



# The Use of Generative Artificial Intelligence for Interpreting Emotions in Corpus-Based Critical Discourse Analysis

Carlos Perrián-Pascual<sup>1</sup>

Received: 30 July 2025 / Accepted: 13 October 2025  
© The Author(s) 2025

## Abstract

Critical discourse analysis has traditionally relied on the dissection of a limited number of texts, leading to criticisms regarding the subjectivity and lack of empirical rigour in its findings. In response, the integration of corpus linguistics provided more robust quantitative methods, but these have primarily focused on the statistical significance of lexical patterns, with limited capacity for deeper interpretive analysis. Recent developments in generative artificial intelligence offer new possibilities for interpreting complex discourse phenomena. This paper addresses how a large language model can interpret emotions within a critical discourse analysis framework. In this regard, our experiment aimed to assess affective polarisation in a corpus of YouTube comments about Brexit. The method primarily involves the identification of emotion-oriented lexical markers through lexical semantic resources, followed by their interpretation and contextualisation via the large language model. The findings suggest that large language models can enhance the analytical process by uncovering latent discourse patterns and offering nuanced interpretations. However, human analysts should still play an active role in guaranteeing the accuracy and validity of results.

**Keywords** Large language model · Affective polarisation · Critical discourse analysis · Emotion · Brexit

## Introduction

Artificial intelligence (AI) can be described as “systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals” (Sheikh et al., 2023: 19). Recent advancements

---

✉ Carlos Perrián-Pascual  
joepas3@upv.es

<sup>1</sup> Universitat Politècnica de València, Valencia, Spain

in generative AI, i.e. large language models (LLMs), have shown potential in processing and analysing unstructured text with prompts (i.e. textual instructions given to the model). LLMs are pre-trained on massive volumes of human-generated text data to produce responses that mimic human reactions. Despite their extraordinary abilities to produce meaningful conversations with humans, LLMs do not have cognitive processes for managing semantic contents (e.g. thinking, reasoning or understanding) similar to those in the human brain. Instead, LLMs operate on the language from a structural perspective, where words are converted into numerical representations and sentences are constructed using probabilistic components of language. In the field of generative AI, various studies have recently demonstrated the use of LLMs in qualitative research (cf. Section “[Generative AI in Qualitative Research](#)”). However, investigating how qualitative analysis, which is based on the interpretation of patterns of meaning, can be conducted by LLMs, which work on structural and probabilistic elements of language, is a research topic that is still in its infancy (De Paoli, 2024). For this reason, Christou (2023) noted that researchers must acknowledge the strengths and weaknesses of generative AI for qualitative analysis.

We chose the research topic of affective polarisation because it has been a growing phenomenon in recent years. Affective polarisation is grounded on social identity theories (Tajfel, 1970; Tajfel & Turner, 1979), which underscore that, under conditions of groups in conflict, there is not only positive sentiment for one’s own group (i.e. in-group) but also negative sentiment towards those identifying with opposing groups (i.e. out-groups). In particular, affective polarisation is related to the intergroup threat theory (Stephan et al., 2015). As explained by Renström et al. (2023), social groups become a significant part of the self-definition of their members. As a group becomes an extension of one’s identity, the desire to protect the group against external threats also increases. As social groups become relevant parts of one’s identity, threats directed at the group are perceived as threats directed at oneself, evoking defensive reactions. Therefore, threats against the in-group may strengthen social identification and increase intergroup differentiation, which is an essential part of affective polarisation. An important aspect of the intergroup threat theory is that threats are perceived, so they may or may not relate to actual threats. However, when threats are perceived, this leads to some form of adverse emotional reaction (e.g. fear, anger, or anxiety), which will finally influence our behaviour and cognition.

Affectively polarised discourse can also be studied in terms of stance-taking, which is defined as “a public act by a social actor, achieved dialogically through overt communicative means, of simultaneously evaluating objects, positioning subjects (self and others), and aligning with other subjects, with respect to any salient dimension of the sociocultural field” (Du Bois, 2007: 163). Based on his stance triangle, the interrelated dimensions of evaluation, positioning, and alignment converge to create antagonistic group identities. For example, in a given utterance, the stancetaker can evaluate the stance object negatively (e.g. through evaluative predicates, metaphors, sarcasm and irony), position the speaker as morally superior (e.g. through pronoun usage and modality markers), and align with other speakers who share this evaluation (e.g. through agreement tokens and repetitions). If the interlocutors agree, they reaffirm in-group solidarity through the shared object, strengthening social identification. If the interlocutors disagree, however, they can generate disalignment. In this case,

the speaker's negative evaluation of the stance object is contested, which reinforces the interlocutors' antagonistic ideological positions, increasing social differentiation with the out-groups. Therefore, stance-taking can help explain the pragmatic process of social identity construction, together with the perceptions of competition between "us" and "them".

To examine affective polarisation, we experimented with a corpus of 631 YouTube comments about the nationwide referendum held on 23 June 2016 with the question "Should the United Kingdom (UK) remain a member of the European Union (EU) or leave the EU?". Despite the relatively small size of the corpus, this study yielded meaningful findings due to its greater emphasis on depth rather than breadth. In particular, the qualitative analysis of 59 different emotion-oriented lexical markers allowed us to get valuable insights into the capability of the LLM to explore affective polarisation in the Brexit debate. Therefore, rather than seeking statistical generalisation, the focus was on in-depth analysis.

It should also be noted that leveraging an LLM to analyse affective polarisation on a corpus is much more than adapting tools to analyse data. It involves integrating a solid methodology in the analysis. In this regard, as we aim to explore social problems related to Brexit (short for *British exit*) by integrating language study and social analysis, we adopted the critical discourse analysis (CDA) approach. Indeed, this approach encompasses a set of theories and methods that deal with the study of social issues through text, which is considered a social product and thus reflects societal values (Fairclough, 1992; Van Dijk, 1995; Dijk, 2003; Fairclough & Wodak, 1997; Wodak, 2001).<sup>1</sup> For our study, we implemented a prototype based on a methodology integrating Fairclough's (1989, 1992) tripartite discourse-analysis model into Mullet's (2018) CDA framework. Therefore, this article resulted from the investigation conducted to answer the following research question:

Can we use an LLM to automatically examine affective polarisation in YouTube comments within the framework of CDA?

To address this question, our research makes several key contributions that bridge generative AI with the critical analysis of affective polarisation. First, we proposed a methodological framework for doing CDA with LLMs. Second, we developed an LLM-aided tool based on such a methodology. Finally, we discussed the implications of using such an automated tool to support a critical approach to digital discourse analysis in affective polarisation. The remainder of this article is organised as follows. Section "[Generative AI in Qualitative Research](#)" describes the most relevant works on using generative AI in qualitative research. Section "[Methodology](#)" presents the proposed methodological framework. Section "[Results and Discussion](#)" examines the results of the experiment. Finally, our conclusions are drawn in Section "[Conclusions](#)".

---

<sup>1</sup> Todolí Cervera et al. (2006) provided a description of the theoretical framework of the main models that have emerged in this approach.

## Generative AI in Qualitative Research

Most studies that have employed LLMs (in particular, GPT) in qualitative research have focused explicitly on Reflexive Thematic Analysis (RTA), a popular qualitative method used in the humanities and social sciences to identify patterns of meaning within text data. Although RTA is not the same as CDA, both methods are two sides of a coin, which can be used to address the same societal concerns with a critical lens (Kogen, 2024). Such studies have demonstrated that RTA is suited for experiments with LLMs, where Braun and Clarke's (2006) thematic-analysis framework was employed. For example, Breazu et al. (2025) employed GPT-4 to analyse hate speech from YouTube comments about Roma beggars in Sweden. Hitch (2024) analysed a newspaper article about COVID-19 in Australia, focusing on the national response to this health problem. Turobov et al. (2024) employed a custom GPT model to analyse United Nations' policy documents. Khan et al. (2024) analysed a small sample of statements (from social media and government reports) concerning the Robodebt controversy, in which the Australian government created an unlawful computer program to recover welfare debts, falsely accusing thousands of citizens. De Paoli (2024) performed RTA with an LLM on two small datasets of interviews: one about the role of video games in culture, economy and education, and the other about using quantitative data to teach undergraduate courses.

As explained by Breazu et al. (2025), LLMs can be employed for RTA in two different manners: (a) inductively, where themes emerge from the data without being influenced by the researcher's preconceptions or theoretical framework (i.e. data-driven approach) and (b) deductively, which starts with a predefined set of themes that the researcher expects to find in the data (i.e. theory-driven approach). Whereas inductive RTA enables researchers to explore how LLMs independently identify themes directly from the dataset, researchers guide the model by introducing theoretical concepts and definitions of specific themes in deductive RTA. Adopting an inductive approach to RTA is a challenging problem since the LLM involves a position of creativity when discovering and presenting the themes directly from the data.

Moreover, Hitch (2024) emphasised two key contributions of LLM-driven qualitative research:

- (a) Analysis augmentation: LLMs can identify patterns in the data that are not immediately apparent to human researchers, leading to potentially novel insights that augment the analysis.
- (b) Analysis review: LLMs can provide an alternative analysis against which human researchers can review their own analysis.

However, most of the above studies highlight two challenges. On the one hand, there is a tendency for the LLM to generate more descriptive than interpretive responses. That is, LLMs can process large corpora to identify broad patterns of meaning but often lack the deeper, inferential reasoning that human analysts apply when interpreting meanings. The challenge lies in improving the synergy between the capabilities of LLMs and human intelligence. On the other hand, analysts should be aware of the potential risks of hallucinations, i.e. when LLMs generate false content, contributing

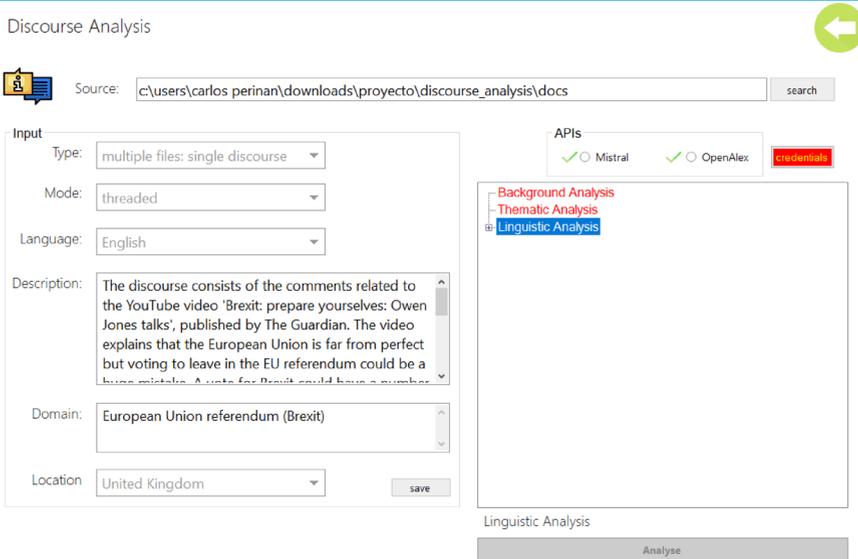
to the spread of disinformation. For these reasons, the importance of manual review for verifying AI-generated data is underscored. Therefore, instead of completely outsourcing the task to AI, LLMs should be considered tools that complement, rather than replace, the critical perspective of analysts, using the LLM as a collaborative research assistant.

## Methodology

CDA takes several different approaches and incorporates a variety of methods. In this context, Mullet (2018) presented a general analytic framework suitable for various disciplines and research problems. The framework was designed for flexibility and simplicity, so it condenses many CDA approaches into seven levels of analysis. However, given our interest in a socio-cultural perspective of discourse analysis, we also considered Fairclough's three-dimensional model (1989, 1992): description, interpretation, and explanation. The first dimension is concerned with the description of the linguistic elements of the text. Assuming that linguistic elements represent a choice made among available options, the second dimension provides an initial interpretation of why these linguistic elements were used in the text. Finally, based on the previous interpretation of the data, the third dimension aims to establish the relationships between the discourse and a specific social, cultural, and political framework, allowing for an analysis that extends from an initial microtextual analysis to a macroanalysis of society with respect to a specific phenomenon. Therefore, we have adapted some of the levels of Mullet's general framework to give greater prominence to the three dimensions of Fairclough's socio-cultural model. As a result, our methodological framework consists of three stages of analysis (i.e. Background Analysis, Thematic Analysis, and Linguistic Analysis), where the last stage is organised into three steps (i.e. Description, Interpretation, and Explanation).

For our experiment, we compiled a small corpus of 631 comments, including 177,803 characters and 33,058 words, from the YouTube video *Brexit: prepare yourselves. Owen Jones talks* published by The Guardian on 21 June 2016.<sup>2</sup> The process of selecting the specific texts to include in the corpus consisted of three steps. First, we compiled a pool of YouTube videos on the Brexit debate using two criteria: they had to be (a) published by a British media outlet (e.g. newspaper or TV channel) and (b) released between 1 January and 22 June 2016. Second, we excluded any comments posted outside this time frame, as we aimed to focus on how people constructed their identities, expressed hostility towards the out-group, and attempted to persuade others before the referendum date. Choosing this sampling window helped us avoid the noise that reactions of celebration, resentment, or blame in the aftermath of the vote would have introduced. Finally, in the third step, we selected the video with the largest number of direct comments, leaving out replies. By focusing on a single video rather than multiple ones, we ensured that all commenters were responding to the same stimulus, which strengthens the validity of the analysis within a unified discursive space. Finally, to facilitate a fine-grained CDA with the LLM, we split

<sup>2</sup> [https://www.youtube.com/watch?v=3FqAaD\\_IsRw](https://www.youtube.com/watch?v=3FqAaD_IsRw).



**Fig. 1** Interface of the prototype

the corpus into five segments, each containing a similar number of characters. In this regard, the process of analysis was structured into three sequences: (a) each segment was analysed individually, (b) results were compared collectively, and (c) a global analysis of the findings was presented. Our methodological framework was applied to sequence (a), where all corpus segments shared the same Background Analysis, but the Thematic and Linguistic Analyses were individually conducted for each segment. It should also be noted that sequence (a) was conducted automatically, whereas sequences (b) and (c) were performed manually.

To guide the AI system to perform a deeper qualitative analysis, the following information about the video was provided through the interface of our prototype, shown in Fig. 1:<sup>3</sup>

**DESCRIPTION:** The discourse consists of the comments related to the YouTube video “Brexit: prepare yourselves. Owen Jones talks”, published by The Guardian. The video explains that the European Union is far from perfect, but voting to leave in the EU referendum could be a huge mistake. A vote for Brexit could have a number of consequences domestically and at a European level. These could include a power grab by the Conservative right, rising anti-immigrant sentiment and fresh attacks on workers’ rights and the NHS. However, British people may also have to take an economic hit and Europe will be unlikely to want to give them a good deal as it seeks to discourage further disintegration of the union. If they vote Leave, they have to be prepared for the consequences.

<sup>3</sup> Most of our description was taken from the original description posted with the video.

DOMAIN: European Union referendum (Brexit).

LOCATION: United Kingdom.

LANGUAGE: English.

In the remainder of this section, we present the methodological framework used to guide the process of CDA, which was computationally implemented as a pipeline of ten tasks:<sup>4</sup>

Task 1: Exploring the background of the corpus by examining the historical, political, economic and social contexts of the discourse with an LLM. Mistral NeMo was the LLM employed in this study. The LLM prompt includes information about the video (e.g. domain and location).

Task 2: Generating up to ten queries for resource retrieval with an LLM prompt that includes information about the video (e.g. domain).

Task 3: For each query automatically generated in Task 2, retrieving bibliographical citations of scholarly publications through OpenAlex (Priem et al., 2022).

Task 4: For each query automatically generated in Task 2, retrieving news items through GDELT, based on the language in which the articles were originally published and the country in which the press outlet is located.

Task 5: Describing the themes and subthemes of the corpus, which are presented along with meaningful descriptions and representative quotations from the texts. Apart from the corpus, the task also requires information about the video (e.g. description, domain, and location) in the LLM prompt.

Task 6: Selecting a specific linguistic feature from our CDA-oriented feature database. In this study, we focused on emotions.

Task 7: Identifying lexical markers in the corpus corresponding to the linguistic feature in Task 6. In the case of emotion-oriented lexical markers, this task consists of two steps. First, the system employs SyntagNet (Maru et al., 2019) to determine the WordNet synset (Fellbaum, 1998) of each lemma in the corpus. Second, the sentiment and emotion associated with each synset are automatically identified using SentiWordNet (Esuli & Sebastiani, 2006) and SentiSense (Carrillo de Albornoz et al., 2012), respectively. We only consider the lemmas whose affective states align with their polarity values; for example, words expressing satisfaction and attraction should be positive, whereas dissatisfaction and repulsion are negative. In affective computing, for example, emotion recognition is considered a subtask of polarity detection, a classification task with outputs such as positive and negative (Cambria, 2016). In this study, a lexical marker is classified as “positive” if its positivity score in SentiWordNet is at least 0.375 and greater than its negativity score, and as “nega-

<sup>4</sup> Appendix 1 includes the specifications of the resources used in our experiment.

**Table 1** Emotion taxonomy

Group	Type	Subtype
Goal-seeking emotions	Attention-grabbing	Surprise
Goal-achievement emotions	Satisfaction	Calmness
		Joy
	Dissatisfaction	Anger
		Fear
		Sadness
Goal-relation emotions	Attraction	Like
		Love
	Repulsion	Disgust
		Hate

tive” if its negativity score is at least 0.375 and greater than its positivity score.<sup>5</sup> This threshold ensures that only words with sufficiently strong sentiment polarity are labelled, thereby reducing ambiguity in polarity detection. On the other hand, our emotion taxonomy is grounded on Appraisal Theory (Martin & White, 2005), particularly on the refinement performed by Benítez-Castro & Hidalgo-Tenorio’s (2019), whose psychologically inspired emotion spectrum is divided into three conceptual groups (i.e. goal-seeking emotions, goal-achievement emotions, and goal-relation emotions), as shown in Table 1.

Task 8: Interpreting each lexical marker identified in Task 7 with an LLM prompt that includes (a) information about the video (e.g., description, domain, and location) and (b) the themes discovered in Task 5.

Task 9: Identifying and explaining patterns of meaning based on the interpretation of lexical markers in Task 8. The task also requires information about the video (e.g. domain and location) in the LLM prompt.

Task 10: Contextualising the explanations of the patterns of meaning provided in Task 9 based on the historical, political, economic and social contexts described in Task 1. The purpose of providing such contextual information is to enhance the accu-

<sup>5</sup> Let  $p, n, o \in [0,1]$  denote the positivity, negativity and objectivity scores in SentiWordNet, respectively, where  $o = 1 - (p + n)$ , and  $t \in [0,1]$  denotes the threshold value. The choice of  $t$  is supported by the tradeoff between coverage and polarity strength. On the one hand, we use “coverage” to refer to the proportions of positively and negatively polarised synsets, based on the following definition of the classification label  $y \in \{ \text{positive, negative, none} \}$ :

$$y = \begin{cases} \text{positive,} & \text{if } p \geq t \text{ and } p > n, \\ \text{negative,} & \text{if } n \geq t \text{ and } n > p, \\ \text{none,} & \text{otherwise.} \end{cases}$$

On the other hand, we use “polarity strength” to reflect the minimum score required for a synset to be considered polarised. In this regard, we noted that coverage is relatively high at lower thresholds (e.g.  $t \leq 0.25$ ), but many of the retained synsets still have large objectivity scores, meaning that they remain predominantly neutral. In contrast, polarity strength increases at higher thresholds (e.g.  $t \geq 0.5$ ), but coverage drops sharply, leaving only a small fraction of synsets available for classification. Therefore, based on the data distributions, we concluded that 0.375 is a reasonable threshold for identifying clearly polarised synsets in SentiWordNet.

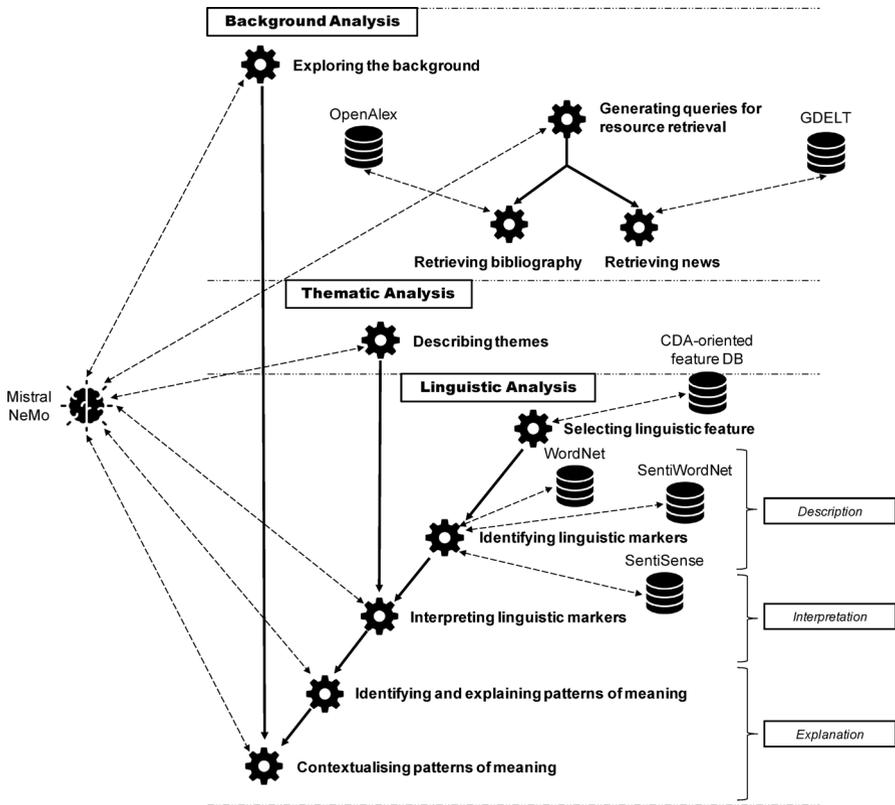


Fig. 2 Pipeline of LLM-driven CDA

racy of qualitative research. The task also requires information about the video (e.g. domain and location) in the LLM prompt.

To conclude, Fig. 2 illustrates the entire process of LLM-driven CDA, which has been deployed in our prototype.

To effectively incorporate the LLM into this pipeline, we provided XML-formatted outputs through format-restricting instructions included in the prompts. In particular, we employed the NL-to-Format technique, which consists in separating “content generation” from “format adherence” (Tam et al., 2024), letting the LLM focus on content richness first and then apply structure later (i.e. wrapping the content into XML), so that the expressiveness of responses is not constrained.

## Results and Discussion

### Background Analysis

#### Historical, political, Economic and Social Contexts

First, the LLM placed the historical context from the mid-20th century to 2016, primarily describing the UK's ambivalence towards European integration through a series of milestones, as seen with the UK's entry into the EEC in 1973, but also with the rise of the UK Independence Party (UKIP) that resulted from the growing Euroscepticism in 1980s and 1990s, among others. In this context, the LLM provided factual information accurately, as in the statement *The referendum was called by Prime Minister David Cameron in 2013 as a means to address the growing Euroscepticism within his Conservative Party.*<sup>6</sup> However, this type of description presents two main shortcomings; it fails to provide a detailed analysis and lacks supporting evidence from scholarly or news sources. For example, the above statement could have been expressed more precisely as:

In 2013, some Conservative MPs were threatening to join UKIP, which had been gaining support with its strong anti-EU stance, a situation that led to Douglas Carswell and Mark Reckless defecting to UKIP in 2014 (Johnston, 2014).

Second, the LLM presented the political context behind the Brexit referendum, tracing its origins back to the early 2000s with the enlargement of the EU (2004), which was followed by two general elections in the UK (2010 and 2015). Although the political context included more factual detail than the historical context, it still lacks citations to sources that would help validate the model's claims, as it also occurred within the economic and social contexts.

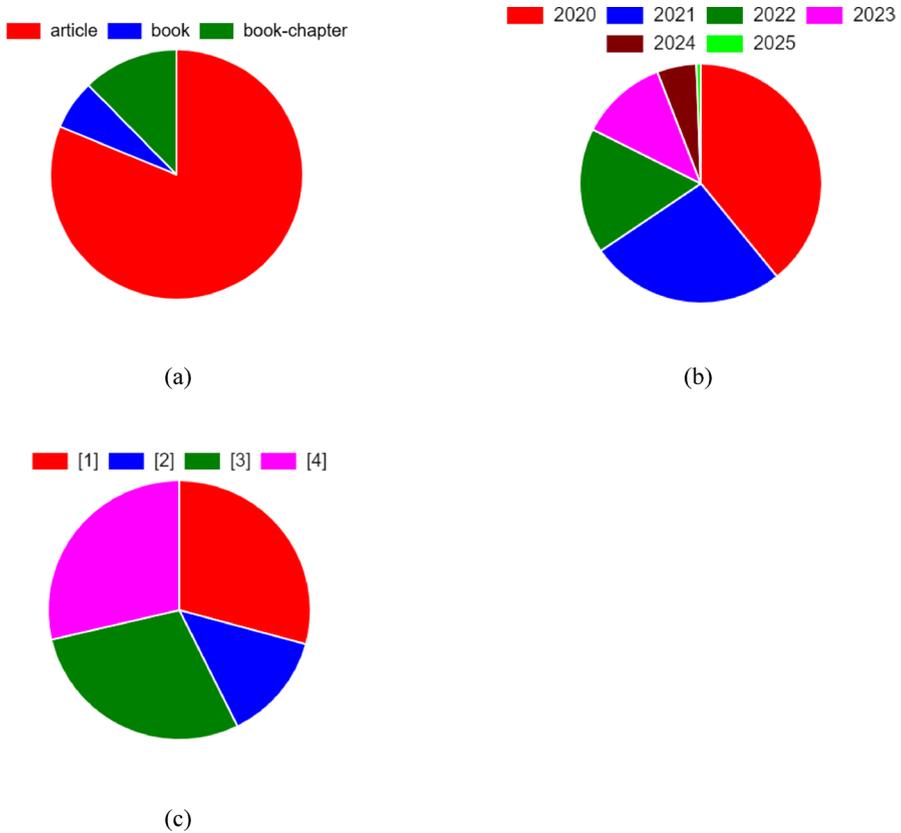
Third, the LLM structured the economic context around the economic arguments for and against Brexit. On the one hand, pro-Leave arguments included the desire for independent trade deals, reduced regulations, and immigration control. On the other hand, pro-Remain arguments emphasised the risks of economic disruption (e.g. potential recession) and the loss of access to the single market.

Fourth, the LLM emphasised the growing discontent among citizens before the referendum within the social context. In particular, Brexit resulted from decades of Euroscepticism, which had become entrenched among some segments of the British population. This sentiment was exacerbated by a variety of factors, including perceptions of excessive immigration and a sense of loss of national sovereignty.

### Complementary Resources

Here are the four LLM-generated queries in the form of comma-separated strings of keywords: [1] "Brexit, impact, economy", [2] "Brexit, timeline", [3] "Brexit, negotiations", and [4] "Brexit, referendum, facts". These queries retrieved 171 bibliographical citations and 289 news items. On the one hand, Fig. 3 shows the distribution of the bibliographical citations.

<sup>6</sup> In this article, italics are used to indicate the LLM-generated text.



**Fig. 3** Distribution of bibliographical citations based on **a** type of publication, **b** year of publication, and **c** query

**Table 2** Publications in bibliographical citations

Publication	Type	#
JCMS Journal of Common Market Studies	Journal	7
Routledge eBooks	Book	7
Edward Elgar Publishing eBooks	Book	6
Oxford University Press eBooks	Book	6
Journal of Contemporary European Studies	Journal	5
SSRN Electronic Journal	Journal	3

To illustrate, we also present the most popular sources of publication in the bibliographical citations, including the number of works cited for each source (Table 2).

On the other hand, the news items retrieved were published between March and June 2025, describing the current situation of Brexit. Therefore, these news items were irrelevant to this study since our goal is to examine the situation before the EU referendum.

## Thematic Analysis

The LLM identified “EU Referendum: Leave vs. Remain” as the only primary theme of the discourse. In contrast, different secondary themes were recognised as the most significant for each corpus segment. To illustrate, here are some of the most frequent secondary themes:

- (a) Arguments for Leave: Commenters argue for leaving the EU.
- (b) Arguments for Remain: Commenters argue for remaining in the EU.
- (c) Media and Polls: Commenters discuss the role of the media in the referendum campaign and the latest polls.
- (d) EU issues: Commenters discuss issues such as (i) the possibility of EU reform and alternative approaches to the UK’s relationship with the EU, (ii) the EU’s role in European politics, and (iii) their opinions about the EU and its institutions.

## Linguistic Analysis

### Identifying Emotion-Oriented Lexical Markers

Table 3 shows the distribution of emotion-oriented lexical markers based on types, subtypes and valence. Following Renström et al. (2023), we evaluated negative and positive emotions to assess affective polarisation.

As we employed natural language processing (NLP) techniques for emotion recognition, the next question is whether we could have used the LLM for this task. To this end, we tried to detect the emotion-oriented lexical markers in the first corpus segment with the LLM on the one hand and with our knowledge-based approach (i.e. WordNet, SentiWordNet, and SentiSense) on the other. We found that our prototype detected 18 words that were not recognised by the LLM (e.g. *complain*, *despise*, *nasty*, and *shameful*), but the LLM detected 23 words that were not recognised by our knowledge-based approach (e.g. *anxious*, *scared*, *upset*, and *worried*). However, the problem with the LLM is that it misidentified *believe* as an emotion-oriented lexical marker and identified five other markers that were not found in this corpus segment. Therefore, the LLM introduces more noise than NLP in emotion recognition. In this regard, we can conclude that LLM-driven CDA can be strengthened by integrating NLP with generative AI, since each offers distinct yet complementary capabilities.

**Table 3** Distribution of emotion-oriented lexical markers

Valence	Emotion		Markers	Percentage		
	Type	Subtype		Subtype	Type	Valence
+	Attraction	Like	93	24.73	31.38	32.44
+	Attraction	Love	25	6.65		
+	Satisfaction	Joy	4	1.06	1.06	
–	Dissatisfaction	Anger	26	6.92	28.46	67.56
–	Dissatisfaction	Fear	54	14.36		
–	Dissatisfaction	Sadness	27	7.18		
–	Repulsion	Disgust	147	39.10	39.10	

**Table 4** Distribution of quotations

		Relevant		Total
		Yes	No	
Verbatim	Yes	26	33	59
	No	38	5	43
	Total	64	38	102

Whereas NLP methods excel at systematically detecting and describing significant linguistic markers in text, generative AI can provide plausible interpretations and explanations of the identified features, helping contextualise patterns of meaning.

### Interpreting Emotion-Oriented Lexical Markers

The interpretation of the emotion-oriented lexical markers was based on quotations taken from the corpus by the LLM. The prompt was devised to analyse each marker considering its emotional content and the themes discovered in the Thematic Analysis.

Regarding the quotations, we also evaluated whether or not they appeared in the corpus (i.e. verbatim) and whether or not they included the markers being analysed (i.e. relevant).<sup>7</sup> Table 4 shows the results after examining the five corpus segments.

In the remainder of this section, we examine the interpretation of the quotations extracted by the LLM (i.e. verbatim: yes). First, we focus on the quotations including the target marker (i.e. relevant: yes). We conclude that 61.54% of these interpretations are correct, apart from contributing to critically analysing affective polarisation, as in (1).<sup>8</sup>

(1) I'm genuinely scared of more immigration from Islamic countries. **Fear** in the face of danger, undoubtedly. That's not fear mongering

*The speaker expresses their anxiety about the potential for increased immigration from Islamic countries, indicating that they feel threatened or worried*

It should be noted that, although fear and anxiety are not expected to increase affective polarisation, these emotions are natural consequences of threat perceptions (e.g. immigration from Islamic countries) from the in-group (Renström et al., 2023). Therefore, the comment is related to affective polarisation. However, this quotation also shows one of the shortcomings of this type of interpretation. In particular, the LLM usually provides shallow interpretations of lexical markers, where such interpretations typically take the form of statements with the pattern “the speaker expresses [negative emotion] towards [target entity]”.

In addition, the most concerning finding is that over one third (38.46%) of the quotations automatically extracted and analysed by the LLM are either irrelevant to affective polarisation (i.e. irrelevant stance object, 15.38%) or misinterpreted (i.e. irony and sarcasm, 19.23%; misassigned stance-taker, 3.85%). For example, many lexical markers that represent positive emotions (e.g. *good, love, perfect*, etc.) do not

<sup>7</sup> In this study, “verbatim” is also used to refer to near-verbatim quotations.

<sup>8</sup> The word in bold within the quotation is the emotion-oriented lexical marker analysed by the LLM.

reflect in-group animosity but a positive evaluation of the video, the channel or the interviewer, as in (2).

(2) **Love** your channel from the USA. Awesome work you do

*This user expresses their attraction towards the channel they are addressing, using the marker "love" to convey their positive feelings towards it*

On the other hand, the interpretations of other comments are misleading. For example, in the case of (3), *love* is not used to express a positive emotion. Instead, the marker is used ironically as a rhetorical device intended to criticise the behaviour of pro-Brexit commenters. Therefore, the real interpretation is that the commenter thinks it is ridiculous how Brexit supporters go to a pro-EU video just to attack it with misinformation.

(3) I **love** how all you leave EU idiots come to a video from a left wing channel that talks about staying in the EU just to hate on it and fill it with your made up Sun newspaper statistics all the major politicians economists and anyone with half a brain cell chooses remain and people like UKIP and BNP are the ones who want to leave Europe use your brain people

*This marker is used to express attraction, specifically love, towards the speaker's political stance. The speaker is expressing love towards the idea that the majority of people with half a brain cell are voting to remain in the EU*

Likewise, the explanation in (4) assumes the commenter is expressing hatred, but in fact they express a perception of others' hatred. The commenter does not necessarily say they hate British selfishness themselves; they say they detect that kind of hatred "from" British selfishness. In other words, the commenter perceives a British attitude characterised by hating everything that is not British, possibly out of selfishness.

(4) I can feel a kind of British selfishness **hatred** towards everything which is not British!

*This marker is used to express repulsion and disgust. The author expresses their hatred towards British selfishness and its perceived hatred towards non-British things*

Second, we examine the interpretation of quotations that do not include the target marker (i.e. relevant: no). In most of these cases, although the target marker is not used in the comment, the marker represents an emotion that is actually expressed by the statement. Indeed, the interpretation of these quotations is correct in 81.82% of the cases, so the comments were meaningful from the perspective of affective polarisation, as shown in (5) for *fear*.

(5) I'm really scared. It's like it's finally sunk in

*This marker is used to convey dissatisfaction and fear. The author expresses their anxiety about the potential consequences of Brexit*

Given the results of this task, we implemented various strategies to improve them. For example, we guaranteed the inclusion of authentic quotations by extracting them directly from the corpus using regular expressions based on the morphological information of the emotion-oriented lexical markers. This approach not only facilitates the faithful reproduction of quotations but also enhances the replicability of the experiment, as the exact quotation associated with a given lexical marker would consistently be retrieved if the study were to be repeated.

Moreover, we relied on prompt engineering to enhance the performance of the LLM in interpreting lexical markers. Particularly, apart from considering the types of emotions and the themes of the discourse, the LLM prompt was optimised to focus the analysis on affective polarisation, for which a definition of this phenomenon was provided at the beginning of the prompt:

Suppose that you are an expert in affective polarisation, which studies how, under conditions of groups in conflict, there is positive sentiment for one's own group and negative sentiment towards those identifying with opposing groups. [...]

In this regard, we noted that the performance of the LLM was not improved when all emotion-oriented lexical markers in a given corpus segment were analysed through a single prompt. In contrast, the LLM generated a more significant interpretation when a single prompt was devised for each marker. To illustrate, (6) and (7) present the new interpretations that correspond to the quotations (1) and (3), respectively.

(6) *This marker of emotion, "fear", is used to express dissatisfaction with the potential consequences of Brexit, specifically the possibility of increased immigration from Islamic countries. The speaker is expressing a genuine fear of this scenario, which they perceive as a real and imminent danger. This emotion is likely driven by a combination of factors, including concerns about cultural identity, security, and the potential strain on resources and services. The use of this marker contributes to the theme of "Immigration and Border Control" in the discourse, as it highlights the emotional stakes involved in this issue. Additionally, it may also be a factor in the theme of "Potential Political Instability", as fear and anxiety can fuel extremism and division. The speaker's use of the phrase "fear in the face of danger" suggests that they believe their fear is a rational response to a real threat, rather than an irrational or baseless emotion.*

(7) *The use of the word "love" in this context is ironic and sarcastic, indicating a strong emotional reaction against those who support leaving the EU. The speaker expresses a sense of affection towards the idea of remain supporters using their intelligence and rationality, contrasting it with the perceived lack of these qualities in leave supporters. This emotional expression contributes to affective polarisation by reinforcing in-group (remain supporters) and out-group (leave supporters) dynamics, with the speaker expressing love towards their in-group and derision towards the out-group. However, in the case of hatred in (4), this strategy did not lead to a better interpretation. In contrast, the result was improved when we provided the LLM with the entire comment, as in (8).*

(8) *The whole brexit debate is only about immigration!!! Do you like foreigners or no !! !It is ridiculous and shameful !! ! I can feel a kind of British selfishness **hatred** towards everything which is not British !! I can't understand why the British want to isolate themselves when the world is becoming closer !! The European union is a great organisation with its defects of course and being in it is like being in a race!!! Why can't Britain compete in this race and take the lead ??*

*This marker of emotion, hatred, is used to express strong repulsion and disgust towards the Brexit debate and the British people who support it. The speaker feels*

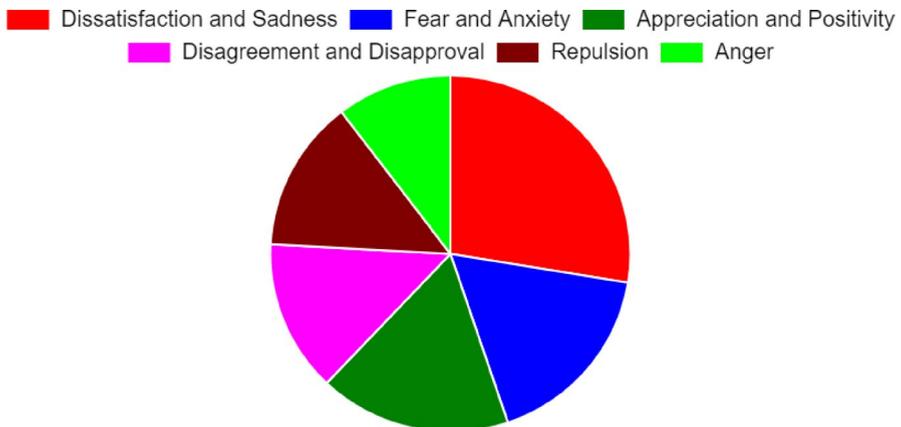
that the debate is solely focused on immigration and displays a hatred towards foreigners and anything that is not British. This emotional response can be seen as a form of affective polarisation, where the speaker has a strong negative sentiment towards the opposing group (those who support Brexit). The use of this marker can be linked to the arguments for Leave, as the speaker perceives the debate as being driven by anti-immigrant sentiment. Additionally, it can be connected to the potential consequences of Brexit, as the speaker fears that a Brexit vote would lead to increased isolation and hatred towards foreigners

Therefore, we conclude that a prompt that includes an extended co-text for the target lexical marker can facilitate the LLM's understanding of the marker, so that the LLM can infer usage patterns based on probabilistic regularities learned from contextual cues. The experiment's findings revealed that it is sufficient with a co-text of 150 characters to the left and right of the target word.

### Identifying and Explaining Patterns of Meaning

The LLM identified six primary patterns of meaning out of the interpretations of the lexical markers in the corpus. The LLM summarised the patterns around a few words that reflect the affective states of commentators: (a) Dissatisfaction and Sadness, (b) Fear and Anxiety, (c) Appreciation and Positivity, (d) Disagreement and Disapproval, (e) Repulsion, and (f) Anger.<sup>9</sup> Fig. 4 shows the relevance of these patterns, which was assessed based on their presence in the five corpus segments.

The LLM also explained these patterns in detail, where most explanations focused on emotional triggers, i.e. entities, events or situations that elicit emotional responses in the individual. As shown in Table 5, we discovered that some of these emotional



**Fig. 4** Distribution of the patterns of meaning

<sup>9</sup> Strictly speaking, disagreement and disapproval are not classified as affective states but as cognitive processes that can evoke emotions. Both of them involve evaluative judgments of another's stance or behaviour so the evaluation is socially situated. Therefore, these cognitive processes are part of the social-appraisal process, as affective states can be shaped by how others think, feel, and act (Manstead & Fischer, 2001).

**Table 5** Emotional triggers in the patterns of meaning

Emotional trigger	Pattern of meaning					
	a	b	c	d	e	f
Potential consequences of Brexit (e.g. economic impact and isolation)	x	x			x	x
Brexit campaign (e.g. behaviour of others)	x			x	x	
Brexit discourse (e.g. lack of constructive dialogue and problem-solving during the referendum debate, particularly in issues such as immigration)	x					x
Political figures (e.g. behaviour and arguments)				x	x	
Perceived threats, which can be external (e.g. the EU's perceived corruption or exploitation) or internal (e.g. the left's treatment of the working class or the spread of false information)					x	x
Brexit process	x					
EU and immigration		x				
EU and work			x			
YouTube commentators			x			
Idea of Brexit				x		

triggers are found in various patterns, with the potential consequences of Brexit as the most recurrent.

To examine whether there was room for improvement through prompt engineering, we instructed the LLM to explain the patterns of meaning from the perspective of affective polarisation. However, no significant improvement was observed in the results.

### Contextualising Patterns of Meaning

In this section, we underscore those aspects of the historical, political, economic and social contexts in the Background Analysis that the LLM relied on to explain the patterns of meaning discovered in the previous task. First, Table 6 shows the historical factors that significantly explain some patterns of meaning.

Although fear and repulsion slightly stand out as emotions associated with the historical context, many explanations do not rely on specific historical events but simply describe generic aspects related to the long, contentious relationship between the UK and the EU.

Second, most of the LLM-generated explanations based on the political context were presented according to the following template:

The political context, with/marked by/characterised by/... [reason], has led to/fueled/... [emotion].

For example:

(9) *The political context, marked by a lack of consensus and acrimonious debate, has fueled disapproval and dislike towards the Brexit campaign and its tactics. Therefore, as illustrated in (9), the political factors are presented as generalisations with no reference to specific events or situations from the Background Analysis. We found only one explanation that refers to a particular political event, but it describes a post-referendum incident:*

**Table 6** Historical factors and the patterns of meaning

Historical factor	Pattern of meaning					
	a	b	c	d	e	f
UK's opt-outs from the Euro, Eurozone and Schengen Area		x				x
Euroscepticism in the UK, particularly in the 1980s				x	x	
Veto of the first application to join the EU and struggle to finally become a member		x				
Strong support for EU membership in the 1975 referendum			x			
Rise of UKIP					x	

**Table 7** Economic factors and the patterns of meaning

Economic factor	Pattern of meaning					
	a	b	c	d	e	f
Immigration on the UK's labour market and public services	x	x			x	x
Wage stagnation or decline	x	x				
Increasing economic inequality	x	x				
Economic uncertainty with Brexit	x	x				
Uneven distribution of the economic benefits of EU membership: sense of economic injustice	x					
Perception that the EU was hindering the UK's economic growth		x				
Economic benefits of EU membership: increased trade, investment and job creation			x			

(10) *The political uncertainty following the referendum, including the resignation of David Cameron and the appointment of Theresa May, fueled fear and anxiety about the political future of the UK*

Third, Table 7 shows the economic factors described in the Background Analysis that can explain some patterns of meaning.

We conclude that dissatisfaction and fear are behind various economic aspects explicitly referenced in the discourse, whereas the perceived impact of immigration on the labour market and public services is the concern that triggers most of the emotions.

Fourth, Table 8 shows the social factors that help explain the patterns of meaning found in the discourse, with dissatisfaction and fear being the most significant again.

To sum up, the last task in our methodological framework consists in instructing the LLM to examine the patterns of meaning through a historical, political, economic

**Table 8** Social factors and the patterns of meaning

Social factor	Pattern of meaning					
	a	b	c	d	e	f
Perception of excessive immigration	x	x		x	x	x
Loss of national sovereignty	x	x		x	x	x
Sense of loss of national identity	x	x		x	x	x
Disenfranchisement from the political establishment	x	x				
Concern about social cohesion and deepening of social divisions	x	x				
Rise in hate crimes	x					
Concern about job security		x				
Sense of belonging to a progressive, multicultural Europe			x			

and social lens. The results in Tables 6, 7 and 8 shed light on the key driving factors that shaped the digital discourse on Brexit. In this regard, economic and social factors were dominant, whereas historical and political dimensions received comparatively limited attention. On the one hand, commentators were emotionally engaged with issues that could immediately impact their daily lives. In the case of Brexit, economic challenges (e.g. wage stagnation and economic inequality) and social conflicts (e.g. immigration and national identity) that were latent before the referendum provided fertile ground for public comments. For example, the first three social factors in Table 8, which triggered most of the affective states of commentators, align with Carter's (2024) study, which concluded that Brexit identity was tied to cultural issues, where immigration and national identity were the primary driving forces behind Leavers. Moreover, the LLM underscored that the economy also played an important role in affective polarisation within Brexit discussions. Similarly, Curtice's (2016) survey on British attitudes towards Europe before the referendum concluded that only cultural considerations were insufficient to persuade people to vote to leave the EU; they also needed to be convinced of the economic disadvantages of EU membership.

On the other hand, discussions around the historical trajectory of the UK's relationship with the EU or the political fragmentation surrounding Brexit are typically framed in more formal registers (e.g. parliamentary speeches) and less overtly referenced in social-media discourse. For example, Brusenbauch Meislová (2021) identified three discursive strategies in political speeches that are in line with the historical and political contexts: the narrative of (a) political failure, where partisan opponents are portrayed as a group that has failed, (b) incompetence, denoting lack of professionalism and limited capacity to govern, and (c) betrayal, resulting in topics such as treachery, sabotage and loyalty. Therefore, these discursive strategies are more closely related to political polarisation than to affective polarisation. Hobolt et al. (2021) explained that affective polarisation in Brexit is grounded on opinion-based groups, where people identify with others based on a shared opinion about a given event, rather than on partisan divisions. Traditionally, partisanship has been a central concept in affective polarisation. However, positive or negative perceptions about Leavers and Remainers were not driven by party identity but by Brexit identity. As noted by Carter (2024: 52), partisan identity was not the cause of affective polarisa-

tion during Brexit, as “internal divisions within the parties show that party loyalty could not have substantially impacted the Brexit debate”. It should be recalled that, except for UKIP, none of the major political parties declared an internal consensus on their stance on this issue in 2016.

Finally, we also tried to refine the LLM prompt so that it could contribute to enhancing the contextualisation of the patterns of meaning. To this end, we instructed the LLM to explain the patterns from the perspective of affective polarisation, together with the contexts in the Background Analysis. However, as in the previous task, the findings did not reveal any substantial improvement.

## Conclusions

Although LLMs have made significant advancements, generative AI for qualitative analysis is still an emerging research area. To improve the performance of a general-purpose LLM in qualitative analysis, we examined affective polarisation in the Brexit debate through the lens of CDA, specifically focusing on emotion-oriented lexical markers. For this purpose, we employed a small corpus of YouTube comments in our experiment, prioritising in-depth analysis over statistical generalisation. Our methodological framework, comprising Background, Thematic, and Linguistic analyses, boosted the interpretive capabilities of the LLM and reduced hallucinations. In particular, the Background and Thematic Analyses provided information about the research topic and the discourse, which guided the AI system towards deeper qualitative analyses in the interpretation and explanation steps. Furthermore, we highlighted the advantages of prompt engineering, with tailored prompts facilitating a more nuanced analysis of affective polarisation. For example, prompt variations enhanced the interpretation of linguistic markers when (a) the prompt included a theoretical definition of affective polarisation, (b) the prompt was aimed at interpreting one particular lexical marker, rather than analysing a set of markers, and (c) the prompt provided larger co-text for the given lexical marker. Nevertheless, LLM-generated claims should have been supported by quotations automatically taken from sources such as scholarly publications or mass media outlets to enhance the trustworthiness of the AI system. Instead, researchers cross-checked the quality of LLM-produced information by comparing data from various sources to ensure the accuracy and reliability of the analysis results. To facilitate this task, the system provided numerous bibliographical citations and news items about the topic of the discourse in the Background Analysis.

In future research, we intend to employ in-context learning to improve the performance of LLM-driven CDA, where the model can learn a novel task through examples and demonstrations included in the prompt itself. In this respect, Retrieval-Augmented Generation can incorporate knowledge into the prompt, which would require a corpus of YouTube comments annotated with interpretations and explanations based on affective polarisation and within the CDA framework.

## Appendix 1: Description of the Resources Used in the Experiment

**Mistral Nemo.**<sup>10</sup> It is a state-of-the-art LLM that performs better than other open-source pre-trained models (e.g. Gemma 2 9B and Llama 3 8B). Mistral NeMo, trained with 12 billion parameters, supports multiple languages, including English and Spanish, making it suitable for multilingual applications. Released as a free model for research in 2024, Mistral NeMo can process up to 500,000 tokens per minute and one billion tokens per month, which is a fairly reasonable amount for research purposes. It is well-known that LLMs select the most suitable tokens based on specific decoding parameters, based on stochastic methods that govern the randomness of the generated text. In our system, after evaluating different combinations of parameter settings, we finally set the temperature to 0.8, top-K to 40 and top-P to 0.8. In this way, we increased response variability and thus reduced repetitive outputs, finding a balance between text diversity and coherence. Moreover, the maximum number of tokens to be processed in a prompt was 4,000, as we aimed to generate long, detailed responses.

**OpenAlex.**<sup>11</sup> It is a catalogue of over 250 million scholarly works (e.g. research papers, books, and conference proceedings) from 250,000 sources. The database can be accessed through a free API service. For each query automatically generated in Task 2 of our methodological framework, the API retrieves a maximum of 50 bibliographical citations within the last five years. For each citation, our system gets the OpenAlex ID, DOI, title, year of publication, type of publication, source of publication, relevance score, and abstract.

**GDELT.**<sup>12</sup> It is an open platform that monitors news media (i.e. broadcast, print, and web news) from around the world in over 100 languages in real time. It extracts information about events, people, organisations, locations, themes, and even the emotional tone of coverage. The database can be accessed through a free API service. For each query automatically generated in Task 2 of our methodological framework, the API retrieves a maximum of 250 news items published within the last three months. For each news item, our system gets the headline, publication date and URL of the full article.

**WordNet.**<sup>13</sup> It is an extensive lexical database of English, where nouns, verbs, adjectives and adverbs are grouped into sets of synonyms (i.e. synsets), each expressing a distinct concept. These synsets are linked through semantic relations, such as antonymy, hyperonymy and meronymy, among others. We downloaded the English WordNet 3.0 and the Spanish WordNet from the Multilingual Central Repository.<sup>14</sup>

**SentiWordNet.**<sup>15</sup> It is a polarity lexicon that resulted from automatically annotating all the English WordNet 3.0 synsets according to their degrees of positivity,

<sup>10</sup> <https://mistral.ai/news/mistral-nemo>.

<sup>11</sup> <https://openalex.org>.

<sup>12</sup> <https://www.gdeltproject.org>.

<sup>13</sup> <https://wordnet.princeton.edu>.

<sup>14</sup> <https://web.archive.org/web/20200314100039/https://adimen.si.edu.es/web/MCR/>.

<sup>15</sup> <https://github.com/aesuli/SentiWordNet>.

negativity, and objectivity, where each dimension was assigned a value from 0 to 1, with 1 being the sum of the three scores for each synset.

**SentiSense.**<sup>16</sup> It is an affective lexicon where more than 2,000 synsets from the English WordNet 3.0 were semi-automatically labelled with an emotion from a set of 14 emotion categories (i.e. ambiguous, anger, calmness, despair, disgust, anticipation, fear, hate, hope, joy, like, love, sadness, and surprise).

**Author Contributions** The author performed all work related to the study.

**Funding** This publication is part of the R&D&I project PID2023-147137NB-I00, funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU, and partially supported by the research project CIP-ROM/2023/29, funded by "Direcció General de Ciència i Investigació" Generalitat Valenciana (Spain)

**Data Availability** No datasets were generated or analysed during the current study.

## Declarations

**Competing Interests** The authors declare no competing interests.

**Ethics Approval** No ethical approval was required for this study.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

## References

- Benitez-Castro, M. A., & Hidalgo-Tenorio, E. (2019). Rethinking Martin & White's AFFECT taxonomy: A psychologically-inspired approach to the linguistic expression of emotion. In Mackenzie, J. L., & L. Alba-Juez (Eds.), *Emotion in discourse* (pp. 301–331). John Benjamins. <https://doi.org/10.1075/pbns.302.12ben>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- Breazu, P., Schirmer, M., Hu, S., & Katsos, N. (2025). Large Language models and the challenge of analyzing discriminatory discourse: Human-AI synergy in researching hate speech on social media. *Journal of Multicultural Discourses*, 1–19. <https://doi.org/10.1080/17447143.2025.2476967>
- Brusenbauch Meislová, M. (2021). Discursive construction of affective polarization in Brexit Britain: Opinion-based identities and out-group differentiation. In Pérez-Escobar, M., & J. M. Noguera-Vivo (Eds.), *Hate speech and polarization in participatory society* (pp. 98–112). Routledge. <https://doi.org/10.4324/9781003109891-9>
- Cambria, E. (2016). Affective computing and sentiment analysis. *IEEE Intelligent Systems*, 31(2), 102–107. <https://doi.org/10.1109/MIS.2016.31>

<sup>16</sup> <https://nlp.uned.es/~jcalbornoz/resources/sentisense>.

- Carter, T. (2024). The united Kingdom and brexit: A case study in affective polarization. *Sigma: Journal of Political and International Studies*, 41(5), 49–59.
- Christou, P. A. (2023). How to use artificial intelligence (AI) as a resource, methodological and analysis tool in qualitative research? *The Qualitative Report*, 28(7), 1968–1980. <https://doi.org/10.46743/2160-3715/2023.6406>
- Curtice, J., & The National Centre for Social Research. (2016). *How deeply does Britain's Euroscepticism run?*. Retrieved June 4, 2020, from <https://whatukthinks.org/eu/wp-content/uploads/2016/02/Analysis-paper-5-How-deeply-does-Britains-Euroscepticism-run.pdf>
- de Carrillo, J., Plaza, L., & Gervás, P. (2012). SentiSense: An easily scalable concept-based affective lexicon for sentiment analysis. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation* (pp. 3562–3567). European Language Resources Association.
- De Paoli, S. (2024). Performing an inductive thematic analysis of semi-structured interviews with a large Language model: An exploration and provocation on the limits of the approach. *Social Science Computer Review*, 42(4), 997–1019. <https://doi.org/10.1177/08944393231220483>
- Du Bois, J. W. (2007). The stance triangle. In R. Englebretson (Ed.), *Stancetaking in discourse: Subjectivity, evaluation, interaction* (pp. 139–182). John Benjamins. <https://doi.org/10.1075/pbns.164.07du>
- Esuli, A., & Sebastiani, F. (2006). SentiWordNet: A publicly available lexical resource for opinion mining. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation* (pp. 417–422). European Language Resources Association.
- Fairclough, N. (1989). *Language and power*. Longman.
- Fairclough, N. (1992). *Discourse and social change*. Polity.
- Fairclough, N., & Wodak, R. (1997). Critical discourse analysis. In Van T. A. Dijk (Ed.), *Discourse studies. A multidisciplinary introduction. Discourse as social interaction* (pp. 258–284). Sage.
- Fellbaum, C. (Ed.). (1998). *WordNet: An electronic lexical database*. The MIT Press. <https://doi.org/10.7551/mitpress/7287.001.0001>
- Hitch, D. (2024). Artificial intelligence augmented qualitative analysis: The way of the future? *Qualitative Health Research*, 34(7), 595–606. <https://doi.org/10.1177/10497323231217392>
- Hobolt, S. B., Leeper, T. J., & Tilley, J. (2021). Divided by the vote: Affective polarization in the wake of the brexit referendum. *British Journal of Political Science*, 51(4), 1476–1493. <https://doi.org/10.1017/S0007123420000125>
- Johnston, C. (2014). *Tory MP Mark Reckless defects to Ukip*. The Guardian. Retrieved May 14, 2025, from <https://www.theguardian.com/politics/2014/sep/27/tory-mp-mark-reckless-defects-ukip>
- Khan, A. H., Kegalle, H., D'Silva, R., Watt, N., Whelan-Shamy, D., Ghahremanlou, L., & Magee, L. (2024). Automating thematic analysis: How LLMs analyse controversial topics. <https://doi.org/10.48550/arXiv.2405.06919>
- Kogen, L. (2024). Qualitative thematic analysis of transcripts in social change research: Reflections on common misconceptions and recommendations for reporting results. *International Journal of Qualitative Methods*, 23, 1–11. <https://doi.org/10.1177/16094069231225919>
- Manstead, A. S. R., & Fischer, A. H. (2001). Social appraisal: The social world as object of and influence on appraisal processes. In Scherer, K. R., Schorr, A., & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, method, research* (pp. 221–232). Oxford University Press.
- Martin, J. R., & White, P. R. R. (2005). *The Language of evaluation: Appraisal in english*. Palgrave Macmillan.
- Maru, M., Scozzafava, F., Martelli, F., & Navigli, R. (2019). SyntagNet: Challenging supervised word sense disambiguation with lexical-semantic combinations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (pp. 3532–38). Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1359>
- Mullet, D. R. (2018). A general critical discourse analysis framework for educational research. *Journal of Advanced Academics*, 29(2), 116–142. <https://doi.org/10.1177/1932202X18758260>
- Priem, J., Piwowar, H., & Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. *ArXiv Preprint*. <https://doi.org/10.48550/arXiv.2205.01833>
- Renström, E. A., Bäck, H., & Carroll, R. (2023). Threats, emotions, and affective polarization. *Political Psychology*, 44(6), 1337–1366. <https://doi.org/10.1111/pops.12899>
- Sheikh, H., Prins, C., & Schrijvers, E. (2023). *Mission AI: The new system technology*. Springer. <https://doi.org/10.1007/978-3-031-21448-6>
- Stephan, W. G., Ybarra, O., & Rios, K. (2015). Intergroup threat theory. In Nelson, T. D. (Ed.), *Handbook of prejudice, stereotyping and discrimination* (pp. 255–278). Psychology.

- Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, 223(5), 96–102. <https://doi.org/10.1038/scientificamerican1170-96>
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In Austin, W. G., & Worchel, S. (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Brooks-Cole.
- Tam, Z. R., Wu, C. K., Tsai, Y. L., Lin, C. Y., Lee, H., & Chen, Y. N. (2024). Let me speak freely? A study on the impact of format restrictions on large language model performance. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track* (pp. 1218–1236). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-industry.91>
- Todolí Cervera, J., Labarta Postigo, M., & Dolón Herrero, R. (2006). What is critical discourse analysis? *Quaderns De Filologia Estudis Lingüistics*, 11, 9–34.
- Turobov, A., Coyle, D., & Harding, V. (2024). Using ChatGPT for RTA. <https://doi.org/10.48550/arXiv.2405.08828>
- Van Dijk, T. A. (1995). Discourse analysis as ideology analysis. In Schäffner, K., & Wenden, A. L. (Eds.), *Discourse and ideologies* (pp. 7–37). Multilingual Matters.
- Van Dijk, T. A. (2003). Critical discourse analysis. In D. Schiffrin, D. Tannen, & H. E. Hamilton (Eds.), *The handbook of discourse analysis* (pp. 352–371). Blackwell. <https://doi.org/10.1111/b.9780631205968.2003.00019.x>
- Wodak, R. (2001). What critical discourse analysis is about: A summary of its history, important concepts and its developments. In Wodak, R., & Meyer, M. (Eds.), *Methods of critical discourse analysis* (pp. 1–13). Sage. <https://doi.org/10.4135/9780857028020.n1>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.