Fine-grained and Dense Annotation of Czech Propaganda Using Large Language Models

Radoslav Sabol and Aleš Horák

NLP Centre, Faculty of Informatics, Masaryk University
Brno, Czechia,
{xsabol, hales}@fi.muni.cz

Abstract. Within the previous project of Czech Propaganda Detection that aimed to recognize manipulative techniques in Czech news articles, the annotation was mostly limited to indicating the presence of a technique in a document. In about 35% of the documents, span-level evidence was also annotated, but only as a support for the document-level labels, resulting in a sparse coverage of the techniques. Thus, the resulting dataset has limitations for training and evaluating more fine-grained propaganda detection models.
In this study, we examine the potential of large language models (LLMs) to generate dense span-level annotations of manipulative techniques in Czech news articles. We designed generation prompts tailored to each technique and experimented with several LLMs to produce annotations for a subset of the Czech Propaganda dataset.
We present the details of the generation process, including the design of the prompts and the selection of models. We also evaluate the generated annotations both quantitatively and qualitatively, including a manual validation and comparison with human annotations.

Keywords: manipulative techniques; propaganda; detection; annotation; large language models; LLM; Czech

1 Introduction

The widespread use of online news and social media has fundamentally transformed the ways people access and interpret information. While this transformation has democratized the creation and distribution of content, it has also facilitated the spread of manipulative and misleading narratives. In the current world-wide situation marked by geopolitical tensions and hybrid warfare, propaganda and disinformation have become pervasive tools of influence. Detecting and countering such manipulative techniques in texts is therefore not only a research challenge but also a societal necessity, with direct implications for media literacy, democratic resilience, and public trust.

Research in computational propaganda detection seeks to automatically identify rhetorical and psychological techniques used to influence audiences – such as fear appeal, name calling, or appeals to authority. A growing number

A. Horák, P. Rychlý, A. Rambousek (eds.): Proceedings of Recent Advances in Slavonic Natural Language Processing, RASLAN 2025, pp. 85–98, 2025.

of datasets and shared tasks have focused on detecting these phenomena in English and other major languages, supporting the development of machine learning models for propaganda detection at the document, sentence, or span level. However, linguistic and cultural differences make it crucial to extend these efforts to smaller and morphologically rich languages, such as Czech, where resources and annotated data remain scarce.

The Czech Propaganda Detection project [5,7] was launched to address this gap by developing a dataset and methodology for identifying manipulative techniques in Czech news articles. The project involved manual annotation of thousands of texts by trained annotators, aiming to capture a diverse set of propaganda strategies and to make the resulting corpus suitable for both research and practical applications. The dataset has already supported several experiments in supervised and semi-supervised detection of propaganda, but its granularity has so far been limited primarily to document-level labels, with only partial span-level evidence available.

In this study, we build upon this foundation and explore the potential of large language models (LLMs) to enrich the dataset with fine-grained and dense span-level annotations. The Propaganda Dataset annotation already included the possibility to mark specific text spans, called *ranges*, as a support for the document-level attribute related to a specific manipulative technique. But this capability was used mainly as an example and evidence of the final annotation. As a result, usually one positive example of e.g. *labelling* or *demonization* is marked in each positively evaluated article but other occurrences of the technique remain unnoticed. We hypothesize that, with carefully designed prompts and model selection, LLMs can identify the manipulative techniques at the span level and thus substantially improve the coverage and utility of existing annotations.

We present a systematic approach to generating these dense annotations, including the design of technique-specific prompts, the choice of suitable LLMs, and the evaluation of their outputs. Our analysis combines quantitative metrics with qualitative inspection and manual validation, providing insights into the reliability and consistency of LLM-generated annotations. The results highlight both the promise and the current limitations of using generative models for fine-grained propaganda detection in Czech and offer directions for future work in combining human expertise with LLM-based assistance.

2 Related Work

In our previous work [7], we have evaluated both deep-learning and traditional document classification methods for 17 techniques annotated in the dataset. We have demonstrated that stylistic cues improve detection performance for many techniques and established strong baselines along with detailed inter-annotator agreement analyses. The average weighted F1-scores of individual techniques ranged from 71% (for *argumentation*) to 96% (for *demonization*). Our planned future work aims at detecting the techniques precisely at the span-level, however, we faced the above mentioned limitations of the Propaganda Dataset. The presented work enables the span annotations to be as dense in the text as possible and thus offers new data for improving the fine-grained span-level manipulation detection.

Abraham et al [1] examine how prompt design affects the accuracy of LLM-based text annotation in social science tasks. They show that annotation quality varies substantially depending on prompt formulation, and propose an automatic prompt optimization method to systematically improve performance. Their browser-based tool illustrates how careful prompt engineering can enhance LLM reliability in labeling tasks.

Hasanain et al [6] investigate both encoder models (AraBERT) and generative models (GPT-4) for fine-grained propaganda span extraction, focusing on Arabic texts. They show that GPT-4 can outperform non-expert annotators, act as an effective consolidator of human labels, and generate annotations that train smaller models to state-of-the-art performance. The work demonstrates that LLMs can be used to cost-effectively expand span-level propaganda datasets.

Ma et al [8] study conversational mental manipulation and propose Intent-Aware Prompting (IAP) to model underlying speaker intentions. Their approach improves detection accuracy, particularly by reducing false negatives in subtle manipulative exchanges. Although targeting a different domain, the work highlights the value of prompt-guided reasoning in capturing nuanced manipulative tactics.

Sprenkamp et al [11] utilized five variations of GPT-3 and GPT-4 in order to explore the effectiveness of multi-label classification on the SemEval-2020 Task 11 dataset. The presented results of GPT-4 are comparable to the state of the art and emphasize both its strengths and its limitations in handling context-dependent propaganda techniques. However, the task was formulated as a multi-label text classification task, not span classification.

3 The Propaganda Dataset

In the following investigations, we use a subset of the Czech Propaganda dataset [7]. For our experiments, we split over 8,000 documents into paragraphs and isolated the paragraphs with relevant range annotations.

Out of 18 annotated document attributes, we have selected the ones that (i) can be considered as manipulative techniques (devices) and (ii) are local attributes that can be expressed using span annotations. These restrictions leave us with ten devices with the following definitions:

– **blaming (vina):** Attributing responsibility for a negative situation to a specific person, group, or event. It must be clear who is to blame and what he is being blamed for.
– **labeling (nálepkování):** Using evaluative, emotionally charged terms with either negative or positive connotations that are intended to elicit a short-sighted attitude without explanation. The label is not a factual description but rather reflects the author's attitude.

- **demonization** (**démonizace**): The entity is depicted in extremely negative terms – mainly as harmful, malicious, irrational, and inhuman. In contrast to standard labeling, demonization emphasizes dehumanization, threat, madness, etc.
- **relativization** (**relativizace**)**:** Belittling the negative actions of some party by comparing them to other negative examples. Serves to divert attention, or to justify the negative deed as necessary. Relativization often takes the form of a suggestive question, creating a diversion.
- **fear-mongering** (**strach**)**:** The statement draws attention to threats, disasters, and a negative foresight into the future. The goal is to evoke fear of future developments. The statements are typically hypothetical and lack further evidence.
- **emotions** (**emoce**)**:** The passage is phrased to evoke a strong emotional response – regret, anger, indignation, or compassion. In contrast with fear-mongering, emotions focus mainly on the present or past events.
- **fabulation** (**fabulace**)**:** Spreading unsubstantiated claims, assumptions, exaggerations, or fabrications. Lacking evidence or a verifiable source.
- **opinion** (**názor**)**:** A strong personal opinion of the author of the article – either directly expressed or indirectly (e.g., through a leading question). Citations of foreign opinions do not fall under this category, unless the author embraces them.
- **source** (**zdroj**): The author refers to another information source – medium, website, document, network, or institution. The range should contain both the name of the source and evidence that the author bases his assumptions on the source.
- **expert** (**odbornik**): The statement is supported by the authority of a specific professional (name, title, profession). The true level of expertise in the given topic is irrelevant. The span structure is otherwise similar to the *source* device.
- **argumentation** (**argumentace**): The author tries to persuade the reader using causal connections or hypothetical scenarios. Argumentation can be logical or emotional.

An overview of examples for each device can be seen in Table 1. For each of the manipulative techniques, we have selected 20 positive and 10 negative paragraphs from the development set for annotation using LLMs. This amount was selected to be feasible for manual evaluation by two annotators.

4 Annotation using Large Language Models

In the following experiment, the LLM is tasked with the following problem: given a paragraph of newspaper text and one of the ten devices, the model must return the original text enriched with extra information about the presence of the given manipulative technique.

We consider two open-source models: GPT OSS 120b [9] and DeepSeek R1 685b [3]. The choice of model was based on decent performance in the Czech language [4], public availability, and its pre-deployment on `e-infra` services.

Table 1: Example annotations for attributes at the span level. Range annotations are highlighted in **bold** (see [10] for further details)

Technique	Example
CS: Argumentace **EN:** Argumentation	byla příliš vylekaná, a proto se neobrátila na policii (she) was too scared, and that's why she didn't contact the police
CS: Vina **EN:** Blaming	tento časopis je pornografický a ponižuje ženy. the magazine is pornographic and it degrades women
CS: Démonizace **EN:** Demonization	Necitlivá, sebestředná a megalomanská práce skladatele Insensitive, self-centered and megalomaniacal work of the composer
CS: Emoce **EN:** Emotions	Former Senator Hlavatý's actions are a fraud on voters The course of action taken by the former senator Hlavatý
CS: Fabulace **EN:** Fabulation	vyloučili ji poté, co přijala rozkaz z Bruselu she was ousted after taking orders from Brussel
CS: Strach **EN:** Fear-Mongering	pokud dojde k dalšímu útoku tohoto druhu, následky budou těžké if another assault of this kind occurs, the consequences will be severe
CS: Nálepkování **EN**: Labeling	nepřemožitelná jaderná hlavice an unstoppable nuclear warhead
CS: Relativizace **EN:** Relativization	každým zlehčováním skutečnosti se společnost ještě více radikalizuje With every trivialization of reality, society becomes even more radicalized
CS: Odborník **EN:** Expert	Sociolog Jan Hartl upozorňuje... Sociologist Jan Hartl points out...
CS: Názor **EN:** Opinion	Nebo je nutné se bát toho, kdo „Topola" do boje vyslal? Or is it necessary to fear the one who sent "Topol" into battle?
CS: Zdroj **EN:** Source	Britské BBC přineslo informace... The British BBC brought information...

```
Input:
    "říká v rozhovoru pro ParlamentníListy.cz europoslanec Tomáš
    Zdechovský, který se tímto tématem několik let zabývá"
Output:
    "<zdroj> říká v rozhovoru pro ParlamentníListy.cz </zdroj>
    europoslanec Tomáš Zdechovský, který se tímto tématem
    několik let zabývá"
```

Fig. 1: An example of input and the expected output for the *source* device

The prompt design was derived from the original definitions of manipulative techniques that were preserved after the human annotation of the Propaganda dataset. These were passed to ChatGPT to guide the formulation of the prompt and to help identify any ambiguities present in the original definitions, as well as to provide some valuable examples that would better guide the LLM. The prompt is heavily templated to focus on a single instance of a manipulative technique at a time due to the overall difficulty of the task. This decision was also guided by our previous observations, which showed that binary, single-label models performed more optimally compared to general, multi-label ones. The prompt, along with an input-output example, can be found in Appendix A.

4.1 Human Evaluation of Generated Annotations

After generating denser annotations for ten manipulative devices using GPT OSS and DeepSeek R1, we have assigned two annotators to examine the correctness of both the original and AI-generated ranges. For the following purpose, we have created a web application that segments the paragraphs into further categories:

- unannotated spans of text (unhighlighted)
- human-annotated segment (highlighted in yellow)
- AI-annotated segment (highlighted in blue)
- overlapping segment annotated by both human and AI (highlighted in red)

For each of the segments, except unannotated spans, each annotator has the option to select one of the following options:

- **Unchecked**: the default state that indicates the need for checking the annotation
- **OK**: the given segment complies with the definition of manipulative technique
- **Not OK**: the given segment cannot be a part of any range annotation due to the lack of compliance with the given technique
- **Too short**: the segment somewhat captures the device, but needs to be longer to be fully in line with the definition
- **Too long**: the segment contains the relevant passage of the manipulative technique, but can be shortened without losing any relevant information
- **Unknown**: the annotator admits his lack of expertise to properly evaluate the given segment, reserved mostly for corner cases

The example of a paragraph with highlighted annotations can be seen in Figure 2. Both annotators have evaluated the ranges of both GPT OSS and DeepSeek R1; however, to have a definitive conclusion, the segments need to be merged back into ranges. To aggregate the partial results, we had to apply the following rules:

1. if either of the segments is labeled as **Not OK**, the entire annotation is **Not OK**
2. if all the segments of an annotation are labeled as **Too short**, the entire annotation is **Not OK**
3. if any segment is **Unknown**, the annotation is omitted from the evaluation
4. all other annotations are labeled as **OK**

5 Results

In this section, we present the evaluation of both generated and pre-existing range annotations. First, we present the inter-annotator agreement for both models between the two annotators. We will show the results of the manual evaluation afterwards.

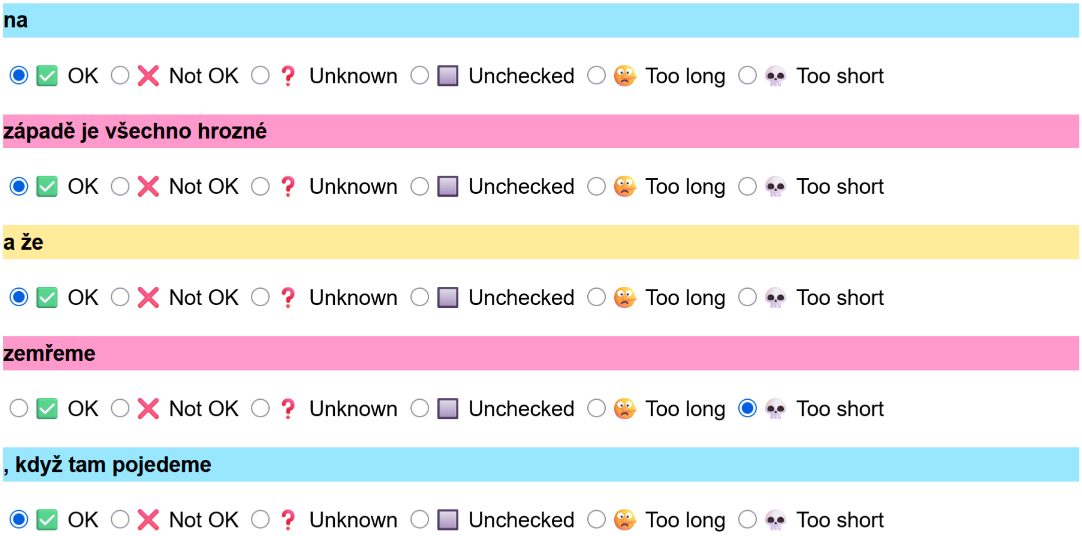

Fig. 2: An example of annotation inside of the web editor. It displays three annotations, two generated by AI, one labeled by human annotator, and two overlaps which are displayed by the red outline.

5.1 Inter-Annotator Agreement

After the human evaluation, we measured Kohen's Kappa [2] metric for each manipulative technique, which is presented in Table 2. Regardless of the used LLM, the scores reflect agreement ranging from slight to strong between the annotators. Lower scores usually reflect the overall difficulty of identifying the technique, such as *argumentation*, where the original definition can have broad ranges of interpretation and is generally difficult to agree on. Most disagreements on the *opinion* device were based on whether all suggestive questions are the author's opinion or not, which is yet to be discussed in future work. In the case of *demonization*, it is sometimes difficult to determine whether the range is indeed demonization, or just very expressive *labelling*. The highest scores are usually attributed to techniques that are generally not difficult to identify, like *expert*, *source*, and *fear-mongering*. To compare the two models, it was moderately easier to agree on the DeepSeek R1 annotations than on the GPT-OSS ones.

5.2 Comparison of Human and AI-generated Annotations

We have collected both **OK** and **not OK** labels across both annotators and summed them up in Table 3. Generally, the percentage of OK annotations is higher for human-annotated ranges. In some cases, namely *emotions*, *opinion*, *fabulation*, and *source*, the ratio of measured **OK** annotations is better than with the expert-annotated ones. The DeepSeek R1 annotations fall behind in quality due to the LLM occasionally altering the format of the output (e.g., extra whitespace characters, unnecessary explanations), which leads to annotations

Technique	DeepSeek R1	GPT OSS
Relativization	33.52	27.97
Labeling	40.98	51.38
Blaming	42.11	31.18
Source	42.86	32.04
Argumentation	22.70	24.75
Expert	51.49	40.95
Fear	62.37	57.91
Opinion	28.57	23.38
Demonization	31.45	23.76
Emotion	38.84	40.16
Fabulation	32.86	24.98
Mean	42.75	37.84

Table 2: Kohen's kappa scores for each manipulative technique

being labeled as **not OK** more often. Further tuning of the prompt specifically for DeepSeek R1 will be necessary in the future. Until this downside is resolved, GPT OSS is a better alternative for annotating manipulative techniques.

Table 3: Combined AI/Human Status Counts by Attribute

	GPT OSS			DeepSeek R1			Human		
Device	not OK	OK	OK %	not OK	OK	OK %	not OK	OK	OK %
argumentation	8	38	82.61	11	27	71.05	5	35	**87.5**
demonization	10	36	78.26	13	43	76.79	6	38	**86.36**
emotions	8	36	**81.82**	11	31	73.81	11	33	75.00
fabulation	3	37	**92.50**	34	16	32.00	4	37	90.24
labelling	12	44	78.57	24	50	67.57	7	35	**83.33**
opinion	2	30	**93.75**	4	18	81.82	3	40	93.02
expert	10	32	76.19	23	21	52.27	6	36	**85.71**
relativization	12	32	72.73	14	24	63.16	9	37	**80.42**
fear-mongering	4	40	90.91	12	36	75.00	2	42	**95.45**
blaming	5	39	88.63	20	28	58.33	4	36	**90.00**
source	9	39	**81.25**	10	38	79.17	10	30	75.00

The amount of overlapping and non-overlapping annotations can be seen in Table 4. From these observations, we can see that AI can cover most of the already existing human annotations, while also annotating new, previously missed occurrences of manipulative techniques. The addition of new annotations is an important property which can allow us to resolve range sparsity in the future.

Table 4: An overview of overlaps between human and generated annotations. **H** stands for non-overlapping human annotations. **AI** stands for non-overlapping AI annotations. **O** stands for overlap of AI and human.

	GPT OSS			DeepSeek R1		
Device	O	AI	H	O	AI	H
argumentation	26	10	7	16	11	12
demonization	35	5	5	33	12	5
emotions	21	11	12	23	13	10
fabulation	16	12	12	24	13	9
labelling	30	12	7	35	19	4
opinion	19	5	11	8	7	16
expert	32	5	5	34	5	4
relativization	16	14	15	21	9	10
fear-mongering	33	5	6	33	7	6
blaming	30	6	6	31	8	5
source	36	6	2	36	6	2

6 Conclusions and Future Directions

In the presented work, we have explored the use of large language models to generate fine-grained and dense span-level annotations of manipulative techniques in Czech news articles. By designing technique-specific prompts and selecting suitable LLMs, we were able to produce annotations that significantly enriched the existing Czech Propaganda dataset.

Our experiments demonstrate that contemporary large language models are capable of producing fine-grained span annotations of manipulative techniques in Czech news articles with a level of consistency that usefully complements existing human-annotated data. Despite differences among models and techniques, all evaluated systems were able to identify a substantial proportion of relevant spans mostly aligned with the definitions used in the Czech Propaganda dataset. These results show that LLMs can effectively mitigate the sparsity of the original corpus, filling in previously unannotated evidence and thus enabling more detailed analyses of propaganda strategies at the paragraph and span level.

The enriched annotations produced in this study form a valuable foundation for training next-generation detection models that require dense supervision, such as sequence-labeling architectures or span-classification frameworks. At the same time, our evaluation highlights several challenges, including occasional over-annotation, inconsistent span boundaries, and technique confusions. These findings point to opportunities for further refinement. Future work will explore strategies for improving annotation reliability, such as iterative human-in-the-loop validation, model agreement filtering, and cross-lingual prompting. We also plan to extend the annotation process to the full dataset and investigate how LLM-augmented supervision influences downstream model performance, thus contributing to more robust tools for analyzing manipulative content in Czech texts.

Acknowledgements. The described study comes from the project "On our own: Opportunities and Risks in the Individualization of Society (PRINS) CZ.02.01.01/00/23_025/0008710", which is co-financed by the European Union.

Computational resources were also provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

References

1. Abraham, L., Arnal, C., Marie, A.: Prompt selection matters: enhancing text annotations for social sciences with large language models. Journal of Computational Social Science **8**(3), 73 (2025)
2. Cohen, J.: A coefficient of agreement for nominal scales. Educational and Psychological Measurement **20**(1), 37–46 (1960). https://doi.org/10.1177/001316446002000104, https://doi.org/10.1177/001316446002000104
3. DeepSeek-AI: Deepseek-v3 technical report (2024), https://arxiv.org/abs/2412.19437
4. Fajcik, M., Docekal, M., Dolezal, J., Ondrej, K., Beneš, K., Kapsa, J., Smrz, P., Polok, A., Hradis, M., Neverilova, Z., Horak, A., et al.: Benczechmark: A czech-centric multitask and multimetric benchmark for large language models with duel scoring mechanism. Transactions of the Association for Computational Linguistics **13**, 1068–1095 (2025)
5. Gregor, M., Mlejnková, P., Pavlíková, M., Šenkýřová, B., Drmola, J., Mareš, M., Kasl, F., Horák, A., Baisa, V., Polčák, R., Hanzelka, J., Syrovátka, J., Herman, O.: Challenging Online Propaganda and Disinformation in the 21st Century. Palgrave Macmillan, Cham (2021). https://doi.org/https://doi.org/10.1007/978-3-030-58624-9
6. Hasanain, M., Ahmad, F., Alam, F.: Large Language Models for Propaganda Span Annotation. In: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (eds.) Findings of the ACL: EMNLP 2024. pp. 14522–14532. Association for Computational Linguistics, Miami, Florida, USA (Nov 2024). https://doi.org/10.18653/v1/2024.findings-emnlp.850
7. Horák, A., Sabol, R., Herman, O., Baisa, V.: Recognition of Propaganda Techniques in Newspaper Texts: Fusion of Content and Style Analysis. Expert Systems with Applications **251**, 124085 (2024). https://doi.org/https://doi.org/10.1016/j.eswa.2024.124085
8. Ma, J., Na, H., Wang, Z., Hua, Y., Liu, Y., Wang, W., Chen, L.: Detecting conversational mental manipulation with intent-aware prompting. In: Rambow, O., Wanner, L., Apidianaki, M., Al-Khalifa, H., Eugenio, B.D., Schockaert, S. (eds.) Proceedings of the 31st International Conference on Computational Linguistics. pp. 9176–9183. Association for Computational Linguistics, Abu Dhabi, UAE (Jan 2025), https://aclanthology.org/2025.coling-main.616/
9. OpenAI: gpt-oss-120b & gpt-oss-20b model card (2025), https://arxiv.org/abs/2508.10925
10. Sabol, R.: Propaganda Detection using Stylometric Text Analysis. Master thesis, Masaryk University, Faculty of Informatics, Brno (2023)
11. Sprenkamp, K., Jones, D.G., Zavolokina, L.: Large Language Models for Propaganda Detection. arXiv preprint arXiv:2310.06422 (2023)

Appendix A

A.1 Prompt Template

```
Jsi jazykový model určený k automatické anotaci manipulačních
technik v textu.

## Úkol
- Najdi v textu všechny výskyty definovaných manipulačních technik.
- Označ je přímo **XML značkami**, bez dodatečného zdůvodnění.
- Použij jen následující XML značky:

{ANNOT_TAGS}

## Podrobné definice technik

{TECHNIQUE_DEFINITIONS}

## Výstup:
- Použij jen XML značky přímo v textu.
- Zachovej původní text a interpunkci,
  jen vlož značky okolo relevantních pasáží.

{N-SHOT EXAMPLES}

## Vstupní text:
"{EXAMPLE}"

## VÝSTUP:
```

A.2 Technique Definitions

Argumentation

```
<argumentace>…</argumentace>

**argumentace**
- Výrok působí jako **přesvědčovací argument** (logický nebo emotivní).
- Nestačí jen popis skutečnosti - musí být cílem ovlivnit čtenáře.
- Příklady: „Kdybychom nezasáhli, dopadli bychom jako Ukrajina."
```

Demonization

```
<demonizace>…</demonizace>

**demonizace**
```

```
- Nepřítel je vykreslen **extrémně negativně**, téměř nelidsky.
- Přisuzují se destruktivní úmysly, zákeřnost, zlo.
- Často silná adjektiva nebo hyperboly.
- Příklady: „Putin chce ovládnout celý svět."
            „Neštítí se ničeho, ani zabíjení dětí."
```

Emotions

```
<emoce>…</emoce>

**emoce**
- Text vyvolává **silnou emocionální reakci u čtenáře**
(strach, rozhořčení, nenávist, úzkost, lítost, vztek).
- Nehodnotí se emoce autora, pouze účinek na čtenáře.
- Často metafory, přehnané popisy nebo emotivní výrazy.
- Příklady: „Západní hyeny čekají na naši slabost."
            „Jsme jen loutky v rukou mocných."
```

Fabulation

```
<fabulace>…</fabulace>

**fabulace**
- Přehánění, nepodložená tvrzení, fámy.
- Často nereálné nebo extrémní tvrzení bez důkazu.
- Příklady: „Miliony uprchlíků už míří do Evropy."
            „Brzy bude chleba stát 100 Kč."
```

Labelling

```
<nalepkovani>…</nalepkovani>

**nalepkovani**
- Krátká, úderná slova nebo spojení, shrnující postoj k osobě/skupině.
- Může být **negativní** nebo **pozitivní**.
- Typicky přídavná jména nebo jednoduché fráze, snadno zapamatovatelné.
- Příklady: „zrádci národa"
            „nezničitelná jaderná hlavice"
```

Opinion

```
<nazor>…</nazor>

**nazor**
- **Osobní názor autora článku**, nikoli citované osoby či jiné entity.
```

```
- Může být vyjádřen přímo („já si myslím…") nebo nepřímo sugestivní otázkou.
- Příklady: „Já si myslím, že vláda zradila národ."
            „Opravdu si někdo myslí, že tohle je náhoda?"
- Názor citované osoby není označován jako `nazor`.
```

Expert

```
<odbornik>…</odbornik>

**odbornik**
- Výrok je podpořen **jmenovaným odborníkem**, ne pouze institucí.
- Může jít i o zdánlivého experta.
- Příklady: „Podle odborníka Petra Nováka…"
            „Jak uvedl analytik Jan Krátký…"
```

Relativization

```
<relativizace>…</relativizace>

**relativizace**
- Zlehčování negativního chování určité entity nebo glorifikace opačné strany.
- Často formou porovnání s jinými případy.
- Příklady: „Ano, Rusko napadlo Ukrajinu, ale Západ také vedl války."
            „A co bombardování Jugoslávie?"
```

Fear-Mongering

```
<strach>…</strach>

**strach**
- Výrazy, které apelují na **budoucí ohrožení, nejistotu nebo kolaps**.
- Cílem je vyvolat pocit hrozby.
- Příklady: „Kvůli sankcím přijdeme o práci."
            „Hrozí kolaps zdravotnictví."
```

Blaming

```
<vina>…</vina>

**vina**
- Svalování viny na konkrétní osobu nebo skupinu.
- Musí být jasně určeno **co je vina** a **kdo je viník**.
- Typicky: „někdo způsobil problém, někdo nese odpovědnost".
- Příklady: „Kvůli politikům jsme přišli o všechno."
            „Západ nese odpovědnost za válku."
```

Source

```
<zdroj>…</zdroj>

**zdroj**
- Výrok je podpořen **externím zdrojem**, např. médium, web,
dokument nebo sociální síť.
- Příklady: „Podle zprávy CNN…"
            „Jak uvádí portál Sputnik…"
```