on social media: A data mining perspective](#). *CoRR*, arXiv:1708.01967.
- Ivan Srba, Olesya Razuvayevskaya, João A. Leite, Robert Moro, Ipek Baris Schlicht, Sara Tonelli, Francisco Moreno García, Santiago Barrio Lottmann, Denis Teyssou, Valentin Porcellini, Carolina Scarton, Kalina Bontcheva, and Maria Bielikova. 2026. [A survey on automatic credibility assessment using textual credibility signals in the era of large language models](#).
- Ben Wu, Olesya Razuvayevskaya, Freddy Heppell, João A. Leite, Carolina Scarton, Kalina Bontcheva, and Xingyi Song. 2023. [SheffieldVeraAI at SemEval-2023 task 3: Mono and multilingual approaches for news genre, topic and persuasion technique classification](#). pages 1995–2008.
- Xinyi Zhou and Reza Zafarani. 2021. [A survey of fake news: Fundamental theories, detection methods, and opportunities](#). *ACM Comput. Surv.*, 53(5):109:1–109:40.

## A User request trends

Figure 9 shows the number of unique requests or inputs submitted by the users during 2025. A clear growing trend can be observed in terms of the increasing number of events.

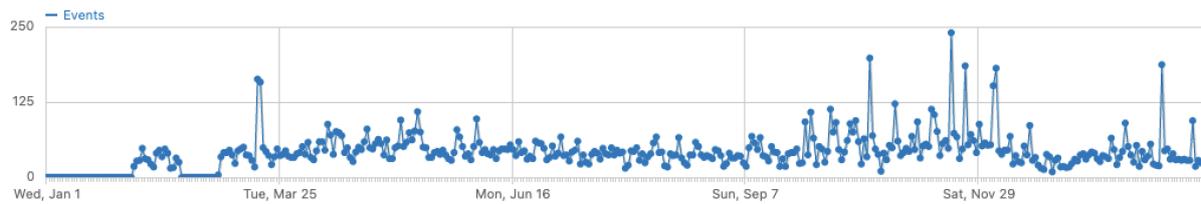


Figure 9: Temporal distribution in the number of events (queries) submitted from the VERIFICATION ASSISTANT page during 2025.

## B Prompt for synthesised YouTube comments

The prompt given to Gemini to generate mock YouTube comments and replies for Figure 7: “I am trying to test a stance classifier. The tool takes a string representing a YouTube comment or twitter post or something and classifies it as either supporting, denying, or questioning a particular statement, or just as a ‘comment’ if it doesn’t support, deny, or question. I need some sample test input for this tool that doesn’t come from anywhere. Please can you generate 100 random inputs with a selection of support, deny, questions and comments. With one input on each line?”