DOI: https://doi.org/10.15276/hait.09.2026.12
UDC 004.89, 004.912

Implicit vs explicit reasoning in strategic narrative detection

Oleg A. Boiko[1)]
ORCID: https://orcid.org/0009-0002-3424-8234; o.a.boiko@kpi.ua
Valeriy Ya. Danylov[1)]
ORCID: https://orcid.org/0009-0000-0875-4868; danilov1950@ukr.net. Scopus Author ID: 7201827051
[1)] National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute".
37, Beresteiskyi Ave. Kyiv, 03056, Ukraine

ABSTRACT

Relevance of the research: The proliferation of state-sponsored propaganda across the global digital ecosystem constitutes a severe threat to international security and democratic stability. As cognitive warfare evolves, adversaries rely on subtle, culturally nuanced strategic narratives to polarize public opinion and alter collective behavior. Defending the information space requires advanced natural language processing systems capable of deep, context-aware reasoning to identify and deconstruct manipulative narratives in real time, moving beyond traditional manual analysis. **Aim and objectives**: This study evaluates the efficacy of advanced artificial intelligence Reasoning Models in classifying strategic narratives. The research investigates the interplay between internal computational reasoning budgets and the enforcement of explicit step-by-step reasoning instructions, seeking to uncover how modern architectures process geopolitical nuances and identify potential interference effects caused by traditional instruction techniques. **Methods used**: To achieve this, the investigation conducted extensive empirical evaluations across a diverse spectrum of large language models utilizing the DIPROMATS 2024 dataset, analyzing social media posts published by diplomatic authorities. The experimental framework isolated variables by varying internal reasoning resources and prompt output instructions. The methodology bypassed complex modular architectures in favor of highly optimized, singular instruction configurations embedded with contextual examples. The assessment rigorously measured exact classification precision, operational latency, and economic efficiency. **Results**: The findings reveal a significant shift in optimal processing strategies. The research identified a critical reasoning paradox, demonstrating that enforcing explicit step-by-step reasoning actively degrades the classification accuracy of some highly capable models by disrupting their optimized latent analytical pathways. Furthermore, a stark architectural divergence emerged regarding computational budgets. While certain models improved with increased reasoning time, others suffered severe performance degradation when allocated maximum computational resources. Conversely, optimized lightweight models demonstrated an unparalleled combination of processing speed and financial efficiency. **Conclusions**: The research concludes that modern cognitive models inherently possess an advanced capacity to natively detect latent propagandistic frameworks, bypassing the need for extensive supervised training data. By exposing the interference effect, the study challenges traditional prompt engineering paradigms and provides the empirical benchmarks required to integrate scalable, real-time strategic narrative detection within comprehensive information security defense systems.

Keywords: Propaganda detection; artificial intelligence; reasoning models; large language models; information security; cognitive warfare

***For citation*: Boiko O. A., Danylov V. Ya. "Implicit vs explicit reasoning in strategic narrative detection". *Herald of Advanced of Information Technology*. 2026; Vol.9 No.2: 168–180. DOI: https://doi.org/10.15276/hait.09.2026.12**

1. INTRODUCTION

The proliferation of propaganda and state-sponsored disinformation across the global digital information ecosystem has become a major threat to international security and democratic stability [1], [2]. Accelerated by the rapid integration of generative artificial intelligence, manipulative content has evolved into a highly advanced form of cognitive warfare. This emergent domain of conflict represents a new frontier where the human mind serves as the primary battleground, specifically aimed at destabilizing cognition and social cohesion through high-velocity information manipulation [3]. Rather than relying solely on the spread of isolated falsehoods, modern disinformation campaigns operate as systemic tools designed to polarize public opinion, influence social identities, and alter collective behavior [1].

This systemic threat is clearly demonstrated by state-directed operations, particularly those driven by the Russian Federation [4]. Following its full-scale invasion of Ukraine, Russia has intensified its psychological and informational campaigns against Western democracies [5]. Using both official state media and covert networks of foreign-language alternative outlets, these operations embed propagandistic narratives deep into the information ecosystems of target nations [5]. To shape foreign media agendas, these adversarial actors actively leverage complex geopolitical framing, such as international relations and armed conflicts [4].

© Boiko O., Danylov V., 2026

This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/deed.uk)

This strategic shift makes the identification of strategic narratives a critical priority for modern information security. Because these storylines and messages are inherently subtle, culturally nuanced, and spread across massive volumes of high-velocity digital content, relying solely on manual human analysis is no longer practical. Protecting the modern information space requires advanced natural language processing systems capable of deep, context-aware reasoning to detect and deconstruct these manipulative narratives in real time.

2. LITERATURE REVIEW AND PROBLEM STATEMENT

The automatic identification of strategic narratives in digital media represents a highly complex computational challenge.

Unlike explicit propaganda detection, narrative identification requires inferring the underlying communication goals and geopolitical framing embedded within a text [6].

Modern frameworks emphasize that these manipulative narratives function as a cornerstone of cognitive warfare, deliberately targeting emotional and cognitive vulnerabilities to induce the behavioral and perceptual shifts necessary to advance an adversary's strategic objectives [3].

Historically, researchers have tackled information manipulation using supervised learning approaches and pre-trained language models. For example, within the context of the DIPROMATS 2024 shared task [7], Cuadrado et al. utilized fine-tuned BERT architectures combined with TF-IDF vectorization to extract structural and lexical features [8]. Similarly, Valls et al. proposed a geopolitically-informed MultiModal BERT approach, attempting to mitigate spurious correlations and concept drift in political tweets by combining text encoding with contextual features [9]. García-Díaz et al. further explored the integration of deep linguistic features and sentence embeddings with transformer models using ensemble learning [10]. Furthermore, to address the inherent class imbalances within these annotated datasets, Fernández et al. demonstrated the efficacy of using LLMs for data augmentation, employing advanced paraphrasing and machine translation to enrich the training corpora prior to model fine-tuning [11].

While these supervised models achieve functional performance, their reliance on extensive, expertly annotated corpora limits their adaptability to rapidly evolving propaganda strategies [12]. To overcome these limitations, recent research has shifted toward utilizing Large Language Models (LLMs) in zero-shot or few-shot paradigms. To handle the complex reasoning required for multiclass narrative classification, state-of-the-art approaches have adopted multi-agent architectures. Notably, Caballero et al. deployed a GPT-4-based multi-agent system comprising a Signal Builder Agent and a Classification Agent, which successfully refined narrative signals and achieved highly competitive results [13].

While multi-agent LLM systems effectively reduce the reliance on annotated datasets, they introduce significant computational overhead, latency, and architectural complexity. The recent emergence of next-generation "Reasoning Models" (also known as Reasoning Language Models or RLMs), which operate by simulating slower, deliberate "System 2" cognitive processes [14] through adjustable internal computational budgets [15], offers a promising alternative. These models are designed to perform stepwise analysis natively, theoretically bypassing the need for external multi-agent signal enhancement [14].

However, the optimal prompt engineering techniques for these advanced architectures are increasingly subject to debate. Historically, Explicit Chain-of-Thought (CoT) prompting has been almost universally prescribed to improve LLM accuracy for complex tasks. Yet, recent cognitive-inspired studies reveal that Explicit CoT can actively reduce performance on tasks heavily reliant on implicit pattern recognition or exception handling [16]. Furthermore, empirical evaluations of modern Reasoning Models (such as the o1 family) have identified an "overthinking" phenomenon, where applying excessive reasoning budgets to certain problems yields diminishing returns or degrades performance [17], while also incurring severe computational costs [18].

Despite these emerging findings, how these phenomena manifest within the highly nuanced domain of strategic narrative classification remains underexplored. It is critical to determine whether enforcing Explicit CoT instructions acts as a necessary cognitive scaffold for geopolitical analysis, or if it actively interferes with the optimized, latent reasoning pathways of advanced models. This study addresses this gap by empirically evaluating the interference effect of Explicit CoT and internal reasoning budgets on classification accuracy, seeking to identify the most efficient and precise configuration for large-scale strategic narrative detection.

3. RESEARCH AIM AND OBJECTIVES

Building upon the recent review of AI-driven propaganda detection [12], this study seeks to evolve the current empirical understanding of next-generation Reasoning Models by evaluating their efficacy in the multiclass classification of strategic narratives. Specifically, the research investigates the interplay between internal reasoning budgets and Explicit Chain-of-Thought (CoT) prompting. By isolating these variables, the study intends to uncover how modern AI architectures process complex geopolitical nuances and identify potential interference effects caused by traditional prompt engineering techniques.

To achieve this aim, the research outlines the following specific objectives:

Performance Benchmarking: To assess the baseline classification accuracy of single-prompt, few-shot configurations across a diverse spectrum of AI models, comparing their capability to map text to predefined strategic narratives against established supervised and multi-agent baseline systems [7], [13].

Analysis of Reasoning Budgets: To examine the impact of internal reasoning budgets across different model architectures, determining whether the allocation of more computational resources universally improves classification accuracy.

Evaluation of the Interference Effect: To investigate how enforcing Explicit CoT instructions (requiring step-by-step reasoning articulation) impacts classification accuracy and performance across different model architectures.

Economic and Operational Assessment: To evaluate the real-world deployment viability of these models for large-scale, high-velocity information monitoring by measuring their economic efficiency (Strict F_1 per USD) and operational latency (processing speed per tweet).

4. METHODOLOGY

4.1. Dataset and Task Definition

The research utilized the official English dataset from the DIPROMATS 2024 shared task (Task 2) [19]. The dataset consists of 800 tweets published by diplomatic authorities representing four major international actors: China, Russia, the European Union, and the United States [19].

The primary objective is a multiclass, multilabel classification task. Each evaluated tweet must be categorized into one or more of 24 predefined strategic narratives (six specific narratives assigned to each of the four actors) [20]. Because a single text can promote multiple concepts simultaneously, or none at all, the classification system must be capable of returning multiple labels or an empty set for each input.

To illustrate the classification challenge, consider the following official example provided by the organizers: “We consider the situation with the rights of the #Russian-speaking population in #Latvia and #Estonia to be discriminatory. Statelessness is a shameful social phenomenon mainly applied to Russian-speaking residents, depriving them of basic democratic rights.” An accurate classification system must process this geopolitical framing and correctly output the corresponding narrative ID, in this case, [“RU1”], which corresponds to the strategic narrative depicting the West and its allies as immoral and hostile actors [20].

4.2. Experimental Setup and Model Selection

To comprehensively evaluate modern AI capabilities on this task, the study conducted a total of 195 experimental runs across 11 distinct models. The selection included 5 models from Google and 6 from OpenAI, ranging from lightweight, highly compressed versions to flagship Reasoning Models. In total, 61 unique configurations were tested to isolate the factors that contribute most significantly to classification accuracy. To assess baseline, out-of-the-box model capabilities, all API calls were executed using the providers' default temperature and top-p settings.

Unlike previous approaches that utilized complex multi-agent architectures, this study adopted a highly optimized single-prompt methodology. The goal was to determine if advanced Reasoning Models could accurately map text to complex narratives using only contextual instructions and an efficient few-shot design.

4.3. Prompt Engineering and Few-Shot Design

The prompt engineering is focused on maximizing evaluation accuracy while minimizing the token budgets. The prompt structure is defined by four core components:

Persona Initialization: The architecture employs an optimized, two-tiered persona initialization. First, a concise system-level instruction assigns the core operational role, establishing the model as an *“experienced international expert in detecting propaganda and strategic narratives”*. Second, the user-level prompt immediately reinforces this framing by initiating the task with the phrase, *“Greetings officer! Today your task is ...”* This direct tonal framing forces the model

ISSN 2663-0176 (Print)
ISSN 2663-7731 (Online)

into a rigorous, security-oriented analytical mindset without relying on bloated, token-heavy context blocks.

Structured Input Context: To clearly delineate task boundaries and prevent context blending, the primary input variables are strictly compartmentalized using explicit XML tags (e.g., <*tweet_to_analyze*>, <*narratives_list*> and <*output_format_instructions*>). This structural clarity ensures the model can easily distinguish the target social media text from the provided classification guidelines.

Embedded Few-Shot Context: Instead of utilizing a separate, traditional "few-shot" block, the context is embedded directly into the narrative definitions. Each narrative is provided with an <*id*>, <*title*>, <*description*>, and a targeted <*example*>. This approach drastically reduces token usage while maintaining high instructional efficiency.

Structured Output: The absolute requirement for the model to return only a valid JSON object is critical for the automated extraction and evaluation pipeline.

4.4. Variables Evaluated: Reasoning Budgets and Explicit Chain-of-Thought

To understand how different AI architectures process complex classification tasks, the experimental framework isolated and controlled two specific variables:

1. Internal Reasoning Budgets: For models that support adjustable internal computing resources, this study adopts the unified term "reasoning budget." This standardizes the terminology across providers, corresponding to what Google designates as a "thinking budget" [21] and what OpenAI defines as "reasoning effort" [22]. The allocated budget was varied across tested configurations, ranging from being completely disabled ("None") for certain models up to the maximum allowable capacities. This approach evaluated whether granting a Reasoning Model more time to internally process the embedded few-shot examples improved classification accuracy.

2. Explicit Chain-of-Thought (CoT) Instructions: The study adjusted the strictness of the requested output format within the XML tags to observe how forcing a model to “think out loud” affected its final decision. Three levels of Explicit CoT were defined:

Strong: The JSON output enforced a strict sequential schema requiring a *reasoning* key to be generated before the *narrative_ids* array. The model was instructed to summarize the tweet and evaluate each narrative step-by-step within this reasoning field. This structural constraint forces the model to articulate its rationale prior to classification.

Light: The JSON output required a reasoning field, but only requested a short, free-form explanation alongside the narrative IDs.

None: The prompt strictly prohibited any explanatory text. The model was instructed to output only the array of applicable narrative IDs.

4.5. Evaluation Metrics and Verification

A critical methodological advantage of utilizing the DIPROMATS 2024 Task 2 dataset is the inherent transparency and reliability of the evaluation process. Because the gold standard test labels are held privately by the organizers, all 195 experimental runs were evaluated directly through the official online evaluation tool [23]. This standardized pipeline ensures that the research results are objectively benchmarked against the exact same criteria as all other participants. Furthermore, it allows the findings to be independently verified on the public leaderboard, providing transparency and guaranteeing a fair, reliable comparison to established baselines and competing systems.

The performance of each configuration was measured using two primary metrics provided by the official evaluation system [24]:

Strict F_1: This metric requires an exact match between the predicted narrative IDs and the primary communicative intention of the tweet (labeled as 'yes' in the gold standard). Partial matches (labeled as 'leaning') are treated as negative cases. This provides a rigorous benchmark for precision, aggressively penalizing models that over-classify.

Lenient F_1: This metric considers both primary (“yes”) and secondary (“leaning”) communicative intentions as positive matches, acting as an upper-bound measure where 'leaning' cases are treated flexibly in the way that best fits the evaluated system's predictions. As shown in Table 1 in Section 5.1, Overall Performance, the evaluated models achieved significantly higher scores under this metric (reaching up to over 0.78), illustrating the models' robust capacity to identify broader, underlying strategic alignments even when they are not the central focus of the text.

4.6. Economic and Operational Speed Assessment

To accurately evaluate the real-world deployment viability of these models for large-scale narrative monitoring, the study tracked both financial cost and operational latency. During the

experimental runs, the evaluation script programmatically recorded the response time for each API call to calculate the average processing speed per tweet (t_{avg}). Concurrently, economic and operational efficiencies were calculated directly within the script utilizing the official API pricing (measured in USD per 10^6 tokens).

The total token cost per run (C_{total}) accounts for the unique billing structure of Reasoning Models, which treat "thinking" or "reasoning" tokens identically to standard output tokens:

$$C_{\text{total}} = \left(\frac{T_{\text{input}}}{10^6} \cdot P_{\text{input}}\right) + \left(\frac{T_{\text{output}}}{10^6} \cdot P_{\text{output}}\right),$$

Where T represents the respective token counts, with T_{output} encompassing both standard completion tokens and any generated reasoning tokens, and P represents the API price per million tokens.

To standardize the economic impact, the average cost per 1,000 tweets (C_{1000}) was calculated based on the total number of evaluated tweets (N_{total}) in the given run:

$$C_{1000} = \left(\frac{C_{\text{total}}}{N_{\text{total}}}\right) \cdot 1000.$$

Finally, to measure the practical trade-off between accuracy, cost, and speed, two efficiency metrics were defined. Economic efficiency, expressed as Strict F_1 points per US Dollar ($F_{1,USD}$), is calculated as:

$$F_{1,\text{USD}} = \frac{F_{1,\text{strict}}}{C_{1000}}.$$

Time efficiency, expressed as operational speed in Strict F_1 points per second ($F_{1,s}$), is calculated as:

$$F_{1,s} = \frac{F_{1,\text{strict}}}{t_{\text{avg}}}.$$

The baseline pricing data was sourced directly from the official documentation for both Google [25] and OpenAI [26] models, with all rates verified as of February 11, 2026. This standardized assessment framework ensures that the resulting cost and speed metrics accurately reflect real-world operational constraints.

4.7. Data and Protocol Availability

To ensure methodological reproducibility, the complete prompt templates, explicit JSON output schemas, and inference parameters utilized in this study are publicly hosted on GitHub [27]. The repository provides comprehensive per-run telemetry, configuration logs, and the JSON outputs for the evaluated models.

5. RESEARCH RESULTS

The study evaluates a total of 195 runs across 11 distinct models (5 models from Google and 6 from OpenAI), validated using the official online evaluation tool [23] under the team name "OB1". The experimental framework tested 61 unique configurations by varying internal reasoning budgets and Explicit Chain-of-Thought (CoT) instructions.

5.1. Overall Classification Performance

The findings indicate a significant shift in the optimal approach for the multiclass classification of strategic narratives. According to the official DIPROMATS 2024 Task 2 leaderboard, the proposed few-shot configurations achieved highly competitive outcomes, ranking 1st overall as of February 28, 2026 [23]. The most successful single run on the leaderboard, achieved by OB1, reached a Strict F_1 score of 0.5568 (using gpt-5.1 with a maximum ("High") reasoning budget and "None" Explicit CoT). However, as shown in Table 1, the variance across multiple independent runs indicates that the top-performing configurations (including gpt-5.1, o4-mini, and gpt-5) yielded statistically comparable results within one standard deviation. Therefore, these models collectively represent a top tier of performance rather than featuring a single absolute leader.

These performance metrics demonstrate that highly optimized single-prompt, few-shot configurations can rival, and even surpass, fine-tuned pretrained transformers and complex LLM-based multi-agent systems [7], [13]. This confirms that the inherent natural language understanding capabilities of modern Reasoning Models can effectively classify nuanced geopolitical themes without relying on extensive supervised training, external multi-agent signal enhancement, or synthetic data generation [13].

To ensure the reliability of the findings, all competitive configurations were evaluated across at least five independent iterations ($n \geq 5$). As demonstrated in Table 1, the Coefficient of Variation (CV) for the top-performing models remained below 1 % (e.g., gpt-5.1 at 0.90 % and gemini-2.5-flash at 0.47 %). While the sample size is relatively compact, this minimal variance confirms that the reasoning models exhibit highly stable and deterministic classification capabilities under the proposed prompt architecture.

ISSN 2663-0176 (Print)
ISSN 2663-7731 (Online)

Table 1. **Top-performing configurations for each evaluated model**

Provider	Model Name	Type	Reasoning budget*	Explicit CoT	# of runs	F_1 Strict	F_1 Lenient	*SD* (σ)	*CV* (%)
OpenAI	gpt-5.1	RLM	Maximum	None	5	0.5485	0.7858	0.0050	0.90
OpenAI	o4-mini	RLM	Default	None	5	0.5484	0.7803	0.0037	0.68
OpenAI	gpt-5	RLM	Minimal	Strong	5	0.5415	0.7812	0.0044	0.81
Google	gemini-2.5-flash	RLM	None	None	5	0.5349	0.7824	0.0025	0.47
OpenAI	gpt-4.1	LLM	None	None	5	0.5339	0.7843	0.0033	0.62
OpenAI	gpt-5-mini	RLM	Medium	None	5	0.5283	0.7782	0.0047	0.89
Google	gemini-2.0-flash	LLM	-	None	5	0.5250	0.7707	0.0014	0.27
Google	gemini-2.5-flash-lite	RLM	None	None	5	0.5203	0.7669	0.0028	0.54
Google	gemini-2.5-pro	RLM	Minimal	None	5	0.5173	0.7779	0.0039	0.75
Google	gemini-2.0-flash-lite	LLM	-	None	5	0.5154	0.7634	0.0033	0.64
OpenAI	gpt-5-nano	RLM	Maximum	Light	5	0.4866	0.7272	0.0035	0.71

Source: **compiled by the authors**

* Here and throughout the paper, for OpenAI models the parameter "Reasoning effort: High" is designated as "Reasoning budget: Maximum", and for Google models the parameter "Thinking budget: 0" is designated as "Reasoning budget: None" to simplify direct comparison across different providers.

5.2. Divergent Impacts of Reasoning Budgets

The research highlights a significant divergence in how different model architectures handle internal reasoning budgets. As detailed in Fig. 1, allocating more computational resources for internal reasoning does not universally improve classification accuracy. In fact, the impact of the reasoning budget is heavily dependent on the specific model's underlying architecture.

For OpenAI gpt-5 family models, increasing the reasoning budget generally yielded better classification results. For instance, the compact gpt-5-nano model saw its Strict F_1 score improve drastically from 0.3633 at a minimal budget to 0.4860 at a maximum budget. Similarly, the gpt-5 model showed steady improvement as the reasoning budget increased, peaking at 0.5385 with a maximum budget configuration. This aligns with the intuitive expectation that more internal thinking time allows the model to better parse the few-shot examples and evaluate the nuanced narrative connections.

In stark contrast, Google's Reasoning Models exhibited a strict inverse relationship. For these models, increasing the reasoning budget consistently degraded performance. The gemini-2.5-flash model achieved its peak Strict F_1 of 0.5349 with the reasoning budget completely disabled (“None”). When the budget was artificially forced to “Maximum”, performance dropped to 0.5054. A similar trend was observed across the gemini-2.5-flash-lite and gemini-2.5-pro models, where the lowest possible reasoning settings consistently outperformed the higher ones.

This divergence suggests profound differences in the training paradigms of these models. For certain architectures, forced over-thinking on a relatively straightforward classification task may lead to over-complicating the context, resulting in narrative hallucination or the misclassification of secondary themes as primary ones. Consequently, maximizing the reasoning budget should not be treated as a default strategy for all modern Reasoning Models.

5.3. The Reasoning Paradox: Internal vs. Explicit Chain-of-Thought

One of the central findings of this research is the identification of a “Reasoning Paradox” regarding the use of Explicit Chain-of-Thought (CoT) prompting. The study tested three Explicit CoT configurations: “Strong” (requiring step-by-step reasoning in a structured JSON format prior to classification), “Light” (a short free-form reasoning), and “None” (direct classification output).

Historically, prompting an AI to articulate its reasoning step-by-step has been a standard method for improving classification accuracy on complex tasks. However, the data reveals that for modern, highly capable Reasoning Models, enforcing

Explicit CoT often degrades performance. As shown

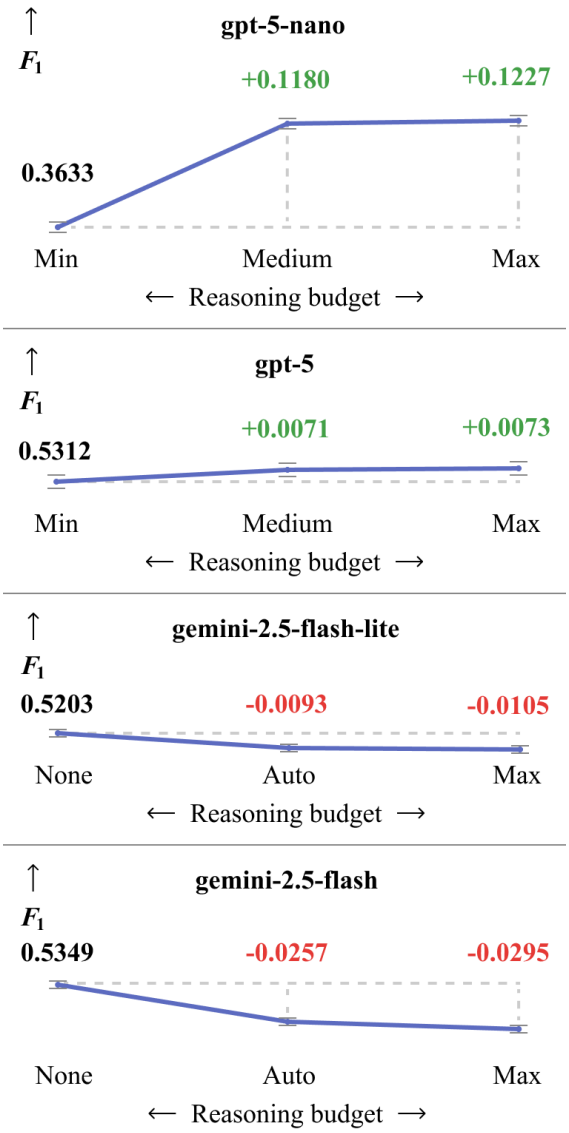


Fig 1. **Reasoning budget impact on Strict F_1**
Source: **compiled by the authors**

in Fig. 2, Google models achieved their highest Strict F_1 scores when Explicit CoT was set to "None". For example, the gemini-2.5-flash model reached a 0.5349 Strict F_1 without Explicit CoT, but experienced a performance drop of 0.0123 when the "Strong" CoT constraint was applied.

Conversely, weaker or highly compressed models still rely heavily on Explicit CoT to function effectively. The compact gpt-5-nano model experienced a dramatic increase in its Strict F_1 score, jumping from 0.3633 without CoT to 0.4287 when forced to use "Strong" CoT. For these smaller models, the requirement to generate intermediate reasoning text acts as a necessary cognitive scaffold.

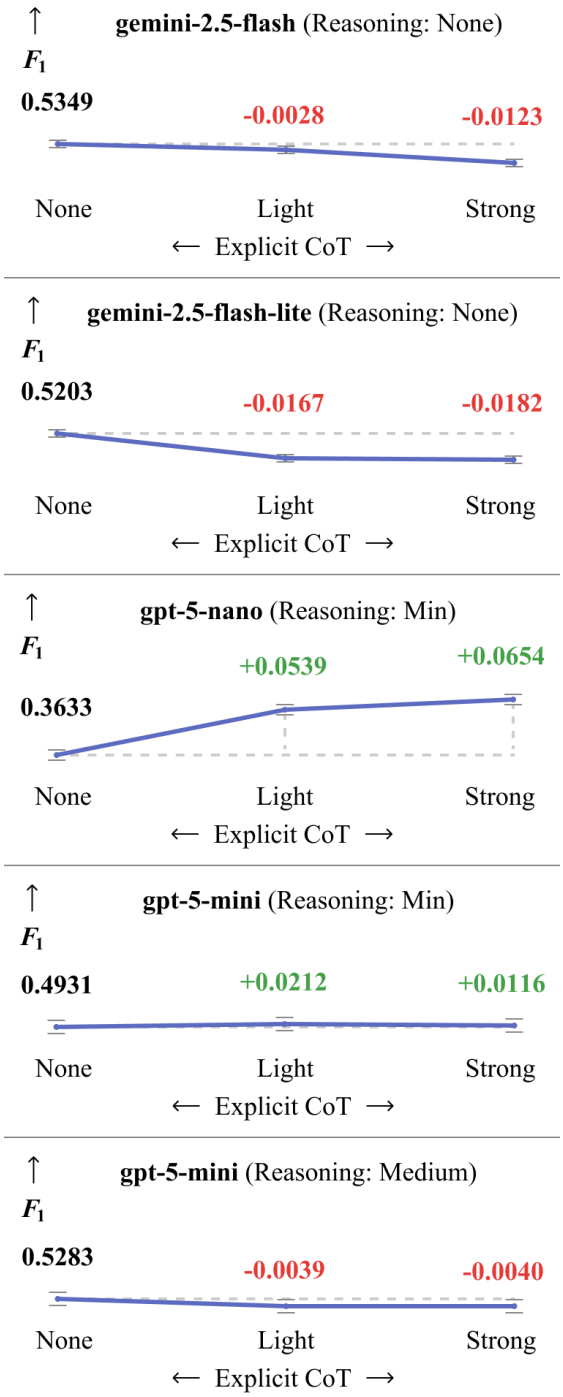


Fig 2. **Impact of Explicit Chain-of-Thought instructions on Strict F_1**
Source: **compiled by the authors**

The interaction between Explicit CoT and internal reasoning budgets further illustrates this paradox. The gpt-5-mini model, when operating on a "minimal" internal reasoning budget, benefited from Explicit CoT (gaining 0.0212 F_1 with "Light" CoT). However, when it's internal budget was increased to "medium", the application of Explicit CoT became detrimental, resulting in a slight performance drop.

ISSN 2663-0176 (Print)
ISSN 2663-7731 (Online)

This dichotomy indicates a fundamental shift in how prompts should be designed for the latest generation of AI. While Explicit CoT remains a vital tool for less capable models, it actively interferes with the optimized, latent reasoning pathways of advanced Reasoning Models. Forcing these models to translate their multidimensional internal logic into human-readable, sequential text output might limit their capacity to correctly identify complex, overlapping strategic narratives.

5.4. Economics and Operational Speed

Beyond pure classification accuracy, the study evaluated the evaluated models for real-world deployment viability based on financial cost and operational latency. As strategic narrative monitoring often involves processing massive volumes of social media data in real time, efficiency is a critical factor.

Table 2 details the economic efficiency of the evaluated models, measured in Strict F_1 points per US Dollar (F_1 per USD). The gpt-5-nano model (operating on a minimal reasoning budget) proved to be the most cost-effective strictly by the numbers, delivering 5.1898 F_1 per USD at a cost of only $0.070 per 1,000 tweets. However, this efficiency comes at the cost of overall accuracy, as its Strict F_1 score was only 0.3633. For a more practical balance of high accuracy and low cost, Google's "lite" models excelled. The gemini-2.0-flash-lite achieved an impressive 5.0529 F_1 per USD (costing $0.102 per 1,000 tweets) while maintaining a competitive Strict F_1 score of 0.5154.

Table 3 illustrates the operational speed of the models, measured in average seconds per tweet and F_1 points per second. Google's gemini-2.5-flash-lite was the fastest model tested, processing an individual tweet in an average of 0.611 seconds and achieving the highest time-efficiency rating of 0.8515 F_1 per second. The gemini-2.0-flash and OpenAI's gpt-4.1 also demonstrated high-speed processing capabilities, both averaging well under one second per tweet.

Conversely, the data reveals that maximizing reasoning budgets or enforcing Explicit Chain-of-Thought (CoT) introduces severe latency bottlenecks. For instance, while the gpt-5-nano model saw accuracy improvements with a maximum reasoning budget, its processing time ballooned to over 16 seconds per tweet. Similarly, enforcing a "Strong" Explicit CoT on the gpt-5 model increased its processing time to 5.622 seconds per tweet.

These metrics confirm that modern, lightweight Reasoning Models are highly capable of powering efficient, real-time propaganda detection pipelines without prohibitive financial or computational overhead. However, attempting to force deeper reasoning processes on these models quickly yields diminishing returns, rendering them impractical for large-scale, high-velocity data monitoring.

Table 2. **Economic efficiency and cost per 1,000 tweets**

Model / Configuration	F_1 Strict	Avg. Price per 1,000 tweets (USD)	$F_{1,USD}$
gpt-5-nano */ Reasoning: Min, CoT: None*	0.3633	$0.070	5.1898
gemini-2.0-flash-lite */ CoT: None*	0.5154	$0.102	5.0529
gemini-2.0-flash */ CoT: none*	0.5250	$0.136	3.8603
gemini-2.5-flash-lite */ Reasoning: None, CoT: None*	0.5203	$0.135	3.8541
gemini-2.5-flash */ Reasoning: None, CoT: None*	0.5349	$0.440	1.2157
gpt-5-mini */ Reasoning: Medium, CoT: None*	0.5283	$0.908	0.5818
gpt-5-nano */ Reasoning: Max, CoT: Light*	0.4866	$0.957	0.5084
gpt-5.1 */ Reasoning: None, CoT: None*	0.5408	$1.750	0.3090
gpt-5 */ Reasoning: Min, CoT: None*	0.5312	$1.759	0.3020
gpt-4.1 */ CoT: None*	0.5339	$2.501	0.2135
o4-mini */ Reasoning: Default, CoT: None*	0.5484	$2.876	0.1907
gpt-5 */ Reasoning: Min, CoT: Strong*	0.5415	$4.132	0.1311
gpt-5.1 */ Reasoning: Max, CoT: None*	0.5485	$4.867	0.1127
gemini-2.5-pro */ Reasoning: Min, CoT: None*	0.5173	$5.206	0.0994

Source: **compiled by the authors**

Table 3. **Operational speed and time efficiency**

Model / Configuration	F_1 Strict	Avg Time per tweet (s)	$F_{1,s}$
gemini-2.5-flash-lite / *Reasoning: None, CoT: None*	0.5203	0.611	0.8515
gemini-2.0-flash / *CoT: none*	0.5250	0.731	0.7182
gpt-4.1 / *CoT: None*	0.5339	0.770	0.6934
gemini-2.0-flash-lite / *CoT: None*	0.5154	0.782	0.6590
gemini-2.5-flash / *Reasoning: None, CoT: None*	0.5349	0.857	0.6241
gpt-5.1 / *Reasoning: None, CoT: None*	0.5408	1.061	0.5097
gpt-5-nano / *Reasoning: Min, CoT: None*	0.3633	0.914	0.3975
gpt-5 / *Reasoning: Min, CoT: None*	0.5312	1.396	0.3805
o4-mini / *Reasoning: Default, CoT: None*	0.5484	3.548	0.1546
gpt-5-mini / *Reasoning: Medium, CoT: None*	0.5283	4.244	0.1245
gemini-2.5-pro / *Reasoning: Min, CoT: None*	0.5173	4.793	0.1079
gpt-5.1 / *Reasoning: Max, CoT: None*	0.5485	5.289	0.1037
gpt-5 / *Reasoning: Min, CoT: Strong*	0.5415	5.622	0.0963
gpt-5-nano / *Reasoning: Max, CoT: Light*	0.4866	16.453	0.0296

Source: **compiled by the authors**

6. DISCUSSION OF RESULTS

6.1. The Interference of Explicit Chain-of-Thought and Overthinking

The empirical results reveal a critical shift in how next-generation AI models process complex, nuanced information. The identified “Reasoning Paradox”, where enforcing Explicit Chain-of-Thought (CoT) actively degrades the performance of advanced models, challenges the established prompt engineering consensus. For flagship models like Google’s gemini-2.5-flash forcing the model to articulate its reasoning step-by-step acts as a cognitive bottleneck rather than a scaffold.

This interference likely occurs because strategic narrative detection is not a strict mathematical or logical deduction task; it requires holistically weighting overlapping, implicit geopolitical contexts. Forcing advanced Reasoning Models to collapse their multidimensional, latent reasoning pathways into a linear, human-readable text output restricts their ability to natively map these complex thematic connections.

Furthermore, the architectural divergence in how these models utilize reasoning budgets exposes a severe "overthinking" vulnerability in specific model families. While OpenAI’s architectures demonstrated a capacity to utilize maximum reasoning budgets to steadily refine their classifications, Google’s models suffered a notable degradation in Strict F_1 scores when granted excess computational time. This suggests that for highly capable models evaluating nuanced social media discourse, prolonged internal processing leads to context over-complication. The model might begin to hallucinate secondary, "leaning" narratives as primary communicative intentions, severely penalizing its strict classification accuracy. Consequently, maximizing test-time compute is not a universally viable strategy for cognitive warfare detection and must be strictly tailored to the underlying model architecture.

6.2. Integrating Reasoning Models into Scalable Detection Systems

The varying computational demands and operational costs observed among the evaluated architectures underscore a clear operational trade-off. While modern Reasoning Models deliver exceptional and consistent classification accuracy, relying on a high-compute, monolithic AI model for real-time cognitive warfare detection is practically and economically inefficient. Furthermore, deploying these high-latency models across fast-paced, multi-domain environments might actively worsen system instability, significantly increasing the risk of a "managerial collapse" during information overload [28]. Conversely, the empirical data demonstrates that specific reasoning models are highly effective for strategic narrative detection at scale. In particular, gemini-2.5-flash-lite delivers an

ISSN 2663-0176 (Print)
ISSN 2663-7731 (Online)

unparalleled combination of speed (highest Strict $F_{1,s}$) and cost-efficiency (exceptional Strict $F_{1,USD}$).

To further maximize this efficiency, detection systems can exploit the scale-free topology of web platforms. Rather than processing every post, monitoring can be selectively focused on network "concentrators" – the highly connected hubs through which most information flows [29]. Armed with this targeted approach, the narrative detection capabilities of optimized reasoning models become a highly valuable addition to broader information security systems. Specifically, these models can be deployed as specialized agents within multi-agent frameworks, such as the "Swarm of Virtual Experts" methodology, which aggregates the outputs of various independent models to evaluate the weight of manipulative signals [30]. By supplying exact performance and economic benchmarks, this research provides the empirical data necessary to seamlessly integrate strategic narrative detection as a scalable, real-time feature within these wider defense ecosystems.

6.3. Limitations of the Study

It is important to acknowledge certain methodological limitations regarding the evaluation pipeline. The study evaluated a predefined grid of 61 configurations across 195 experimental runs using the public DIPROMATS 2024 online evaluator. While no iterative prompt tuning occurred between runs, the optimal hyperparameters (internal reasoning budgets and Explicit CoT levels) for each model were ultimately identified based on their performance on this single test set. Because the test set effectively functioned as a validation set for selecting these optimal configurations post-hoc, there is a theoretical risk of mild adaptive overfitting. Consequently, while the reported Strict F_1 scores accurately reflect the models' capabilities on this specific dataset under optimal conditions, their absolute performance might slightly vary when generalized to entirely unseen data streams.

Furthermore, the central finding regarding the "Reasoning Paradox", where Explicit Chain-of-Thought (CoT) degrades performance in advanced models, was observed using a highly specific prompt architecture. In this design, few-shot examples were embedded directly into the narrative definitions rather than presented in a separate block. This embedded structure inherently provides a strong contextual framework. Therefore, the observed performance drop might partially stem from a structural conflict between the inherently structured definitions and the forced sequential reasoning of Explicit CoT. While the trend was consistent across different models from different providers, the robustness of this interference effect across varying prompt formats remains an important area for future research.

6.4. Ethical Considerations and Dual-Use Risk

The automation of strategic narrative detection presents inherent dual-use risks. While this architecture is designed to defend information ecosystems against state-sponsored manipulation, automated narrative classification systems could theoretically be repurposed by authoritarian actors for mass censorship or the suppression of legitimate dissenting discourse. To mitigate these risks, the deployed architecture is strictly intended for analytical monitoring and flagging. Furthermore, our evaluation prioritizes the Strict F_1 metric, which aggressively penalizes over-classification, thereby structurally reducing the risk of false-positive censorship in real-world deployments.

7. CONCLUSIONS

This study empirically evaluated the efficacy of next-generation Reasoning Models in the complex natural language processing task of strategic narrative detection. By systematically isolating the variables of internal reasoning budgets and Explicit Chain-of-Thought (CoT) prompting, the research investigated how advanced AI architectures process nuanced, state-sponsored geopolitical manipulation. The findings confirm that modern System 2 models possess a highly advanced capacity to detect latent propagandistic frameworks without the need for extensive supervised training data. This represents a significant leap forward in cognitive warfare defense.

Critically, the research uncovered a profound "Reasoning Paradox" regarding prompt engineering and test-time compute. The empirical data demonstrates that enforcing Explicit Chain-of-Thought instructions, while benefiting compact models like gpt-5-nano, often introduces a severe "interference effect" by disrupting the optimized, latent reasoning pathways of more advanced architectures. Furthermore, a distinct divergence was observed regarding internal reasoning budgets: while OpenAI's models mostly benefited from increased compute, Google's gemini-2.5-flash and gemini-2.5-pro experienced notable performance degradation and "overthinking" when allocated maximum reasoning budgets.

Finally, the operational analysis of these models highlights important nuances in deployment strategies. While flagship models achieve the highest

Strict F_1 scores, their associated latency and financial costs render them impractical for real-time digital space monitoring. Conversely, highly optimized models, particularly gemini-2.5-flash-lite, demonstrated an unparalleled combination of classification speed and cost-efficiency. These empirical benchmarks provide the precise data required to integrate Reasoning Models within advanced information security architectures, ensuring both cognitive depth and real-time resilience in the modern digital space.

8. ACKNOWLEDGMENTS

The authors express their deepest gratitude to the Security and Defense Forces of Ukraine (Сили безпеки та оборони України) for the opportunity to conduct this research.

REFERENCES

1. Plikynas, D., Rizgelienė, I. & Korvel, G. “Systematic review of fake news, propaganda, and disinformation: Examining authors, content, and social impact through machine learning”. *IEEE Access*. 2025; 13, https://www.scopus.com/pages/publications/85216006685.
DOI: https://doi.org/10.1109/ACCESS.2025.3530688.

2. Ahmad, P. N., Guo, J., AboElenein, N. M., ul Haq, Q. M., Ahmad, S., Algarni, A. D. & Ateya, A. A. “Hierarchical graph-based integration network for propaganda detection in textual news articles on social media”. *Scientific Reports*. 2025; 15: 1827, https://www.scopus.com/pages/publications/85215547874. DOI: https://doi.org/10.1038/s41598-024-74126-9.

3. Paziuk, A., Lande, D., Shnurko-Tabakova, E. & Kingston, P. “Decoding manipulative narratives in cognitive warfare: a case study of the Russia-Ukraine conflict”. *Frontiers in Artificial Intelligence*. 2025; 8: 1566022, https://www.scopus.com/pages/publications/105017058174.
DOI: https://doi.org/10.3389/frai.2025.1566022.

4. Yang, Y., McCabe, S. & Hindman, M. “Does russian propaganda lead or follow? Topic coverage, user engagement, and RT and Sputnik's agenda influence on US media”. *The International Journal of Press/Politics*. 2025; 30 (4): 956–979, https://www.scopus.com/pages/publications/85202849284. DOI: https://doi.org/10.1177/19401612241271074.

5. Beseler, A. & Toepfl, F. “Conduits of the Kremlin's informational influence abroad? How German-Language alternative media outlets are connected to Russia's ruling elites”. *The International Journal of Press/Politics*. 2025; 30 (3): 659–678, https://www.scopus.com/pages/publications/85186411409. DOI: https://doi.org/10.1177/19401612241230284.

6. Fraile-Hernández, J. M., Peñas, A. & Moral, P. “Automatic identification of narratives: evaluation framework, annotation methodology, and dataset creation”. *IEEE Access*. 2024; 13, https://www.scopus.com/pages/publications/85207153103.
DOI: https://doi.org/10.1109/ACCESS.2024.3475579.

7. Moral, P., Fraile, J. M., Marco, G., Peñas, A. & Gonzalo, J. “Overview of DIPROMATS 2024: detection, characterization and tracking of propaganda in messages from diplomats and authorities of world powers”. *Procesamiento del Lenguaje Natural*. 2024; 73: 347–358, https://www.scopus.com/pages/publications/85201636504. DOI: https://doi.org/10.1109/10.26342/2024-73-26.

8. Cuadrado, J., Martinez, E., Cuadrado, J., Martinez-Santos, J. C. & Puertas, E. “VerbaNex AI at DIPROMATS 2024: enhancing propaganda detection in diplomatic tweets with fine-tuned BERT and integrated NLP techniques”. *CEUR Workshop Proceedings*. 2024; 3756, https://www.scopus.com/pages/publications/85204400912.

9. Valls, A., Komander, B. & Cerquides, J. “Geopolitically-informed multimodal BERT for propaganda detection in political tweets”. *CEUR Workshop Proceedings*. 2024; 3756, https://www.scopus.com/pages/publications/85204408926.

10. García-Díaz, J. A., Pan, R., Lázaro, A. R., Cristancho, C. & Valencia-García, R. “UMUTEAM at DIPROMATS 2024: Feature integration for detecting fine-grained propaganda and narrative”. *CEUR Workshop Proceedings*. 2024; 3756, https://www.scopus.com/pages/publications/85204377063.

11. Fernández, M., Ojeda, M., Guevara, L., Varela, D., Mendoza, M. & Barrón-Cedeño, A. “VICTOR VECTORS @ DIPROMATS 2024: Propaganda detection with LLM paraphrasing and machine translation”. *CEUR Workshop Proceedings*. 2024; 3756, https://www.scopus.com/pages/publications/85204375548.

12. Boiko, O. A. “Modern AI methods for detecting propaganda in text”. *Registration, Storage and Processing of Data*. 2025; 27 (1): 120–131. DOI: https://doi.org/10.35681/1560-9189.2025.27.1.336147.

13. Caballero, A., Centeno, R. & Rodrigo, Á. “LLM-based multi-agent models for multiclass classification of strategic narratives”. *CEUR Workshop Proceedings*. 2024; 3756, https://www.scopus.com/pages/publications/85204364077.

14. Li, Z.-Z., Zhang, D., Zhang, M.-L., et al. “From System 1 to System 2: A survey of reasoning large language models”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2025; 1–20, https://www.scopus.com/pages/publications/105023055770.
DOI: https://doi.org/10.1109/TPAMI.2025.3637037.

15. Snell, C., Lee, J., Xu, K. & Kumar, A. “Scaling LLM test-time compute optimally can be more effective than scaling model parameters” *arXiv*. 2024, https://www.scopus.com/pages/publications/105010225173. DOI: https://doi.org/10.48550/arXiv.2408.03314.

16. Liu, R., Geng, J., Wu, A. J., Sucholutsky, I., Lombrozo, T. & Griffiths, T. L. “Mind Your Step (by Step): Chain-of-thought can reduce performance on tasks where thinking makes humans worse”. *arXiv*. 2024, https://www.scopus.com/pages/publications/105023636323.
DOI: https://doi.org/10.48550/arXiv.2410.21333.

17. Chen, X., Xu, J., Liang, T., et al. “Do NOT think that much for 2+3=? On the overthinking of o1-like LLMs”. *arXiv*. 2025. DOI: https://doi.org/10.48550/arXiv.2412.21187.

18. McGinness, L. & Baumgartner, P. “Large language models imitate logical reasoning, but at what cost?” *arXiv*. 2025, https://www.scopus.com/pages/publications/105023512352. DOI: https://doi.org/10.48550/arXiv.2509.12645.

19. “Task 2 Data”. *Zenodo. Test set*. 2024. DOI: https://doi.org/10.5281/zenodo.12663310.

20. “Set of Narratives”. – Available from: https://sites.google.com/view/dipromats2024/task-2/set-of-narratives. – [Accessed: Feb. 16, 2025].

21. “Gemini thinking”. *Gemini API Documentation*. – Available from: https://ai.google.dev/gemini-api/docs/thinking. – [Accessed: Feb. 11, 2025].

22. “Reasoning Models”. *OpenAI API Documentation*. – Available from: https://platform.openai.com/docs/guides/reasoning. – [Accessed: Feb. 11, 2025].

23. “Official DIPROMATS 2024 Task 2 Leaderboard”. *Hugging Face Spaces*. – Available from: https://huggingface.co/spaces/NLP-UNED/dipromats2024-task2-leaderboard. – [Accessed: Feb. 28, 2025].

24. “Evaluation & Baselines”. – Available from: https://sites.google.com/view/dipromats2024/task-2/evaluation-baselines. – [Accessed: Feb. 25, 2025].

25. “Gemini API Pricing”. – Available from: https://ai.google.dev/gemini-api/docs/pricing. – [Accessed: Feb. 11, 2025].

26. “Pricing”. – Available from: https://platform.openai.com/docs/pricing. – [Accessed: Feb. 11, 2025].

27. Boiko, O. A., Danylov, V.Ya. “dipromats2024-task2-ob1. Supporting materials for a research study on strategic narrative detection. (Team OB1)”. *GitHub*. – Available from: https://github.com/BoyKot/dipromats2024-task2-ob1. – [Accessed: Mar. 1, 2025].

28. Bazarnyi, S., Husak, Y., Voitko, T., Aliew, F. & Yevseiev, S. “Mathematical model of multi-domain interaction based on game theory”. *Advanced Information Systems*. 2025; 9 (3): 22–31, https://www.scopus.com/pages/publications/105012905062. DOI: https://doi.org/10.20998/2522-9052.2025.3.03.

29. Uhryn, D. I., Karachevtsev, A. O., Tomka, Y. Ya., Zakharov, M. M., & Troianovska, Y. L. “Information system for analyzing public sentiment in web platforms based on machine learning”. *Herald of Advanced Information Technology*. 2024; 7 (2): 199–212, https://www.scopus.com/pages/publications/105024800269. DOI: https://doi.org/10.15276/hait.07.2024.14.

30. Lande, D., Alekseichuk, L., Svoboda, I. & Strashnoy, L. “Methodology of a swarm of virtual experts for evaluating the weight of connections in networks”. *Theoretical and Applied Cybersecurity*. 2024; 6 (2): 25–33. DOI: https://doi.org/10.20535/tacs.2664-29132024.2.319946.

Conflicts of Interest: The authors declare that they have no conflict of interest regarding this study, including financial, personal, authorship or other, which could influence the research and its results presented in this article

Received 05.01.2026
Received after revision 26.02.2026
Accepted 09.03.2026

DOI: https://doi.org/10.15276/hait.09.2026.12
УДК 004.89, 004.912

Явне та неявне міркування у розпізнаванні стратегічних наративів

Бойко Олег Анатолійович[1)]
ORCID: https://orcid.org/0009-0002-3424-8234; o.a.boiko@kpi.ua
Данилов Валерій Якович[1)]
ORCID: https://orcid.org/0009-0000-0875-4868; danilov1950@ukr.net. Scopus Author ID: 7201827051
[1)] Київський політехнічний інститут імені Ігоря Сікорського, пр. Берестейський, 37. Київ, 03056, Україна

АНОТАЦІЯ

Актуальність дослідження: Поширення державної пропаганди в глобальній цифровій екосистемі становить серйозну загрозу міжнародній безпеці та демократичній стабільності. Еволюція когнітивної війни призводить до того, що супротивники все більше покладаються на тонкі, культурно нюансовані стратегічні наративи для поляризації громадської думки та зміни колективної поведінки. Захист інформаційного простору вимагає передових систем обробки природної мови, здатних до глибокого, контекстно-орієнтованого міркування для виявлення та деконструкції маніпулятивних наративів у режимі реального часу, виходячи за межі традиційного ручного аналізу. **Мета та завдання**: Це дослідження оцінює ефективність передових моделей штучного інтелекту з функцією міркування у класифікації стратегічних наративів. Дослідження вивчає взаємозв'язок між внутрішніми обчислювальними бюджетами на міркування та застосуванням явних покрокових інструкцій для міркування, з метою з'ясувати, як сучасні архітектури обробляють геополітичні нюанси, та виявити потенційні ефекти інтерференції, спричинені традиційними методами інструктування. **Методи дослідження**: Для досягнення цього в рамках дослідження було проведено масштабні емпіричні оцінювання широкого спектра великих мовних моделей з використанням набору даних DIPROMATS 2024, аналізуючи публікації в соціальних мережах, зроблені дипломатичними відомствами. У межах експерименту змінні були ізольовані шляхом варіювання внутрішніх ресурсів міркування та інструкцій щодо виводу у запитах. Методологія уникла використання складних модульних архітектур на користь високооптимізованих одиничних конфігурацій інструкцій, доповнених контекстними прикладами. Оцінювання ретельно вимірювало абсолютну точність класифікації, операційну затримку та економічну ефективність. **Результати**: Отримані результати свідчать про значний зсув в оптимальних стратегіях обробки даних. Дослідження виявило критичний парадокс міркування, продемонструвавши, що примусове явне покрокове міркування активно знижує точність класифікації деяких високопродуктивних моделей, порушуючи їхні оптимізовані приховані аналітичні шляхи. Крім того, виявилася різка архітектурна розбіжність щодо обчислювальних бюджетів. У той час як певні моделі покращували свої результати зі збільшенням часу на міркування, інші зазнавали серйозного погіршення продуктивності при виділенні максимальних обчислювальних ресурсів. Натомість оптимізовані полегшені моделі продемонстрували неперевершене поєднання швидкості обробки та фінансової ефективності. **Висновки**: У дослідженні робиться висновок, що сучасні когнітивні моделі за своєю природою володіють розвиненою здатністю нативно виявляти приховані пропагандистські структури, усуваючи потребу у великих обсягах даних для керованого навчання. Викриваючи ефект інтерференції, дослідження кидає виклик традиційним парадигмам інженерії запитів та надає емпіричні орієнтири, необхідні для інтеграції масштабованого виявлення стратегічних наративів у режимі реального часу в комплексні системи захисту інформаційної безпеки.

Ключові слова: виявлення пропаганди; штучний інтелект; мовні моделі з функцією міркування; мовні моделі з логічним міркуванням; великі мовні моделі; інформаційна безпека; когнітивна війна.

ABOUT THE AUTHORS



Oleg A. Boiko - PhD Student, Department of Artificial Intelligence. National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”. 37, Beresteiskyi Ave. Kyiv, 03056, Ukraine
ORCID: https://orcid.org/0009-0002-3424-8234; email: o.a.boiko@kpi.ua
***Research field*:** Artificial Intelligence

Бойко Олег Анатолійович - студент аспірантури кафедри Штучного інтелекту. Національний технічний університет України “Київський політехнічний інститут імені Ігоря Сікорського”, пр. Берестейський, 37. Київ, 03056, Україна



Valeriy Ya. Danylov - Doctor of Engineering. Sciences, Professor, Laureate of Borys Paton National Prize of Ukraine, Department of Artificial Intelligence. National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”. 37, Beresteiskyi Ave. Kyiv, 03056, Ukraine
ORCID: https://orcid.org/0009-0000-0875-4868; danilov1950@ukr.net. Scopus Author ID: 7201827051
***Research field*:** System Analysis; Neural Networks; Deep Learning in Medicine and Socio-Economics; Hydroacoustic Signal Processing

Данилов Валерій Якович - доктор технічних наук, професор, лауреат Національної премії України імені Бориса Патона, професор кафедри Штучного інтелекту. Національний технічний університет України “Київський політехнічний інститут імені Ігоря Сікорського”, пр. Берестейський, 37. Київ, 03056, Україна