



OPEN

DATA DESCRIPTOR

HALT-PROP: Human-Annotated Lithuanian Textual Corpus for Propaganda Narratives and Techniques

Ieva Rizgeliene¹✉, Vilma Zubaitienė², Nerijus Maliukevičius³ & Virginijus Marcinkevičius¹

In the contemporary technological landscape, propaganda has become one of the most pervasive tools in information warfare. Social media platforms and entire media ecosystems are leveraged to disseminate hostile propaganda aimed at polarizing societies, destabilizing states, and eroding longstanding democratic processes. Malign propaganda is not only common in widely-spoken languages but also targets less-spoken languages to maximize its reach and influence. While progress has been made in developing models capable of detecting propaganda, most advances have focused on high-resource languages. In contrast, low-resource languages continue to face significant limitations, the most critical being the scarcity of annotated datasets. In many regions and countries, such resources are entirely absent. To address this gap, we present the HALT-PROP dataset, the first human-annotated Lithuanian textual propaganda corpus. The corpus comprises two complementary datasets: (1) 2,870 news articles manually labeled by five experts at the article level to identify the presence of propaganda; and (2) a subset of 1,000 articles annotated for specific propaganda techniques and narratives using a cross-annotation approach.

Background & Summary

With the rapid advancement of artificial intelligence, increasingly sophisticated natural language processing (NLP) methods are being developed to address a wide range of tasks, including the detection of propaganda and disinformation. Due to abundant linguistic data, the main progress has been achieved in rich-resource languages, where machine learning models demonstrate high levels of accuracy in identifying fake news¹ and propaganda². Progress is also made in low-resource languages³, which is mostly related to the development of multilingual transformer-based models^{4,5}. These models, pre-trained simultaneously across multiple languages, leverage cross-lingual knowledge transfer to effectively capture both universal and language-specific textual characteristics.

Despite recent advancements, applying machine learning models to propaganda detection in low-resource languages still faces significant limitations. These challenges are common across all NLP tasks for low-resource languages, including limited access to digital data and resources, complex linguistic structures, and a lack of annotated datasets⁶. There remains a significant need for new data resources in low-resource languages. In tasks such as propaganda and disinformation detection, this scarcity is even more pronounced: entire regions and countries often lack any available datasets at all.

Related work. Although textual corpora specifically tailored for propaganda analysis remain relatively scarce, considerable progress has been made, particularly with English-language resources. The pioneering *TSHP-17* corpus comprises news articles labeled as trusted, satire, hoax, or propaganda, based on the credibility of their originating news outlets⁷. Subsequently, the *Proppy* corpus was introduced, distinguishing between propaganda and trustworthy content⁸. Numerous datasets also explore related aspects of propaganda, including fake news, misinformation, rumors, and veracity. Notable datasets for fake news detection encompass the *LIAR* dataset,

¹Vilnius University Institute of Data Science and Digital Technologies, Vilnius, Lithuania. ²Vilnius University Faculty of Philology Institute of Applied Linguistics, Vilnius, Lithuania. ³Vilnius University Institute of International Relations and Political Science, Vilnius, Lithuania. ✉e-mail: ieva.rizgeliene@mif.stud.vu.lt

Dataset (Year)	Article-Level Labels	Propaganda Techniques	Narratives or Topics	Pro-Russian Propaganda	Language(s)	Corpus Size
TSHP-17 ⁷ (2017)	Yes	No	No	No	English	22,580 articles
Proppy ⁸ (2019)	Yes	No	No	No	English	51,294 articles
PTC-SemEval20 ¹⁸ (2020)	No	Yes	No	No	English	536 articles
Chang <i>et al.</i> ²⁴ (2021)	No	Yes	No	No	Mandarin	9,950 tweets
WANLP ²² (2022)	No	Yes	No	No	Arabic	930 tweets
H-Prop ²⁵ (2022)	Yes	No	No	No	Hindi	28,630 articles
H-Prop-News ²⁵ (2022)	Yes	No	No	No	Hindi	5,500 articles
SemEval-2023 Task 3 Dataset ¹⁹ (2023)	No	Yes	No	No	English, French, German, Italian, Polish, Russian	2,000 articles Mbzuai-nlp
/propaganda-codeswitched-text ²³ (2023)	No	Yes	No	No	English, Roman Urdu	1,030 Social media posts
ZenPropaganda ²¹ (2024)	No	Yes	No	No	Russian	125 texts
MU-NLPC/Propaganda ³ (2024)	No	Yes	Yes	Yes	Czech	8,646 articles
SemEval-2024 Task 4 ²⁰ (2024)	No	Yes	No	No	English, Bulgarian, North Macedonian, Arabic	10,000 memes

Table 1. Publicly available propaganda datasets.

which utilizes six distinct truthfulness labels derived from fact-checking sources⁹; the *Political News dataset* with three categories, satire, fake and real news¹⁰; the comprehensive *FakeNewsNet* repository integrating articles with rich metadata¹¹; and the *ISOT* dataset¹². Additionally, datasets are available for rumor detection and assessing claim veracity, such as *PHEME*¹³, *CREDBANK*¹⁴, *Emergent*¹⁵, *FEVER*¹⁶, and *FNC-1*¹⁷. All the aforementioned datasets provide labels at the article level, meaning that labels apply to entire texts rather than at more granular levels such as tokens, words, or sentences. Consequently, these datasets do not facilitate an in-depth analysis of specific linguistic aspects, such as the use of propaganda techniques.

In 2020, the SemEval organization, known for its series of international workshops on NLP, introduced propaganda technique detection as Task 11¹⁸. This task presented the PTC-SemEval20 corpus, an English-language dataset comprising 536 articles collected from 13 propaganda and 36 non-propaganda news media outlets, annotated for 14 distinct propaganda techniques. This corpus was the first human-annotated dataset specifically addressing propaganda techniques at a more granular level, focusing on specific fragments or spans of text rather than labeling entire articles, thus enabling a more detailed and nuanced analysis.

Progress has also been achieved in developing propaganda corpora for languages other than English. SemEval has extended its original 2020 propaganda detection task twice: first in 2023, by introducing the Multilingual Propaganda Detection Dataset, featuring span-level annotations of 23 persuasion techniques, framing annotations, and news genre labels across nine languages (English, French, German, Italian, Polish, Russian, Georgian, Greek, Spanish)¹⁹; and subsequently in 2024, with a multilingual and multimodal dataset comprising approximately 10,000 English-language memes, supplemented by smaller test sets in Bulgarian, North Macedonian, and Arabic, annotated for 22 persuasion techniques grouped into Ethos, Pathos, and Logos categories²⁰.

Regarding pro-Russian propaganda, a comprehensive Czech-language corpus was developed containing 8,646 news articles from four outlets previously identified as distributors of Russian propaganda. This dataset includes fine-grained annotations for manipulative propaganda techniques along with document-level attributes such as genre, sentiment, topic, location and additional textual attributes³.

Additional datasets annotated specifically for propaganda techniques include a Russian-language corpus focused on COVID-19 propaganda²¹, an Arabic dataset of 930 tweets annotated for 20 techniques²², a multilingual English Roman Urdu dataset comprising 1,030 social media posts annotated for 20 techniques²³, and a Mandarin corpus consisting of 9,950 annotated tweets for 21 techniques²⁴. Furthermore, datasets lacking granular annotations, providing only article-level labels for propaganda, include H-Prop and H-Prop-News for Hindi²⁵, as well as various fake news datasets in Spanish²⁶, Chinese²⁷, and Urdu²⁸.

Table 1 summarizes existing propaganda corpora. Since our dataset includes article-level labels to identify the presence of propaganda, as well as span-level annotations for propaganda techniques and labels for expressed narratives, we focused on datasets that contain at least one of these features. Given that our corpus was developed for Lithuanian, a low-resource language, we prioritized corpora created for other low-resource languages. However, due to the limited availability of such resources, we also considered datasets annotated for propaganda features in languages other than English, or in multilingual settings that included English alongside other languages. For completeness, we also include foundational English-language corpora, such as TSHP-17, Proppy, and PTC-SemEval-2020. The columns in Table 1 indicate whether each dataset includes article-level labels, annotations of propaganda techniques, and information related to specific narratives or topics, either as explicit labels or identifiable through annotation features. Additionally, we assessed whether the datasets explicitly address pro-Russian propaganda.

The analysis reveals that while several corpora exist for low-resource languages annotated for propaganda techniques, none are dedicated specifically to the languages of the Baltic countries, or even to those of Russia's neighboring states, regions frequently targeted by Russian propaganda. Furthermore, we found that only one existing corpus explicitly incorporates propaganda narratives or topics, despite their central role in the dissemination of propaganda messages. Regarding pro-Russian propaganda, only a single dataset, focused on the Czech language, directly addresses this issue. These findings highlight the uniqueness and significance of our

No.	Technique	Explanation
1	Emotional Expression	Intentionally uses emotionally charged language (e.g., fear, anger, pride, sympathy) to provoke strong feelings and influence audience beliefs or actions. Often avoids logical reasoning and relies on exaggeration, personal attacks, or vague but positive terms to shape perception.
2	Whataboutism/ Red Herring/Straw Man	Distracts from the main issue by shifting blame or criticism to others (Whataboutism), introducing irrelevant information or arguments (Red Herring), or misrepresenting an opponent's view by exaggerating, distorting, or oversimplifying it to attack a weaker version (Straw Man). These strategies serve to divert or deflect attention from the core argument.
3	Simplification	Deliberately reduces complex issues to overly simple explanations by attributing blame or responsibility to a single cause or group, framing problems as having only two opposing options, using stereotyped phrases that shut down deeper thinking, and relying on short, catchy slogans that appeal more to emotions than logic. These strategies limit critical analysis and obscure the true complexity of issues.
4	Intentional Vagueness (Obfuscation)	Uses ambiguous or imprecise language to obscure meaning, allowing multiple interpretations and helping avoid accountability or direct scrutiny.
5	Appeal to Authority	Refers to perceived authoritative figures or institutions to legitimize a claim, implying it is true based solely on the authority's status, often without supporting evidence.
6	Flag-Waving	Promotes a position by invoking patriotism or national pride, suggesting the idea serves the country's interests, even in the absence of a clear rationale or evidence.
7	Bandwagon	Encourages alignment with a belief or action by implying widespread acceptance. Leverages social pressure and the desire to conform to persuade individuals to adopt the majority view.
8	Doubt/Smears	Seeks to undermine credibility by casting suspicion or attacking character—either subtly (Doubt) or directly through baseless accusations or insinuations (Smears)—without presenting concrete evidence.
9	Reduction ad Hitlerum/Stalinum	Discredits a person, idea, or group by associating them with historically vilified figures (e.g., Hitler, Stalin), appealing to emotion rather than addressing the actual argument.
10	Repetition	Reinforces a message through frequent repetition. Over time, repeated statements may appear more familiar and thus more believable, even in the absence of evidence—a psychological effect known as the “illusion of truth.”

Table 2. Classification and explanation of common propaganda techniques.

Lithuanian propaganda corpus. It is the first annotated dataset developed for Lithuanian, and, more broadly, for any Baltic or Russia-bordering language. Uniquely, it includes both propaganda technique annotations and explicit labels for propaganda narratives, addressing a critical but largely neglected dimension in existing resources.

Motivation. Despite notable progress in developing propaganda-related textual corpora for low-resource languages, resources specifically targeting Russian propaganda in the native languages of Russia's immediate neighbors and other post-communist states remain limited. For low-resource languages in this region, textual datasets focused explicitly on Russian propaganda are scarce, with the notable exception of a comprehensive Czech-language corpus³. It is also important to note that no dedicated propaganda corpora currently exist for the languages of the Baltic countries (Lithuanian, Latvian, Estonian), nor for other languages spoken in Russia's neighboring states, such as Ukrainian, Kazakh, Azerbaijani, or Georgian.

This gap is especially striking, given that Russian information operations have long prioritized targeting the so-called “near abroad”²⁹. Numerous studies have highlighted this pattern³⁰. This Russian approach is rooted in doctrines such as the “humanitarian dimension of foreign policy”³¹, which frame the neighboring countries not only as spheres of privileged interest but also as ideological and cultural extensions of Russia. Although substantial research has examined the Russian malign influence toolbox, focusing on actors, platforms, and tactics^{32–34}, few efforts have been sought to systematically document or annotate propaganda texts in the native languages of the targeted societies. This lack of native-language corpora hinders the development of AI-driven detection tools, automated narrative analysis, and multilingual disinformation monitoring systems. This is particularly problematic in the Baltic region, where Lithuania, Latvia, and Estonia remain on the frontlines of Russian information warfare, but lack dedicated linguistic resources for the computational analysis of propaganda content. Addressing this gap is critical for advancing both academic understanding and practical measures against Russian malign influence in vulnerable language ecosystems.

Methods

The aim in creating this dataset was to create a dataset that is suitable to train/fine-tune large language models for propaganda identification in media outlets. To successfully achieve this aim, we did a deep analysis in order to correctly define and describe what propaganda is, what propaganda techniques are, and what topics are most sensitive and popular for propaganda in an information war context.

Preparation process. In this preparation process, we define propaganda, propaganda narratives and techniques. Later, these definitions were used by annotators to uniquely identify them.

Definition of propaganda. Propaganda can be defined in many ways: it can be described as ideas or statements that may be false or present only one side of an argument that is used to gain support for a political leader, party, etc³⁵, a systematic effort to shape perceptions and behavior³⁶, messaging that exploits emotional rather than rational responses³⁷, or methods employed by organized groups to psychologically unify and mobilize individuals³⁸. Broadly, propaganda involves deliberate communication strategies designed to influence public opinion through specific topics, narratives, and techniques.

In propaganda analysis, a topic denotes the core issue or theme being discussed, while a narrative refers to the specific framing or interpretation applied to that topic, combining facts, values, and emotions into a coherent message. In our study, focused on pro-Russian propaganda aimed at polarizing and radicalizing Lithuanian society, we classify an article as propagandistic if it explicitly presents at least one of our identified propaganda narratives, and employs at least one recognized propaganda technique.

Classification and explanation of propaganda techniques. This classification groups related propaganda techniques into broader categories based on shared rhetorical strategies and persuasive functions, reflecting their conceptual similarities. Building on SemEval datasets, which show that grouping similar techniques improves annotation accuracy and clarity, this framework simplifies the approach by consolidating various emotional appeals under “Emotional Expression” and merging “Doubt” with “Smears.” Clear definitions accompany each technique to support conceptual clarity and interdisciplinary use. Prioritizing usability over fine-grained differentiation, it identifies ten distinct techniques, mainly for qualitative and interdisciplinary analysis, rather than automated processing.

In contrast, the SemEval-2023 Task 3 framework by Piskorski *et al.*¹⁹ offers a more detailed scheme with 23 techniques, distinguishing closely related strategies such as “Whataboutism,” “Red Herring,” and “Straw Man” separately, and splitting emotional appeals into “Loaded Language,” “Appeal to Fear-Prejudice,” and “Exaggeration-Minimisation.” Some categories in the original classification, like “Reductio ad Hitlerum/Stalinum,” are absent from SemEval-2023, which includes other categories merged or missing in the original. These differences reflect contrasting priorities: simplicity and broad applicability for qualitative analysis, versus detailed differentiation for computational annotation in multilingual contexts.

To maintain clarity, this simplified classification integrates SemEval-exclusive techniques into broader categories: Slogans fall under Simplification as catchy phrases simplifying complex issues; Conversation Killer aligns with Doubt/Smears or Intentional Vagueness due to its role in disrupting discussion and evading scrutiny; Obfuscation-Vagueness-Confusion is grouped under Intentional Vagueness; Loaded Language and Exaggeration-Minimisation belong to Emotional Expression, with the latter also fitting under Simplification for selective emphasis.

The SemEval-2024 techniques “Thought-Terminating Cliché,” “Glittering Generalities,” and “Presenting Irrelevant Data” are incorporated as follows: “Thought-Terminating Cliché” fits under Simplification for shutting down critical thinking; “Glittering Generalities” is part of Emotional Expression due to its vague, positive language evoking emotion; and “Presenting Irrelevant Data” aligns with Doubt/Smears or Intentional Vagueness, as it distracts from the main issue²⁰. This grouping simplifies detailed distinctions for clearer qualitative analysis.

By situating these techniques within broader categories, the classification balances accessibility and usability for qualitative and interdisciplinary analysis, while acknowledging SemEval’s emphasis on fine-grained distinctions. The table below presents the ten propaganda techniques identified here, with explanations to support clear understanding and application.

Alongside defining and categorizing propaganda techniques, it is equally important to examine how these techniques manifest within specific narratives.

Selection of propaganda narratives. The next stage involved identifying specific Russian-originated narratives directed at Lithuania. Drawing on existing research and institutional assessments, a list of key propaganda narratives attributed to Russia was compiled. This list served as the basis for the annotation process and enabled systematic analysis and categorization of propaganda narratives in the dataset.

While the form of storytelling, how messages are structured and delivered, is the focus of propaganda technique analysis, this study adopted Catherine Kohler Riessman’s thematic narrative analysis to examine the content of propaganda messaging. This dual-level approach allows distinguishing between the deeper narrative structures and the rhetorical techniques used to reinforce them. Thematic analysis, as defined by Riessman, emphasizes “what” is said rather than “how” it is said, aiming to identify recurring conceptual groupings and thematic units across cases³⁹. Propaganda narratives can be analyzed not merely as collections of false claims, but as coherent meaning-making structures that organize information into politically consequential storylines. Patterson and Monroe describe narrative as “the ways in which we construct disparate facts in our own worlds and weave them together cognitively in order to make sense of our reality”⁴⁰. Accordingly, identifying Russian propaganda narratives against Lithuania involves tracing how malign actors selectively assemble events, symbols, and moral judgments into strategic narratives that seek to normalize anti-Lithuanian worldviews.

Accordingly, existing research on Russian propaganda in Lithuania and the broader region was reviewed in order to extract already-identified narrative themes. These served as the basis for a systematic annotation scheme that enabled the classification of messages into recurring ideological and thematic categories. The 2018 report by NATO StratCom COE “Russia’s Footprint in the Nordic-Baltic Information Environment” identified a total of 29 narratives promoted by Russia in regard to the Nordic-Baltic countries⁴¹. The 2020 follow-up report “Information Influence Campaigns in the Baltic States” further refined the typology, e.g. some narratives were added, and some were removed from the list⁴². Among the Russian propaganda narratives targeting the wider Baltic-Nordic region, several were specifically adapted to the Lithuanian context. These could be summed up as following:

- NATO as a threat: portraying NATO’s presence in Lithuania as provocative and destabilizing.
- Historical revisionism: denying the Soviet occupation of Lithuania, questioning the legitimacy of Lithuanian independence.
- Failed state: depicting Lithuania as a failed and irrelevant state, manipulated by foreign powers.
- Russophobia: alleging systemic discrimination against Russian-speakers.

- Western moral decay: contrasting Western immorality with Russian traditional values.
- Geopolitical victimhood: presenting Russia as encircled by enemies (Lithuania, NATO), shifting the blame for regional tensions away from Russian actions.
- Migration and refugees: framing Lithuania's policy in this sector as a threat to national identity, security, and social cohesion.
- Questioning rearmament and defence spending: criticising Lithuania's increased defence budget and military modernisation as unnecessary or externally imposed.
- The Baltics as geopolitical playground: portraying Lithuania as a puppet of global political-economic elites, particularly George Soros and the Washington establishment.

In addition to the propaganda narratives already identified in prior research, it was essential to incorporate the assessments provided by Lithuanian state institutions. Equally important was the need to contextualize the observed narratives within the broader political and social environment, namely, the key events that occurred during the period covered by the annotated corpus. As a first step, the annual National Threat Assessments, published by Lithuania's State Security Department and the Defence Intelligence and Security Service, were analyzed. These reports offer a perspective on how state national security institutions evaluate Russian propaganda narratives targeting Lithuania. The 2023 report outlined two key narratives that focused on the war in Ukraine: the first one targeted Western sanctions against Russia; the second Ukrainian refugees⁴³. The 2024 report added three more anti-Ukraine related narratives: war in the Baltic states is imminent; support for Ukraine is pointless; the war in Ukraine is of no importance to the West⁴⁴. Building on the previous analysis, the 2025 report further underscored that, alongside anti-Ukrainian propaganda, Russia had revived its traditional narratives directed against Lithuania: first, Russia sought to discredit Lithuania by accusing it of rewriting history, promoting Nazism, and fostering Russophobia, framing Lithuania as one of the most aggressively anti-Russian countries in Europe; second, historical revisionist narratives accused the Baltic states of distorting the history of the Second World War, glorifying Nazi collaborators, and erasing Soviet contributions; the third narrative claimed systematic discrimination against Russian speakers in Lithuania, aiming to attract criticism from international human rights institutions, and to equate Lithuania's domestic policies with historical atrocities⁴⁵. Moreover, when assessing the broader geopolitical disruptions that shaped the information environment during this period, developments in Belarus must be considered. These include the refugee crisis at the Lithuanian border, orchestrated by the Lukashenko regime, and the mass protests following the fraudulent presidential elections. The migration wave was accompanied by coordinated propaganda campaigns, while the subsequent flight of Belarusians from the Lukashenko regime created new societal tensions in Lithuania. Russian and Belarusian propaganda actors have exploited these developments, particularly by reviving historical themes related to "Litvinism", to sow mistrust in Lithuania towards Belarusians⁴⁶.

These narrative themes provided an academically and contextually grounded foundation for the annotation process and were synthesized for the empirical coding phase as follows, to reflect both continuity and evolution in hostile messaging patterns:

1. Disinformation about the war in Ukraine (i.e. spreading false narratives to justify Russia's aggression and delegitimize Ukrainian resistance);
2. Delegitimization of the Lithuanian State (i.e. slandering the Republic of Lithuania as a failed or artificial "project," questioning its sovereignty and historical foundations);
3. Undermining the Lithuanian Armed Forces (i.e. attacks on military funding, modernization efforts, and NATO deployments, aiming to portray Lithuania as militaristic or provocatively anti-Russian);
4. Erosion of Trust in Lithuanian Institutions (i.e. promoting narratives that depict state authorities as corrupt, incompetent, or unrepresentative);
5. Attacks on Western Institutions and Alliances (i.e. discrediting the EU, NATO, and other multilateral bodies, framing them as exploitative, ineffective, or morally bankrupt);
6. Decline of Western Civilization (i.e. spreading claims about Western moral decay, often emphasizing themes like gender ideology, LGBT rights, or secularism, to contrast it with "traditional values");
7. Authoritarian Model Promotion (i.e. highlighting regimes like Moscow, Minsk, or Beijing as examples of stability, efficiency, and sovereign governance in contrast to Western "chaos");
8. Narratives of US Decline and "Washington Hegemony" (i.e. framing the United States as a waning imperial power and suggesting the emergence of a multipolar or "New World Order" led by Russia and China);
9. Geopolitical Reordering and the "New World Order" (i.e. promoting conspiracy-laden ideas about a global realignment that replaces liberal democratic systems with alternative authoritarian alliances);
10. Weaponization of Migration and Refugees (i.e. amplifying fears around migrant flows and depicting them as tools of hybrid warfare or existential threats to national identity and security);
11. Revival of "Litvinism" (i.e. exploiting historical revisionism to push the narrative that parts of Lithuania historically belonged to Belarus, undermining the Lithuanian national identity and territorial integrity).

Data collection process. In the initial stage of selecting media outlets for articles for the dataset, we referenced findings from research⁴⁷ conducted by Lithuanian National Radio and Television (LRT) and the independent disinformation analysis centre DebunkEU.org, which provided a detailed analysis of Lithuanian online networks disseminating pro-Kremlin disinformation, particularly during the initial phase of Russia's invasion of Ukraine. Starting with a list of potential sources for the dataset, we analyzed these media outlets to assess their current accessibility, as several were shut down during 2024. We also investigated the presence and activity of

Source/Media outlet	Number of topics	Number of items collected	Number of items after the filtration
Bükimevieningi.lt	50	7,409	6,787
Ekspertai.eu	132	13,052	10,367
Infa.lt	82	6,996	6,491
Komentaras.lt	68	1,893	1,403
Ldiena.lt (ldiena.com)	50	18,416	15,583
Lrt.lt	54	139,845	133,197
Minfo.lt	94	15,880	15,530
Total:	203,491	186,376	

Table 3. Media outlets chosen for the dataset formation.

these media outlets on social media to identify the most influential ones. During this investigation, we found that some sources are just republishing information from other sources. Therefore, they were excluded from the candidate list. All sources that were ranked and chosen for our dataset were identified by this investigation as key distributors of misleading narratives, conspiracy theories, and propaganda intended to distort public perceptions, legitimise Russian military actions, and undermine trust in Western democratic institutions. In Table 3, we present all media outlets chosen as the text basis for the dataset. The texts were collected during the period from 1 January 2018 to 1 October 2024. This period was chosen to see how propaganda techniques and narratives changed before and during the war. Not all media outlets' texts could consist of propaganda narratives. Therefore, we chose specific categories (topics) for the data collection. For future analysis, we also chose one media outlet (Lrt.lt) that is sponsored by the Lithuanian government; therefore, we assume it should not have any propaganda texts, but could be very valuable to validate our annotation results. For the same reason, we also did not remove news articles based on topics (e.g. sport, leisure).

At the next stage, the collected data was filtered to remove articles that have the same name or the same content, or have a content length shorter than 200 characters or longer than 10,000 characters. Short texts do not contain any information or enough information to identify them as propaganda or not. On the other hand, the very long pages are hard to analyse and annotate.

Data selection process. Items from the collected dataset (see Table 3) were selected randomly in batches of 50 to 100 articles for each media outlet individually for the first stage of the annotation process (see Fig. 1). There were no other criteria to select articles for the annotation.

Annotation process. Since our goal was to create a propaganda corpus, we needed to annotate articles not only for specific techniques and narratives, but also to identify and filter out propagandistic content. Therefore, our collected articles underwent a two-phase annotation process, as illustrated in Fig. 1. Five annotators were involved in this annotation process. Annotator 1 participated throughout the entire annotation period, while Annotators 2 and 3 contributed during the earlier stages and were later replaced by Annotators 4 and 5 in the later stages. In total, six distinct annotator pairs were formed: three pairs (comprising Annotators 1, 2, and 3) were active in the earlier phase, and three different pairs (comprising Annotators 1, 4, and 5) annotated the corpus during the later phase.

In the first phase, annotators individually assessed whether articles contained propagandistic content by assigning each article one of four possible labels: “yes,” “no,” “propaganda citation,” or “unclear.” An article was labelled “propaganda” (“yes”) if it contained at least one identified propaganda narrative and employed at least one propaganda technique within the text. An article was labelled “not propaganda” (“no”) if neither a propaganda narrative nor a propaganda technique was identified. Two additional labels were also available: “unclear” and “propaganda citation.” The “unclear” label was assigned if an article contained only propaganda techniques without an identifiable propaganda narrative, or, vice versa, only a propaganda narrative without the presence of any propaganda techniques. The “propaganda citation” label was used when propagandistic statements were directly quoted or cited from other sources. This initial phase primarily served to select articles suitable for subsequent detailed annotation. During this phase, each annotator individually reviewed and labelled approximately 50 articles per week, resulting in around 150 articles annotated weekly by the three annotators. The main objective was to identify articles suitable for detailed cross-annotation, aiming for at least 30 articles positively identified as propagandistic (“yes”) per week. If fewer than 30 articles were labelled “yes” from the initial batch of 150, annotators continued reviewing additional articles until the threshold of 30 “yes” articles was achieved.

In the second phase, each week a batch of 30 articles identified as propagandistic in the first phase was selected for propaganda techniques and narratives annotation. These 30 articles were distributed among the annotators so that each annotator was assigned 20 articles. Each article was annotated independently by two different annotators (ten articles overlapping between annotators). After individual annotations were completed, the annotations for each article were merged, and pairs of annotators held weekly discussions to resolve any conflicts, thereby preparing a finalized, consensus-based annotation for each article.

The entire annotation process was conducted in Label Studio, an open-source data labeling platform. We used its built-in tools for span annotation and article-level labeling, and adapted the interface to align with our annotation methodology. Figure 2 presents a gold-standard annotation example for a single article in Label Studio.

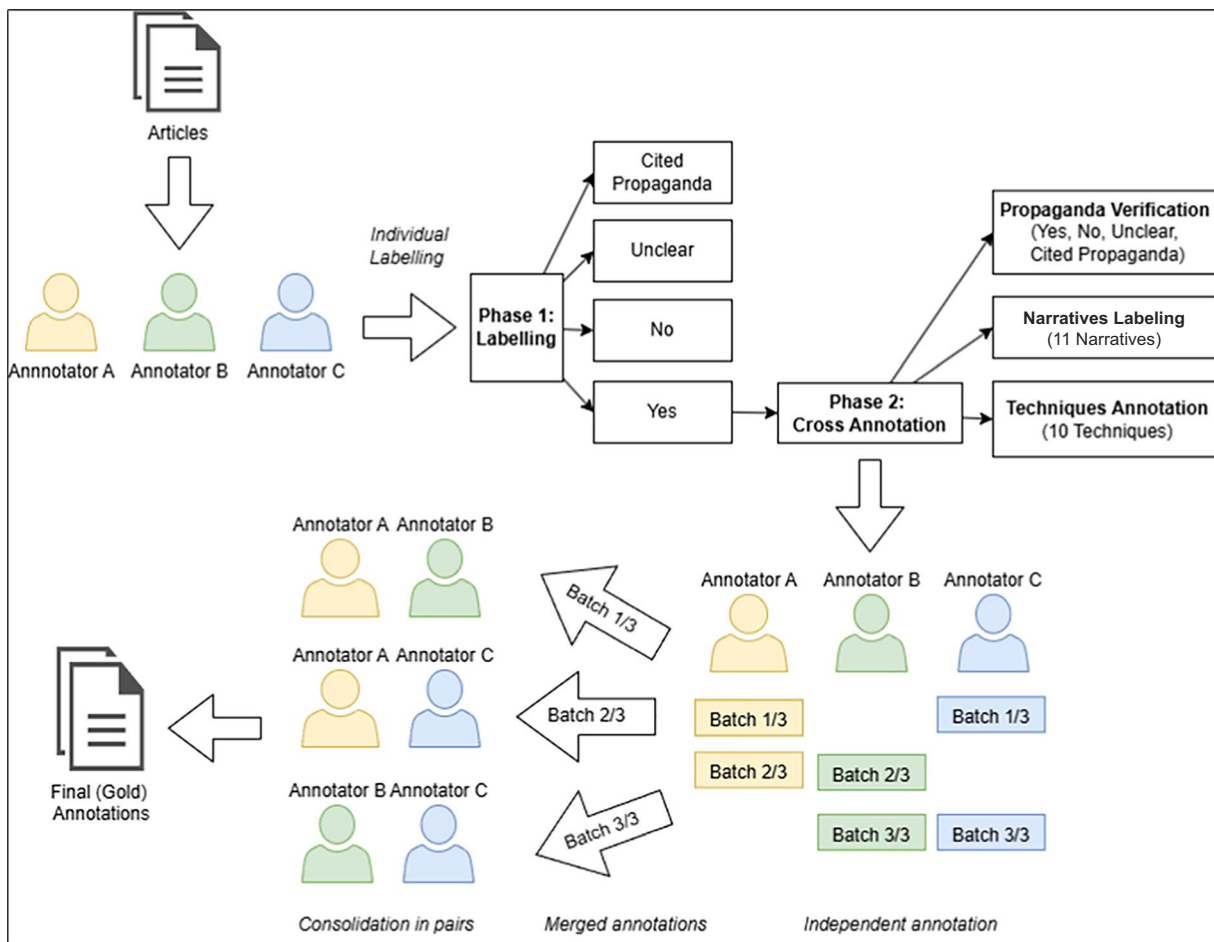


Fig. 1 The annotation process.

Ar šis tekstas yra propaganda?

Ne¹¹

Taip²¹

Propagandos citavimas¹¹

Nejmanoma nustatyti¹¹

Pasirinkite tekste minimas PD temas:

- Pabėgėlių/migrantų krizės eskalavimas¹⁴
- Karo Ukrainoje dezinformacija¹⁴
- Lietuvos respublikos 'projekto' šmeižimas¹⁷
- Lietuvos ginkluotųjų pajėgų finansavimo šmeižimas¹⁷
- Nepasitikėjimas Lietuvos institucijomis¹⁹
- Pasitikėjimo Vakarų institucijomis mažinimas²⁰
- Vakarų civilizacijos pabaigos idėjos skatinimas (vertybių išsigimimas, LGBT...)²⁴
- Efyktyvaus valdymo pavyzdžio (pvz., Minskas, Maskva, Pekinas) skatinimas²⁴
- 'Vašingtono hegemonijos' pabaigos teorija²⁴
- 'Naujosios pasaulio tvarkos' idėjos propagavimas²¹
- 'Litvinizmas'¹⁴

Pažymėkite tekste atpažintas propagandos technikas:

1. Emocinė raiška s
2. Whataboutism_Red Herring_Straw Man d
3. Supaprastinimas f
4. Neapibrėžtumas g
5. Apeliavimas į autoritetą z
6. Mojavimas vėliava x
7. Sekimas iš paskos c
8. Abejojimas v
9. Reductio ad hitlerum b
10. Pakartojimas y

Lietuva vietoj „aklo lojimo“ ant kaimynų, turėtų pasakyti didelį ačiū už mūsų regiono saugumą Baltarusijai. Lietuvos politikai galimai patys neišdrįsta to padaryti, nes trūksta politinės valios, savarankiškumo ir šeiminių (JAV-ČŽV) leidimo.

Baltarusijos sienos apsaugos tarnybos saugo regioninį ir tarptautinį saugumą. Apie tai pranešė Baltarusijos prezidentas Aleksandras Lukašenka.

„Svarbu, kad šandien baltarusių sienos apsaugos tarnybos iš tiesų saugotų ne tik nacionalinius šalies interesus, bet ir regioninį bei tarptautinį saugumą“, – sakė Aleksandras Lukašenka.

„Mūsų kariai sugebėjo gerokai sumažinti nelegalių migrantų į ES valstybės narės sistemingai identifikuoti pasienyje asmenims, kurie tarptautiniu mastu yra paieškomi. Kartu su kitų šalių teisėsaugos institucijomis, reguliariai užkirsti kanalus prekyba narkotikais ..“, – sakė prezidentas.

Aleksandras Lukašenka pareiškė, kad šiandien pasaulyje vis dar nerimą kelia ir nenuspėjama, karinė konfrontacija kuri vyksta beveik visuose žemynuose, įskaitant tiesiogiai prie Baltarusijos sienų.

Fig. 2 Gold-standard span- and article-level annotation for a single example in Label Studio.

Cross-annotation methodology. As one of the primary goals in developing our propaganda corpus was to minimise the risk of annotator bias, for the second phase of the annotation, a cross-annotation method was selected, as this approach ensures that each document is annotated independently by multiple annotators⁴⁸. As was previously mentioned, five different annotators participated in the annotation process, from which six

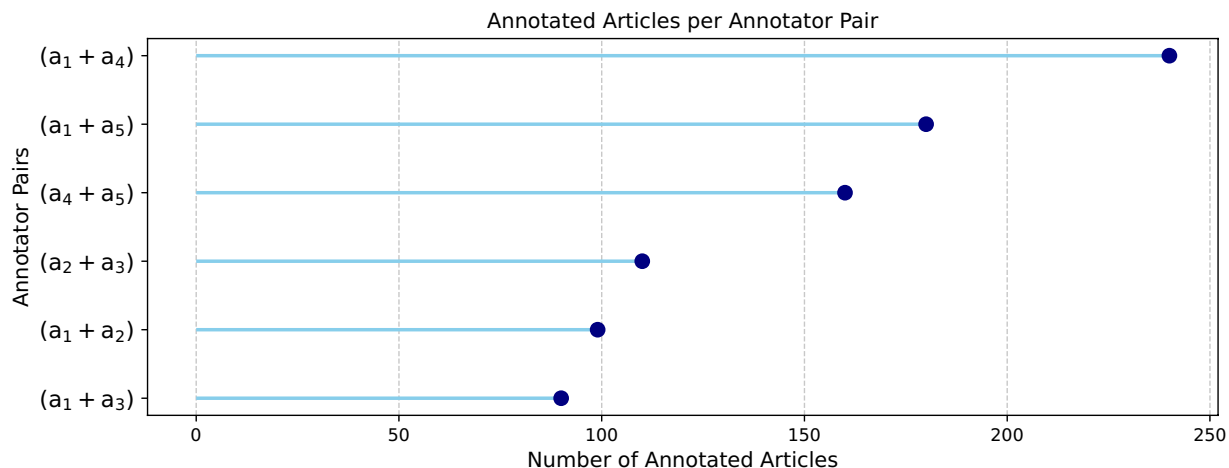


Fig. 3 Annotations distribution among the pairs.

Agreement between annotators			Agreement between the annotator and the final golden annotation			
Annotations of the annotator pair	Spans (γ_s)	+labels (γ_d)	Annotations		Spans (γ_s)	+labels (γ_d)
$(a_1 + a_2)$	0.32	0.22	a_1	c_1	0.52	0.56
			a_2	c_1	0.59	0.69
$(a_1 + a_3)$	0.1	0.07	a_1	c_2	0.61	0.67
			a_3	c_2	0.31	0.43
$(a_2 + a_3)$	0.23	0.19	a_2	c_3	0.69	0.76
			a_3	c_3	0.35	0.46
$(a_1 + a_4)$	0.35	0.25	a_1	c_4	0.64	0.68
			a_4	c_4	0.54	0.6
$(a_1 + a_5)$	0.29	0.16	a_1	c_5	0.67	0.71
			a_5	c_5	0.45	0.53
$(a_4 + a_5)$	0.36	0.25	a_4	c_6	0.64	0.71
			a_5	c_6	0.48	0.55

Table 4. Gamma (γ) inter-annotator agreement between annotators spotting spans alone (spans) and spotting spans + labelling (+labels).

distinct annotator pairs were formed. Below, we provide the formal definition of the cross-annotation process used in our study.

Let $A = \{a_1, a_2, a_3, a_4, a_5\}$ denote the set of five annotators involved in the annotation task. The annotators were systematically grouped into six distinct annotation pairs, defined as follows:

$$P = \{(a_1, a_2), (a_1, a_3), (a_2, a_3), (a_1, a_4), (a_1, a_5), (a_4, a_5)\}$$

Each pair $p_i \in \{1, \dots, 6\}$ annotated a subset of articles. The annotation procedure consisted of two distinct stages:

- **Stage 1 (Independent annotation):** Each annotator within pair $p_i = (a_j, a_k)$ independently annotated the assigned articles, producing individual annotation sets U_j and U_k , respectively.
- **Stage 2 (Consolidation):** Annotators a_j and a_k within each pair p_i convened to reconcile differences from their independent annotations. This process resulted in a final consolidated annotation (golden standard) c_i , where c_i corresponds uniquely to a pair p_i .

Thus, for each annotation pair $p_i \in P$, the final consolidated annotation (gold annotation) is denoted by: c_i for $i \in \{1, 2, 3, 4, 5, 6\}$, where each c_i serves as the reference annotation for evaluating the annotations of annotators a_j and a_k within the respective pair p_i .

Only golden annotations were included in the final propaganda corpus, and, as previously mentioned, these golden annotations were derived from the annotations of two annotators. Figure 3 illustrates the distribution of annotated articles among all annotator pairs.

Technique	Spans (γ_s)	+labels (γ_{sl})
Emotional Expression	0.23	0.23
Whataboutism/Red Herring/Straw Man	0.13	0.14
Simplification	0.21	0.21
Intentional Vagueness (Obfuscation)	0.23	0.21
Appeal To Authority	0.37	0.37
Flag-Waving	0.43	0.42
Bandwagon	0.43	0.38
Doubt/Smears	0.30	0.30
Reduction ad Hitlerum/Stalinum	0.37	0.39
Repetition	0.50	0.49

Table 5. Gamma (γ) inter-annotator agreement between annotators for each technique spotting spans alone (spans) and spotting spans + labelling (+labels).

	a_1	a_2	a_3	a_4	a_5
a_1	—	0.565 ± 0.46	0.3 ± 0.43	0.566 ± 0.46	0.54 ± 0.43
a_2	0.565 ± 0.46	—	0.377 ± 0.45		
a_3	0.3 ± 0.43	0.377 ± 0.45	—		
a_4	0.566 ± 0.46			—	0.556 ± 0.42
a_5	0.54 ± 0.43			0.556 ± 0.42	—

Table 6. Average values and standard deviations of the Jaccard similarity metric between annotators based on narrative labelling.

Data Records

The HALP-PROP corpus 53 is divided into two files:

- Annotations.csv has information about the article and deep annotation information (whether it is propaganda, narratives, techniques) (see details in Table 7)
- Primary_filter.csv contains information about article (heading, content) and an answer to the question “is it propaganda?” (see details in Table 7)

Both CSV files have UTF-8 encoding and semicolon (“;”) delimiters for column separation. For convenience, we included an XLSX version of the files, too.

Technical Validation

Validation of the propaganda technique annotations. To evaluate the quality of the propaganda technique annotations, the γ inter-annotator agreement metric⁴⁹ was employed. This metric is designed for tasks where both spans and their labels, in our case, propaganda techniques, must be identified. Moreover, it can effectively handle overlaps in annotations, which frequently occur in propaganda technique annotation, as multiple different techniques may be assigned to the same textual unit (token, word, sentence, etc). Below, we briefly introduce the methodology and mathematical definitions underlying the γ metric, which was selected to validate the quality of the propaganda technique annotations in our newly created corpus.

Definition of annotation units. An annotation unit u identified by annotator a can be defined as: $u = ([u_{start}, u_{end}], u_{label})$, where u_{start} and u_{end} denote the start and end character offsets of the annotated span within the text, respectively, and u_{label} denotes the specific propaganda technique label assigned to this span.

Calculation of observed disagreement (D_o). Observed disagreement measures how different the annotations provided by annotators are. It is computed by aligning annotation units from two annotators to minimize the total disagreement:

- **Spans-only disagreement (D_o^s):** Positional discrepancies between two spans u^j (annotator a_j) and u^k (annotator a_k) were calculated as follows:

$$d_{pos}(u^j, u^k) = \left(\frac{|u_{start}^j - u_{start}^k| + |u_{end}^j - u_{end}^k|}{(u_{end}^j - u_{start}^j) + (u_{end}^k - u_{start}^k)} \right)^2$$

Given two annotators a_j and a_k , with aligned annotations, the observed disagreement for spans-only is:

File	Field name	Description
Annotations.csv (1000 records)	index	An article number
	heading	Media outlet article headline
	content	Media outlet article content, without HTML formatting in the Lithuanian language
	is_propaganda	Label for the content. It could be one of the following values: <ul style="list-style-type: none"> • “yes” for text marked as propaganda • “no” for text that is not marked as propaganda • “propagandaCitation” for an article that consists only of propaganda citation • “nonDeterminable”/“Unclear” for an article that could not be defined as propaganda or not propaganda
	narratives	Narratives found in the content of the article. A detailed description of each narrative can be found in the section “Selection of Propaganda Narratives”. The content could have a few narratives from the list: <ol style="list-style-type: none"> 1. “disinformationAboutTheWarInUkraine” 2. “delegitimizationOfTheLithuanianState” 3. “underminingTheLithuanianArmedForces” 4. “erosionOfTrustInLithuanianInstitutions” 5. “attacksOnWesternInstitutionsAndAlliances” 6. “declineOfWesternCivilization” 7. “authoritarianModelPromotion” 8. “U.S.DeclineAndWashingtonHegemony” 9. “geopoliticalReorderingAndTheNewWorldOrder” 10. “weaponizationOfMigrationAndRefugees” 11. “revivalOfLitvinism”
	custom_narratives	It could be a list of other narratives in Lithuanian language identified by the annotators
	techniques	The list of techniques found in the article content. Each item on the list has a structure as follows: <ul style="list-style-type: none"> • “start” – the beginning of the technique • “end” – the end of the technique • “technique” – abbreviation of the technique found in Table 2 • “text” – the part of article content corresponding to the technique
Primary_filter.csv (2870 records)	index	An article number
	heading	Media outlet article headline
	content	Media outlet article content, without HTML formatting in the Lithuanian language
	is_propaganda	Has the same values as in Annotations.csv

Table 7. HALP-PROP corpus structure.

$$D_o^s = \frac{1}{|U|} \sum_{(u^j, u^k) \in \text{alignemnt}} d_{pos}(u^j, u^k),$$

where $|U|$ represents the average number of annotated units per annotator, and alignment refers to the optimal pairing of annotation units from annotators a_j and a_k .

- **Span and label agreement (D_o^{sl}):** combines positional discrepancies with a categorical label difference, calculated by the combined dissimilarity d_{comb} :

$$d_{comb}(u^j, u^k) = \alpha \cdot d_{pos}(u^j, u^k) + \beta \cdot d_{cat}(u_{label}^j, u_{label}^k),$$

with weights set as $\alpha = 1$ (for positional discrepancies) and $\beta = 2$ (for categorical discrepancies), where categorical dissimilarity is defined as:

$$d_{cat}(u_{label}^j, u_{label}^k) = \begin{cases} 0, & \text{if } u_{label}^j = u_{label}^k \\ 1, & \text{otherwise} \end{cases}$$

Thus, the observed disagreement combining spans and labels is:

$$D_o^{sl} = \frac{1}{|U|} \sum_{(u^j, u^k) \in \text{alignemnt}} d_{comb}(u^j, u^k)$$

Calculation of expected disagreement (D_e). Expected disagreement (D_e) represents the disagreement that would occur by random chance. It is estimated by generating random annotations (with the same distribution of span lengths and counts as observed annotations) and averaging their disagreement. Expected disagreement for span-only case (D_e^s) and for span-and-label case (D_e^{sl}) were computed by simulation, using random annotation samples $D_e = E[D_e^{random}]$.

Final Gamma Calculation (γ). The final gamma (γ) inter-annotator agreement measure is computed as the normalized difference between observed and expected agreements:

- **Span-only gamma (γ_s):**

$$\gamma_s = 1 - \frac{D_o^s}{D_e^s}$$

- **Span-and-label gamma (γ_{sl}):**

$$\gamma_{sl} = 1 - \frac{D_o^{sl}}{D_e^{sl}}$$

Inter-annotator agreement analysis results. Table 4 presents the results of the gamma (γ) inter-annotator agreement analysis. Span-only agreement between annotators ranged predominantly from 0.10 to 0.36, suggesting considerable disagreement regarding the exact textual boundaries of annotated propaganda instances. When considering both spans and labels, agreement further dropped to a range of 0.07 to 0.25, revealing substantial additional discrepancies in labeling propaganda techniques. However, when comparing individual annotators' results against consolidated gold annotations, agreement notably improved. Span-only agreement with gold annotations ranged from 0.31 to 0.69, demonstrating that consolidation significantly enhanced clarity in identifying spans. Agreement on both spans and labels with gold annotations ranged from 0.43 to 0.76. Considering that the annotations were produced entirely by human annotators, some variability and discrepancies were expected. Similar patterns and agreement levels were observed in the SemEval propaganda techniques corpus⁵⁰, where initial inter-annotator agreement values for spans ranged from 0.30 to 0.34, and for spans and labels from 0.24 to 0.28. Their consolidated agreement with gold annotations ranged from 0.42 to 0.76 for spans only, and from 0.39 to 0.74 when labels were considered as well. The results from our corpus are therefore comparable to the SemEval corpus. Although our initial annotator agreement values exhibit a wider range, this is primarily due to the greater complexity of our annotation setup: five annotators formed six different annotation pairs in our study, whereas the SemEval corpus involved four annotators organized into only two pairs. This more extensive pairing structure in our study contributes to increased diversity in annotations, and further reduces potential annotator bias.

Table 5 presents γ -inter annotator agreement results for each propaganda technique. The highest agreement occurred for “repetition,” “following behind,” and “waving the flag,” as these techniques have clear indicators, like repeated phrases, encouragement of group conformity, and patriotic appeals. The combined technique Whataboutism/Red Herring/Straw Man yielded the lowest inter-annotator agreement, reflecting the difficulty of consistently identifying topic diversion or misrepresentation. It was relatively rare in our corpus (3.74% of spans), so even small inconsistencies could noticeably lower agreement. Moreover, this challenge is not unique to our study: in SemEval-2020 Task 11, participating teams reported among the lowest classification scores for this technique, indicating broader difficulties for automated systems⁵¹. Moderate agreement was observed for “appeal to authority” and “reductio ad Hitlerum.” Although these techniques involve clear identification markers (references to experts or institutions, or associations with Nazism or communism), annotators differed in marking span boundaries. Lower agreement levels appeared in “doubt,” “uncertainty” (or “intentional vagueness”), “emotional expression,” and “simplification.” Despite being individually straightforward to identify, annotators frequently disagreed, due to overlapping characteristics, leading to annotations marked as overlapping. “Uncertainty,” involving vague language, similarly presents identification challenges akin to “whataboutism/red herring/straw man.”

Narratives labelling validation. We have also evaluated a match between annotators $\{a_1, a_2, a_3, a_4, a_5\}$ narratives for each article using the Jaccard similarity metric. The average values of the metric and its standard deviations are presented in Table 6. The results indicate that average pairwise Jaccard similarity scores ranged from 0.30 to 0.57, suggesting a generally moderate level of agreement. However, all annotator pairs exhibited substantial variability across articles, with standard deviations consistently exceeding 0.42. This high variability can be attributed to the multi-label nature of the task, where annotators were permitted to assign multiple narratives to each article. Instances in which one annotator labeled a narrative that another did not initially identify were addressed during the final discussion phase of the cross-annotation methodology. In most cases, annotators reached a consensus to retain multiple narrative labels in the consolidated annotation⁵².

Data availability

The corpus is available on National Open Access Research Data Archive (MIDAS) <https://doi.org/10.18279/MIDAS.260057>.

Code availability

No custom code was used.

Received: 24 July 2025; Accepted: 27 November 2025;

Published online: 03 December 2025

References

- Hashmi, E., Yayilgan, S. Y., Yamin, M. M., Ali, S. & Abomhara, M. Advancing Fake News Detection: Hybrid Deep Learning With FastText and Explainable AI. *IEEE Access* **12**, 44462–44480 (2024).
- Liu, X., Ma, K., Ji, K., Chen, Z. & Yang, B. Graph-based multi-information integration network with external news environment perception for Propaganda detection. *International Journal of Web Information Systems* **20**, 195–212 (2024).
- Horák, A., Sabol, R., Herman, O. & Baisa, V. Recognition of propaganda techniques in newspaper texts: Fusion of content and style analysis. *Expert Syst Appl* **251**, 124085 (2024).
- Kuntur, S., Wróblewska, A., Paprzycki, M. & Ganzha, M. Under the Influence: A Survey of Large Language Models in Fake News Detection. *IEEE Transactions on Artificial Intelligence* **6**, 458–476 (2025).
- Wang, X., Zhang, W. & Rajtmajer, S. Monolingual and Multilingual Misinformation Detection for Low-Resource Languages: A Comprehensive Survey. (2025).
- Pakray, P., Gelbukh, A. & Bandyopadhyay, S. Natural language processing applications for low-resource languages. *Natural Language Processing* **31**, 183–197 (2025).
- Rashkin, H., Choi, E., Jang, J. Y., Volkova, S. & Choi, Y. Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking. in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* 2931–2937, <https://doi.org/10.18653/v1/D17-1317> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2017).
- Barrón-Cedeño, A., Jaradat, I., Da San Martino, G. & Nakov, P. Propopy: Organizing the news based on their propagandistic content. *Inf Process Manag* **56**, 1849–1864 (2019).
- Wang, W. Y. 'Liar, Liar Pants on Fire': A New Benchmark Dataset for Fake News Detection. in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* 422–426, <https://doi.org/10.18653/v1/P17-2067> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2017).
- Horne, B. & Adali, S. This Just In: Fake News Packs A Lot In Title, Uses Simpler, Repetitive Content in Text Body, More Similar To Satire Than Real News. *Proceedings of the International AAAI Conference on Web and Social Media* **11**, 759–766 (2017).
- Shu, K., Mahudeswaran, D., Wang, S., Lee, D. & Liu, H. FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotemporal Information for Studying Fake News on Social Media. *Big Data* **8**, 171–188 (2020).
- Ahmed, H., Traore, I. & Saad, S. Detection of Online Fake News Using N-Gram Analysis and Machine Learning Techniques. in 127–138, https://doi.org/10.1007/978-3-319-69155-8_9 (2017).
- Zubiaga, A., Wong Sak Hoi, G., Liakata, M. & Procter, R. PHEME dataset of rumours and non-rumours. [Dataset]. Preprint at (2016).
- Mitra, T. & Gilbert, E. CREDBANK: A Large-Scale Social Media Corpus With Associated Credibility Annotations. *Proceedings of the International AAAI Conference on Web and Social Media* **9**, 258–267 (2021).
- Ferreira, W. & Vlachos, A. Emergent: a novel data-set for stance classification. in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* 1163–1168, <https://doi.org/10.18653/v1/N16-1138> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2016).
- Thorne, J., Vlachos, A., Christodoulopoulos, C. & Mittal, A. FEVER: a large-scale dataset for Fact Extraction and VERification. (2018).
- Riedel, B., Augenstein, I., Spithourakis, G. P. & Riedel, S. A simple but tough-to-beat baseline for the Fake News Challenge stance detection task. (2017).
- Martino, G. D. S., Barrón-Cedeño, A., Wachsmuth, H., Petrov, R. & Nakov, P. SemEval-2020 Task 11: Detection of Propaganda Techniques in News Articles (2020).
- Piskorski, J., Stefanovitch, N., Da San Martino, G. & Nakov, P. SemEval-2023 Task 3: Detecting the Category, the Framing, and the Persuasion Techniques in Online News in a Multi-lingual Setup. in *Proceedings of the The 17th International Workshop on Semantic Evaluation (SemEval-2023)* 2343–2361, <https://doi.org/10.18653/v1/2023.semeval-1.317> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2023).
- Dimitrov, D. et al. SemEval-2024 Task 4: Multilingual Detection of Persuasion Techniques in Memes. in *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)* 2009–2026, <https://doi.org/10.18653/v1/2024.semeval-1.275> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2024).
- Anton Chernyavskiy, S. et al. 17795–17807 (ELRA and ICCL, Torino, Italia, 2024).
- Alam, F., Mubarak, H., Zaghouani, W., Martino, G. D. S. & Nakov, P. Overview of the WANLP 2022 Shared Task on Propaganda Detection in Arabic. (2022).
- Salman, M. U., Hanif, A., Shehata, S. & Nakov, P. Detecting Propaganda Techniques in Code-Switched Social Media Text. (2023).
- Chang, R.-C., Lai, C.-M., Chang, K.-L. & Lin, C.-H. Dataset of Propaganda Techniques of the State-Sponsored Information Operation of the People's Republic of China. (2021).
- Chaudhari, D. & Pawar, A. V. Barrón-Cedeño, A. H-Prop and H-Prop-News: Computational Propaganda Datasets in Hindi. *Data (Basel)* **7**, 29 (2022).
- Posadas-Durán, J.-P., Gómez-Adorno, H., Sidorov, G. & Escobar, J. J. M. Detection of fake news in a new corpus for the Spanish language. *Journal of Intelligent & Fuzzy Systems* **36**, 4869–4876 (2019).
- Hu, Y., Huang, H., Chen, A. & Mao, X.-L. Weibo-COV: A Large-Scale COVID-19 Social Media Dataset from Weibo. (2020).
- Farooq, M. S., Naseem, A., Rustam, F. & Ashraf, I. Fake news detection in Urdu language using machine learning. *PeerJ Comput Sci* **9**, e1353 (2023).
- Götz, E. Near Abroad: Russia's Role in Post-Soviet Eurasia. *Eur Asia Stud* **74**, 1529–1550 (2022).
- Winnerstig, M. *Tools of Destabilization. Russian Soft Power and Non-Military Influence in the Baltic States*. <https://www.foi.se/en/foi/reports/report-summary.html?reportNo=FOI-R-3990-SE> (2014).
- Pelnēns, G. et al. *The 'Humanitarian Dimension' of Russian Foreign Policy toward Georgia, Moldova, Ukraine, and the Baltic States*. (Centre for East European Policy Studies, 2010).
- Nerijus Maliukevičius. *Rusijos informacijos geopolitikos potencialas ir sklaida Lietuvoje*. (Vilniaus universiteto leidykla, Vilnius, 2008).
- Mantas Martišius. *(Ne)Akivaizdus Karas: Nagrinėjant Informaciją Karą*. (Versus Aureus, 2010).
- Denisenko, V. *Propagandos apsupty*. (Vilniaus universiteto leidykla, Vilnius, 2021).
- Oxford Learner's Dictionary. Propaganda. <https://www.oxfordlearnersdictionaries.com/us/definition/english/propaganda> (2025).
- Jowett, G. S. & O'Donnell, V. *Propaganda & Persuasion*. (SAGE, 2018).
- Stanley, J. *How Propaganda Works*. (Princeton University Press, 2016).
- Ellul, J. *Propaganda: The Formation of Men's Attitudes*. (Vintage, 2021).
- Riessman, C. K. Narrative Analysis. in *Narrative, Memory & Everyday Life* 1–7 (University of Huddersfield, Huddersfield, 2005).
- Patterson, M. & Monroe, K. R. Narrative in political science. *Annual Review of Political Science* **1**, 315–331 (1998).
- Bērziņa I., Cepurītis M., Kaljula D. & Juurvee I. *Russia's Footprint in the Nordic-Baltic Information Environment 2016/2017*. <https://stratcomcoe.org/publications/russias-footprint-in-the-nordic-baltic-information-environment-20162017/138> (2018).
- Cepurītis, M. et al. *Russia's Footprint in the Nordic-Baltic Information Environment. Report 2019/2020*. (NATO Strategic Communications Centre of Excellence, 2020).
- National Threat Assessment*. https://www.vsd.lt/wp-content/uploads/2023/03/ENG-2023-Gresmes-ENG-el_-1.pdf (2023).
- National Threat Assessment*. <https://www.vsd.lt/wp-content/uploads/2025/03/2024-03-07-VGV-EN.pdf> (2024).

45. *National Threat Assessment*. <https://www.vsd.lt/en/reports/national-threat-assessment-2025/> (2025).
46. Hyndle-Hussein, J. & Klysiński, K. A difficult legacy. Tensions over how to interpret the shared past of Lithuanian and Belarusian peoples. <https://www.osw.waw.pl/en/publikacje/osw-commentary/2024-08-07/a-difficult-legacy-tensions-over-how-to-interpret-shared-past> (2024).
47. LRT tyrimai. LRT tyrimas. Lietuvos “penktoji kolona”: Rusijos propagandą platina šeimos gynėjai, sektos ir knygu apie Staliną leidėjai. <https://www.lrt.lt/naujienos/lrt-tyrimai/5/1700792/lrt-tyrimas-lietuvas-penktoji-kolona-rusijos-propaganda-platina-seimos-gynejai-sektos-ir-knygu-apie-stalina-leidejai> (2022).
48. Stollenwerk, F. *et al.* Text Annotation Handbook: A Practical Guide for Machine Learning Projects. (2023).
49. Mathet, Y., Widlöcher, A. & Métivier, J.-P. The Unified and Holistic Method Gamma (γ) for Inter-Annotator Agreement Measure and Alignment. *Computational Linguistics* **41**, 437–479 (2015).
50. Da San Martino, G., Yu, S., Barrón-Cedeño, A., Petrov, R. & Nakov, P. Fine-Grained Analysis of Propaganda in News Article. in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* 5635–5645, <https://doi.org/10.18653/v1/D19-1565> (Association for Computational Linguistics, Stroudsburg, PA, USA, 2019).
51. Da San Martino, G., Barrón-Cedeño, A., Wachsmuth, H., Petrov, R. & Nakov, P. SemEval-2020 Task 11: Detection of Propaganda Techniques in News Articles. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, 1377–1414, Barcelona (online). International Committee for Computational Linguistics (2020).
52. HALT-PROP: Human-Annotated Lithuanian Textual Corpus for Propaganda Topics and Techniques. MIDAS <https://doi.org/10.18279/MIDAS.260057> (2025).

Acknowledgements

We are grateful for the financial support provided by the Lithuanian Government Priority Research Program “Building Societal Resilience and Crisis Management in the Context of Contemporary Geopolitical Developments” (implemented through the Lithuania Research Council) under grant number S-VIS-23-8. Project title: ‘Propaganda and Disinformation Research: Machine Learning Based Automatic Detection, Impact and Societal Resilience’. Special thanks to Juratė Ruzaite and Mantas Martišius for initial annotations, topic and techniques identification. Thank you to Paulius Zaranka for contributing to the implementation and support of the environment for annotations and partial data gathering. Special thanks to the project manager Darius Plikynas for securing funds for the project.

Author contributions

Background and Summary: Related Work I.E., Motivation N.M. (with I.E.). Methods: Preparation Process, Data Collection Process & Data Selection Process V.M., Propaganda Definition I.E., Definitions of Propaganda Techniques V.Z., Selection of Propaganda Narratives N.M., Annotation Process I.R. Technical Validation: Validation of Propaganda Techniques annotation I.R., Narratives Labelling Validation V.M. (with I.E.). Data Records: V.M. I.E., V.Z. and N.M. were the main annotators of the propaganda texts. V.M. was responsible for the development of the annotation methodology, data gathering (with I.E.) and checking. All the authors contributed to the improvement of the manuscript and approved the final version.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to I.R.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025