ANNALES

Anali za istrske in mediteranske študije
Annali di Studi istriani e mediterranei
Annals for Istrian and Mediterranean Studies
Series Historia et Sociologia, 36, 2026, 2

UDK 009 Annales, Ser. hist. sociol., 36, 2026, 2, pp. 163-336, Koper 2026 ISSN1408-5348

UDK 009

ISSN 1408-5348
e-ISSN 2591-1775

ANNALES

Anali za istrske in mediteranske študije
Annali di Studi istriani e mediterranei
Annals for Istrian and Mediterranean Studies

Series Historia et Sociologia, 36, 2026, 2

KOPER 2026

ISSN 1408-5348
e-ISSN 2591-1775

UDK 009

Letnik 36, leto 2026, številka 2

Glavni urednik/*Redattore capo/ Editor in chief:* Darko Darovec

Odgovorni urednik/*Redattore responsabile/Responsible Editor:* Salvator Žitko

Uredniki/*Redattori/Editors:* Urška Lampe, Boštjan Udovič, Žiga Oman, Veronika Kos

Prevajalka/*Traduttrice/Translator:* Cecilia Furioso Cenci (it.)

Oblikovalec/*Progetto grafico/ Graphic design:* Dušan Podgornik, Darko Darovec

Tisk/*Stampa/Print:* Založništvo PADRE d.o.o.

Založnika/*Editori/Published by:* Zgodovinsko društvo za južno Primorsko - Koper / *Società storica del Litorale - Capodistria*© / Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja / *Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment / Istituto IRRIS di ricerca, sviluppo e strategie della società, cultura e ambiente*©

Sedež uredništva/*Sede della redazione/ Address of Editorial Board:* SI-6000 Koper/*Capodistria,* Garibaldijeva/*Via Garibaldi 18*
e-mail: annaleszdjp@gmail.com, ***internet:*** https://zdjp.si

Redakcija te številke je bila zaključena 30. 06. 2026.

Sofinancirajo/*Supporto finanziario/ Financially supported by:* Javna agencija za znanstvenoraziskovalno in inovacijsko dejavnost Republike Slovenije (ARIS)

Annales - Series Historia et Sociologia izhaja štirikrat letno.

Maloprodajna cena tega zvezka je 11 EUR.

Naklada/*Tiratura/Circulation:* 300 izvodov/copie/copies

Revija Annales, Series Historia et Sociologia je vključena v naslednje podatkovne baze / *La rivista Annales, Series Historia et Sociologia è inserita nei seguenti data base / Articles appearing in this journal are abstracted and indexed in*: Clarivate Analytics (USA): Arts and Humanities Citation Index (A&HCI) in/and Current Contents / Arts & Humanities; IBZ, Internationale Bibliographie der Zeitschriftenliteratur (GER); Sociological Abstracts (USA); Referativnyi Zhurnal Viniti (RUS); European Reference Index for the Humanities and Social Sciences (ERIH PLUS); Elsevier B. V.: SCOPUS (NL); Directory of Open Access Journals (DOAJ).

ANNALES · Ser. hist. sociol. · 36 · 2026 · 2

Anali za istrske in mediteranske študije - Annali di Studi istriani e mediterranei - Annals for Istrian and Mediterranean Studies

UDK 009 Volume 36, Koper 2026, issue 2 ISSN 1408-5348
e-ISSN 2591-1775

VSEBINA / *INDICE GENERALE / CONTENTS*

received: 2026-02-25 DOI 10.19233/ASHS.2026.11

MAPPING CONTESTED CULTURAL MEMORY: AN LLM-SUPPORTED APPROACH TO ANALYSING NARRATIVE STRUCTURES, DISCURSIVE MODES AND DISCOURSE FUNCTIONS

Marjan HORVAT
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: marjan.horvat@irris.eu

Jure KORAŽIJA
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: jure.korazija@irris.eu

Jan BABNIK
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: jan.babnik@irris.eu

Tadej ŠKVORC
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: tadej.skvorc@irris.eu / tadej.skvorc@fri.uni-lj.si

Darko DAROVEC
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: darko.darovec@irris.eu

Žiga OMAN
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: ziga.oman@irris.eu

Urška LAMPE
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: urska.lampe@irris.eu

Angelika ERGAVER
Institute IRRIS for Research, Development and Strategies of Society, Culture and Environment, Čentur 1f, 6273 Marezige, Slovenia
e-mail: angelika.ergaver@irris.eu

Marko ROBNIK-ŠIKONJA
University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000 Ljubljana, Slovenia
e-mail: marko.robniksikonja@fri.uni-lj.si

ABSTRACT

This article develops and evaluates an exploratory LLM-supported methodology for analysing contested cultural memory in hybrid public spheres. The methodology combines three analytical dimensions: Greimasian actantial extraction for identifying narrative structures, antagonism–agonism–deliberation classification for assessing discursive modes, and epistemic, ethical, democratic, and ideological scoring for analysing discourse functions. These outputs are examined through semantic clustering, statistical analysis, and expert-based validation. The pipeline is demonstrated through Slovenia's Day of Resistance Against the Occupier, comparing newspaper articles and publicly visible X posts. The findings show that the approach can identify meaningful relations between narrative configurations, discursive modes, and discourse functions, especially in longer texts, while X posts remain exploratory due to reliability limits.

Keywords: social media, cultural memory, deliberative democracy, Greimas actant model, large language models (LLMs), prompt engineering, statistical analysis, political discourse

MAPPARE LA MEMORIA CULTURALE CONTESA: UN APPROCCIO SUPPORTATO DA LLM ALL'ANALISI DELLE STRUTTURE NARRATIVE, DELLE MODALITÀ DISCORSIVE E DELLE FUNZIONI DISCURSIVE

SINTESI

Il presente articolo sviluppa e valuta una metodologia esplorativa supportata da LLM per l'analisi della memoria culturale contesa nelle sfere pubbliche ibride. La metodologia combina tre dimensioni analitiche: l'estrazione attanziale greimasiana per identificare le strutture narrative, la classificazione antagonismo–agonismo–deliberazione per valutare le modalità discorsive, e l'attribuzione di punteggi epistemici, etici, democratici e ideologici per analizzare le funzioni discorsive. Tali risultati sono ulteriormente esaminati attraverso clustering semantico, analisi statistica e validazione basata sul giudizio di esperti. La pipeline viene applicata al caso della Giornata slovena della resistenza contro l'occupatore, attraverso il confronto tra articoli di giornale e post pubblicamente visibili su X. I risultati mostrano che l'analisi attanziale supportata da LLM può identificare configurazioni narrative significative, soprattutto nei testi più lunghi e ricchi di contesto. I risultati relativi ai brevi post su X sono invece utilizzati soltanto come indicatori esplorativi, a causa della minore affidabilità e della limitata disponibilità di informazioni contestuali.

Parole chiave: social media, memoria culturale, democrazia deliberativa, modello attanziale di Greimas, grandi modelli linguistici (LLM), prompt engineering, analisi statistica, discorso politico

INTRODUCTION

Cultural memory provides communities with shared historical reference points (Nora, 1989), through which collective identities, political orientations, and normative attachments are articulated (Assmann & Czaplicka, 1995).[1] Such reference points include monuments, rituals, commemorations, formative historical events, and collective traumas such as the Holocaust (Winter, 2008). They can stabilise public discourse by creating common interpretive horizons, particularly in nation-building processes (Anderson, 1983; Hobsbawm, 1983). Yet cultural memory is never politically neutral. The same commemorative reference can also become a site of exclusion, ideological struggle, and competing interpretations of the past (Rothberg, 2009). When mobilised strategically, especially by populist or nationalist actors, memory conflicts can reinforce social divisions and political polarisation (Wodak & Richardson, 2013; Tveskov & Bissonnette, 2023). Contemporary "memory wars" are therefore not confined to reinterpretations of 20th-century history, but continue to shape present political conflicts and geopolitical tensions (Rutten et al., 2013; Pshenychnykh et al., 2024).

Digital media have made these dynamics more visible, faster, and more difficult to stabilise. Cultural memory continues to be shaped by state institutions, museums, schools, official commemorations, traditional media, as well as by family narratives, personal testimonies, and other forms of everyday remembrance. What has changed is the speed, visibility, and scale at which these practices circulate and interact. Social media posts, hashtags, memes, and affective reactions

1 The authors acknowledge funding from the European Union's Horizon Europe programme under grant agreement No. 101094752, *Social Media for Democracy – Understanding the Causal Mechanisms of Digital Citizenship* (SoMe4Dem). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or Horizon Europe. Neither the European Union nor the granting authority can be held responsible for them. This work was also supported by the European Union's Horizon Europe programme under grant agreement No. 101186647, *Centre of Excellence in Artificial Intelligence for Digital Humanities* (AI4DH) and by COST Action CA22149, *Research Network for Interdisciplinary Studies of Transhistorical Deliberative Democracy* (CHANGECODE), funded by COST, European Cooperation in Science and Technology. The article is also the result of research supported by the Slovenian Research and Innovation Agency (ARIS) within the research programme P6-0435, *Practices of Conflict Resolution Between Customary and Statutory Law in the Area of Today's Slovenia and Its Neighbouring Lands*; and research projects P6-0411, *Language Resources and Technologies for Slovene; V5-2297, Media Landscape in Slovenia between Pluralisation and Homogenisation*; L2-50070, *Embeddings-based Techniques for Media Monitoring Applications; GC-0002, Large Language Models for Digital Humanities*; and J6-60100, *Divided Memory and Collectivisation of Individual Memories in Slovenia between 1941 and 1996*.

have expanded the arenas in which memory is produced, contested, and recontextualised. This transformation has been described as part of the mediatisation and digitalisation of memory (van Dijck, 2007) and as a "connective turn" in which public and private memory practices increasingly overlap and blur (Hoskins, 2011).

These developments raise a methodological challenge. If collective memory is increasingly articulated across hybrid media environments (Chadwick, 2017) researchers need tools capable of analysing not only what is remembered, but also how mnemonic claims are narratively structured, how conflict is discursively enacted, and what democratic functions such discourse performs. Existing approaches in memory studies have provided rich theoretical accounts of memory regimes (Olick et al., 2023; Levy & Sznaider, 2002), digital remembrance, polarisation, and public-sphere transformation (Hoskins, 2011; Rutten et al., 2013; Richardson-Little et al., 2022). However, it remains difficult to operationalise these concepts systematically across larger corpora while preserving interpretive depth.

This article addresses this challenge by developing and testing an LLM-supported methodology for analysing cultural memory across hybrid media environments. The approach builds on a theoretical-conceptual framework (Horvat, 2026), which shifts the focus from memory regimes to the discursive modes through which cultural memory is enacted in hybrid public spheres. The methodology of the study combines three main analytical dimensions and examines their empirical applicability through an LLM-supported workflow. First, texts are examined through Greimasian actantial extraction, which reconstructs narrative structures by identifying subjects, objects, senders, receivers, helpers, and opponents. Second, discursive modes are examined through antagonism–agonism–deliberation (AAD) classification, which assesses how these narrative configurations are publicly performed. Third, discourse functions are examined through epistemic, ethical, democratic, and ideological scoring, which evaluates whether texts mainly contribute to factual clarification, ethical recognition, democratic inclusion, or ideological positioning. The outputs of these three analytical dimensions are then linked through semantic clustering and statistical analysis in order to identify recurring narrative and discursive patterns and are subsequently assessed through expert-based validation. The full methodological pipeline is presented in section Case selection, data and method.

The empirical aim is to compare how the same commemoration is represented in two communicative environments: newspaper articles and publicly visible X posts. More specifically, the article examines how narrative structures, AAD classifications, and discourse functions can be computationally linked to identify areas of conflict and patterns of agonistic and deliberative potential. The comparison is used to demonstrate the proposed pipeline. Accordingly, findings based on X posts are interpreted cautiously: short social media texts provide limited contextual information, are difficult to classify reliably, and should be treated as publicly visible discursive traces rather than as a representative deliberative debate.

The empirical case is Slovenia's Day of Resistance Against the Occupier, commemorated on April 27. The holiday refers to the founding of the Liberation Front in 1941 and to the broader memory of anti-fascist resistance, the Partisan movement, the communist revolution, post-war violence, and the socialist past. Since Slovenian independence, it has remained politically contested and has been mobilised in disputes over national identity, anti-fascism, anti-communism, reconciliation, and historical responsibility (Kranjc, 2013; Prunk, 1992; Pirjevec, 1995; Godeša, 1995; Hribar & Hribar, 2021). This makes it a suitable case for testing how an LLM-supported methodology can map contested memory discourse across media.

The article addresses four research questions:

RQ1: How can LLM-supported actantial analysis be used to identify recurring narrative structures in contested memory discourse?
RQ2: How do these narrative structures relate to antagonistic, agonistic, and deliberative modes of discourse?
RQ3: How do newspaper articles and publicly visible X posts differ in their actantial configurations, AAD modes, and discourse-function profiles?
RQ4: What are the methodological limits of LLM-supported analysis when applied to longer news articles and shorter X posts?

The article's main contribution is methodological. It demonstrates how actantial narrative analysis, AAD classification, and discourse-function analysis can be combined in a scalable LLM-supported pipeline for analysing contested cultural memory, while also specifying the scope conditions of this approach.

The paper proceeds in several steps. First, we introduce the background for the three analytical dimensions used in the study. We then present the case selection, data, methodological pipeline, clustering, and validation strategy. Next, we examine the empirical results of the application. Finally, we discuss the methodological and substantive implications of the findings, including scope conditions and limitations, and outline directions for future validation and application.

THEORETICAL BACKGROUND AND RELATED WORK

Because the broader theoretical-conceptual framework is developed in detail elsewhere (Horvat, 2026), this article does not attempt to reconstruct the full debates on cultural memory, deliberative democracy, agonistic theory, or hybrid public spheres. Instead, it focuses on the methodological development and empirical testing of an LLM-supported pipeline for analysing contested memory discourse.

The study distinguishes between three interlinked analytical dimensions. First, narrative structures capture how cultural memory is organised through relations between actors, goals, sources of legitimation, beneficiaries, helpers, and opponents. Second, discursive modes capture how these narrative structures are enacted in the hybrid public spheres: antagonistically, agonistically, or deliberatively. Third, discourse functions capture what kind of epistemic, ethical, democratic, or ideological work commemorative claims perform.

These dimensions are operationalised through actantial extraction, AAD classification, and discourse-function scoring. Table 1 summarises the three analytical dimensions of the study and their methodological implementation. The following subsections briefly introduce the background required for each, while the next section presents the implementation of the LLM-supported pipeline: annotation, semantic clustering, statistical analysis, interpretation, and validation.

Narrative Structures: Greimas' Actantial Model

The first analytical component concerns the narrative structuring of commemorative claims. Commemorative discourse organises the past into narrative configurations that assign agency, responsibility, victimhood, legitimacy, and opposition, thereby linking historical interpretation to present political identities and future orientations (Meyer, 2008).

To operationalise this dimension, we draw on Greimas' actantial model. Building on structural analyses of myth and narrative (Lévi-Strauss, 1955/1963; Propp, 1968), Greimas (1966/1983) proposed six basic actantial roles that structure narratives: subject, object, sender, receiver, helper, and opponent. These roles describe the relational architecture of a narrative: who pursues what, on whose behalf, for whose benefit, with whose support, and against which obstacles.

For cultural memory research, the actantial model is useful because commemorative conflicts often involve competing assignments of these roles. In the case of Slovenia's Day of Resistance Against the Occupier, the Liberation Front and the Partisan movement may be positioned as subjects of anti-fascist liberation, fascist occupation as the opponent, and national freedom as the object. In alternative framings, however, the same commemoration may be reorganised around anti-communism, post-war violence, national reconciliation, or competing claims to historical legitimacy, with the Liberation Front or the Partisan movement positioned as opponents rather than subjects of national liberation. Actantial analysis therefore helps identify the internal narrative structure of competing memory claims.

Greimas' model has been applied across legal, media, literary, visual, and political discourse analysis (Jiang, 2017; Aarva & Pakarinen, 2006; Langer, 2000; Salama, 2021; Wang & Roberts, 2005; Schoor, 2021). Computational applications remain more limited, although recent studies have begun to apply actantial or actantial scheme inspired narrative-structural methods to larger text corpora and social media discourse (Tangherlini et al., 2020; Willaert, 2023; Elfes, 2024). Our approach builds on this emerging literature but treats actants not as surface-level entities or subject–verb–object units, but as inferred narrative functions. The same actantial role can be realised through different actors, characters, or formulations across texts. Actantial analysis thus serves as a bridge between cultural-memory theory and computational analysis by identifying where conflict is located in memory structure.

Antagonism, Agonism, and Deliberation as Discursive Modes

The second analytical dimension of the study concerns the discursive mode through which

Table 1: Analytical dimensions and their methodological implementations.

Analytical dimension	Analytical question	Operationalisation	Methodological implementation
Narrative structures	How is cultural memory narratively structured?	Greimasian actantial roles: subject, object, sender, receiver, helper, opponent	LLM-supported actantial extraction; semantic clustering of extracted actants
Discursive modes	How is conflict over memory enacted?	Antagonism, agonism, deliberation	LLM-supported AAD classification
Discourse functions	What work does the discourse perform?	Epistemic, ethical, democratic, ideological function scores	LLM-supported function scoring linked to narrative configurations

commemorative narratives are articulated. We distinguish between antagonism, agonism, and deliberation. These modes (AAD) do not refer to different memory regimes, but to different ways in which conflict over memory is publicly performed.

Antagonism refers to conflict organised around mutual delegitimation, exclusion, and friend–enemy distinctions. In commemorative debates, antagonism appears when actors frame the past as a struggle between "our truth" and "their lies," refusing to recognise opponents as legitimate participants in public discourse. In digital environments, such dynamics may be intensified by homophily, echo chambers, and affective polarisation (McPherson et al., 2001; Pariser, 2011; Cinelli et al., 2021; Törnberg, 2022).

Agonism, by contrast, preserves conflict while transforming enemies into legitimate adversaries. Drawing on Mouffe (1999), the democratic challenge is not to eliminate conflict, but to transform antagonism into agonism. In memory studies, agonistic approaches criticise both exclusionary national memory and overly consensual models of remembrance (Bull & Hansen, 2016). They stress that contested pasts cannot always be resolved through rational consensus and that attempts to impose a single reconciliatory narrative may marginalise alternative memories or generate backlash (Nienass, 2023).

Deliberation refers to modes of discourse oriented toward mutual justification, reason-giving, responsiveness, and inclusion. It is studied at the level of deliberative systems (Mansbridge & Parkinson, 2012) or through more specific indicators of discourse quality, such as the Deliberative Quality Index and its later adaptations (Steenbergen et al., 2003; Bächtiger et al., 2010; Beauchamp, 2020; Fournier-Tombs & MacKenzie, 2021; Oswald, 2022). In the theoretical-conceptual framework developed by Horvat (2026), deliberation does not mean the elimination of conflict or the production of consensus, but the opening of contested memory claims to justification, response, correction, and possible transformation.

For the purposes of this article, AAD is used as an operational distinction between three ways of enacting commemorative conflict: antagonism as exclusionary conflict, agonism as legitimate adversarial contestation, and deliberation as reason-giving, responsive, and potentially transformative discursive engagement. The framework (Horvat, 2026) does not treat these modes as a linear scale from conflict to consensus, but as analytically distinct ways of organising conflict over memory. This allows us to examine not only whether commemorative discourse is conflictual, but also how conflict is articulated and whether contested memory claims are opened to justification, response, and reinterpretation.

Discourse Functions in Cultural Memory Discourse

The third dimension of the study concerns the functions performed by commemorative discourse. Memory debates do not only transmit information about the past. They also justify political identities, recognise or deny harm, include or exclude voices, and legitimise particular interpretations. To capture these differences, we adapt the deliberative functions identified by Mansbridge et al. (2012), adding an ideological function.

The epistemic function concerns informed judgement, factual clarification, and responsible engagement with historical evidence. In cultural memory debates, this includes whether claims about the past are open to correction and grounded in verifiable information rather than denialism, selective distortion, or unsupported assertion. The ethical function concerns recognition, respect, and the treatment of others as legitimate interlocutors. This is especially important in memory conflicts, where the right to remembrance may come into conflict with democratic norms, particularly when memories glorify authoritarianism, exclusion, or violence. As Wüstenberg (2023) argues, not all forms of remembrance are democratically valid. The democratic function concerns inclusion: whose memories become visible, whose experiences are marginalised, and who is allowed to participate in defining the public meaning of the past. Commemorations may broaden recognition and participation, but they may also reproduce exclusionary hierarchies of memory. The added ideological function captures how memory is mobilised to stabilise political identities, naturalise particular interpretations, or support strategic political claims. This function is not necessarily illegitimate, since ideological articulation is part of democratic contestation. It becomes problematic when it overrides epistemic robustness, ethical recognition, and democratic reciprocity.

These four functions are not mutually exclusive. A single text may combine epistemic, ethical, democratic, and ideological elements. The purpose of this analytical dimension is therefore not to classify each text into one fixed category, but to assess which functions are more or less salient and how they interact with actantial structures and AAD modes.

LLMs and Computational Analysis of Political Discourse

The methodological motivation for using LLMs lies in the difficulty of analysing complex narrative and discursive structures at scale. Traditional NLP and machine-learning approaches have enabled large-scale analysis of social media and political discourse, but they often require task-specific annotated datasets and are less suitable for interpretive tasks where categories are complex, context-dependent, or theoretically derived.

Large language models, including GPT-based models, Llama, and DeepSeek, have opened new possibilities for semantic and multilingual text analysis (Brown et al., 2020; OpenAI et al., 2023; Grattafiori et al., 2024; DeepSeek-AI et al., 2024). Their strength lies in their ability to process context, infer implicit meaning, and apply complex annotation instructions without task-specific fine-tuning. This renders them potentially useful in the analysis of cultural memory discourse, where narrative roles, discursive modes, and discourse functions cannot easily be reduced to keyword dictionaries. In the empirical analysis below, we use GPT-4o as one implementation of this broader class of models.

At the same time, LLM-supported annotation has notable limitations. Existing computational research on political discourse remains strongly oriented toward English-language, US-centred, Twitter/X, and election-related data (AlDayel & Magdy, 2021; Kawintiranon & Singh, 2022; Németh, 2023). These limitations are particularly relevant for less-resourced languages such as Slovene and for complex socio-semiotic tasks such as actantial and deliberative discourse analysis.

For this reason, the present article treats LLM-supported annotation as an experimental method rather than as a fully validated measurement instrument. The method is used to test whether theoretically derived categories can be operationalised at scale, while the validation section evaluates where the approach is more reliable and where its outputs require cautious interpretation.

CASE SELECTION, DATA AND METHOD

This section presents the empirical design and methodological pipeline. The aim is to test whether a theory-guided LLM-supported workflow can identify recurring narrative and discursive patterns in a large corpus of commemorative discourse. Figure 1 summarises the main steps of the pipeline, which are described in detail in the following subsections.

The first subsection introduces Slovenia's Day of Resistance Against the Occupier as a case of contested memory. The second subsection describes the two corpora: newspaper articles and publicly visible X posts. The third subsection explains the LLM-supported annotation procedure, including actantial extraction, AAD classification, and discourse-function scoring. The fourth subsection describes semantic clustering and statistical analysis. The final subsection presents the validation strategy used to assess the reliability and plausibility of the outputs.

Case Selection: Slovenia's Day of Resistance as Contested Memory

Slovenia's Day of Resistance Against the Occupier, commemorated on April 27, provides a suitable empirical case with which to test the proposed LLM-supported methodology because it condenses several unresolved layers of collective memory and divided memories in

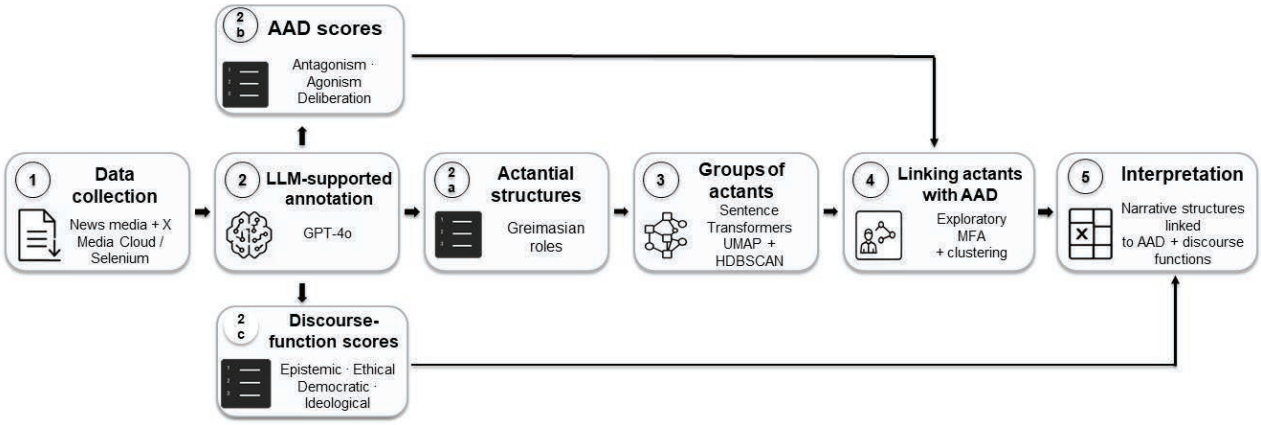

Figure 1: Overview of the LLM-supported methodological pipeline. After data collection from newspaper articles and publicly visible X posts (1), GPT-4o was used to generate three types of annotations: actantial structures (2a), AAD scores (2b), and discourse-function scores (2c). Extracted actants were embedded with SentenceTransformers, projected into a lower-dimensional space with UMAP, and clustered with HDBSCAN to identify recurrent groups of actants (3). Exploratory Multiple Factor Analysis with clustering was conducted to uncover relationships between actantial narrative structures and AAD scores. (4). The final step involved qualitative interpretation of narrative structures in relation to AAD and discourse-function patterns (5). The outputs were subsequently assessed through expert-based validation, including reliability checks for AAD and discourse-function scores and plausibility assessment of actantial extraction. Source: authors' own illustration based on the methodological pipeline developed in this study.

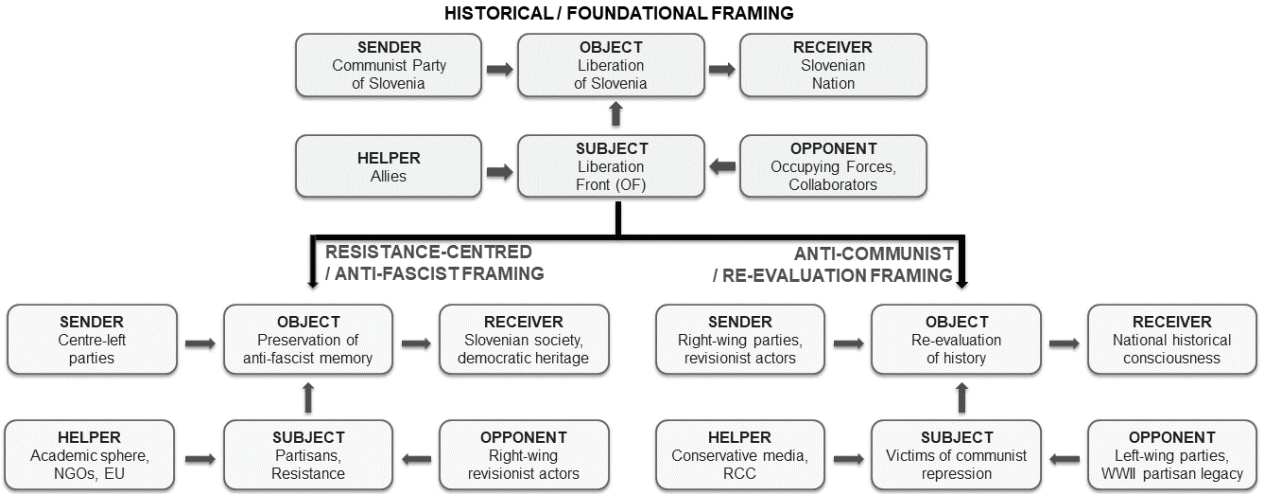

Figure 2: Heuristic actantial framings of Slovenia's Day of Resistance Against the Occupier. The figure schematically illustrates dominant narrative configurations of the commemoration. It is not used as a coding template or as a map of the full empirical field, but as a heuristic example of the polarising narrative structure that the empirical analysis seeks to examine, complicate, and compare with more fine-grained patterns in the corpus (Source: authors' own illustration).

Slovenia, especially related to the Second World War (Lampe & Bajc, 2025). The holiday refers to the founding of the Liberation Front in 1941 and is officially associated with the National Liberation Struggle, anti-fascist resistance, the Partisan movement, and national liberation. At the same time, this memory remains politically contested because of the entanglement of resistance with revolutionary transformation, wartime collaboration and anti-communism, post-war extrajudicial violence, socialist heritage, national reconciliation, and the political instrumentalisation of history (Godeša, 1995; Prunk, 1992; Pirjevec, 1995; Kranjc, 2013; Hribar & Hribar, 2021).

The case is analytically useful because it cannot be reduced to a simple opposition between an anti-fascist/Partisan narrative and an anti-communist/reconciliation narrative. While this opposition remains politically important, public discourse around the commemoration also includes institutional, ethical, democratic, historical, generational, and ideological framings. A discourse-analytical approach makes it possible to examine not only which interpretation of the past is present, but also how historical agency, legitimacy, responsibility, victimhood, and opposition are narratively distributed.

At the narrative level, the commemoration raises questions about who is positioned as the legitimate subject of resistance, what object the commemoration is said to pursue, who authorises or legitimises this pursuit, who benefits from it, and who is presented as helper or opponent. At the level of discursive modes, it allows us to examine whether conflict is performed antagonistically, agonistically, or deliberatively. At the level of discourse functions, it facilitates an assessment of whether references to the past primarily clarify historical facts, recognise suffering, defend democratic values, or stabilise ideological positions.

Figure 2 provides an illustration of such actantial framings of the Day of Resistance. It shows how the commemoration is commonly structured through two broad and politically recognisable narrative orientations: a resistance-centred/anti-fascist framing and an anti-communist/re-evaluation framing. However, the figure should not be understood as a map of the full empirical field or as a predefined coding template.

We use the comparison between newspaper articles and publicly visible X posts to examine how the dominant polarising framing of the commemoration is reproduced, modified, or complicated across different communicative environments. This perspective is important because public representations of the Day of Resistance often rest on a binary opposition that can perpetuate political divisions and reinforce polarisation. The aim is therefore not to demonstrate the impact of social media on cultural memory per se or simply to confirm the existence of polarisation. Rather, our analysis tests whether the proposed LLM-supported pipeline can move beyond this dominant binary image and identify more fine-grained patterns within public discourse: how conflict is structured, how agency and legitimacy are assigned, whether disagreement is enacted antagonistically, agonistically, or deliberatively, and which discourse functions become salient.

Corpus and Data Collection

The empirical analysis is based on two corpora referring to Slovenian commemorations of the Day of Resistance Against the Occupier: publicly visible X posts and online news articles. The two corpora were selected to compare narrative and discursive patterns across two communicative environments: shorter, platform-mediated social media posts and longer news texts.

The X corpus covers the period from April 2023 to April 2024 and was collected by Sciences Po within the SoMe4Dem project. The collection focused on Slovenian-language commemorative discussions and comprises 753 publicly visible posts. For publication and data-sharing purposes, the original dataset was transformed into an anonymised, post-level analytical dataset. The version used in this article includes the text of individual posts together with non-identifying metadata, such as language, timestamp, external links, and variables indicating whether a post is a retweet or a quote. The analysis was conducted at the level of individual posts. Retweets and quotes were treated as textual units only insofar as they appeared in the dataset; the analysis did not examine retweet networks, diffusion patterns, or user-level interaction dynamics. The corpus was compiled using the query terms "Dan upora proti okupatorju" and "Dan upora", with special-character normalisation to improve retrieval. These keywords were selected to capture posts directly referring to the Day of Resistance and its broader commemorative context in Slovenia.

The news corpus was collected using Media Cloud, an open-source platform for collecting and analysing online news content across media ecosystems, developed by the Berkman Klein Center for Internet & Society at Harvard University (Roberts et al., 2021). The Slovenian database was queried using 14 case-sensitive keywords: "dan upora", "dnevu upora", "dan OF", "dneva OF", "proti okupatorju", "državna proslava", "državne proslave", "državni proslavi", "dan spomina", "dnevu spomina", "osvobodilna fronta", "osvobodilne fronte", "protiimperialistična fronta", and "protiimperialistične fronte". Additional material was collected through links found in the X dataset and through manual retrieval from three Slovenian publications: *Delo*, *Demokracija*, and *Mladina*. The final news corpus includes relevant articles published over three consecutive years, from 2022 to 2024. After removing irrelevant and duplicate items, the dataset was reduced from 308 to 164 articles.

The two corpora are not directly equivalent in genre, length, or communicative context. News articles are longer, more explicitly contextualised, and produced within traditional media environments. X posts are shorter, more fragmented, and often dependent on platform conventions, implicit references, quotes, retweets, or external links. For this reason, comparisons between the two corpora are interpreted cautiously. In particular, the X corpus should be understood as a collection of publicly visible discursive traces around the commemoration, not as a representative sample of public opinion or as a deliberative debate.

LLM-Supported Annotation Procedure

To carry out this analysis, we built a computational workflow around GPT-4o. All annotations were generated through the API using the default temperature of 0.7. During prompt development, we reviewed trial outputs and decided to keep this setting. Outputs at temperature 0, which would have allowed full replicability, felt overly rigid and simply didn't work well for the kinds of interpretive tasks this study required. Once the prompts were finalised, we applied the same structure and settings consistently across the full corpus for each analytical component. We did not formally test whether submitting the same texts multiple times would yield identical or near-identical outputs. The validation section therefore focuses on how well the final annotations agree with expert judgement and whether they seem plausible, rather than on whether the model would produce the same results if run again. For this reason, we treat the LLM-generated outputs as annotations produced under a specific prompting and parameter setup, not as fixed, repeatable measurements.

Actantial extraction

This section explains how Greimas's actantial narrative analysis was operationalised for the two corpora of news articles and X posts. Greimas's model identifies six actantial roles that structure the narrative logic of a text: subject, object, sender, receiver, helper, and opponent. Each role was analysed at three levels: actant as the abstract functional role in the narrative, actor as the entity manifesting that role, and character as the specific textual manifestation of the actor.

To operationalise this analysis, we implemented a computational workflow using the GPT-4o API. The system was programmed to process each text as a single unit of analysis by applying a standardised prompt designed to identify and classify actantial roles according to the three levels described above: actant, actor, and character (cf. Appendix A for full prompt (GITHUB, 2026a)). Actantial extraction was therefore used to map narrative configurations within commemorative discourse.

Rather than treating actors as fixed categories, the analysis focuses on how their narrative roles change across texts and how these configurations structure the meaning of the commemoration. This is particularly important in contested memory debates, where the same historical actor, event, or collective identity may be positioned differently depending on the narrative logic of the text.

The results of the analysis were compiled into a tabular dataset in the results section. This dataset includes both the raw textual output from the language model and a structured parsing of the actantial roles, enabling further qualitative and quantitative analysis. The structured output categorises data into fields corresponding to each actantial role and its three analytical levels — actant, actor, and character — ensuring clarity, consistency, and transparency in subsequent analysis. These structured outputs were then used for semantic clustering, statistical analysis, and qualitative interpretation of recurring narrative configurations.

AAD classification

The second analytical component of the methodology concerns Antagonism–Agonism–Deliberation (AAD) classification. This analysis was used to identify how commemorative discourse is articulated through three discursive modes: antagonism, agonism, and deliberation. Each mode was defined by specific linguistic, rhetorical and tonal features, enabling a systematic evaluation of how conflict, disagreement, and justification appeared in the dataset (cf. Appendix B for the full prompt (GITHUB, 2026b)).

For the news articles, GPT-4o evaluated each text by detecting discursive features associated with antagonistic, agonistic, and deliberative communication and assigning a six-point score ranging from strong antagonism to strong deliberation. For X posts, in contrast, due to the brevity and contextual limitations, a simplified three-category classification was applied, categorising discourse as primarily "Antagonistic," "Agonistic," or "Deliberative". If no clear dominance of any category was detected or insufficient information was available, the model assigned the label "N/A".

The distinction between the two procedures reflects the different properties of the corpora. News articles are longer and usually provide more contextual information, allowing the model to detect multiple discursive features within a single text. X posts are shorter, often implicit, and may depend on hashtags, quotations, retweets, external links, or shared background knowledge. For this reason, AAD outputs for X posts are interpreted as exploratory indicators of discursive orientation rather than as directly equivalent to the scores assigned to news articles.

The results were compiled into a structured tabular dataset in which each textual unit — news article or X post — was recorded together with its assigned AAD classification and, where applicable, the identified presence of antagonistic, agonistic, and deliberative features. These AAD variables were later used to examine whether particular narrative configurations were associated with more antagonistic, agonistic, or deliberative forms of engagement.

Discourse-function scoring

The third analytical component of the methodology is the detection of four key discourse functions: epistemic, ethical, democratic, and ideological. For each text, GPT-4o assigned a 0–100 score to each dimension, indicating its relative prominence in the text (cf. Appendix C for full prompts (GITHUB, 2026c)). The epistemic function addresses the clarity and accessibility of fact-based knowledge; the ethical function examines fairness, justice, and moral responsibility; the democratic function captures inclusion, participation, and public accountability; and the ideological function captures the extent to which the text mobilises memory to legitimise particular positions, identities, or political interpretations.

These scores were interpreted as indicators of functional salience rather than as precise interval-level measurements. A higher score does not necessarily mean that a text is more important or more influential, but that a given function is more strongly present in the text according to the annotation criteria. Since commemorative discourse often performs several functions at the same time, the analysis does not assign each text to a single dominant function only. Instead, it examines how epistemic, ethical, democratic, and ideological functions combine within individual texts and how these functional profiles relate to narrative configurations and AAD classifications.

Table 2 illustrates the outputs of the three analytical components using one example from a newspaper article and one example from an X post. The first example is taken from a newspaper article that challenges the official meaning of the Day of Resistance Against the Occupier and frames the commemoration as a divisive rather than unifying national holiday. The second example is taken from an X post by Slovenian Prime Minister Robert Golob, who presents the holiday affirmatively as "a holiday of remembrance and reflection." The examples are used only to illustrate the structure of the LLM-generated outputs; they are not treated as representative of the full corpus.

Table 2: Illustrative examples of LLM-generated outputs for one newspaper article and one X post. The table shows how the pipeline records actantial roles, AAD classification, and discourse-function scores. The examples are used for methodological illustration and are not treated as representative of the full dataset.

Analytical output	News Article	X post
Subject	*National identity discourse*	*Commemoration*
Object	*Historical narrative*	*National unity*
Sender	*Ideological beliefs*	*Historical memory*
Receiver	*Slovenian public*	*National identity*
Helper	*Historical revisionism*	*Public discourse*
Opponent	*Established historical narratives*	*Forgetfulness*
AAD classification	*Mild antagonism*	*Deliberation*
Epistemic Function	*62*	*20*
Ethical Function	*65*	*15*
Ideological Function	*78*	*10*
Democratic Function	*58*	*5*

Semantic Clustering and Statistical Analysis

After the LLM-supported annotation, the extracted actantial outputs were processed through semantic clustering and statistical analysis. The purpose of this step was to identify recurrent semantic patterns across the corpus and to examine how these patterns relate to AAD classifications and discourse-function scores. The clustering procedure was primarily used to support interpretation, not to produce definitive categories.

To compare actants across texts, group semantically similar outputs, and generalise the findings, we used neural text representation methods from natural language processing (NLP). These methods require textual units to be represented as numerical vectors that approximate their semantic properties. Each text in the dataset was first annotated with six actantial roles: Subject, Object, Sender, Receiver, Helper, and Opponent. Many of the extracted actants overlapped semantically, differing only in wording or level of abstraction. To capture and consolidate these variations, each extracted actant string was represented as a high-dimensional vector using the SentenceTransformer model all-mpnet-base-v2. The embeddings were then subjected to dimensionality reduction via UMAP (Uniform Manifold Approximation and Projection). For each actantial role separately, the resulting low-dimensional projections were grouped using HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) in order to identify clusters of semantically similar actants.

Our code was designed to test several parameter settings for dimensionality reduction and HDBSCAN minimum cluster size. Rather than treating the highest Silhouette score as a purely automatic optimisation criterion, parameter selection was guided by a combination of cluster diagnostics and expert inspection of cluster coherence. HDBSCAN was selected because it can discover clusters of varying densities without requiring a predetermined number of clusters. Points that lacked sufficient density in the vector space were labelled as outliers.

Descriptions in each non-outlier cluster were then generalised by assigning them a concise descriptive label using GPT-4o via API. We provided the model with actant names from each cluster; in response, it generated a short one- to three-word designation reflecting their shared semantic theme. After naming the clusters, we prompted GPT-4o again to detect and merge clusters that appeared semantically identical, thus reducing redundancy and improving clarity in subsequent analyses. Outliers were re-examined through a second prompt, where GPT-4o attempted to integrate them into existing clusters based on semantic alignment. Only if ten or more outliers coalesced around a distinct conceptual theme would a new cluster be created, ensuring that minor variations did not fragment the data. All GPT-generated cluster labels, proposed merges, and outlier assignments were treated as interpretive aids and inspected by the researchers before being used in subsequent analysis. For illustrative purposes, Table 3 presents an example of

Table 3: Illustrative example of an Opponent actant cluster.

Original Opponent Actants	*Final Opponent Actant Cluster*
Oppressive Forces	**Oppression & Threats**
External Aggression	Oppression & Threats
Historical oppression	Oppression & Threats
Foreign oppression	Oppression & Threats
Modern societal threats	Oppression & Threats
Ideological threats	Oppression & Threats
Structural Challenges	Oppression & Threats
Modern distractions	Oppression & Threats
Threats to freedom	Oppression & Threats
Modern threats	Oppression & Threats

a final Opponent actant cluster with original actants assigned by GPT-4o. The example is included to illustrate the consolidation procedure and should not be read as a complete or representative list of all clusters.

In the next step, after consolidating actantial clusters, we explored how these clusters relate to AAD classifications. Having consolidated the actantial data from Greimas' model and established AAD scores for each text, we examined how the six clustered actant roles — Subject, Object, Sender, Receiver, Helper, and Opponent — corresponded with AAD levels. To do this, we employed Multiple Factor Analysis (MFA), followed by HDBSCAN clustering.

First, we compiled all relevant variables into a single dataset: six columns corresponding to the clustered actant roles and one numeric column representing the AAD score. MFA was conducted using the FactoMineR package, with one block of variables corresponding to the six actant roles and a second block corresponding to the AAD score. The purpose was to identify common dimensions that captured both actantial configuration and variation in AAD classification. Following the MFA, we extracted the factor coordinates reflecting the combined actantial configuration and AAD score. HDBSCAN was then applied to identify clusters of texts with similar narrative structures and AAD orientations, while observations that did not belong to a dense region were labelled as outliers.

Because HDBSCAN can flag a subset of observations as outliers, we performed a second pass of HDBSCAN on those outliers. This two-stage procedure allowed us to detect whether some outliers formed coherent subclusters. Each text's first- and second-pass cluster assignments were saved in tabular form to facilitate further interpretation. Graphical outputs, including scree plots of variance explained and two-dimensional scatterplots of factor coordinates coloured by cluster membership, were generated to visualise how AAD patterns align with or diverge from particular actantial configurations. Through this combination of MFA and iterative HDBSCAN clustering, the analysis identified recurrent configurations in which particular actantial structures co-occurred with more antagonistic, agonistic, or deliberative modes of engagement.

For each cluster, we then calculated the mean epistemic, ethical, democratic, and ideological scores. The same logic was used to support interpretation of discourse-function scores. This made it possible to identify whether particular narrative clusters were associated with stronger epistemic clarification, ethical recognition, democratic framing, or ideological mobilisation.

All parts of the methodology described above were implemented within a documented computational workflow, which we make freely available. This design allows for a streamlined process: once the initial dataset is uploaded, the scripts sequentially execute each stage. By structuring the methodology as a coherent pipeline, we increase transparency and reproducibility, as each step is automated and documented. However, the workflow should not be understood as fully automated interpretation. The outputs of semantic clustering, MFA, and HDBSCAN were used as guides for qualitative interpretation, not as self-standing proof of stable narrative categories or causal relationships.

Validation Strategy

The reliability and plausibility of the LLM-supported outputs were evaluated through comparison with human expert annotations. The validation was designed as an initial expert-based assessment of the proposed pipeline rather than as a definitive validation of a general-purpose measurement instrument. Its purpose was to examine whether GPT-4o outputs were sufficiently consistent with expert judgement to support large-scale mapping of commemorative discourse. The validation covered the three main analytical components of the pipeline: AAD classification, discourse-function scoring, and Greimasian actantial extraction. The validation sample included 18 news articles and 30 X posts. Since the two corpora differ substantially in length, genre, and contextual density, results are reported separately for news articles and X posts.

Validation of AAD classification and discourse-function scoring

First, we evaluated reliability for AAD classification and discourse-function scoring by comparing expert annotations with GPT-4o outputs. Three measures were used: absolute agreement, Krippendorff's alpha, and intraclass correlation coefficients.

Table 4: Inter-rater reliability statistics for discourse functions and AAD analysis for news articles and X posts. A = news articles; X = X posts; Kα = Krippendorff's alpha; ICC = intraclass correlation; AA = absolute agreement. Significance levels: *** p < 0.001, ** p < 0.01, * p < 0.05.

Category	A - Kα	A - ICC	A - AA	X - Kα	X – ICC	X - AA
Epistemic f.	0,45	0,51**	83,3%	0,13	0,22	80%
Ethical f.	0,60	0,61**	83,3%	-0,11	0,11	60%
Ideological f.	0,63	0,66***	88,8%	0,16	0,14	50%
Democratic f.	0,58	0,57**	77,7%	-0,08	0,08	70%
AAD score	0,69		66,60%	0,41		43,3%

The four discourse-function parameters were originally annotated by GPT-4o on a 0–100 scale and then normalised to a 10-point scale for human annotation. Absolute agreement was calculated as the percentage of cases in which the GPT-4o rating matched the human rating within a ±2 margin. Given the interval-like nature of the data, this margin was used to capture near-equivalent ratings while still distinguishing larger discrepancies. Absolute agreement therefore provides an intuitive measure of how often LLM scores approximate expert ratings within an acceptable range.

We also employed intraclass correlation coefficients and Krippendorff's alpha. ICC was used because it captures consistency between two sets of continuous ratings and is therefore suitable for interval-like data. Following Koo and Li (2016), we used a two-way mixed-effects model, ICC (3,1), with a consistency definition, since the aim was to assess whether the LLM reproduced the pattern of expert judgements rather than exact numerical identity. Krippendorff's alpha was calculated as a complementary measure, since it is widely used in content analysis and accounts for chance agreement.

Table 4 shows a clear difference between news articles and X posts. For news articles, the reliability statistics indicate moderate alignment between GPT-4o and expert judgement. ICC values for discourse functions range from 0.51 to 0.66, Krippendorff's alpha ranges from 0.45 to 0.69, and absolute agreement reaches up to 88.8%. These results suggest that GPT-4o performs reasonably well when applied to longer, more context-rich texts.

For X posts, reliability is substantially lower. Alpha and ICC values indicate weak agreement for several dimensions, even though absolute agreement remains between 50% and 80% for some discourse functions. This discrepancy is important. It shows that tolerance-based agreement may capture cases where LLM and expert ratings are relatively close, but low chance-corrected agreement suggests that GPT-4o does not reliably reproduce expert judgement at the item level for short posts. The AAD score follows a similar pattern, with stronger agreement for news articles and weaker agreement for X posts.

These findings suggest that GPT-4o more reliably approximates expert judgement in longer, more formal texts and performs less consistently on semantically sparse X posts. This is not surprising, since short social media posts are difficult to label even for simpler tasks such as sentiment analysis, where human inter-annotator agreement has been reported between 0.42 and 0.67, while self-agreement ranges between 0.46 and 0.84 (Mozetič et al., 2016). The X corpus is therefore treated as a source of indicators of discursive orientation rather than as a basis for robust item-level measurement comparable to news articles.

Validation of Greimasian actantial extraction

The validation of Greimasian actantial extraction required a different procedure. Exact string matching was not appropriate because actants are abstract narrative functions and can be expressed through multiple semantically similar formulations. For example, an opponent may be named differently across texts while performing the same narrative function.

Two human expert annotators were therefore presented with the actants and characters generated by GPT-4o alongside the original texts. They were asked to indicate their level of agreement with the LLM-generated output on a 1–5 scale, where 1 meant complete disagreement and 5 meant complete agreement. This approach assessed the plausibility and coherence of the actantial extraction rather than exact identity between human and machine labels.

For news articles (Table 5), both annotators generally rated GPT-4o outputs favourably, with mean scores often exceeding 4. This suggests that the

Table 5: Agreement scores between human expert annotators and GPT-4o for Greimasian actantial roles in news articles. Values indicate means and standard deviations. MAD = mean absolute difference between human annotators.

Actant Role	Human 1	Human 2	MAD
Subject Actant	4,35 ± 0,63	4,94 ± 0,27	0,57
Subject Character	3,86 ± 0,86	4,29 ± 0,61	1,00
Object Actant	4,86 ± 0,36	4,57 ± 0,51	0,43
Object Character	4,50 ± 0,76	4,07 ± 0,83	0,86
Sender Actant	3,93 ± 0,83	4,79 ± 0,58	1,00
Sender Character	4,29 ± 0,91	3,50 ± 1,02	1,64
Receiver Actant	4,29 ± 0,47	4,93 ± 0,27	0,64
Receiver Character	4,14 ± 0,66	4,64 ± 0,63	0,64
Helper Actant	4,29 ± 0,73	4,50 ± 0,52	0,21
Helper Character	4,14 ± 0,95	4,21 ± 0,80	0,93
Opponent Actant	4,64 ± 0,63	4,64 ± 0,50	0,57
Opponent Character	4,64 ± 0,75	4,07 ± 0,83	1

Table 6: Agreement scores between human expert annotators and GPT-4o for Greimasian actantial roles in X posts. Values indicate means and standard deviations. MAD = mean absolute difference between human annotators.

Actant Role	Human 1	Human 2	MAD
Subject Actant	2,80 ± 0,66	3,97 ± 1,22	1,43
Subject Character	3,00 ± 3,43	3,43 ± 1,36	1,3
Object Actant	3,23 ± 0,77	3,77 ± 1,31	1,2
Object Character	2,93 ± 0,91	3,43 ±1,31	1,1
Sender Actant	3,17 ± 0,70	4,00 ± 1,29	1,37
Sender Character	2,96 ± 1,00	3,53 ± 1,22	1,23
Receiver Actant	3,23 ± 0,57	4,40 ± 0,77	1,23
Receiver Character	2,97 ± 0,56	4,33 ± 0,88	1,57
Helper Actant	2,80 ± 0,96	3,93 ± 1,41	1,33
Helper Character	2,70 ± 1,01	3,70 ± 1,37	1,13
Opponent Actant	3,10 ± 1,19	3,93 ± 1,23	1,5
Opponent Character	3,17 ± 1,21	3,70 ±1,24	1,27

extracted actantial roles were usually considered narratively coherent and textually plausible. The strongest convergence appears for Helper Actant, where the mean absolute difference between annotators is 0.21. The largest discrepancy appears for Sender Character, where the MAD reaches 1.64. This indicates that some roles, especially more abstract or interpretively demanding ones, remain difficult to identify consistently.

For X posts (Table 6), agreement scores are lower and more variable. Annotator 1's scores range from 2.70 to 3.23, while Annotator 2's scores range from 3.43 to 4.40. The higher MAD values indicate greater divergence between expert assessments. Discrepancies are especially visible for Receiver Character and Opponent Actant. These results suggest that extracting complete actantial structures from short social media posts is considerably more difficult than from news articles.

Implications of the validation

Taken together, the validation results indicate that the proposed LLM-supported pipeline can produce analytically useful outputs, especially for longer, context-rich texts such as news articles. In the news corpus, GPT-4o showed moderate alignment with expert judgement in discourse-function scoring and AAD classification, while actantial outputs were generally evaluated as plausible and coherent. This supports the use of the pipeline as a scalable exploratory tool for mapping narrative and discursive patterns in commemorative discourse.

The results for X posts were more limited. Short posts are often semantically compressed, context-dependent, and difficult to interpret even for human annotators. Lower chance-corrected reliability, weaker ICC values, and greater variability in actantial assessment show that X-based outputs should be treated cautiously. Although absolute agreement within the ±2 margin remained relatively high for some categories, this does not imply reliable item-level reproduction of expert judgement.

Overall, the validation results call for a cautious approach to using GPT-4o as an annotation tool within our exploratory pipeline. For news articles, the LLM-generated annotations showed partial agreement with expert judgement. Reliability was moderate for AAD classification and discourse-function scoring, while actantial extractions

were generally considered plausible. For X posts, however, the results were considerably weaker, suggesting that the LLM-generated annotations are not reliable enough to support solid item-level measurement. We therefore retain the X corpus only for exploratory comparison of communicative environments, and treat X-based results as indicative patterns rather than validated quantitative findings. The empirical results of this exploratory application are presented in the next section.

APPLYING THE PIPELINE: NARRATIVE CONFIGURATIONS AND DISCURSIVE MODES

This section presents the application of the LLM-supported pipeline to the two corpora on Slovenia's Day of Resistance Against the Occupier. The aim is to show how the pipeline can identify recurrent actantial configurations and examine how these configurations relate to AAD classifications and discourse-function scores. The analysis identified recurrent narrative configurations in both corpora. In the news article corpus, 12 configurations were identified. In the X corpus, 21 configurations were identified. Because the validation results showed stronger reliability for longer and more context-rich texts, the news configurations are presented in greater detail. X-based results are retained as an exploratory contrast, but are interpreted more cautiously due to the lower reliability of short-text annotation and the fragmentary nature of platform-mediated discourse.

Narrative Configurations in News Articles

The news article corpus reveals a more differentiated field of commemorative discourse than a simple binary opposition between pro-resistance and anti-communist narratives would suggest (Table 7).

The table shows three important patterns. First, the most frequent configuration (No. 6), which account for 32 % of the dataset, is indeed antagonistic, but only mildly so. Its main narrative structure centres on historical identity / preservation and historical acknowledgement, while the opponent is constructed as historical distortion. The high ideological score (76.1) indicates that this configuration is strongly involved in stabilising a particular interpretation of the past. However, its relatively high epistemic and ethical scores suggest that this ideological positioning is not epistemically empty, but is combined with claims to historical clarification and moral responsibility. Second, the news corpus is not dominated by antagonism alone. Several configurations are classified as strong agonism, strong agonism / mild deliberation, or even strong deliberation. Configurations centred on national unity, resistance remembrance, national liberation, cultural heritage preservation, historical acknowledgement, and collective solidarity indicate that contested memory discourse in the news corpus contains multiple modes of engagement. Third, high epistemic function is not limited to deliberative configurations. Some agonistic configurations have very high epistemic scores, especially configurations 7, 10, 3, 8, and 11. At the same time, democratic scores are generally lower than epistemic, ethical, and ideological scores. This indicates that the commemoration is more often articulated through historical clarification, moral evaluation, and political identity positioning than through explicit claims about inclusion or democratic accountability.

Narrative Configurations in X Posts

The X corpus was analysed using the same logic as the news corpus, but the results are interpreted more cautiously. As shown in the validation section, X posts are shorter, more implicit, and more dependent on platform context than news articles. The configurations identified in the X corpus should therefore be read as indicators of discursive orientation rather than as robust item-level measurements. The full tabular overview of all configurations is provided in Appendix E2 (GITHUB, 2026d). In the main text, we report the aggregate AAD distribution, the most frequent configuration, and the relationship between AAD modes and discourse-function scores.

The analysis identified 21 narrative configurations in the X corpus. At the level of identified narrative configurations and their associated instances, antagonistic configurations account for 50.8% of the X corpus, while deliberative configurations account for 30.3%. Strong agonistic configurations represent 7.4%, and mild agonistic configurations 2.3%. One configuration, accounting for 9.1%, could not be assigned a clear AAD classification. This distribution suggests that antagonistic framings are prominent in the X corpus. A substantial share of the corpus is classified as deliberative, and smaller agonistic configurations are also present.

The discourse-function scores further clarify how AAD modes differ internally. The main distinction between AAD modes is not simply the presence or absence of ideology. Ideological salience appears across antagonistic, agonistic, and deliberative configurations. The more important difference lies in how ideological positioning is combined with epistemic, ethical, and democratic functions (cf. Table 8).

Table 7: Narrative configurations in the news article corpus. Note: C = Cluster number; Subject, Object, Sender, Receiver, Helper, and Opponent refer to Greimasian actantial roles; AAD = antagonism–agonism–deliberation classification; % = share of instances in the corpus assigned to the configuration; Ep. = epistemic function; Et. = ethical function; Id. = ideological function; De. = democratic function.

C	Subject	Object	Sender	Receiver	Helper	Opponent	AAD	%	Ep.	Et.	Id.	De.
4	Historical Identity Preservation	Cultural Heritage Preservation	Historical Significance	National Awareness	Cultural / Collective Memory	Historical Distortion	**Strong Deliberation**	5,8	74,5	60,8	49,8	36,9
1	Resistance Remembrance	National Commemoration	Historical Significance	National Identity	Cultural / Collective Memory	Societal Division	**Strong Agonism / Mild Deliberation**	4,4	70,0	65,2	56,7	43,3
2	National Unity	X	Historical Significance	National Identity	Cultural / Collective Memory	Societal Division	**Strong Agonism / Mild Deliberation**	13,1	71,4	69,2	67,8	52,2
12	Historical Identity Preservation	Freedom and Rights	Historical Significance	Slovenian Social Future /National Awareness	Collective Solidarity	Oppression & Threats	**Strong Agonism / Mild Deliberation**	2,9	65,0	70,0	72,5	48,8
3	National Resistance	National Liberation	Historical Interpretation	National Identity / Slovenian Social Future	Collective Solidarity	Foreign Occupation	**Strong Agonism / Mild Deliberation**	10,9	78,4	65,7	68,3	44,7
8	Historical Identity Preservation / National Resistance	Historical Acknowledgment	National Identity Crisis	National Awareness	Cultural / Collective Memory	Societal Division	**Strong Agonism**	5,1	77,1	66,4	72,9	47,1
11	National Unity / Resistance Remembrance	National Unity	Historical Significance	Slovenian Social Future	Cultural / Collective Memory	Societal Division	**Strong Agonism**	12,4	76,9	72,8	71,8	56,7
5	Historical Identity Preservation	Historical Acknowledgment/ Freedom and Rights	Ethical Duty	National Awareness	Supportive Alliances	Societal Division	**Strong Agonism**	3,6	65,0	48,0	44,2	30,0
7	Historical Identity Preservation	Historical Acknowledgment	Ethical Duty	National Awareness	Historical Documentation	Historical Distortion	**Strong Agonism**	2,2	83,7	71,0	68,3	43,3
9	Historical Identity Preservation	National Commemoration	Historical Significance	National Awareness	Supportive Alliances	Historical Distortion / Oppression & Threats	**Mild Agonism**	3,6	71,2	48,0	64,2	42,0
10	Historical Identity Preservation	x	Historical Injustice	National Awareness	Cultural / Collective Memory	Societal Division	**Mild Agonism**	3,6	81,0	76,0	75,0	58,0
6	Historical Identity Preservation	Historical Acknowledgment	Ideological Agenda / Historical Interpretation	National Awareness / Slovenian Social Future	Cultural/ Collective Memory / Historical Documentation	Historical Distortion	**Mild antagonism**	32,1	67,4	62,5	76,1	43,7

Table 8: Discourse-function scores by AAD classification in the X corpus. Shares indicate the summed size of configurations within each AAD category. Discourse-function scores are weighted averages calculated by multiplying each configuration's score by its share of the X corpus and dividing by the total share of the corresponding AAD category.

AAD classification	Share of X corpus	Epistemic function	Ethical function	Ideological function	Democratic function
Antagonism	50.8%	18.2	30.2	48.3	17.6
Deliberation	30.3%	27.6	30.4	39.9	20.3
Strong agonism	7.4%	33.0	19.8	41.6	15.1
Mild agonism	2.3%	14.3	33.0	49.3	24.1
No clear classification	9.1%	8.1	8.1	13.4	6.8

Note: *The X posts were originally labelled using three categories, antagonism, agonism, and deliberation. Once narrative patterns were identified, each pattern was given an AAD label based on how the posts within it were distributed across those three categories. Patterns with a mix of labels got intermediate designations like "strong agonism" or "mild agonism". The table aggregates configurations with the same derived AAD orientation. Shares indicate the summed size of these configurations within the X corpus.*

Antagonistic configurations have the highest ideological score among the main AAD categories, but relatively low epistemic and democratic scores. This suggests that antagonistic X corpus tends to mobilise the commemoration as a political or identity marker, with limited historical elaboration and weak democratic framing. Deliberative configurations show a more balanced profile, with higher epistemic and democratic scores, while strong agonistic configurations display the highest epistemic score.

The most frequent configuration (No. 4), accounting for 20.7% of the corpus, is classified as antagonism and centres on resistance and independence, national sovereignty, national unity, collective memory, and threats to sovereignty (Table 9). Its epistemic and democratic scores are relatively low, while its ideological score is close to the overall X average. This configuration is therefore best understood as a condensed oppositional mnemonic signal rather than as an elaborated historical or democratic argument.

Overall, the X results suggest that the commemoration appears in a more compressed and platform-mediated form than in news articles, but not as a uniformly antagonistic discourse. Because of the lower validation reliability for X posts, these findings are treated as exploratory patterns rather than as a basis for strong conclusions about platform effects.

DISCUSSION: CONTRIBUTION, SCOPE, AND LIMITATIONS

This article examined whether a theory-guided, LLM-supported pipeline can operationalise an analytical approach to contested cultural memory discourse by linking three analytical dimensions: actantial narrative structure, discursive modes and discourse functions. The results show that this approach can identify recurrent narrative and discursive patterns across larger corpora, especially in longer and more context-rich texts. At the same time, the findings also clarify the limits of the method. The pipeline should be understood as an exploratory mapping procedure rather than as a fully automated measurement instrument.

Mapping narrative structures in contested memory discourse

With regard to RQ1, the findings show that LLM-supported actantial analysis can identify recurring narrative structures in contested memory discourse. In the news corpus, the extracted configurations reveal how the commemoration is organised around recurring relations between subjects, objects, senders, receivers, helpers, and opponents. These configurations show that the Day of Resistance Against the Occupier is not articulated through one stable narrative, but through several competing assignments of agency, legitimacy, opposition, and collective purpose.

Table 9: Most frequent narrative configuration in the X corpus (No. 4). The table shows the actantial role clusters, AAD classification, discourse--function scores, and share of instances for this X configuration.

Analytical component	X corpus
Subject	Resistance and Independence
Object	National Sovereignty and Independence
Sender	Historical Themes
Receiver	National Unity
Helper	Collective Memory / Collective Empowerment
Opponent	Sovereignty Threats
AAD classification	Antagonism
Share of instances	20.7%
Epistemic function	14.0
Ethical function	25.5
Ideological function	42.1
Democratic function	15.1

This is the main added value of actantial analysis. It does not merely identify topics or named actors. It reconstructs how texts organise the past into narrative relations: who is presented as acting, what goal is pursued, who benefits, who legitimises the action, who supports it, and who is constructed as an obstacle. This makes it possible to move beyond a simple binary between pro-resistance and anti-communist framings and to identify more specific narrative configurations within public discourse.

At the same time, this component involves an important limitation. Actants are inferred narrative functions, not simple named entities. The same actor may occupy different actantial positions across texts, and the same narrative function may be expressed through different formulations. We addressed this limitation by treating actantial outputs as structured interpretive data rather than as directly observable textual facts.

Linking narrative structures to AAD modes and discourse functions

With regard to RQ2, the results show that narrative conflict and discursive modes are related, but not identical. The presence of an opponent in an actantial structure does not automatically imply antagonism. Some configurations contain clear opponents but are classified as agonistic or deliberative, especially in the news corpus. This distinction is important because contested memory discourse should not be equated with polarisation as such.

The combined use of actantial analysis, AAD classification, and discourse-function scoring shows not only where conflict is narratively located, but also how it is publicly performed and what kind of work it does. In the news corpus, several configurations combine ideological salience with relatively strong epistemic or ethical scores. This may suggest that ideological positioning does not necessarily exclude historical clarification or moral reflection. In the X corpus, antagonistic configurations are more prominent, but deliberative and agonistic configurations are also present. The AAD dimension therefore helps distinguish different forms of conflict, rather than simply marking texts as conflictual or non-conflictual.

This distinction is important for the interpretation of the Slovenian case. The public discourse around the Day of Resistance is clearly conflictual and polarising but the analysis shows that conflict is not always organised in the same way. Some configurations are exclusionary and antagonistic, while others preserve disagreement in more agonistic or deliberative forms.

Comparing newspaper articles and X posts

With regard to RQ3, the comparison between newspaper articles and X posts shows that the same commemoration appears differently across communicative environments. News articles provide more elaborated and context-rich narrative configurations. They allow for more explicit historical explanation, moral evaluation, and institutional framing. The X corpus, by contrast, contains more compressed and context-dependent configurations. The dominant X configuration functions more as a condensed mnemonic and political signal than as an elaborated historical or democratic argument.

These differences should not be interpreted as evidence of a platform effect. Newspaper articles and X posts differ in genre, length, institutional setting, sampling logic, communicative function,

and validation reliability. The comparison therefore indicates how the pipeline behaves across two contrasting communicative environments, rather than demonstrating causal differences between media systems.

Methodological contribution and scope conditions

With regard to RQ4, the validation results indicate that the pipeline is more reliable for longer and more context-rich texts than for short X posts. For news articles, GPT-4o showed moderate alignment with expert judgement in AAD classification and discourse-function scoring, while actantial outputs were generally evaluated as plausible and coherent. For X posts, reliability was lower, reflecting the brevity, ambiguity, and contextual dependence of social media discourse.

This does not make the method unusable, but it defines its proper scope. The pipeline is best understood as a scalable interpretive tool. It helps detect patterns that would be difficult to identify manually across larger corpora, but it does not replace close reading, expert knowledge, or historical contextualisation. Semantic clustering further supports this function by grouping similar actants and identifying recurrent configurations, but clusters should be treated as heuristic aids rather than self-evident empirical categories.

We addressed these methodological risks by combining automated annotation with validation, clustering, and qualitative interpretation. The article therefore does not present LLM outputs as autonomous findings, but as structured inputs into an interpretive research process. The central methodological point is therefore straightforward: LLM-supported analysis can help operationalise complex theoretical concepts in cultural memory research, but only when its outputs are validated, contextualised, and interpreted with caution.

CONCLUSION

This article developed and evaluated an exploratory LLM-supported pipeline for analysing contested cultural memory discourse. Building on a broader theoretical-conceptual model for linking cultural memory and democratic practices, the study was organised around three analytical dimensions: narrative structures, discursive modes, and discourse functions. These dimensions were implemented through Greimasian actantial extraction, antagonism–agonism–deliberation (AAD) classification, and discourse-function scoring. The resulting outputs were integrated through semantic clustering and statistical analysis, interpreted qualitatively, and assessed through expert-based validation. By combining narrative-structural, discursive-mode, and discourse-function analysis with semantic clustering, statistical analysis, and expert-based validation, the article showed how theoretically complex concepts can be operationalised for larger-scale discourse analysis without reducing them to simple keywords, sentiment categories, or topic labels.

The empirical case of Slovenia's Day of Resistance Against the Occupier demonstrates the usefulness of this approach. The analysis showed that contested memory discourse is not exhausted by a binary opposition between pro-resistance and anti-communist narratives. In the news corpus, antagonistic, agonistic, and deliberative configurations coexist, while ideological positioning is often combined with epistemic and ethical claims. In the X corpus, the commemoration appears in a more compressed and platform-mediated form, with stronger antagonistic tendencies but also internal variation across deliberative and agonistic configurations. The substantive finding is therefore not that one medium is simply more polarised than the other, but that the same commemoration is organised through different narrative and discursive profiles across the two corpora. These findings support the article's main methodological claim: the proposed pipeline can help identify where conflict is narratively located, how it is discursively enacted, and what functions it performs.

At the same time, the article does not present LLM-supported analysis as a substitute for human interpretation. The validation results indicate that the method performs more reliably on longer, context-rich news articles than on short X posts. The findings based on X should therefore be treated as exploratory indicators rather than as robust item-level classifications or evidence of platform effects. More broadly, the pipeline should be understood as a scalable interpretive tool that can help structure and extend qualitative analysis, but it cannot replace close reading, historical contextualisation, or expert judgement.

Future research should further test the robustness and portability of the approach. This includes expanding expert validation samples, comparing different LLMs and prompting strategies, testing repeated-run stability, and applying the pipeline to other commemorations, languages, media systems, and political contexts. Social media analysis would also benefit from incorporating thread-level, reply-level, quotation, and external-link context rather than relying only on individual posts. Such work would help determine whether the narrative-discursive patterns identified in this case also

appear in other contested memory debates, and whether LLM-supported methods can contribute more broadly to the study of cultural memory in hybrid public spheres.

DATA AVAILABILITY

The data collection was carried out within the project CIVICA Research – The European University of Social Sciences: Research and Innovation, funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101017201. The datasets supporting the findings of this study are publicly available through the CLARIN.SI repository under the title Slovenian Day of Resistance X & news corpus 1.1: http://hdl.handle.net/11356/2216. The repository contains two CSV files: an anonymised, feature-based analytical dataset of publicly visible X posts and a dataset of traditional media articles. The X dataset includes post-level entries with derived analytical features, including Greimasian actantial coding, actant clusters, actor and character fields, author stance, antagonism scores, discourse-function indicators, HDBSCAN-based cluster information, and average cluster scores. For privacy and platform-compliance reasons, the published X dataset does not contain user-identifying information. The news dataset contains the traditional media corpus used in the analysis.

AUTHORS' CONTRIBUTION STATEMENT

Conceptualisation: MH
Methodology: MH, JK, JB, TŠ
Formal analysis: JK, JB, MH, TŠ, MRŠ
Investigation: MH, JB, JK, UL, DD, ŽO, AE
Data curation: JK
Software: JK, TŠ, JB
Visualisation: JK, MH
Writing – original draft: MH, JK, JB, TŠ
Writing – review & editing: MH, JK, JB, TŠ, MRŠ, UL, DD, ŽO, AE

KARTIRANJE SPORNEGA KULTURNEGA SPOMINA: LLM-PODPRTI PRISTOP K ANALIZI NARATIVNIH STRUKTUR, DISKURZIVNIH MODUSOV IN DISKURZIVNIH FUNKCIJ

Marjan HORVAT
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: marjan.horvat@irris.eu

Jure KORAŽIJA
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: jure.korazija@irris.eu

Jan BABNIK
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: jan.babnik@irris.eu

Tadej ŠKVORC
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: tadej.skvorc@irris.eu

Darko DAROVEC
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: darko.darovec@irris.eu

Žiga OMAN
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: ziga.oman@irris.eu

Urška LAMPE
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: urska.lampe@irris.eu

Angelika ERGAVER
Inštitut IRRIS za raziskave, razvoj in strategije družbe, kulture in okolja, Čentur 1f, 6273 Marezige, Slovenija
e-mail: angelika.ergaver@irris.eu

Marko ROBNIK-ŠIKONJA
Univerza v Ljubljani, Fakulteta za računalništvo in informatiko, Večna pot 113, 1000 Ljubljana, Slovenija
e-mail: marko.robniksikonja@fri.uni-lj.si

POVZETEK

Članek razvija in ovrednoti eksplorativni raziskovalni pristop, podprt z velikimi jezikovnimi modeli (LLM), za analizo spornega kulturnega spomina v hibridnih javnih sferah. Izhaja iz širšega teoretsko-konceptualnega okvira, ki težišče premakne od spominskih režimov k diskurzivnim modusom, skozi katere se kulturni spomin artikulira (Horvat, 2026). Metodološki postopek študije povezuje tri analitične dimenzije: Greimasov aktantski model za analizo narativnih struktur, klasifikacijo antagonizem–agonizem–deliberacija za analizo diskurzivnih modusov ter ocenjevanje epistemskih, etičnih, demokratičnih in ideoloških funkcij diskurza. Rezultati teh analiz so nato povezani s semantičnim gručenjem, statistično analizo in ekspertno validacijo. Pristop je preizkušen na primeru slovenskega praznika Dan upora proti okupatorju, ki ostaja pomembno polje nasprotujočih si interpretacij antifašističnega odpora, revolucije, povojnega nasilja, socialistične dediščine in nacionalne sprave. Članek primerja časopisne članke in javno dostopne objave na omrežju X ter preverja, ali lahko LLM-podprt postopek prepozna ponavljajoče se narativne konfiguracije in pokaže, kako se te povezujejo z različnimi diskurzivnimi modusi in diskurzivnimi funkcijami. Ugotovitve kažejo, da pristop omogoča prepoznavanje pomenljivih narativnih in diskurzivnih vzorcev, zlasti v daljših in kontekstualno bogatejših besedilih. Pri kratkih objavah na X so rezultati manj zanesljivi, zato jih članek obravnava predvsem kot eksplorativne kazalnike. Prispevek ponuja empirično preverjeno metodološko operacionalizacijo analize spornega spomina, hkrati pa jasno opredeli njene omejitve in potrebe po nadaljnji validaciji.

Ključne besede: kulturni spomin, hibridna javna sfera, družbena omrežja, deliberativna demokracija, Greimasov aktantski model, diskurzivni modusi, veliki jezikovni modeli (LLM), semantično gručenje, politični diskurz

SOURCES AND LITERATURE

Aarva, Pauliina & Marja Pakarinen (2006): Studying the Striving and Opposing Forces in Newspaper Journalism: The Actantial Model of Health Promotion. Health Promotion International, 21, 2, 160–168.

AlDayel, Abeer & Walid Magdy (2021): Stance Detection on Social Media: State of the Art and Trends. Information Processing & Management, 58, 4, 102597.

Anderson, Benedict (1983): Imagined Communities: Reflections on the Origin and Spread of Nationalism. London, Verso.

Assmann, Jan & John Czaplicka (1995): Collective Memory and Cultural Identity. New German Critique, 65, Spring-Summer, 125–133.

Bächtiger, André, Shikano, Susumu, Pedrini, Seraina & Mirjam Ryser (2010): Measuring Deliberation 2.0: Standards, Discourse Types, and Sequenzialization. Proceedings of the ECPR General Conference. Potsdam, September 2009.

Beauchamp, Nick (2020): Modeling and Measuring Deliberation Online. In: Welles Foucault, Brooke & Sandra González-Bailón (eds.): The Oxford Handbook of Networked Communication. Oxford, Oxford University Press, 321–349.

Brown, Tom et al. (2020): Language Models Are Few-Shot Learners. Advances in Neural Information Processing Systems, 33, 1877–1901.

Bull, Anna Cento & Hans Lauge Hansen (2016): On Agonistic Memory. Memory Studies, 9, 4, 390–404.

Chadwick, Andrew (2017): The Hybrid Media System: Politics and Power. 2nd edition. New York, Oxford University Press.

Cinelli, Matteo, De Francisci Morales, Gianmarco, Galeazzi, Alessandro & Michele Starnini (2021): The Echo Chamber Effect on Social Media. Proceedings of the National Academy of Sciences, 118, 9, e2023301118.

DeepSeek-AI et al. (2024): DeepSeek-V3 Technical Report. arXiv preprint, 2412.19437v2.

Elfes, Jan (2024): Mapping News Narratives Using LLMs and Narrative-Structured Text Embeddings. arXiv preprint, 2409.06540v1.

Fournier-Tombs, Élise & Michael K. MacKenzie (2021): Big Data and Democratic Speech: Predicting Deliberative Quality Using Machine Learning Techniques. Methodological Innovations, 14, 2.

GITHUB (2026a): REFMF / Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach. Appendix A: Prompt Structure for Greimas's Narrative Analysis. https://github.com/REFMF/Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach-/blob/main/Appendix%20A%3A%20Prompt%20Structure%20for%20Greimas%E2%80%99s%20Narrative%20Analysis (last access: 2026-06-16).

GITHUB (2026b): REFMF / Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach. Appendix B: Prompt Structure for AAD Analysis. https://github.com/REFMF/Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach-/blob/main/Appendix%20B%3A%20Prompt%20Structure%20for%20AAD%20Analysis (last access: 2026-06-16).

GITHUB (2026c): REFMF / Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach. Appendix C: Prompt Structure for Discourse Function Analysis. https://github.com/REFMF/Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach-/blob/main/Appendix%20C%3A%20Prompt%20Structure%20for%20Discourse%20Function%20Analysis (last access: 2026-06-16).

GITHUB (2026d): REFMF / Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach. Appendix E2-Tabular Overview of Narrative Structures on X.pdf. https://github.com/REFMF/Cultural-Memory-from-Antagonism-to-Deliberation-in-Social-Media-AI-Approach-/blob/main/Appendix%20E2-Tabular%20Overview%20of%20Narrative%20Structures%20on%20X.pdf (last access: 2026-06-16).

Godeša, Bojan (1995): Kdor ni z nami, je proti nam: Slovenski izobraženci med okupatorji, Osvobodilno fronto in protirevolucionarnim taborom. Ljubljana, Cankarjeva založba.

Grattafiori, Aaron et al. (2024): The LLaMA 3 Herd of Models. arXiv preprint, 2407.21783v3.

Greimas, Algirdas Julien (1966/1983): Structural Semantics: An Attempt at a Method. Translated by Daniele McDowell, Ronald Schleifer & Alan Velie. Lincoln and London, University of Nebraska Press.

Hobsbawm, Eric (1983): The Invention of Tradition. Cambridge, Cambridge University Press.

Horvat, Marjan (2026): From Memory Regimes to Discursive Modes: A Theory-driven Framework for Analysing Cultural Memory in Hybrid Public Spheres. Zenodo. DOI: 10.5281/zenodo.18787238.

Hoskins, Andrew (2011): Media, Memory, Metaphor: Remembering and the Connective Turn. Parallax, 17, 4, 19–31.

Hribar, Spomenka & Tine Hribar (2021): Slovenski razkol in slovenska sprava. Ljubljana, Ciceron.

Jiang, Yicun (2017): The Logic of Legal Narrative. Language and Semiotic Studies, 3, 4, 17–31.

Kawintiranon, Kornraphop & Lisa Singh (2022): PoliBERTweet: A Pre-trained Language Model for Analyzing Political Content on Twitter. Proceedings of the Thirteenth Language Resources and Evaluation Conference, 7360–7367.

Koo, Terry K. & Mae Y. Li (2016): A Guideline of Selecting and Reporting Intraclass Correlation Coefficients for Reliability Research. Journal of Chiropractic Medicine, 15, 2, 155–163.

Kranjc, Gregor Joseph (2013): To Walk with the Devil: Slovene Collaboration and Axis Occupation, 1941–1945. Toronto, University of Toronto Press.

Lampe, Urška & Gorazd Bajc (2025): Contested Memories and the Dynamics of Remembrance in Slovenia related to the Second World War. Balkanistic Forum, 34, 2, 261–284.

Langer, Roy (2000): Germans Welcome? Reflections on Media Discourse Analysis in Practice. International Journal of Applied Linguistics, 10, 1, 7–24.

Lévi-Strauss, Claude (1955/1963): Structural Anthropology. Translated by Claire Jacobson & Brooke Grundfest Schoepf. New York, Basic Books.

Levy, Daniel & Natan Sznaider (2002): Memory Unbound: The Holocaust and the Formation of Cosmopolitan Memory. European Journal of Social Theory, 5, 1, 87–106.

Mansbridge, Jane & John Parkinson (eds.) (2012): Deliberative Systems: Deliberative Democracy at the Large Scale. Cambridge, Cambridge University Press.

Mansbridge, Jane, Bohman, James, Chambers, Simone, Christiano, Thomas, Fung, Archon, Parkinson, John, Thompson, Dennis F. & Mark E. Warren (2012): A Systemic Approach to Deliberative Democracy. In: Parkinson, John & Jane Mansbridge (eds.): Deliberative Systems: Deliberative Democracy at the Large Scale. Cambridge, Cambridge University Press, 1–26.

McPherson, Miller, Smith-Lovin, Lynn & James M. Cook (2001): Birds of a Feather: Homophily in Social Networks. Annual Review of Sociology, 27, 1, 415–444.

Meyer, Erik (2008): Memory and Politics. In: Erll, Astrid & Ansgar Nünning (eds.): Cultural Memory Studies: An International and Interdisciplinary Handbook. Berlin, Walter de Gruyter, 173–180.

Mouffe, Chantal (1999): Deliberative Democracy or Agonistic Pluralism? Social Research, 66, 3, 745–758.

Mozetič, Igor, Grčar, Miha & Jasmina Smailović (2016): Multilingual Twitter Sentiment Classification: The Role of Human Annotators. PLOS ONE, 11, 5, e0155036.

Németh, Renáta (2023): A Scoping Review on the Use of Natural Language Processing in Research on Political Polarization: Trends and Research Prospects. Journal of Computational Social Science, 6, 1, 289–313.

Nienass, Benjamin (2023): Between Conflict and Consensus. In: Gutman, Yifat & Jenny Wüstenberg (eds.): The Routledge Handbook of Memory Activism. London, Routledge, 454–458.

Nora, Pierre (1989): Between Memory and History: Les Lieux de Mémoire. Representations, 26, Spring, 7–24.

Olick, Jeffrey K., Sierp, Aline & Jenny Wüstenberg (2023): Introduction: Taking Stock of Memory Studies. Memory Studies, 16, 6, 1399–1406.

OpenAI et al. (2023): GPT-4 Technical Report. arXiv preprint, 2303.08774v6.

Oswald, Lisa (2022): Automating the Analysis of Online Deliberation? A Comparison of Manual and Computational Measures Applied to Climate Change Discussions. SocArXiv Papers Preprint.

Pariser, Eli (2011): The Filter Bubble: What the Internet Is Hiding from You. New York, Penguin Press.

Pirjevec, Jože (1995): Jugoslavija: Nastanek, razvoj ter razpad Karadjordjevićeve in Titove Jugoslavije. Koper, Lipa.

Propp, Vladimir (1928/1968): Morphology of the Folktale. Translated by Laurence Scott. Austin, University of Texas Press.

Prunk, Janko (1992): Slovenski narodni vzpon, narodna politika (1768–1992). Ljubljana, Državna založba Slovenije.

Pshenychnykh, Anastasiya, Pfoser, Alena & Sabina Mihelj (2024): When Monument Battles Go Digital: Russian–Ukrainian Conflicts over Material Heritage on Telegram. Social Media + Society, Online First.

Roberts, Hal et al. (2021): Media Cloud: Massive Open Source Collection of Global News on the Open Web. Proceedings of the International AAAI Conference on Web and Social Media, 15, 1034–1045.

Rothberg, Michael (2009): Multidirectional Memory: Remembering the Holocaust in the Age of Decolonization. Stanford, Stanford University Press.

Rutten, Ellen, Fedor, Julie & Vera Zvereva (eds.) (2013): Memory, Conflict and New Media: Web Wars in Post-Socialist States. London, Routledge.

Salama, Amir H. Y. (2021): Towards a Discourse-Semantic Approach to Visual Narrative Analysis: The Case of the BBC News Story of Islamic State (IS). Pragmatics and Society, 12, 2, 189–223.

Schoor, Carola (2021): Probing into Populism's Core: An Analysis of the Deep Semio-Linguistic Structure Underlying Populism. Critical Discourse Studies, 18, 2, 226–244.

Steenbergen, Marco R., Bächtiger, André, Spörndli, Markus & Jürg Steiner (2003): Measuring Political Deliberation: A Discourse Quality Index. Comparative European Politics, 1, 1, 21–48.

Tangherlini, Timothy R., Shahsavari, Shadi, Shahbazi, Behnam , Ebrahimzadeh, Ehsan & Vwani Roychowdhury (2020): An Automated Pipeline for the Discovery of Conspiracy and Conspiracy Theory Narrative Frameworks: Bridgegate, Pizzagate and Storytelling on the Web. PLOS ONE, 15, 6, e0233879.

Törnberg, Petter (2022): Echo Chambers and Viral Misinformation: Modeling Fake News as Complex Contagion. PLOS ONE, 13, 9, e0203958.

Tveskov, Mark Axel & Ashley Ann Bissonnette (2023): Conflict Archaeology, Historical Memory, and the Experience of War: Beyond the Battlefield. Foreword by Paul A. Shackel. Gainesville, FL, University Press of Florida.

van Dijck, José (2007): Mediated Memories in the Digital Age. Stanford, CA, Stanford University Press.

Wang, Yong & Carl W. Roberts (2005): Actantial Analysis: Greimas's Structural Approach to the Analysis of Self-Narratives. Narrative Inquiry, 15, 1, 51–74.

Willaert, Tom (2023): A Computational Analysis of Telegram's Narrative Affordances. PLOS ONE, 18, 11, e0293508.

Winter, Jay (2008): Sites of Memory and the Shadow of War. In: Erll, Astrid & Ansgar Nünning (eds.): Cultural Memory Studies: An International and Interdisciplinary Handbook. Berlin, Walter de Gruyter, 61–76.

Wodak, Ruth & John E. Richardson (2013): Analyzing Fascist Discourse: European Fascism in Talk and Text. London, Routledge.

Wüstenberg, Jenny (2023): Memory and Democracy: A Complex Relationship. Observing Memories, 7. European Observatory on Memories.