

FORSEE

| Forging Successful AI Applications
| for European Economy and Society

Project Information

Project Number: 101177579

Project Title: Forging Successful AI Applications for European Economy and Society - FORSEE

Funding Scheme: HORIZON-CL2-2024-TRANSFORMATIONS-01-06

Project Start Date: February 1st, 2025

Deliverable Information

Title: "D4.1 Social Media Analysis of AI applications"

Work Package: "WP 4 - Researching societal acceptance and criteria of success - Identifying controversies and convergence"

Lead Beneficiary: University of Toulouse

Due Date: 31/01/2026

Revision Number: Fx.1

Authors: Charis Papaevangelou, Lucie Loubère, Nikos Smyrnaiois, Pierre Ratinaud

Dissemination Level: Public

Deliverable Type: Report

Overview

This task examines how public discourse on social media shapes, negotiates, and contests notions of success and failure surrounding Artificial Intelligence, with particular attention to generative AI. Recognising social media platforms as central arenas for public sense-making beyond formal institutional and journalistic settings, the analysis focuses on how diverse users express their expectations and anxieties of, as well as aspirations for AI



Funded by
the European Union

FORSEE has received funding from the EU's Horizon Europe research and innovation programme under Grant Agreement 101177579

in a variety of contexts. The research draws on a comparative analysis of Facebook, TikTok, and YouTube across four national contexts (France, Spain, Germany, and Ireland), combining lexicometric analysis with discourse and, where relevant, network analysis. Across platforms and countries, the analysis reveals persistent discursive tensions around AI, including the swaying between promises of efficiency and experiences of failure, economic prospects and labour pressures, and all within the context of broader existential and (geo)political concerns related to AI.

Revision History				
Version	Changes implemented by	Revision Date	Reviewed by:	Brief description of changes
<Fx.1>	Charis Papaevangelou, Lucie Loubère, Nikos Smyrnaio, Pierre Ratinaud	<19/01/2026>	Bernhard Rieder, Eugenia Siapera	Revised according to internal peer-review

Approval Procedure			
Version	Deliverable Name	Approved by	Approval Date
<Fx.x>	<.....>	<Name>	<dd/mm/yyyy>

Table of Acronyms	
Acronym	Definition
AI	Artificial Intelligence
API	Application Programming Interface
CDA	Critical Discourse Analysis
DHC	Descending Hierarchical Classification
CTS	Characteristic Text Segment
DHC	Descending Hierarchical Classification
DSA	Digital Services Act
EU	European Union
FORSEE	Forging Successful AI Applications for European Economy and Society
GenAI	Generative Artificial Intelligence

IRaMuTeQ	Interface de R pour les Analyses Multidimensionnelles de Textes et de Questionnaires
LLM	Large Language Model
SoE	Sociology of Expectations
TS	Text Segment
WP	Work Package



Table of Contents

Table of Contents	4
List of Figures	5
List of Tables	6
Executive Summary	7
1. Introduction	8
2. Main Report Content	8
2.1. Literature Review	8
2.1.1. The societal shaping of Artificial Intelligence.....	8
2.1.2. Social media as an arena of contestation.....	9
2.1.3. Three recurring tensions in public imaginaries.....	9
2.1.4. Methodological implications.....	10
3. Methodology	11
3.1. Platform selection strategy.....	11
3.2. Lexicometric and Discourse Analysis.....	12
4. Analysis of social media platforms	14
4.1. Facebook.....	14
4.2. TikTok.....	21
4.3. YouTube.....	27
5. Conclusions	36
6. References	38
7. Appendix	43



List of Figures

Figure 1 - Histogram of common themes in market-oriented and application-based discourse for the Facebook corpus (results as a percentage of each corpus)	14
Figure 2 - Histogram of common themes in societal and human-centred discourse for the Facebook corpus (results as a percentage of each corpus)	17
Figure 3 - Histogram of common themes in discourse on GenAI as a content-creation force for the TikTok corpus (results as a percentage of each corpus)	21
Figure 4 - YouTube comments network graph of the corpus of Germany	30
Figure 5 - YouTube comments network graph of the corpus of Spain	30
Figure 6 - YouTube comments network graph of the corpus of France	31
Figure 7 - YouTube comments network graph of the corpus of Ireland	31
Figure 8 – Facebook: Dendrogram of the clustering on the corpus of France	40
Figure 9 – Facebook: Dendrogram of the clustering on the corpus of Spain	40
Figure 10 – Facebook: Dendrogram of the clustering on the corpus of Ireland	40
Figure 11 – Facebook: Dendrogram of the clustering on the corpus of Germany	41
Figure 12 - TikTok: Dendrogram of the clustering on the corpus of France	42
Figure 13 – TikTok: Dendrogram of the clustering on the corpus of Spain	42
Figure 14 – TikTok: Dendrogram of the clustering on the corpus of Ireland	42
Figure 15 – TikTok: Dendrogram of the clustering on the corpus of Germany	43
Figure 16 - YouTube: Dendrogram of the clustering on the videos from the corpus of Germany	43
Figure 17 - YouTube: Dendrogram of the clustering on the videos from the corpus of Spain	43
Figure 18 - YouTube: Dendrogram of the clustering on videos from the corpus of France	44
Figure 19 - YouTube: Dendrogram of the clustering on the videos from the corpus of Ireland	44
Figure 20 - YouTube: Dendrogram of the clustering on the comments from the corpus of Germany	44
Figure 21 - YouTube: Dendrogram of the clustering on the comments from the corpus of Spain	45
Figure 22 - YouTube: Dendrogram of the clustering on the comments from the corpus of Ireland	45
Figure 23 - YouTube: Dendrogram of the clustering on the comments from the corpus of France	45



List of Tables

Table 1 - Facebook corpus statistics by country	14
Table 2 - Statistics on the number of comments of TikTok videos by country	19
Table 3 - Statistics on the number of views of TikTok videos by country	20
Table 4 - Statistics on the number of comments of YouTube videos by country	25
Table 5 - Statistics on the number of views of YouTube videos by country	25



Executive Summary

This report analyses how social media users in France, Spain, Germany, and Ireland frame the “success” and “failure” of Artificial Intelligence. Drawing on user-generated content from Facebook, TikTok, and YouTube, it identifies core societal tensions, controversies, and patterns of conditional acceptance of AI, primarily Generative AI (GenAI). The report is part of FORSEE Work Package 4, which examines societal acceptance of AI applications and research on controversies and points of tension.

The report uses a mixed-methods approach combining lexicometric discourse analysis (Descending Hierarchical Classification using IRaMuTeQ/Reinert method) and network analysis (Gephi). Across platforms, users frame “success” through an optimistic lens that relate to economic growth, including productivity gains, content creation, employability, and competitiveness. “Failure”, on the other hand, is framed as a form of social disruption and loss of control, including job displacement, deepfakes and manipulation, unreliable outputs, accountability gaps, and AI-fuelled warfare and surveillance.

We also find that platform affordances and demographics shape these framing. Specifically, our report shows how certain Facebook users, primarily with institutional and semi-professional standing, dominate the discussion with content ranging from news to promotion of professional events around a wide gamut of topics, including financial markets, geopolitics, and creativity; TikTok users in our dataset treat GenAI as a tool to be harnessed to increase visibility and content monetisation, or to find “shortcuts” and “life hacks” for work and school assignments. Last, our YouTube data demonstrate how the platform functions as a key arena for better elaborated debates, wherein user communities coalesce around issues spanning political economy and labour implications and renewed quasi-existential questions about humanity in the age of GenAI.

While cross-country themes are evident, local nuances persist. Our data from France and Spain feature more politically charged, contextually grounded accounts of AI, particularly regarding armed conflict and artistic legitimacy. Discourse from Ireland seem to focus more on labour market pressures and the professionalisation of content creation, while German users display a more sceptical tone regarding GenAI’s capacities, emphasising on the qualitative distinctions between humans and machines.

By and large, the findings show that societal acceptance of AI is conditional and instrumental, that is, users seem to embrace GenAI’s promises when they can leverage it in a way that is beneficial to them with little concern or thought about the broader implications. In that sense, we see that many users adopt GenAI to “keep up” with platform economies or workplace expectations rather than because they endorse its broader trajectory. The most salient points of controversy concern distribution (who benefits/who loses), accountability (who is responsible for harms), and power (who deploys AI in policing, warfare, and information control).



1. Introduction

Work package 4 of FORSEE-project focuses on societal acceptance of AI applications and research on controversies and points of tension. As AI technologies continue to permeate various aspects of society, understanding how they are received and perceived by different stakeholders becomes crucial. Firstly, research concentrates on the public sphere; news organisations and social media play a pivotal role in shaping societal perceptions of AI by disseminating information, influencing public discourse, and framing narratives around AI applications. Then, points of tension are further explored through litigation analysis; court cases involving AI applications serve as a critical barometer, revealing societal tensions and concerns, providing a legal lens through which ethical, accountability and transparency issues are scrutinized and addressed.

This report (D4.1 Social Media Analysis of AI applications, Task 4.1) specifically aims to determine how social media users across four countries, France, Spain, Germany and Ireland, frame the “success” and “failure” of AI, by identifying patterns and themes, contradictions or tensions that emerge from different forms of discourse (posts, videos and comments). By examining narratives from user-generated content, D4.1 performs a foundational analysis of public discourse on AI (complementary to the one of D4.2, which focuses on mainstream press).

2. Main Report Content

2.1. Literature Review

2.1.1. The societal shaping of Artificial Intelligence

Technological innovation is not driven by technical feasibility alone but is shaped by socioeconomic interests and political contestation, which determine which futures gain momentum and which lose support. Drawing on the Sociology of Expectations (Bakker et al., 2011; Brown & Michael, 2003; Pinch & Bijker, 1984), this report frames AI not as an inevitable trajectory but as a socially positioned debate. In that sense, anticipatory or future-driven narratives are strategic: they attract investment, legitimise policy, and normalise specific uses. As Van Lente (2012) argues, these expectations function as organising devices, coordinating actors’ behaviour before a technology is fully realised.

A central political challenge follows: AI adoption requires “social legitimacy”, that is, a shared consensus that the technology is desirable, appropriate, and properly governed (Deng & Ahmed, 2025; Suchman, 1995). Public discourse is the arena where this legitimacy is tested. It is where citizens negotiate the definitions of “success” and “failure” and assign responsibility for harms (Coeckelbergh, 2022). These contestations provide the analytical



lens for Task 4.1, which maps how social media users in France, Spain, Germany, and Ireland perceive and construct the value and risks, as well as potential successes and failures, of AI.

2.1.2. Social media as an arena of contestation

Social media platforms offer a vantage point distinct from the mainstream press coverage analysed in Deliverable 4.2. Driven by algorithmic recommendations and engagement incentives, these platforms accelerate the circulation of content that aligns with the business models of these companies (Poell et al., 2021; Siapera, 2018). At the same time, they allow users to bypass traditional gatekeepers, amplifying controversies or challenging industry narratives (Carlson, 2018; Poell et al., 2023; Smyrnaio & Rebillard, 2019). This environment is well-suited for Critical Discourse Analysis (CDA) (Mogashoa, 2014), as it reveals the raw negotiation of power and meaning outside institutional filters (Bouvier & Machin, 2020). As a result, social media discourse often reveals tensions and resistance that remain muted or stabilised in institutional settings.

Existing research on AI discourse has understandably relied to a large extent on text-based platforms, such as Twitter/X and Reddit (e.g., Heaton et al., 2025; Xu et al., 2024), where large-scale data access was relatively easier. This focus has generated valuable insights into how AI is discussed in textual form, particularly within technologically engaged communities such as those on Reddit. At the same time, recent restrictions on platform APIs and data access (Hendrix, 2023) have further shaped what kinds of empirical research is possible, reinforcing a reliance on platforms and formats where text remains more readily available. Parallel to this, public debate around AI increasingly unfolds across multiple environments, where audiovisual content plays an increasingly important role in shaping perceptions (Brewer et al., 2025). Rather than positioning one type of platform or format against another, this report brings these strands together by examining discourse across Facebook, TikTok, and YouTube. It allows us to analyse how AI-related narratives circulate across video-centred online communities with strong textual communicative elements, to situate platform-specific dynamics within a broader, cross-platform understanding of public discourse around AI (Qi et al., 2024; Tsimpoukis, 2025).

2.1.3. Three recurring tensions in public imaginaries

Across the literature on AI expectations, public risk perception, and social media discourse, a set of recurring and interconnected tensions emerges that structure how AI is framed as a success or a failure. Drawing in particular on studies of sociotechnical imaginaries, platform discourse, and early public reactions to generative AI, we identify three such tensions that consistently organise public debate on social media, which we briefly lay out next. Moreover, framing our findings through the lens of the Sociology of Expectations reveals that public discourse negotiates the “images of the future” through which “actions, reactions and decisions” are framed, determining how anticipatory narratives meet social reality (van Lente, 2012, p. 772).

The first tension we find pertains to AI narratives that oscillate between utopian promises and visible failures. Works on early public discourse, especially with regard to the release of ChatGPT, often present generative AI (GenAI) as a form of “panacea”, with success defined as amplifying productivity in various professional activities like coding and digital marketing (Abdullah et al., 2022; Ng & Chow, 2024). However, this sentiment is particularly volatile, often shifting rapidly from excitement to anxiety, especially when it relates to professionals like journalists who feel both threatened by and pressured to adopt GenAI services (Lewis et al., 2025). Moreover, as scholars studying expectations around emerging technologies have noted, optimistic framings around an emerging ‘hyped’ technology may lead to sharp public disappointment when limitations emerge (Kerr et al., 2020; Van Lente, 2012). For instance, GenAI users may prematurely celebrate efficiency but quickly pivot to scepticism when systems produce “hallucinations” or unreliable information, which may have significant negative consequences in crucial sectors like (public) health (De Angelis et al., 2023). Thus, “success” is not a fixed technical metric or dimension, but a framing that is contested with every new controversy and among actors with different political-economic interests.

The second tension is economic. While managerial and industry stakeholders emphasize efficiency, business optimization, and sustained growth resulting from GenAI adoption (Grover et al., 2022), public concern consistently focuses on the displacement of human labour (Qi et al., 2024). Indeed, as seen in the aforementioned example of journalists, the fear of job displacement is more prominently voiced by those who are more likely to be affected by the widespread adoption of GenAI (Deng & Ahmed, 2025). Moreover, research indicates that similar fears of job obsolescence are particularly acute outside specialist technical communities, including in non-Western countries such as China (Mao & Shi-Kupfer, 2023; Qi et al., 2024). In this context, what industry actors frame as a “success” (automation) is frequently interpreted by the public as a “failure” of social protection (Brantner & Saurwein, 2021). This persistent tension between the promise of technological efficiency and the reality of economic disruption for specific user groups highlights unresolved questions of distributive justice: who gets to benefit from the mass deployment of GenAI in industrial settings?

The third tension concerns governance. Public discourse frequently focuses on opacity, bias, and the credibility of AI outputs (Sarisakaloğlu, 2021). These concerns are complicated by ambiguity regarding AI’s agency. Users often anthropomorphise systems, describing them as actors that “know” or “create” (Heaton et al., 2025). However, when harms occur, linguistic patterns often shift blame onto “the system”, obscuring the human developers and institutions behind it (Saurwein et al., 2025). This creates a perceived responsibility gap that undermines trust and fuels demands for regulation (Brewer et al., 2022), especially regarding who must be held accountable when things go wrong mainly, at the consumer-facing level of the GenAI value chain.



2.1.4. Methodological implications

To capture the tensions identified in the literature, this report employs a methodology designed to map how these abstract conflicts manifest as concrete discursive structures and community dynamics within the public sphere. Existing research has broadly identified recurring tensions, such as the oscillation between utopian promises and technological failure, the conflict between managerial efficiency and labour displacement, and the gaps in governance and accountability. Building upon this strand, our study adds significant empirical depth by examining the extent to which these tensions occur across diverse social media environments. By moving, thus, beyond previous studies on well-studied platforms like X/Twitter or Reddit, this research analyses how these discursive tensions are refracted through the specific affordances of Facebook, TikTok, and YouTube across four distinct European national contexts: France, Spain, Germany, and Ireland. This approach allows us to determine if these theoretically identified tensions are universal or if they are intensified by local sociopolitical realities and platform-specific cultures.

The present study further advances the literature by mapping how themes falling under these tensions are distributed across different types of actors and user communities. Rather than treating the “public” as a monolith, our mixed-methods approach, combining lexicometric analysis and network analysis, identifies which specific groups of social media users drive particular narratives. Therefore, our methodological framework serves as a benchmark to measure the “social legitimacy” of AI by observing how it is tested in the daily interactions of European citizens. It identifies the specific conditions under which acceptance is granted or withheld, moving the discussion from technical metrics to social contestation. What is more, the study and its cross-country scope highlight how local nuances, such as concerns in our social media data from France over digital sovereignty or anxieties expressed in our data from Ireland regarding labour market pressures, reshape the broader global tensions identified thus far in the literature. This granular level of analysis provides a foundational understanding of the governance pressure points that emerge when AI is integrated into varied social and professional realities.

3. Methodology

3.1. Platform selection strategy

We selected Facebook, TikTok, and YouTube for this report based on empirical relevance and methodological feasibility, as well as for the variety of socio-technical affordances of the platforms. We further based our strategy on key findings of The Reuters Institute Digital News Report (2025) regarding our report’s four national contexts, which indicate substantial usage rates: Facebook ranges from 42% (Germany) to 61% (France), YouTube from 51% (Germany/Ireland) to 59% (Spain), and TikTok from 15% (Germany) to 31% (Spain). Put simply, these figures show how the selected platforms function as key spaces for the public’s engagement with digital content and information and, therefore, are relevant for observing how the discourse around AI is shaped.



Beyond aggregate usage, the three platforms differ meaningfully in terms of user practices and content. Facebook, despite a continuing decline among younger audiences, remains a key arena for public-facing communication by organizations and political leaders. YouTube, one of the oldest online platforms, spans a broader demographic, offering a wide range of content entertainment to education (Rieder et al., 2020). TikTok, whose usage is predominantly driven by younger populations, operates through short-form videos shaped by its recommendation algorithm and by users' widespread practice of reusing and remixing popular songs and other audiovisual content. These platform differences were used not to generalize about users' identities or motivations but to enable a comparative analysis of how AI discourse circulates through distinct technical and cultural formats (Siapera, 2018).

This platform selection strategy was further constrained by practical limitations such as access to data. Specifically, despite multiple attempts, securing access to the Application Programming Interface (API) for X (formerly Twitter) proved unsuccessful, a difficulty attributed to the severe restrictions implemented following its acquisition by Elon Musk; furthermore, the European Commission has taken enforcement action against X under the Digital Services Act (DSA) in relation to its compliance obligations, including data access for vetted researchers.¹ As a result, we found it more feasible and preferable to utilise the platforms' provided means to conduct our research. Moreover, to the extent possible, we wished to gain insights into the discourse around AI from varied perspectives that would also centre on general users, that is, not fully professionals or prominent figures. However, the quality and exhaustiveness of the data obtained are uneven across platforms, which unavoidably affected our subsequent methodology. For instance, the Meta Content Library retrieves only posts published by Pages and public Profiles meeting specific visibility/reach criteria.

As a result, we combined the data sources to approximate this diverse perspective, using the Meta Content Library, the TikTok Research API, the YouTube Data Tools (Rieder, 2015) and the YouTube API for comments. In doing so, we believe that our approach captures each platform's nuances, which could be summed as follows: Facebook's page posts for institutional and semi-professional discourses; TikTok's videos for a more hype-driven and creator-oriented discourse; and YouTube's comments, and communities formed based on them, a more educational and political one.

3.2. Lexicometric and Discourse Analysis

To analyse the discourse emerging from our diverse dataset comprising Facebook posts, TikTok video content, and YouTube comments, we primarily rely on lexicometric analysis. This methodological choice follows precedents in the field, such as the approach employed by numerous international scholars of applied social sciences studying public debates across different media environments over a long period (e.g., Smyrniaios & Ratinaud, 2017; Ramos et al., 2018; de Rosa et al., 2023; Tsimpoukis, 2025; Mieles et al., 2026).

¹ https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2934.



In our methodology, discourse analysis is computationally supported by IRaMuTeQ (Ratinaud, 2024), an open-source software that algorithmically dissects textual corpora into thematic clusters based on lexicometric principles and is broadly considered part of the recent “computational turn in the humanities and social sciences” (Rogers, 2019, p. 4). This process, known as Descending Hierarchical Classification (DHC) (also known as the Reinert method) (Reinert, 1993), allows us to explore and compare the discourse employed in different textual corpora (Chaves et al., 2017; Ratinaud et al., 2019). The subsequent interpretation of the DHC-produced clusters remains theory-led according to the overarching conceptual framework of WP4 and FORSEE. To be clear, the concept of discourse analysis is quite malleable but, in general, refers to the investigation of written or spoken texts, that is, “the analysis of language-in-use, along with the structures and institutions in which it is shaped, which it shapes and in which it takes place” (Ali, 2019, p. 405). By doing so, we can understand both how social media users frame AI and how these framings intersect with broader sociopolitical and economic debates.

There are several reasons for choosing the Reinert method. First, the project team has implemented IRaMuTeQ-based DHC across varied and large corpora (e.g. press discourse, platform-related texts, and grey literature) in prior research, which supports methodological continuity and enables cautious comparison with earlier results. Second, while alternative semi-automated discourse analysis techniques (e.g. topic modelling, supervised classification) offer different affordances, existing research suggests that the results produced by these methods are broadly comparable when applied to similar corpora (Alboni et al., 2023). The Reinert method was therefore selected as a well-tested and appropriate methodological approach for inductively identifying recurring lexical and discursive patterns even in heterogeneous textual corpora like the ones analysed in this report and across FORSEE. The choice of a word bag classification, without resorting to embedding, is also based on the type of data we have. This work focuses on written traces found in specific contexts such as TikTok videos or YouTube comments. This particularity implies variations in narrative forms and vocabulary used, such as slang, informal language, and platform-specific vernaculars. These variations can pose problems when using vector representations based on training corpora that are far from our sample (novels, press, encyclopaedias, etc.).

In our textual analysis of the social media corpora, IRaMuTeQ segments the corpus into Text Segments (TSs) and partitions them into thematic classes based on their vocabulary distribution; the outcome is visualised as a dendrogram that spatialises lexical convergences and divergences. Each lexical class produced by DHC can be understood as a frame of discourse, since the statistically significant co-occurrences that constitute a class describe the meaning universe within which a word takes shape. This approach, known as “frame mapping” (Ledouble & Marty, 2019), offers an inductive method for identifying frames without predetermining categories, in contrast to deductive frame coding. As Solberg and Kirchoff (2024) write with regard to topic modelling, which is not very dissimilar to what we are doing with our research, “provides an empirically grounded



method for automated coding of manifest content in large, domain-specific text corpora sampled from public discourse” that “when combined with close readings of latent meanings in exemplary texts” (like our characteristic TSs) we can identify prominent topics in public discourse (p. 656).

Moreover, the way clusters are visualised in the generated dendrogram reflects their semantic vicinity. Put simply, if two clusters belong to the same sub-branch, this suggests that their TSs mobilise similar vocabularies or closely related contexts; thus, this also allows us to observe and interpret clusters’ relations. It should also be mentioned that this analysis is agnostic to provenance at the clustering stage. In other words, the DHC classes are formed on the entire dataset, without regard to variables such as country, outlet, or date; those variables are only reintroduced later, via chi-squared (χ^2) over/under-representation checks, to examine how modalities (e.g., metadata) distribute across classes.

The chi-squared association test underpins IRaMuTeQ’s classification procedure by measuring how particular words or text segments are associated with a specific lexical class rather than distributed randomly across the corpus (Souza et al., 2018). In practice, the test compares the observed frequency of each term within a class to the expected frequency if the distribution were even. A high χ^2 value, therefore, indicates that a word occurs far more often in that class than chance would predict. This same test can be extended to external variables such as platform, date, or language, revealing when certain lexical worlds become statistically overrepresented. This allows us to dive into specific themes and explore such dimensions in a more nuanced fashion.

Following the lexicometric analysis, each lexical, or DHC, class produced by IRaMuTeQ is defined by a set of characteristic TSs, that is short excerpts statistically representative of that class. For each class, we examined several hundred characteristic text segments (CTSs) to gain a comprehensive view of the themes discussed but ended up retrieving quotes (when necessary) only from the top 50. Effectively, these segments provide important context to the DHC clusters, allowing us to critically analyse the constructed discourse. Last, it is worth mentioning that our analysis is also combined with network analysis—only in the case of YouTube in the form of co-commenting networks—which we thought to be crucial to explore whether patterns of engagement, i.e., users commenting on the same set of videos, corresponded to similar patterns of discourse; this is further elaborated on in YouTube’s subsection.

4. Analysis of social media platforms

4.1. Facebook

4.1.1. *Data and methods*

The Facebook dataset consists of four independent corpora, one for each national context examined in WP4: France, Spain, Ireland, and Germany. All data was extracted using the Meta Content Library, Meta’s public transparency interface that replaced CrowdTangle



allowing researchers to search for and export posts published by eligible Facebook Pages and Profiles. The dataset (Table 1) was constructed using keyword-based searches in the Content Library’s “Page Posts” search function, applying the extraction constraints imposed by the tool. Specifically, the tool only retrieves posts published to:

- Pages
- Groups
- Events
- Profiles with $\geq 1,000$ followers
- Profiles with a verified badge

Also, the tool does not allow exact-phrase searches (i.e., no quoted strings). Searches must therefore be performed on individual terms. For methodological consistency, the extraction was carried out using the terms *intelligence* and *artificial* (including their equivalents in French, German, Spanish, and English). The inability to search for the exact phrase “artificial intelligence” raised the possibility of retrieving irrelevant content. However, initial qualitative verification of sampled posts showed that the retrieved material overwhelmingly referred to the technological meaning of AI, rather than to unrelated uses of “intelligence” or “artificial”. To ensure corpus homogeneity and limit noise from groups, which often host heterogeneous or even spam content, the extraction was restricted to Pages and eligible Profiles. We found our decision necessary after early tests revealed that group-based results downgraded the quality of our corpus, primarily due to noise and irrelevance.

Publicly available evidence indicates that content published by Facebook Pages accounts for a relatively small share of the material displayed to users in the News Feed. Meta’s Widely Viewed Content² reports suggest that posts originating from followed Pages represent approximately 15-20% of the content seen by users, with the remainder largely coming from personal profiles and groups. In parallel, Facebook does not release official statistics on the distribution of followers across users; however, based on well-documented heavy-tailed attention distributions observed across social platforms, it is reasonable to estimate that only a small minority of users (on the order of 1-5%) have more than 1,000 followers. These figures should be interpreted with caution, as they rely on indirect measurements and modelling rather than complete platform-wide data. Consequently, a dataset composed of posts from Pages and from profiles with more than 1,000 followers captures a minority but structurally influential subset of Facebook content, characterised by higher public visibility and reach, and not a representative sample of typical user activity.

Moreover, another key constraint of the tool is that country targeting, i.e., geolocation filtering, works only for Pages and Profiles, not Groups. We therefore restricted our extraction to Pages and eligible Profiles with an identified posting location to ensure that the corpus corresponds as closely as possible to the four countries of our scope. In fact, this step became critical after initial exploration of the French-language corpus revealed large volumes of content originating from French-speaking countries outside the EU

² <https://transparency.meta.com/reports/widely-viewed-content-report/>.

(particularly in sub-Saharan Africa); despite our attempt to minimise the ‘noise’, we were unable to completely eliminate it and, as such, caveat wherever the findings accordingly. Last, the exported data contained numerous identical posts that were cross-posted across multiple accounts, or slightly modified versions circulating as “near duplicates”. The identification of these duplicates is complicated by minimal changes (for instance, differences in punctuation or URL parameters); to the extent, which was possible, we deduplicated the corpus, but some repetition persists. In the end, we ended up with a dataset of 218,838 Facebook posts, with Spanish being by far the largest sub-corpus (125,224) and Irish the smallest (8,568), which warrants our attention as regards the comparison of findings between countries but also invites us to look at the specificities of each country individually. Our entire dataset was then parsed with IRaMuTeQ for subsequent lexicometric analysis, which was conducted on approximately 600,000 Text Segments (Table 1). Below, we present an overview of the DHC analysis for the four countries (Figures 8-11).

Country	Query	Number of Posts	Number of Text Segments
France	intelligence artificielle	59,047	198,405
Spain	inteligencia artificial	125,224	304,151
Ireland	artificial intelligence	8,568	19,929
Germany	künstliche intelligenz / künstliche + intelligenz	25,999	76,880
Total		218,838	599,365

Table 1 - Facebook corpus statistics by country

4.1.2. Findings

The analysis of Facebook discourse across the French, Spanish, Irish and German corpora reveals a heterogeneous yet patterned communicative environment. Although each national corpus reflects specific sociopolitical and media contexts, the thematic structure produced by the DHC exhibits cross-country consistencies that can be grouped into three distinct framings: a market –and application–oriented discourse; a technical, institutional and regulatory discourse; and a broader societal and human-centric dimension.

4.1.2.1. Market-oriented and application-based discourse

A significant portion of the discourse across all four countries frames AI as a tool for economic growth, professional development, and specific utility, mirroring the application-oriented trends observed on other platforms of our sample (Figure 1).



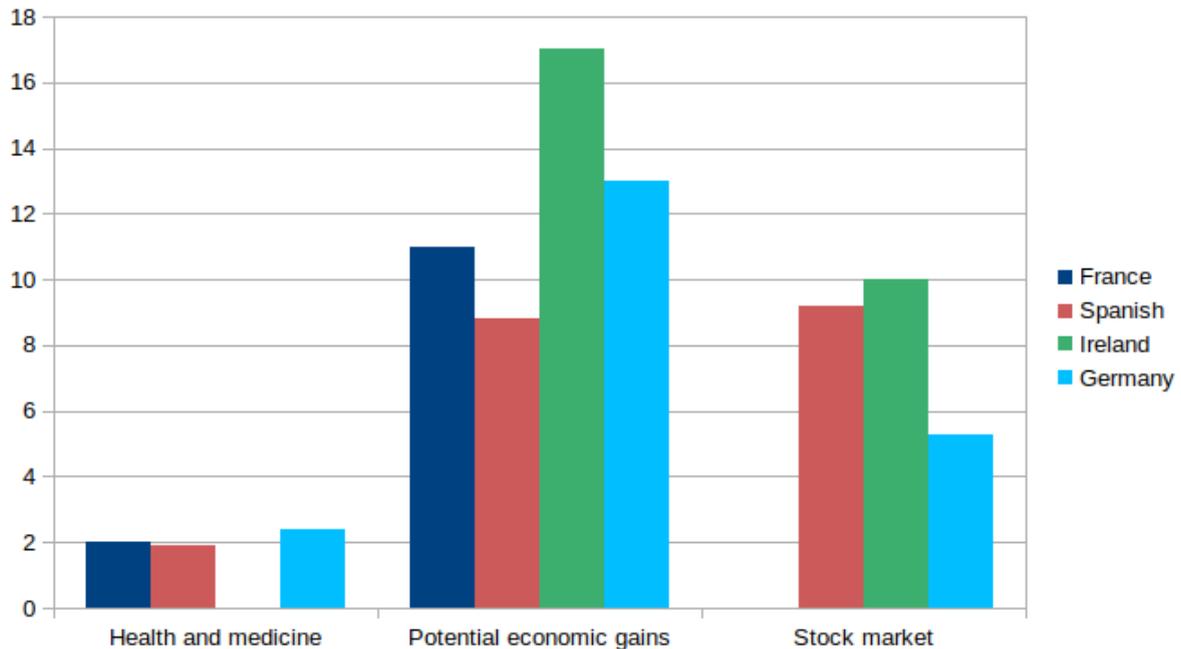


Figure 1 - Histogram of common themes in market-oriented and application-based discourse for the Facebook corpus (results as a percentage of each corpus)

In the corpus of France (Figure 8), this domain is represented by Class 7 (12.9%), which focuses on AI use cases, especially in commercial and entrepreneurial contexts, including productivity tools, digital marketing applications, and customer service automation. Similarly, Class 6 (11%) focuses on AI development oriented towards developing countries, particularly in French-speaking Africa such as Senegal and Cameroon, presenting AI as a lever for transformation and progress across key sectors of their economies. Furthermore, Class 10 (2.03%) addresses AI in health and medicine, including posts about early detection of neurological conditions and applications in patient monitoring or diagnostics.

The corpus of Spain (Figure 9) exhibits a broad thematic distribution in this domain. Class 7 (12.9%) concerns training courses and advertising of professional courses and technological upskilling, while Class 11 (11.6%) highlights public programmes, often not exclusively AI-related, that promote digital tools as instruments of sectoral development. Class 10 (8.82%) addresses AI's technical capabilities and potential productivity gains, and Class 3 (9.2%) focuses on corporate competition, takeovers, and AI-related announcements. Specific applications are also found in Class 5 (1.89%), detailing AI-related advances in healthcare and medicine, and Class 4 (1.41%), which features AI-generated predictions especially for sports analytics and sports betting. Finally, the two smallest clusters, Class 2 (0.85%) and Class 1 (0.73%), consist of posts ending with "written using AI" and primarily contain AI-generated cooking recipes.

In the corpus of Ireland (Figure 10), broader societal implications of AI are captured in Class 10 (17%), where AI is framed as a market booster and a driver of scientific and economic

progress, with discourse that includes concerns about job displacement and optimism about new employment opportunities and innovation.

This economic focus is echoed in Class 6 (6.3%), which foregrounds the latest developments as regards the stock market news related to AI, and Class 5 (3.7%), which contained references to key stock market indexes like NASDAQ. The corpus of Germany follows this pattern with its largest cluster, Class 7 (24.2%), encompassing talk about applications across health (including early detection and diagnosis), business management, sustainability, and education. Smaller but thematically significant clusters include Class 1 (5.3%), which is centred on economic news, stock-market trends, mergers, and fundraising activity in the tech sector.

4.1.2.2. Technical and institutional discourse

Distinct from user-centric technical advice, the technical discourse on Facebook is deeply institutional, dominated by references to Big Tech, academia, and governmental regulation.

The corpus of France (Figure 8) displays a distinct technical domain with a sector-specific and professional orientation. Class 11 (16.8%), the largest cluster in the technical domain, cites to academics, institutional actors, research programmes, and governmental initiatives, frequently invoking national-level political figures such as French President Emmanuel Macron and situating AI within broader digitalisation agendas. Class 4 (7%) contains news from the AI sector, emphasising corporate competition, technological announcements, and developments involving large corporate tech actors such as Google, Apple, Microsoft, and OpenAI. Additionally, Class 8 consists of posts announcing AI events, conferences, and workshops, which often serve as dissemination points for institutional priorities.

In the corpus of Spain (Figure 9), enterprise and institutional discourse are strongly represented. Class 12 (11.6%) compiles institutional profiles, CVs, and academic conferences. The corpus of Ireland (Figure 7) similarly foregrounds this dimension in Class 7 (13.7%), which is centred on key corporate and state actors and foregrounds the geopolitical dimension and global competition of AI. Class 2 (3.8%) centres on regulation and contains references to the EU AI Act, its legislative process, and, broadly, the EU's regulatory agenda. Furthermore, Class 4 (12.1%) outlines the technical foundations of AI, such as machine learning, and repeatedly emphasises data processing as a core capability across multiple fields.

The corpus of Germany (Figure 11) reflects this institutional weight in Class 6 (18%), which highlights exchanges and events involving academics and companies, foregrounding the same institutional and corporate actors we saw in the other countries. Class 2 (15.6%) focuses on the substantial public funding schemes in the EU and US devised to boost AI development, as well as references to the invasion of Ukraine by Russia and the broader geopolitical developments that concern Europe, and that also relate to the so-called global



AI race. Sporadic references to AI-related research and academic projects are also found in Class 5 (7.4%).

4.1.2.3. Societal and human-centric discourse

Finally, a complex social domain emerges across all corpora, addressing the geopolitical, ethical, philosophical, and creative implications of AI (Figure 2).

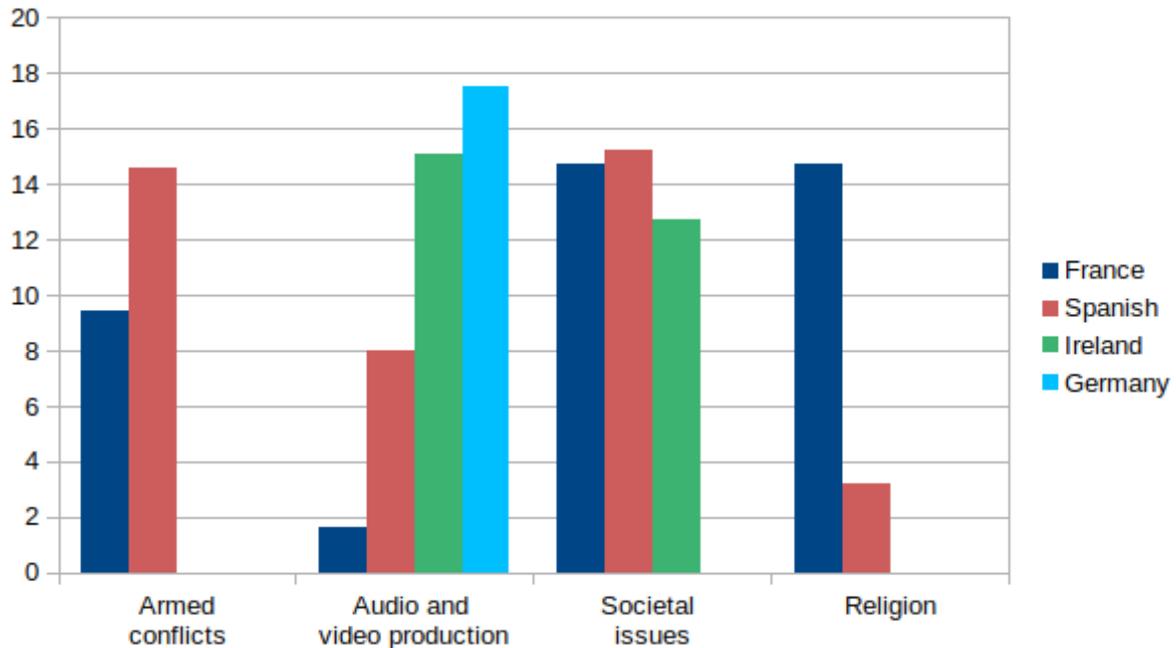


Figure 2 - Histogram of common themes in societal and human-centred discourse for the Facebook corpus (results as a percentage of each corpus)

The corpus of France (Figure 8) separates this material into distinct social fields. Class 2 (9.41%) is one of the densest and most thematically consolidated clusters, centred on armed conflicts, particularly Israel's war on Gaza and the Russian invasion of Ukraine, and includes strong criticism of AI-fuelled warfare and surveillance. Class 1 (1.61%) contains audiovisual news segments, typically reposted materials summarising AI-related developments circulated by media outlets. Regarding creativity, Class 5 (11.6%) focuses on AI-generated images, ranging from casual experimentation with generative tools to explicit debates concerning artistic legitimacy and concerns about deepfakes, while Class 9 (14.76%) contains personal stories and commentary that reflect anxieties or hopes regarding AI's social and emotional implications, including reflections on authenticity and the nature of human creativity. Uniquely, Class 3 (7.94%) comprises religious-like discourse in which AI is juxtaposed with, or even associated with, divine consciousness and spirituality; this class we identified is mostly from African French-speaking users.

The corpus of Spain (Figure 9) mirrors these themes with large clusters concentrated in narrative and geopolitical discussions. Class 6 (14.6%) centres on ongoing wars, particularly Gaza and Ukraine, and includes explicit criticism of AI-fuelled military practices, especially in the context of Israel's war on Gaza. Class 13 relates to various stories (15.2%) that

comprise narrative segments addressing societal transformations and the future implications of AI. Class 8 (8%) covers AI-generated audio and visual content, including posts reflecting on the artistic legitimacy of AI and on what constitutes art broadly, while smaller clusters, like Class 9 (3.2%), juxtapose divine or quasi-religious themes with AI.

In the corpus of Ireland (Figure 10), Class 9 (3.4%) contains references to deepfakes and AI-generated or amplified mis- and disinformation and is closely connected to Class 8 (15.1%), which contains references to—not exclusively related to AI-generated—audiovisual and cultural production, including Hollywood films. While the corpus of Ireland necessitated specific preprocessing due to URL-heavy posts (clustering noisy data into Class 1, 12.7%), the remaining discourse aligns with these societal concerns.

The corpus of Germany (Figure 11) addresses these human-centric themes in Class 3 (17.5%), which covers both the ease of producing texts, images, and poems using AI and, to a certain extent, the ongoing debate over the limits of AI in performing complex or emotion-laden tasks. Class 4 (11%) also contains references to media content.

4.1.3. Concluding remarks

The analysis of Facebook discourse across the French, Spanish, Irish and German corpora reveal a heterogeneous yet patterned communicative environment in which AI circulates through news recirculation, institutional messaging, audiovisual content, promotional material, and varied user-generated posts. Although each national corpus reflects specific sociopolitical and media contexts, the thematic structure produced by the DHC exhibits some cross-country themes. Across all four countries, AI discourse consistently organizes around six domains: creativity, AI-generated content, manipulation and deepfakes; economic competition, market activity, and professional training; applications in health, science and public management; geopolitical implications and military uses of AI; institutional, academic and regulatory discourse; and philosophical, existential, or even spiritual framings. Moreover, Big Tech firms function as dominant reference nodes, shaping discussion around competition and investment (e.g., FR Classes 4, 7; ES Classes 3, 7; IE Class 7; DE Classes 1, 2). Notably, economic and labour narratives recur across corpora, linking AI to productivity, automation, entrepreneurship, and the need for upskilling, while scientific and medical framings emphasise progress in diagnostics, data analysis, and operational efficiency.

At the same time, though, risk narratives display sharper national variation. By way of illustration, France and Spain host the most contextualised and ethically charged accounts of AI, especially in relation to armed conflict (FR Class 2; ES Class 6) and manipulation through deepfakes or synthetic media. Ireland and Germany exhibit more diffuse risk framings: Irish discourse focuses on audiovisual manipulation and labour concerns (IE Classes 6, 4), while the corpus of Germany treats risk primarily through scepticism regarding machine vs. human cognition and the limits of generative systems (DE Class 3). Surveillance, privacy breaches and misuse appear most prominently in France and Ireland



but remain secondary themes when compared to the pervasive economic and institutional discourse.

4.2. TikTok

4.2.1. Data and methods

The TikTok dataset consists of the four national corpora corresponding to the same four countries examined in WP4: France, Spain, Ireland and Germany. All data were collected using the TikTok Research API, which provides controlled access to public content for academic research. According to TikTok's API documentation, keyword queries are applied to the video description field. Other elements of the video object, though (e.g., hashtags, on-screen text, audio metadata), are not indexed for keyword retrieval through the Research API. Usernames and account identifiers in the API results were retained in the dataset but not used as analytical variables, as the focus of this task was on discursive patterns rather than influencer or account-level analysis. The resulting dataset, therefore, reflects not only random user-generated content but also TikTok's own internal ranking and indexing mechanisms, which constitute methodological constraints for working with this platform.

For methodological consistency with the other social-media corpora, we tested several possible query terms. However, acronyms such as "AI", "IA", or "KI" (the abbreviated forms of Artificial Intelligence in the languages covered) were excluded from the search queries. Preliminary tests showed that these acronyms generated large volumes of irrelevant material. For instance, "KI" may be used outside a technological context, including as an onomatopoeic expression to resemble the sound of birds (e.g., crows) or as a common function word in other languages, which substantially reduced the thematic precision of some of the retrieved results. The final dataset was therefore built using only the terms "Artificial Intelligence" and "ChatGPT", which yielded the most relevant and least ambiguous results during the collection period. We believe that our decision was further warranted due to TikTok's own affordances, whereby the platform's content is often steered by popularity, or even hype cycles of specific phenomena that are translated via hashtags; "ChatGPT", thus, functioned as the dominant proxy to discuss AI during the concerned period. In the same vein, we also tested other Generative-AI-related keywords (e.g., Gemini, Copilot, DeepSeek); they either returned no new relevant content or returned earlier videos with little relevance to our enquiry. Quickly, then, we understood that for many TikTok users, only a small set of hashtags, chiefly "ChatGPT", served as their 'interface' to discuss and tether their content to AI-related content, further fuelling the hype of OpenAI's famous Gen-AI product. As a result, we stuck with "Artificial Intelligence" and "ChatGPT" as our two queries.

We gathered 244,784 videos in total. We then removed unrelated videos based on their descriptions (approx. 32% of the initial dataset, or 7,883 videos) and deduplicated the dataset. Still, the dataset contained significant noise, and after some initial tests, our analyses did not yield meaningful results. We reduced the dataset for transcription by selecting the 500 most-viewed and 500 most-commented videos per country (after deduplication). Thus, the sample of videos per country is between 500 and 1000,

21



Funded by
the European Union

depending on the number of those that were both in the 100 most-viewed and the 100 most-commented lists (figures described in Tables 2 and 3).

	Country	N	Mean	SD	Minimum	Maximum	Percentiles		
							25th	50th	75th
comment_count	France	741	680.59	1301.4	0	19030	185.000	387.000	716.00
	Germany	767	816.72	1819.2	8	28586	194.500	426.000	763.50
	Ireland	735	81.75	538.1	0	10495	4.000	9.000	19.50
	Spain	755	1304.20	2001.4	0	27995	369.500	756.000	1385.00

Table 2 - Statistics on the number of comments of TikTok videos by country

	Country	N	Mean	Median	SD	Minimum	Maximum	Percentiles		
								25th	50th	75th
view_count	France	741	645914	366811	1.111e+6	0	14040098	167682	366811	670965
	Germany	767	653113	307909	1.511e+6	0	28231745	120442	307909	708102
	Ireland	735	88626	5747	639190	0	11966676	1952	5747	14899
	Spain	755	1.617e+6	1072435	2.161e+6	0	34631177	524393	1.072e+6	1.893e+6

Table 3 - Statistics on the number of views of TikTok videos by country

This decision was taken based on computational and logistical constraints: audio-to-text transcription using OpenAI Whisper is resource-intensive, and processing the full dataset was not feasible within the available time and hardware capacity. The reduced sample, therefore, represents the most visible and most interacted-with TikTok content related to GenAI during the period. It is also worth noting that Whisper constantly mis-transcribes key terms like ChatGPT or produces non-words, and that accuracy of transcription was not equivalent across languages, owing to different accents and the quality of the recorded audio. A further recurring issue was that a portion of the retrieved videos contained no intelligible speech, despite being tagged with AI-related hashtags. This phenomenon reflects common TikTok user practices, in which hashtags are used not to describe content but to exploit algorithmic visibility cues. These transcription errors introduce noise into the corpus. While the lexicometric analysis can mitigate some of this (i.e., clusters often stabilise even when key terms are distorted), transcription inconsistencies remained in our dataset despite systematic efforts to mitigate them to the best of our capabilities. Each sampled video was further annotated with engagement metrics (i.e., likes, shares, comments, views), with their views divided into quarters. Finally, all transcribed material was formatted for DHC analysis using IRaMuTeQ.

4.2.2. Findings

Across countries, the DHC analysis consistently revealed significant heterogeneity in the discourse regarding AI on TikTok, which we identified to revolve primarily around three main axes: GenAI as a content-creation force, GenAI as a tool to be mastered, and GenAI as a complex societal issue. These themes appear across all countries, but each corpus foregrounds different aspects depending on local concerns, platform cultures, and the



narratives users choose to amplify. These variations underscore that TikTok is less a single “AI discourse space” than a set of parallel publics shaped by national contexts, platform norms, and the immediate utility users seek from generative systems. At the same time, the three axes also reveal shared concerns and expectations. Users across all four countries treat AI as something they must learn to handle, whether to create content, remain employable, or simply optimise and manage daily tasks. They express a mixture of enthusiasm and unease: enthusiasm in the form of “optimising” workflows, scaling creative output, or generating playful synthetic media; unease in the form of anxiety about human replacement, personal data, institutional reactions, or the direction of global competition. Some societal concerns emerge in key domains, ranging from automated traffic enforcement (Spain) and environmental costs (France) to issues pertaining to the labour-market (Ireland). These discourses do not express a single, coherent position on AI “successes” or “failures”, but rather show that public orientations fluctuate between opportunity, experimentation, and anxiety. In the next paragraphs, we unpack each of these axes, while situating them in their respective national contexts, which further nuance our findings.

4.2.2.1. GenAI as a content-creation force

This axis centres on discourses in which GenAI is primarily framed as capable of (re)producing textual or audiovisual material with a significant emphasis on the outcome rather than the process or mechanics of producing that content (Figure 3).

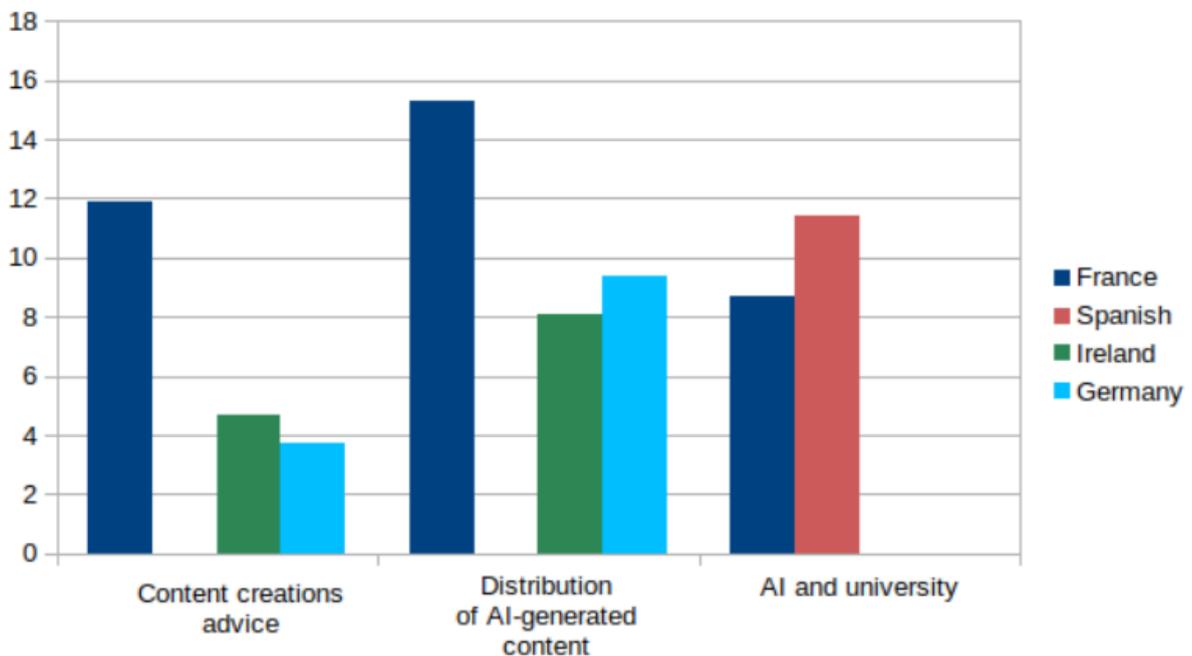


Figure 3 - Histogram of common themes in discourse on GenAI as a content-creation force for the TikTok corpus (results as a percentage of each corpus)

The corpus of France (Figure 12) foregrounds concerns about academic integrity alongside a seeming concern regarding the impact of GenAI on creative production. Class 13 (8.7%), defined by terms related to research (e.g., research, dissertation), seems to situate AI within

debates regarding higher education and student use of GenAI tools. Characteristic Text Segments (CTSs), for instance, refer to debates on how universities respond to ChatGPT, e.g., how Sciences Po, an elite French educational institution, announced a ban on its use for assignments, with possible sanctions ranging up to student exclusion. In parallel, Class 9 (8.1%) and Class 10 (3.2%) are dominated by content-creation terms (e.g., create, generate, image, etc.), and illuminate a more instrumentalist, or even entrepreneurial, approach to creativity focused on scaling content generation. For example, CTSs illustrate this by promoting 'optimal use' of GenAI chatbots via prompts to "instantly generate" hundreds of reels to flood social media (Video, 20/04/2024) or through the circulation of AI-generated trailers for non-existent sequels such as Titanic 2, blurring the boundary between conventional fan-made creations and fully automated outputs

Gen-AI content creation in the corpus of Spain (Figure 13) mostly revolves around commercialisation thereof and the pursuit of algorithmic visibility. Class 14 (15.3%), the largest cluster in this domain, contains songs related to audio generation (song, hear, popular, etc.), referring to the use of GenAI for remixes and voice cloning, including for the 'reanimation' of deceased artists (e.g., the release of a new song of The Beatles using John Lennon's voice). CTSs include examples of GenAI-created advertisements, such as a KFC commercial described as being made "100% with AI" and produced without physical recording. Class 3 (7.8%) further reflects this development, containing references to AI-generated surreal content, such as videos of animals performing tongue-twisters—what the Internet slang has dubbed "brain rot" content created solely to capture attention and generate buzz. We also noted here a prevalent discourse about students using GenAI applications like ChatGPT (Class 7).

The corpus of Ireland (Figure 14) further emphasises the professionalisation of GenAI content production, structured into two clusters. First, Class 6 (8.1%) contains more instructions and prompts to generate viral content, leveraging the platform's affordances and recommender systems. However, we also noticed instances where users promoted the use of GenAI to produce professional commercial content, from birthday flyers to song lyrics and other sellable assets. Second, Class 3 (4.7%) includes a complementary strategy, where GenAI can be used to optimise content copy and reach, with CTSs describing workflows for using generative systems to "reverse engineer viral videos" and rewrite them for a particular "niche" audience. It is also worth mentioning that in Class 7 (17.46%), we found a relatively prevalent discussion around the question of copyright, or rather, the reproduction of artists' material by GenAI systems without artists' consent or remuneration.

Last, creative discourse in the corpus of Germany (Figure 15) is roughly split between an embrace of the absurd and the surreal, and a more structured, strategic approach to GenAI content creation. To illustrate, Class 3 (9.4%) contains CTSs that refer to the "Noga Boga" trend, in which users intentionally generated absurd stories, following the broader "brain rot" phenomenon. (e.g., one video describes a character who "ate Paris because they thought it was a croissant". In contrast, Class 7 (3.7%) indicates a more professionalised use of GenAI, offering guidance on audience segmentation and editorial planning (a trend that is even



more pronounced in the next axis). In one example, a user posted a prompt for GenAI chatbots to “think thoroughly about the problems, misunderstandings, and dream goals of my target audience” as a content-planning strategy. We also noted mentions of GenAI to help users with everyday chores such as preparing a shopping list (Class 12)

4.2.2.2. AI as a tool to be mastered

On this axis, AI is perceived and discussed on TikTok as a tool that requires technical skills and training, with a strong focus on prompting strategies. Parts of this axis were also found in the previous one but, in this case, they focus predominantly on the technical training aspect, rather than just optimising the output and the subsequent instrumentalisation thereof. Certainly, as with all content, the videos themselves operate within the same logic and platform affordances and thus reproduce what they also seek to impart to other users.

In the corpus of France (Figure 12), we observed the dynamic of this axis primarily in one DHC cluster related to mastering consumer-facing tools. First, we find Class 7 (7.6%) corresponds to a “tutorial” discourse, defined by words indicating instructions (e.g., click, press, choose, download, etc.), while CTSs contain “click-by-click” workflows to automate specific tasks such as generating animations or presentations. Notably, there is also a significant volume of content dedicated to “humanising” GenAI text to bypass plagiarism detectors, especially for academic use.

In the corpus of Spain (Figure 13), we see significant focus on professional upskilling, particularly, in the programming sector. Class 4 (8.4%) contains terms like Python and coding, and videos promoting “full-stack development bootcamps” to learn Python using GenAI-supported tutorials. Further, we see two other two clusters relating to the mastering of GenAI as a tool: Class 7 (11.4%), which, as in most other countries, contains numerous references to the use of AI as a shortcut for schoolwork, and Class 5 (4.2%), which includes various tutorials for users to leverage general-purpose tools for specific tasks.

Our data on Ireland (Figure 14) highlights the practical application of AI as a support tool for employment and job search. In particular, Class 2 (8.9%) is clearly oriented towards the job market, containing terms like CV, job, resume, etc. This is supported by CTSs that serve as step-by-step guides, where users share specific GenAI chatbot prompts to rewrite CVs to better match job descriptions or to prepare for interviews. Also, Class 7 (17.5%) invokes the need for people to master (Gen)AI to commercial ends.

Lastly, the corpus of Germany (Figure 15) frames AI as a tool to managing everyday life tasks. Class 12 (9.6%), for instance, includes examples of using ChatGPT to organise daily life, such as generating weekly meal plans or shopping lists based on specific budgets, highlighting how GenAI can be used to ‘optimise’ mundane tasks in a supposedly individualised manner.



4.2.2.3. *AI as a complex societal issue*

The final axis shifts from individual practices to collective imaginaries, treating AI as a socio-economic or geopolitical object.

The French cluster (Figure 12) also encompassed a part of the discourse that predominantly related to the country's strategy for digital sovereignty. Class 12 (9%) and Class 5 (9.6%) situate AI within a narrative of national competitiveness and investments in AI infrastructure. For instance, we see references to the "France 2030" plan and a "€109 billion investment" to compete in the global AI race, while certain users raise awareness of the physical footprint of the technology (e.g., the environmental impact of data centres). Interestingly, we also noticed a sizable part of the discussion in Class 2 (8.3%), focuses on religious matters with users asking GenAI chatbots existential, or even esoteric questions.

The Spanish data (Figure 13) also highlights the AI applications in public and civil administration. For example, Class 8 (3.2%) includes discussions about AI in the context of traffic management, referring to the use of AI to automatically fine drivers for violations such as not wearing seatbelts or using mobile phones. Beyond this rather narrow focus, we also see broader societal issues captured in Class 10 (11.6%) and Class 6 (4%), which discuss the replacement or even the extinction of humans by machines, invoking well-known science fiction themes.

Our data on Irish TikTok users (Figure 14) seem to mobilise this axis through a political-economic and historical lens. To illustrate, Class 8 (17.8%) situates the current AI boom within a broader historical context, comparing it to the rise of the commercialised Internet, while we also notice critiques of techno-solutionist ideas and worries over data privacy.

Finally, the social aspect appears in our corpus of Germany (Figure 15) as heavily shaped by geopolitical competition and developments in the global AI race, as we also observed in the French data. For instance, Class 5 (9.9%) is dominated by discussions of DeepSeek and China, as well as the broader discourse on China's approach to AI. As we saw in our Irish data earlier, we identified a discourse related to science fiction (Class 6; 9%), in which users reflect on the future "intentions" of AI systems and their broader impact on humanity.

4.2.2.4. **Concluding remarks**

The comparative analysis of the four national corpora shows that TikTok primarily serves as a space where AI is approached in highly practical, use-oriented ways. Rather than hosting abstract debates about AI policy or long-term societal impacts, TikTok content largely frames GenAI as a set of tools to be applied immediately, whether to improve visibility on the platform, simplify schoolwork, or enhance professional competitiveness. Across countries, AI is discussed less as a technological breakthrough in itself and more as something that needs to be learned and used efficiently to keep up with others.



This more application-focused orientation produces a shared dynamic across the corpora. Tutorials, prompts, and demonstrations of AI use are widespread, particularly around content production, productivity, and monetisation. At the same time, these practices are accompanied by recurring concerns about falling behind. In several cases, users explicitly frame AI mastery as necessary to avoid exclusion from creative, educational, or labour markets. Alongside these utilitarian aspects, we also observe the circulation of deliberately absurd or surreal AI-generated content, particularly in Germany and Spain. These forms of “low-effort” or humorous experimentation suggest that TikTok is also used as a space to test the limits of generative systems and engage with AI in playful or ironic ways.

Despite these common patterns, national specificities remain visible. In France, discussions around AI are more frequently linked to institutional concerns, including education, authorship, and state-level strategy. In Ireland, AI use is strongly connected to labour market pressures and the professionalisation of content creation. Spanish discourse places greater emphasis on automation, commercial simulation, and surveillance-related concerns, while the corpus of Germany combines everyday uses of AI with more sceptical or observational commentary. These differences indicate that TikTok does not homogenise public discourse; instead, it provides a shared format for local concerns to be expressed.

4.3. YouTube

4.3.1. Data and methods

We designed our enquiry into YouTube, aiming to capture the part of public debate that sits between the short, fast-moving and engagement-focused content of TikTok, and the more ‘institutionalised’, public-figure-led communication found on Facebook. Video collection was carried out using YouTube Data Tools (Rieder, 2015), using keyword searches in the relevant languages as on other social media platforms (e.g., “intelligence artificielle” for French). The initial retrieval targeted the most-viewed and most-commented videos to focus on content that was highly visible and likely to generate discussion. Searches were run in consecutive six-month intervals, both to work within retrieval limits and to reflect that the videos attracting attention change over time. This work was carried out during July 2025 and focused on videos published between 1 January 2022 and 30 June 2025. We used the country filter offered by YouTube Data Tools (YTDT) for selecting the videos, but as the platform has no geographical limits, the comments are international.

Across all countries, the initial queries yielded 24,184 videos. Some English-language technical terms (e.g., machine learning, deep learning) were excluded from later rounds because they returned large volumes of English-speaking content even when a country filter was applied. After deduplication within each national corpus, we narrowed the material to the 100 most-viewed and the 100 most-commented videos per country and time slice. Thus, the initial sample of videos per country is between 100 and 200, depending on the number of those that were both in the 100 most-viewed and the 100 most-commented lists (figures described in Tables 4 and 5).



	Country	N	Missing	Mean	Median	SD	Minimum	Maximum	Percentiles		
									25th	50th	75th
commentCount	France	145	0	901.7	611	1498.3	0	12702	344.0	611.0	935.0
	Germany	132	4	643.2	419.0	702.4	1	4734	271.8	419.0	833.0
	Ireland	161	1	6966.1	5605	6819.9	0	38708	2084.0	5605.0	8364.0
	Spain	148	0	4502.2	2317.0	8358.0	0	68819	1224.5	2317.0	4438.0

Table 4 - Statistics on the number of comments of YouTube videos by country

	Country	N	Mean	Median	SD	Minimum	Maximum	Percentiles		
								25th	50th	75th
viewCount	France	145	557078	310557	852577	8163	6838042	213955	310557	555012
	Germany	136	407792	208750	495255	20306	2975105	138734	208750	516934
	Ireland	162	1.348e+7	7.406e+6	1.993e+7	189589	139000532	3.616e+6	7.406e+6	1.314e+7
	Spain	148	8.200e+6	3.489e+6	1.408e+7	290733	118350245	1.595e+6	3.489e+6	8.257e+6

Table 5 - Statistics on the number of views of YouTube videos by country

We then screened the videos for relevance, retaining only those that contained substantive references to AI (in the spoken content, visual explanation, and, where needed the video description). Videos were excluded when AI appeared mainly as a production technique without any discussion (e.g., AI-generated music or visuals posted as standalone outputs), when “AI” functioned primarily as a buzzword without meaningful engagement with the topic, or when content was unrelated despite keyword matches. Where possible, we also limited very short videos (including YouTube Shorts) that provided little context for the discussion in the comment section, and we removed items that primarily promoted competitions or similar calls without meaningful content. Certainly, this process introduces an element of researcher judgment, or even bias, shifting the material selected towards primarily explanatory, commentary, news, and “how-to” videos. Regardless, we believe this trade-off aligns with the purpose of the YouTube corpus in this deliverable, which focuses on identifying perceptions of success and failure around AI, and also due to the platform’s data collection limitations necessitate a certain level of abstraction and qualitative methodological decisions. Consequently, such an abstraction also means that, as with the rest of the analyses in this report, our findings and conclusions draw upon and concern the sample analysed here; while we can make some generalisations to a certain extent, the findings should in principle read through this prism and rather limited scope. At any rate, the final corpus consists of 502 videos, distributed as follows:

- **France:** 89
- **Spain:** 136
- **Germany:** 92
- **Ireland (EN):** 185

We first conducted a DHC analysis of the transcriptions of these 502 videos, obtained using Google’s Notebook LLM (Figures 16-19) to confirm they met the above criteria. Subsequently, we retrieved the comments associated with these videos, specifically targeting the authors within the co-commenting network (users commenting on multiple



videos). This process yielded a raw set of **1,244,311 comments**. After cleaning the data, we retained a final corpus of **1,046,821** comments, distributed as follows:

- **France:** 132,552
- **Spain:** 263,548
- **Germany:** 74,654
- **Ireland (EN):** 576,067

The information was cross-referenced to identify the co-presence of accounts commenting on the same video. This data was used to produce an undirected graph in the open-source software Gephi (Bastian et al., 2009), where the nodes represented the accounts, and the links indicated the number of times two accounts commented on the same video. The produced graphs were cleaned by deleting accounts that appeared fewer than three times and/or had links with an edge weight lower than three. Then, we deployed a community detection algorithm (Blondel et al., 2008) to identify groups of accounts that frequently commented on the same videos: communities that represented more than 3% of the remaining accounts were annotated as distinct “communities” in the comment metadata; clusters below this threshold were labelled “other”, while accounts not belonging to any cluster after filtering were labelled “none”. It is also important to note that these clusters should not be interpreted as “communities” in the sociological sense (i.e., exhibiting shared identity) but rather as groups exhibiting common patterns in engagement (i.e., commenting) practices. Put simply, they represent user groups that tend to appear together around similar types of videos (e.g., technical tutorials and political commentary). Finally, we visualised the network of communities using ForceAtlas2 (Jacomy et al., 2014).

Our final step of the analysis combined two approaches. First, as with the previous analyses, the DHC, or lexicometric analysis, was applied to the annotated comments to examine how these clusters of communities distribute across different DHC classes. In addition, we constructed co-commenting networks using Gephi to examine how audiences cluster around different types of content, which we visualised as a network graph.

To accurately interpret the network visualisations (Figures 4-7) presented in the findings:

- **Thematic labels (black text):** The distinct theme of the DHC class in which a user community is overrepresented is indicated in black text next to the cluster. In cases where a community is associated with multiple themes, they are sorted in descending order by on their statistical association (χ^2).
- **Class numbers (red text):** The DHC class number corresponding to the theme is indicated in red.
- **Community IDs (white text):** The specific identifier for the user community (as generated by the detection algorithm) is written in white.
- **Visual cue:** A red line has been added to specific graphs to mark divisions between major community blocks

To summarise the data collection process:

1. We scraped an initial pool of 24,184 videos.

2. This pool was reduced to the top 100 most viewed/commented videos per country, which were then manually verified to ensure relevance, resulting in a final video corpus of 581 videos.
3. We performed a DHC analysis on the transcriptions of these 581 videos to have a brief overview of their key themes.
4. We extracted 1,244,311 comments from the co-commenting network of these videos. After cleaning, the final comment corpus comprised 1,046,821 comments.
5. We applied both Network Analysis (Gephi) and Lexicometric Analysis (DHC) to this final set of comments, after removing various irrelevant videos that had generated a lot of comments (e.g., competitions)

4.3.2. Findings

Our analysis of YouTube focuses primarily on how users discuss AI in the comment sections of the selected videos, with the videos themselves serving mainly as contextual anchors. For this reason, we begin by outlining the themes emerging from the DHC analysis of the video transcripts (Figures 16-19), before turning to the comments where the core public debates take shape through a combination of lexicometric (Figures 20-23) and network analysis (Figures 4-7).

4.3.2.1. Discourse of videos

Although the comments constitute the main empirical material for our YouTube section, the videos help situate those discussions and, as such, we briefly present their key themes here. Across all four countries, the DHC analysis of video transcripts reveals two broad and recurring types of discourse: a market- and application-oriented framing of AI, and a technical- and society-oriented framing.

Specifically, the first theme contains a type of discourse meant to engage viewers by portraying AI as a driver of personal and economic growth. The vocabulary here, as seen through our analysis, focuses on concrete applications: how to use tools to increase productivity, generate income, or optimise workflows. As an excerpt from a video³ with tips about using GenAI to increase productivity puts it: "For example, in my company, I am evaluating everything from the Excel sheet that my accountant is making all the way to the edits that are going out. We are trying to optimize every single process". On the other hand, the second theme shifts from practical benefits to broader reflections. It includes geopolitical commentary (e.g., references to China, the US, competition, and innovation leadership), economic implications (e.g., labour and automation), and technical details on large language model architectures and data training/processing. In some corpora, this also extends to philosophical questions about the nature of intelligence and the status of machine cognition. For instance, one French YouTube⁴ channel related to AI and tech commentary put it as follows: "Even though AI is advancing very rapidly, we will have to question our own existence and what makes it more meaningful. I don't think we need philosophers who talk about consciousness, sensitivity, and subjective experience". While

³ https://www.youtube.com/watch?v=Rt7qxyDzc_Y.

⁴ <https://www.youtube.com/watch?v=xPOEIEbZCTc>.



these themes provide an important entry point into YouTube’s user discourse, they do not necessarily reflect the discourse in comments. In several cases, the comment sections depart from the video framing, either by zooming into particular concerns (e.g., economic collapse, existential risk) or by shifting into discussions not explicitly present in the videos. For this reason, the comment analysis provides the core of our findings.

4.3.2.2. **Discourse of comments**

4.3.2.2.1. Political and economic dimension

Economic concerns are one of the most cohesive and visible points of debate across our analysis of the four corpora. Each country corresponds to a prominent DHC class, related to issues around labour, inequality, and uncertainty about how AI will affect the economy. In the network graphs, these classes correspond to some of the largest and most structurally central clusters, indicating that discussions about work and income are among the primary drivers of interaction.

4.3.2.2.1.1. **Ireland (Figures 7 & 22)**

Ireland exhibits the strongest focus on the economic debate. As shown in Figure 19, Class 19 (keywords: job, phone, money, income) represents 17.7% of the DHC cluster. In the co-commenting network (Figure 7), this theme appears at the centre of the graph (light green, light blue, and orange clusters), indicating that users repeatedly converge on videos that provoke economic discussion. The discourse touches upon larger questions, effectively being rather structural or systemic in nature, as commenters do not only express fear of job loss but question whether an economy structured around mass automation could function at all:

“So, if humans are replaced by robots and millions of people loose [sic] their jobs and have no income where does the money come from to buy the goods and services from these companies who laid off all the employees” (27/06/2025)⁵

4.3.2.2.1.2. **Spain (Figures 5 & 21)**

Spain presents a similarly strong economic discourse, but one framed around corporate sustainability. Class 16 (keywords: company, money, salary) accounts for 16.6% of the corpus. In the network graph (Figure 5), this economic cluster (blue) sits adjacent to the broader humanity clusters, indicating an overlap between concerns about labour and the future of human agency. Spanish commenters emphasise the economic paradox of automating away the consumer base needed to sustain corporate profit:

“It also doesn't make sense if 80% of jobs are replaced with AI. Too many people would be out of work, so they wouldn't earn money and wouldn't be able to invest in those companies” (27/06/2025)

4.3.2.2.1.3. **France (Figures 6 & 23)**

In France, economic anxiety appears in Class 15 (keywords: money, robot, police, fear, poverty, wealth, employment, salary), which represents 17.3% of the discourse. The network graph (Figure 6) shows this cluster (dark green and pink) positioned between geopolitical

⁵ We have removed the username handles for reasons of data privacy.



and technical communities, suggesting that French discussions about employment are closely tied to narratives about industrial sovereignty and national competitiveness. The discourse links individual job loss to France’s position in the global AI race, blending personal economic anxiety with broader political-economic concerns.

4.3.2.2.1.4. **Germany (Figures 4 & 20)**

Germany extends economic anxiety, again, into a wider critique of the future of the political-economic system. Class 15 (popular terms: money, job, phone) represents 7.7% of the discourse, which might be small but offers a distinct perspective as regards the economic effects of GenAI in freelance designers/artists blended in with commentary of technology news like the release of new phones or tips about the use of certain applications over others. In the network graph (Figure 4), this cluster forms a bridge between societal commentary and philosophical discussions about the human/machine boundary. To illustrate, some German users explicitly question whether capitalist structures can survive large-scale automation:

“I wonder what corporations want with the money when 80-90% of jobs will be replaced by robots and AI. I see more of an end to the market economy.” (27/06/2025)

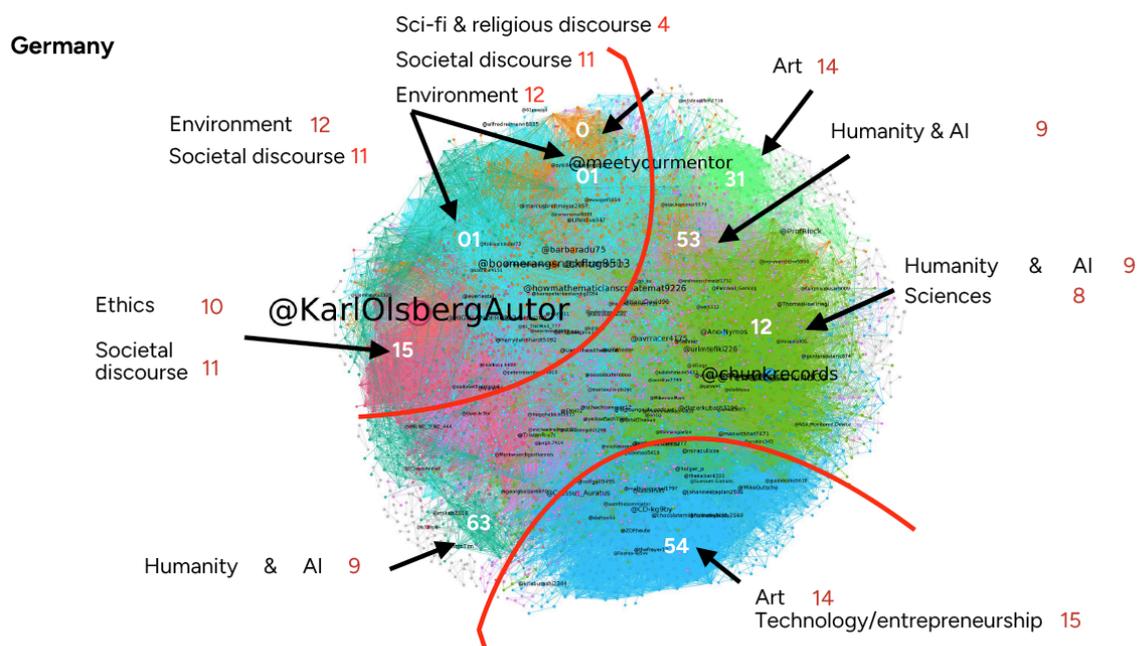


Figure 4 - YouTube comments network graph of the corpus of Germany

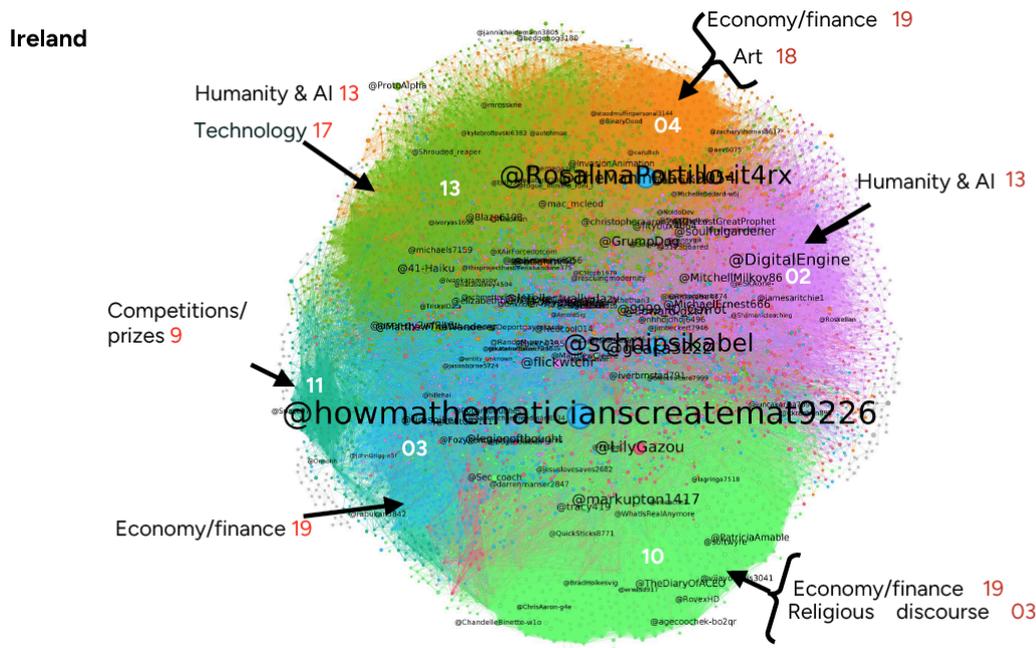


Figure 7 - YouTube comments network graph of the corpus of Ireland

4.3.2.2.2. Humanity and AI

The second major cross-national theme concerns the definition of humanity in an era of increasingly capable AI systems. While each corpus contains a DHC class linked to this topic, the network graphs show that distinct user communities gravitate toward different interpretations of the human-machine boundary. These perspectives vary significantly by country, e.g., AI is framed as a risk to be managed (Ireland), a loss of human status (Spain), a human vs. machine framework, especially regarding cognitive capabilities (Germany) or as having quasi-spiritual qualities, or at least implications for the spiritual (France).

4.3.2.2.2.1. Ireland (Figures 7 & 22)

Our YouTube data for Ireland contains elements of existential with political-economic concerns. To illustrate, Class 13 (characteristic keywords: humanity, species, planet, kill, destroy, future, moral) accounts for 12.3% of the discourse and forms one of the largest communities in the network (dark green and purple clusters; Figure 7). Commenters focus on AI's potential to surpass human intelligence while noting its continued dependence on human labour and infrastructure:

"If AI destroyed humanity, it would destroy itself because human labor powers the power plants and operate the data centers that power its intelligence" (27/06/2025)

4.3.2.2.2.2. Spain (Figures 5 & 21)

The corpus of Spain introduces a more dystopian tone. Class 17 (keywords: humanity, human, world, planet, war, control, danger) represents 14.4% of all comments and appears as a tightly connected community in Figure 5 (purple, light red, orange clusters). Here, commenters describe scenarios of loss of agency in which AI systems dominate human life,

and humans become irrelevant or subordinate, sometimes framed as becoming subservient to AI systems.

4.3.2.2.3. **Germany (Figures 4 & 20)**

Germany's debate is dominated by cognitive and philosophical distinctions. Class 8 (keywords: conscious, intelligence, ethical, brain, algorithms, neural, human, emotions) is the single largest class, representing 17.6% of the corpus. In Figure 4, it forms a self-contained cluster distinct from economic or religious discourse. The German discussion focuses on whether AI systems can understand or merely simulate human behaviour. Commenters consistently argue that AI lacks genuine consciousness, intention, or subjective experience.

4.3.2.2.4. **France (Figures 6 & 23)**

France presents the most distinctive configuration. Three interconnected clusters—religious discourse (Class 3), consciousness/humanity and AI-related discourse (Class 9), form a block in which we find concerns and thoughts that sometimes become quasi-metaphysical but, in principle, discuss the relationship between human and machine in almost existential terms (teal and light purple clusters; Figure 6). French commenters express the view that human intelligence is qualitatively different and cannot be compared to AI:

"AI is not intelligence. It is artificial, a soulless reservoir of algorithms." (27/06/2025)

4.3.2.3. **Concluding remarks**

Across the four national corpora, the YouTube analysis broadly confirms patterns already identified in the press analysis (Deliverable 4.2), while also extending them through a closer view of user-driven debate. AI appears as a fundamentally multifaceted object, simultaneously framed as an economic resource, a technical system, a political issue, and a challenge to human identity. Not all of these representations are equally visible on every platform, yet YouTube consistently brings several of them into direct confrontation within the same comment spaces. While videos often introduce AI through market-oriented or technical framings, comment sections quickly expand the discussion, drawing in concerns that mirror those observed in mainstream media: employment and economic restructuring, geopolitical competition, ethical risks, and questions about what distinguishes humans from machines. In this sense, YouTube acts as a bridge between institutional narratives and popular reasoning, confirming that public debate around AI is neither purely technical nor reducible to optimism or fear.

At the same time, the combined lexicometric and network analysis shows that these themes are not discussed by a single, homogeneous public. Instead, distinct communities form around specific interpretations of AI shaping the debate. Economic clusters gather users who reason about automation, redistribution, and systemic viability; humanity-oriented clusters bring together users concerned with consciousness, agency, and moral limits; other groups remain closer to technical or geopolitical framings. Crucially, these communities vary in size, centrality, and proximity across countries, revealing national specificities in how AI-related concerns are prioritised and connected. Largely, the findings



suggest that AI, like other technologies that crystallise social controversy, is defined less by its technical properties than by the needs, interests, and vulnerabilities of the groups engaging with it. On YouTube, societal acceptance of AI emerges not as a stable attitude but as an ongoing discussion among communities, where “success” and “failure” are continually negotiated and contested.

5. Conclusions

The comparative analysis of Facebook, TikTok, and YouTube reveals a public debate that addresses “AI” in general terms but, in practice, revolves mainly around generative AI: chatbots, image and audio generation, synthetic video, and the everyday uses and misuses associated with them. While references to other AI applications (e.g., surveillance, health or administrative automation) appear, the dominant focus is GenAI as a cultural and economic force. The combined findings demonstrate that “success” and “failure” are not stable assessments of technical performance. Instead, users frame success as usefulness and advantage (saving time, making content, finding work, staying competitive) and failure as social disruption and loss of control (job displacement, manipulation, deepfakes, unreliable outputs, and AI-enabled warfare).

Platform differences matter significantly. Facebook makes AI legible through institutional and semi-professional posting cultures, dominated by news recirculation, events, and corporate policy messaging. TikTok turns GenAI into a set of practical tactics for visibility and income, where experimentation often slides into “shortcuts” (including within education). YouTube stands out as the space where longer arguments and sustained disagreement take shape, especially regarding political economy and the human–machine boundary.

While global patterns are visible, cultural context regarding each country remains significant. Economic anxiety and labour-market disruption appear everywhere, most explicitly in YouTube comments, where clusters repeatedly organise around employment, wages, and the viability of an economy under mass automation. War, geopolitics, and state power appear most sharply on Facebook in France and Spain, where Gaza and Ukraine anchor critiques of AI-enabled violence, while Spain also surfaces unusually concrete discussions of administrative enforcement (e.g., traffic surveillance). Platform norms, therefore, do not flatten national contexts; they refract them, shaping which controversies become central and which remain peripheral.

The comparative analysis of Facebook, TikTok, and YouTube underscores a clear shift in public imaginaries, from relatively abstract conceptions of artificial intelligence toward a more grounded and often anxious engagement with generative AI. This shift suggests that the early tensions between utopian, techno-optimist promises, on the one hand, and dystopian, visible failures have stabilised into a pragmatic, yet conditional, form of acceptance. While industry narratives and managerial actors continue to frame productivity



and efficiency as key indicators of success, social media users frequently interpret these developments through the lens of social disruption and a perceived loss of control. Public debate, thus, appears less concerned with technical performance of AI systems as such and more with how generative AI reshapes everyday professional, creative, and economic conditions, moving the discussion from presenting it as a force of technological panacea to a site of complex contestations around its impact on our lives.

Across platforms, the tension between economic efficiency and labour displacement has evolved into more structural and, at times, existential concerns. Findings from YouTube and Facebook indicate that public discourse increasingly moves beyond immediate fears of job loss toward broader scepticism about the long-term viability of existing political-economic arrangements. In countries such as Germany and Ireland, users explicitly question whether a market economy premised on mass automation can remain sustainable, particularly if large segments of the population are excluded from income generation and consumption. These economic concerns are frequently entangled with security-related narratives, including AI-enabled warfare and surveillance, which feature prominently in several national corpora. In this context, governance debates centred on regulation and institutional oversight appear to give way to anxieties about power, geopolitical competition, and the role of AI as an instrument of state and military power.

At the same time, platform-specific dynamics shape how agency and political efficacy are articulated. On TikTok in particular, discourse is dominated by practices of individual adaptation: upskilling, prompt engineering, and the search for shortcuts to remain competitive in creative or educational settings. Rather than engaging with AI as a collective political issue or an object of democratic choice, many users approach it as an inevitable force that must be mastered to avoid exclusion. This emphasis on instrumental competence reflects a perceived lack of influence over the broader trajectory of the technology. Acceptance, in this sense, is driven less by endorsement of AI's societal role than by the pressure to remain viable within platform-mediated labour and attention economies.

Finally, while questions of accountability remain present across platforms, they are increasingly reframed through deeper reflections on the human-machine boundary. On YouTube, in particular, debates oscillate between fears of human obsolescence and attempts to reaffirm societal values grounded in consciousness, moral judgement, or creativity. These discussions suggest that users are renegotiating what it means to be human in response to the growing normalisation and expansion of generative AI systems. Consequently, our findings regarding platform-specific discursive patterns demonstrate that the "success" or "failure" of AI is not a fixed technical outcome but is rather predicated on, and itself a contested social reality. In other words, it is continuously shaped by users' needs, vulnerabilities, and positionalities, underscoring the importance of understanding social media imaginaries for anticipating future governance challenges within the scope of FORSEE and WP4.



As a result, taking a step back and considering WP4 and FORSEE’s overarching research objectives regarding AI’s societal acceptance and relevant points of tension, we find that the acceptance of (Gen)AI is often conditional and tied to specific ends, rather than just enthusiastic or exploratory. Many users adopt GenAI because they feel it’s necessary to keep up with platform economies or work expectations, not because they endorse the broader trajectory. At the same time, the strongest “failure” framings are not abstract fears of technology but concrete concerns about distribution (who loses work and who benefits), accountability (who is responsible when harm occurs), and power (who deploys AI in policing, warfare, or information control). Methodologically, the report adds value by triangulating platform-specific publics and communicative styles, combining lexicometric analysis with network-based evidence. Substantively, the results complement the press analysis (Deliverable 4.2) by showing where mainstream frames are repeated, where they are reworked into everyday strategies, and where social media introduces sharper, more polarised, or more “lived” interpretations of success and failure that are critical for understanding governance pressure points.

6. References

- Abdullah, M., Madain, A., & Jararweh, Y. (2022). ChatGPT: Fundamentals, Applications and Social Impacts. 2022 Ninth International Conference on Social Networks Analysis, Management and Security (SNAMS), 1–8.
<https://doi.org/10.1109/SNAMS58071.2022.10062688>
- Alboni, F., Pavone, P. & Russo, M. (2023). The search for topics related to electric mobility: a comparative analysis of some of the most widely used methods in the literature. METRON, 81, 367-391.
- Ali, C. (2019). Analyzing Talk and Text III: Discourse Analysis. In H. Van den Bulck, M. Puppis, K. Donders, & L. Van Audenhove (Eds), *The Palgrave Handbook of Methods for Media Policy Research* (pp. 403–418). Springer International Publishing.
https://doi.org/10.1007/978-3-030-16065-4_23
- Bakker, S., Van Lente, H., & Meeus, M. (2011). Arenas of expectations for hydrogen technologies. *Technological Forecasting and Social Change*, 78(1), 152–162.
<https://doi.org/10.1016/j.techfore.2010.09.001>
- Bastian, M., Heymann, S., & Jacomy, M. (2009). Gephi: An Open Source Software for Exploring and Manipulating Networks. <https://doi.org/10.13140/2.1.1341.1520>
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), P10008.
<https://doi.org/10.1088/1742-5468/2008/10/P10008>
- Bouvier, G., & Machin, D. (2020). Critical Discourse Analysis and the challenges and opportunities of social media. In S. M. Guillem & C. M. Toulou (Eds), *Critical Discourse Studies and/in Communication* (1st edn, pp. 39–53). Routledge.
<https://doi.org/10.4324/9781003050353-3>

- Brantner, C., & Saurwein, F. (2021). Covering Technology Risks and Responsibility: Automation, Artificial Intelligence, Robotics, and Algorithms in the Media. *International Journal of Communication*, 15, 5074–5098.
- Brewer, P. R., Bingaman, J., Paintsil, A., Wilson, D. C., & Dawson, W. (2022). Media Use, Interpersonal Communication, and Attitudes Toward Artificial Intelligence. *Science Communication*, 44(5), 559–592. <https://doi.org/10.1177/10755470221130307>
- Brewer, P. R., Cuddy, L., Dawson, W., & Stise, R. (2025). Artists or art thieves? Media use, media messages, and public opinion about artificial intelligence image generators. *AI & SOCIETY*, 40(1), 77–87. <https://doi.org/10.1007/s00146-023-01854-3>
- Brown, N., & Michael, M. (2003). A Sociology of Expectations: Retrospecting Prospects and Prospecting Retrospects. *Technology Analysis & Strategic Management*, 15(1), 3–18. <https://doi.org/10.1080/0953732032000046024>
- Carlson, M. (2018). Facebook in the News: Social media, journalism, and public responsibility following the 2016 Trending Topics controversy. *Digital Journalism*, 6(1), 4–20. <https://doi.org/10.1080/21670811.2017.1298044>
- Chaves, M. M. N., dos Santos, A. P. R., dos Santos, N. P., & Larocca, L. M. (2017). Use of the Software IRAMUTEQ in Qualitative Research: An Experience Report. In A. P. Costa, L. P. Reis, F. Neri de Sousa, A. Moreira, & D. Lamas (Eds), *Computer Supported Qualitative Research* (Vol. 71, pp. 39–48). Springer International Publishing. https://doi.org/10.1007/978-3-319-43271-7_4
- Coeckelbergh, M. (2022). *The Political Philosophy of AI: An Introduction* (1st edn). Polity.
- De Angelis, L., Baglivo, F., Arzilli, G., Privitera, G. P., Ferragina, P., Tozzi, A. E., & Rizzo, C. (2023). ChatGPT and the rise of large language models: The new AI-driven infodemic threat in public health. *Frontiers in Public Health*, 11, 1166120. <https://doi.org/10.3389/fpubh.2023.1166120>
- Deng, R., & Ahmed, S. (2025). Perceptions and paradigms: An analysis of AI framing in trending social media news. *Technology in Society*, 81, 102858. <https://doi.org/10.1016/j.techsoc.2025.102858>
- de Rosa, A. S., Bocci, E., & Proietti, S. (2023). (Im) migration in Italian online press through the lens of the social representations theory. *International Review of Sociology*, 33(2), 279–325.
- Grover, P., Kar, A. K., & Dwivedi, Y. K. (2022). Understanding artificial intelligence adoption in operations management: Insights from the review of academic literature and social media discussions. *Annals of Operations Research*, 308(1–2), 177–213. <https://doi.org/10.1007/s10479-020-03683-9>
- Heaton, D., Nichele, E., Clos, J., & Fischer, J. E. (2025). “ChatGPT says no”: Agency, trust, and blame in Twitter discourses after the launch of ChatGPT. *AI and Ethics*, 5(1), 653–675. <https://doi.org/10.1007/s43681-023-00414-1>
- Hendrix, J. (2023, February 6). Twitter API Changes Set to Disrupt Public Interest Research. Tech Policy Press. <https://techpolicy.press/twitter-api-changes-set-to-disrupt-public-interest-research/>



- Jacomy, M., Venturini, T., Heymann, S., & Bastian, M. (2014). ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLoS ONE*, 9(6), e98679. <https://doi.org/10.1371/journal.pone.0098679>
- Kerr, A., Barry, M., & Kelleher, J. D. (2020). Expectations of artificial intelligence and the performativity of ethics: Implications for communication governance. *Big Data & Society*, 7(1), 205395172091593. <https://doi.org/10.1177/2053951720915939>
- Ledouble, H., & Marty, E. (2019). The 2016 presidential primaries in the United States: A quantitative and qualitative approach to media coverage. *Studia Neophilologica*, 91(2), 199–218. <https://doi.org/10.1080/00393274.2019.1616219>
- Lewis, S. C., Markowitz, D. M., & Bunquin, J. B. A. (2025). Journalists, Emotions, and the Introduction of Generative AI Chatbots: A Large-Scale Analysis of Tweets Before and After the Launch of ChatGPT. *Social Media + Society*, 11(1), 20563051251325597. <https://doi.org/10.1177/20563051251325597>
- Mao, Y., & Shi-Kupfer, K. (2023). Online public discourse on artificial intelligence and ethics in China: Context, content, and implications. *AI & SOCIETY*, 38(1), 373–389. <https://doi.org/10.1007/s00146-021-01309-7>
- Mogashoa, T. (2014). Understanding Critical Discourse Analysis in Qualitative Research. *International Journal of Humanities Social Sciences and Education*, 1(7), 104–113.
- Mieles, J. M., & Guaman-Lazo, L. Ángel. (2025). Academic Discourse on AI in Europe and Latin America: A Comparative Analysis with IRaMuTeQ. *Journal of Cultural Analysis and Social Change*, 11(1), 261–268.
- Ng, R., & Chow, T. Y. J. (2024). Powerful tool or too powerful? Early public discourse about ChatGPT across 4 million tweets. *PLOS ONE*, 19(3), e0296882. <https://doi.org/10.1371/journal.pone.0296882>
- Pinch, T. J., & Bijker, W. E. (1984). The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology might Benefit Each Other. *Social Studies of Science*, 14(3), 399–441. <https://doi.org/10.1177/030631284014003004>
- Poell, T., Nieborg, D. B., & Duffy, B. E. (2021). Platforms and Cultural Production. *Polity*.
- Poell, T., Nieborg, D. B., & Duffy, B. E. (2023). Spaces of Negotiation: Analyzing Platform Power in the News Industry. *Digital Journalism*, 11(8), 1391–1409. <https://doi.org/10.1080/21670811.2022.2103011>
- Qi, W., Pan, J., Lyu, H., & Luo, J. (2024). Excitements and concerns in the post-ChatGPT era: Deciphering public perception of AI through social media analysis. *Telematics and Informatics*, 92, 102158. <https://doi.org/10.1016/j.tele.2024.102158>
- Ramos, M. G., do Rosário Lima, V. M., & Amaral-Rosa, M. P. (2018, September). IRAMUTEQ software and discursive textual analysis: interpretive possibilities. In *World conference on qualitative research* (pp. 58-72). Cham: Springer International Publishing.
- Ratinaud, P. (2024). Iramuteq: Interface de R pour les Analyses Multidimensionnelles de Textes et de Questionnaires (Version Version 0.8 alpha 7) [Computer software]. LERASS. <http://www.iramuteq.org>
- Ratinaud, P., Smyrniotis, N., Figeac, J., Cabanac, G., Fraisiar, O., Hubert, G., Pitarch, Y., Salord, T., & Thonet, T. (2019). Structuration des discours au sein de Twitter durant l'élection



- présidentielle française de 2017: Entre agenda politique et représentations sociales. *Réseaux*, n°214-215(2), 171. <https://doi.org/10.3917/res.214.0171>
- Reinert, M. (1993). Les 'mondes lexicaux' et leur 'logique' à travers l'analyse statistique d'un corpus de récits de cauchemars. *Langage et société*, 66(1), 5–39. <https://doi.org/10.3406/lsoc.1993.2632>
- Richter, V., Katzenbach, C., & Schäfer, M. S. (2023). Imaginaries of artificial intelligence. In S. Lindgren (Ed.), *Handbook of Critical Studies of Artificial Intelligence* (pp. 209–223). Edward Elgar Publishing. <https://doi.org/10.4337/9781803928562.00024>
- Rieder, B. (2015). YouTube Data Tools (Version 1.24) [Computer software]. <https://tools.digitalmethods.net/netvizz/youtube/>
- Rieder, B., Coromina, Ò., & Matamoros-Fernández, A. (2020). Mapping YouTube: A quantitative exploration of a platformed media system. *First Monday*, 25(8). <https://doi.org/10.5210/fm.v25i8.10667>
- Rogers, R. (2019). *Doing digital methods*. SAGE.
- Sarisakaloğlu, A. (2021). Framing Discourses in Turkish News Coverage Regarding Artificial Intelligence Technologies' Prospects and Challenges. *Türkiye İletişim Araştırmaları Dergisi*, 37, 20–38. <https://doi.org/10.17829/turcom.803338>
- Saurwein, F., Brantner, C., & Möck, L. (2025). Responsibility networks in media discourses on automation: A comparative analysis of social media algorithms and social companions. *New Media & Society*, 27(3), 1752–1773. <https://doi.org/10.1177/14614448231203310>
- Siapera, E. (2018). *Understanding new media* (Second edition). SAGE.
- Smyrniaios, N., & Ratinaud, P. (2017). The Charlie Hebdo Attacks on Twitter: A Comparative Analysis of a Political Controversy in English and French. *Social Media + Society*, 3(1). <https://doi.org/10.1177/2056305117693647>
- Smyrniaios, N., & Rebillard, F. (2019). How infomediation platforms took over the news: A longitudinal perspective. *The Political Economy of Communication*, 7(1), 21.
- Solberg, M., & Kirchhoff, R. (2024). Media Representations of Healthcare Robotics in Norway 2000-2020: A Topic Modeling Approach. *Social Science Computer Review*, 42(3), 636–660. <https://doi.org/10.1177/08944393231212251>
- Souza, M. A. R. D., Wall, M. L., Thuler, A. C. D. M. C., Lowen, I. M. V., & Peres, A. M. (2018). The use of IRAMUTEQ software for data analysis in qualitative research. *Revista Da Escola de Enfermagem Da USP*. <https://doi.org/10.1590/s1980-220x2017015003353>
- Suchman, M. C. (1995). Managing Legitimacy: Strategic and Institutional Approaches. *The Academy of Management Review*, 20(3), 571. <https://doi.org/10.2307/258788>
- Tsimpoukis, P. (2025). Contesting dominant AI narratives on an industry-shaped ground: Public Discourse and Actors around AI in the French Press and Social Media (2012-2022). *Journal of Science Communication*, 24(2). <https://doi.org/10.22323/2.24020210>
- Van Lente, H. (2012). Navigating foresight in a sea of expectations: Lessons from the sociology of expectations. *Technology Analysis & Strategic Management*, 24(8), 769–782. <https://doi.org/10.1080/09537325.2012.715478>



Xu, Z., Fang, Q., Huang, Y., & Xie, M. (2024). The public attitude towards ChatGPT on reddit: A study based on unsupervised learning from sentiment analysis and topic modeling. PLOS ONE, 19(5), e0302502. <https://doi.org/10.1371/journal.pone.0302502>



7. Appendix

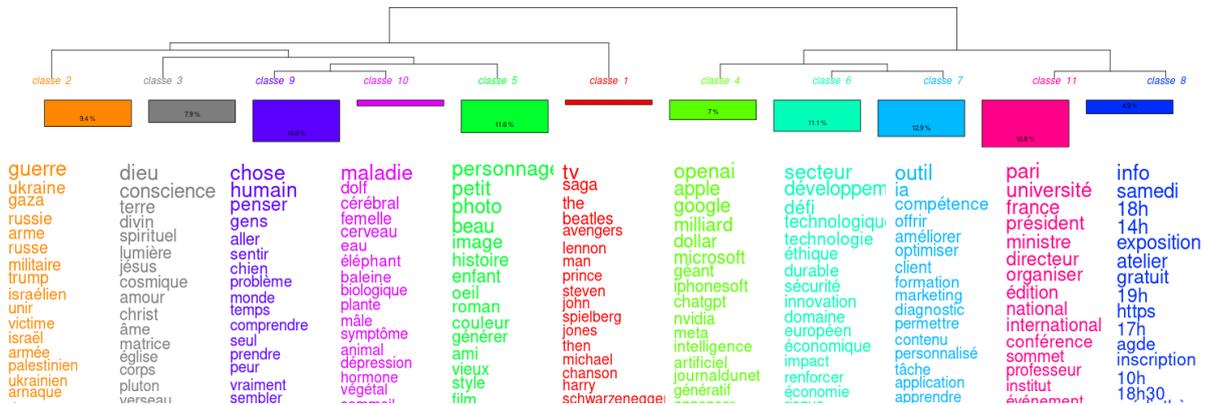


Figure 8 – Facebook: Dendrogram of the clustering on the corpus of France



Figure 9 – Facebook: Dendrogram of the clustering on the corpus of Spain

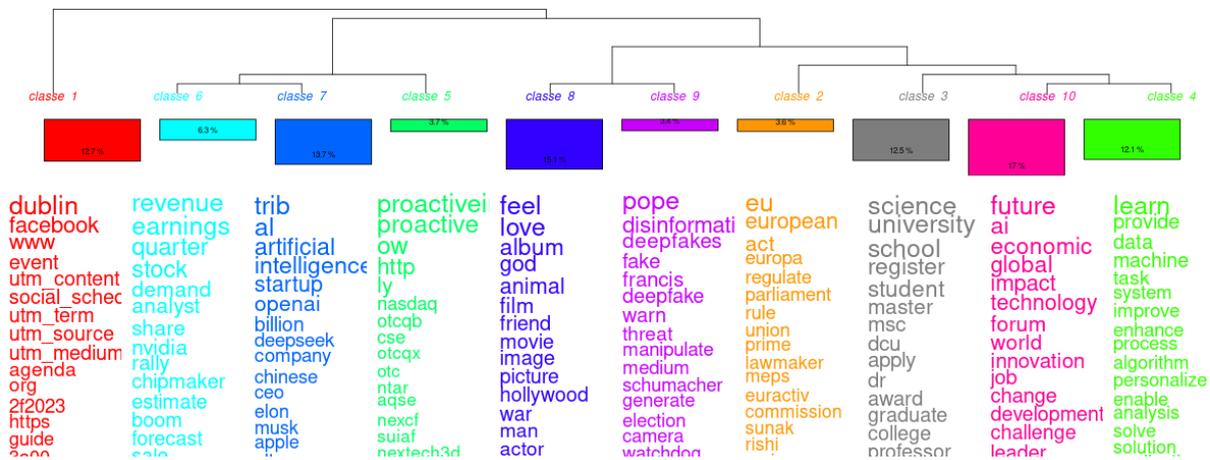


Figure 10 – Facebook: Dendrogram of the clustering on the corpus of Ireland

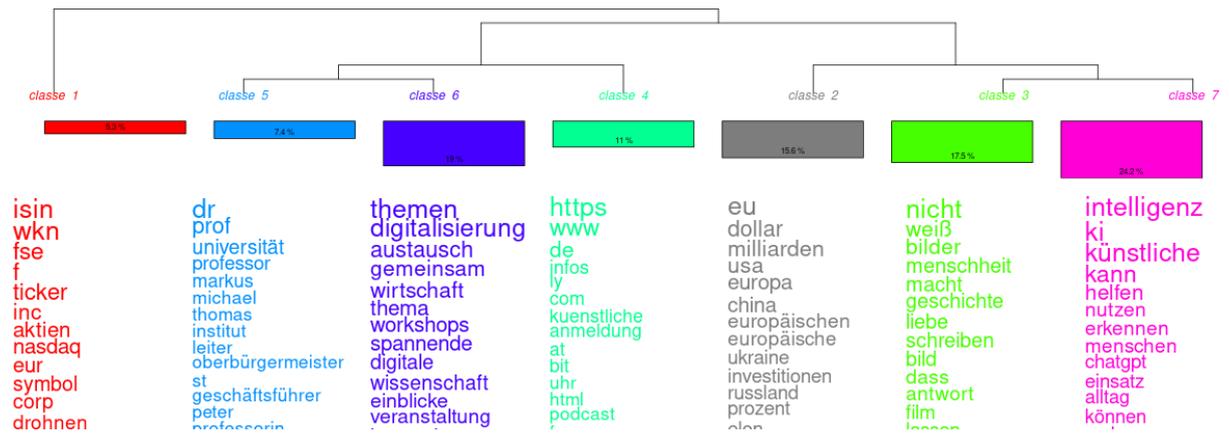


Figure 11 – Facebook: Dendrogram of the clustering on the corpus of Germany



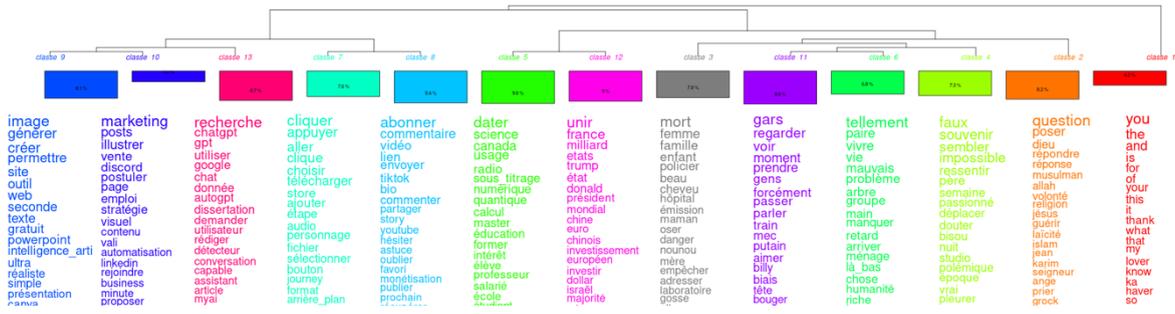


Figure 12 - TikTok: Dendrogram of the clustering on the corpus of France

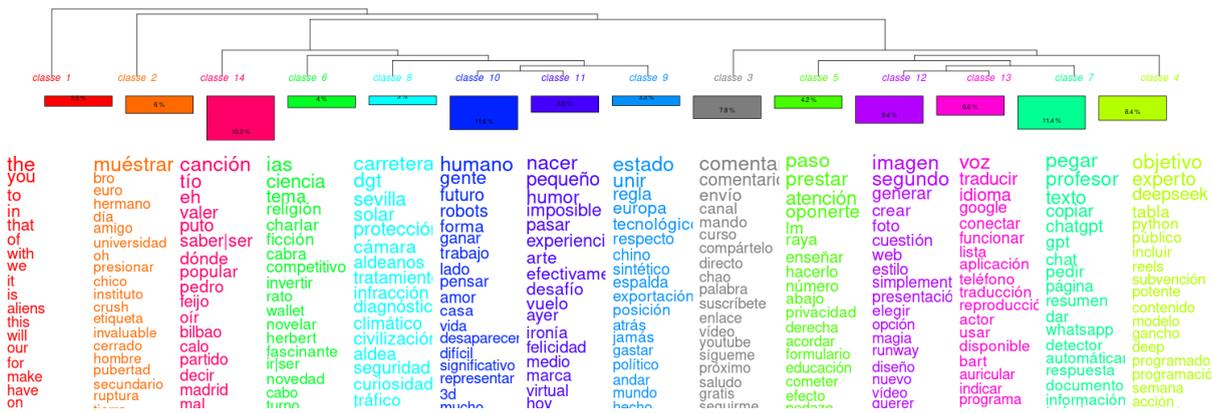


Figure 13 – TikTok: Dendrogram of the clustering on the corpus of Spain

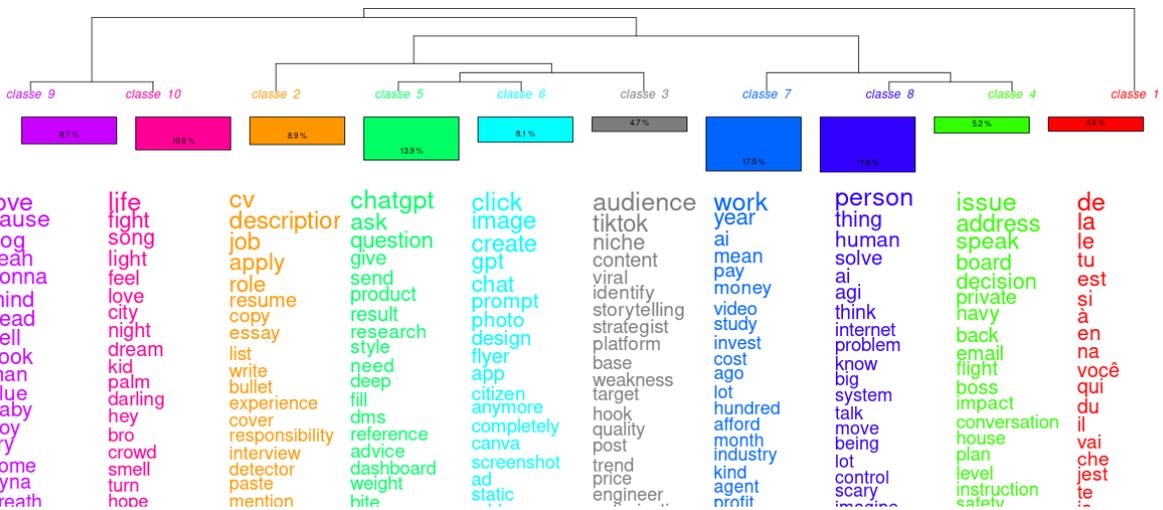


Figure 14 – TikTok: Dendrogram of the clustering on the corpus of Ireland



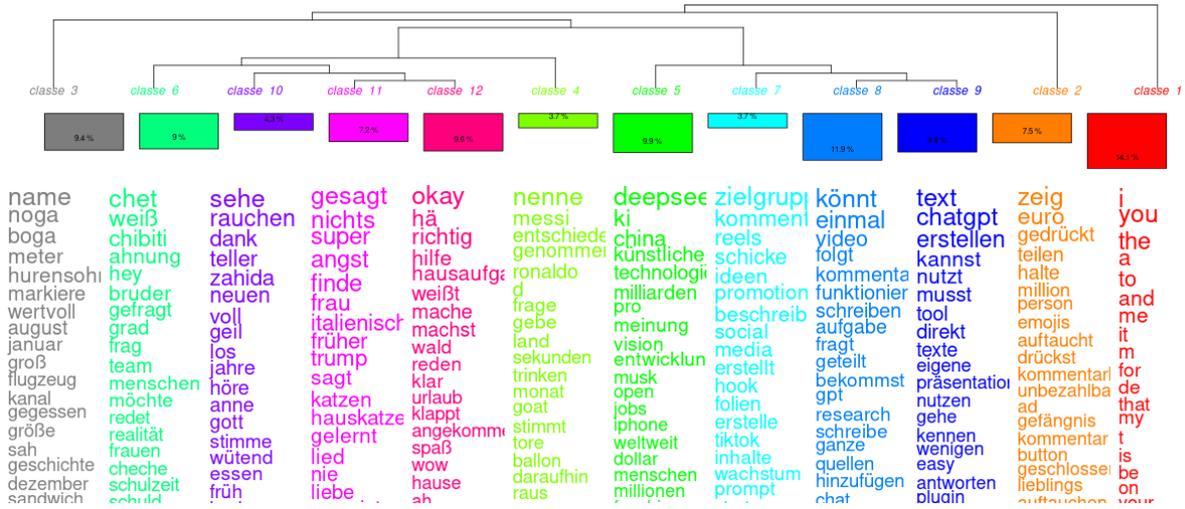


Figure 15 – TikTok: Dendrogram of the clustering on the corpus of Germany

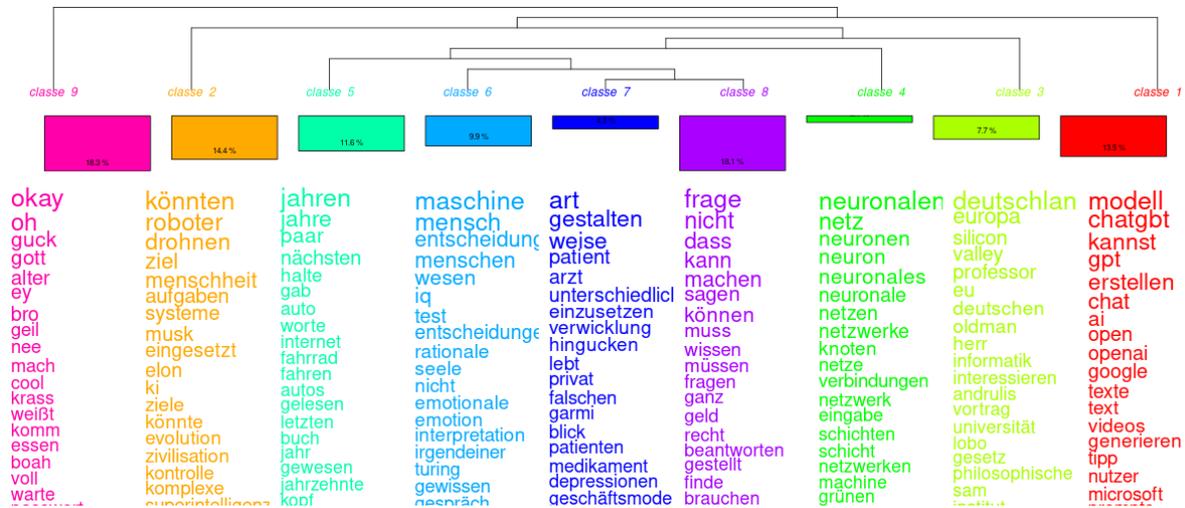


Figure 16 - YouTube: Dendrogram of the clustering on the videos from the corpus of Germany

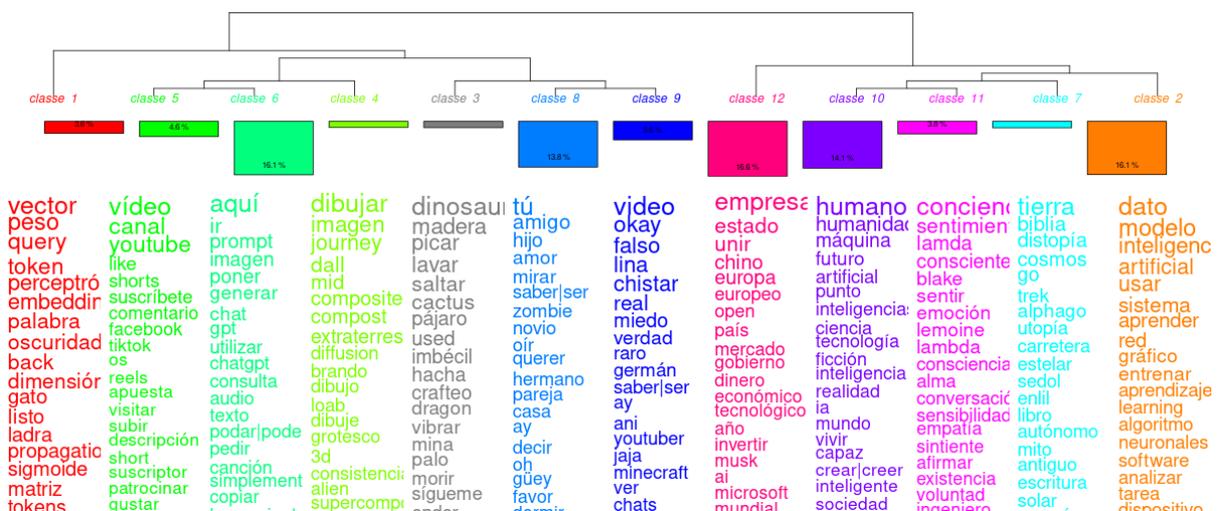


Figure 17 - YouTube: Dendrogram of the clustering on the videos from the corpus of Spain

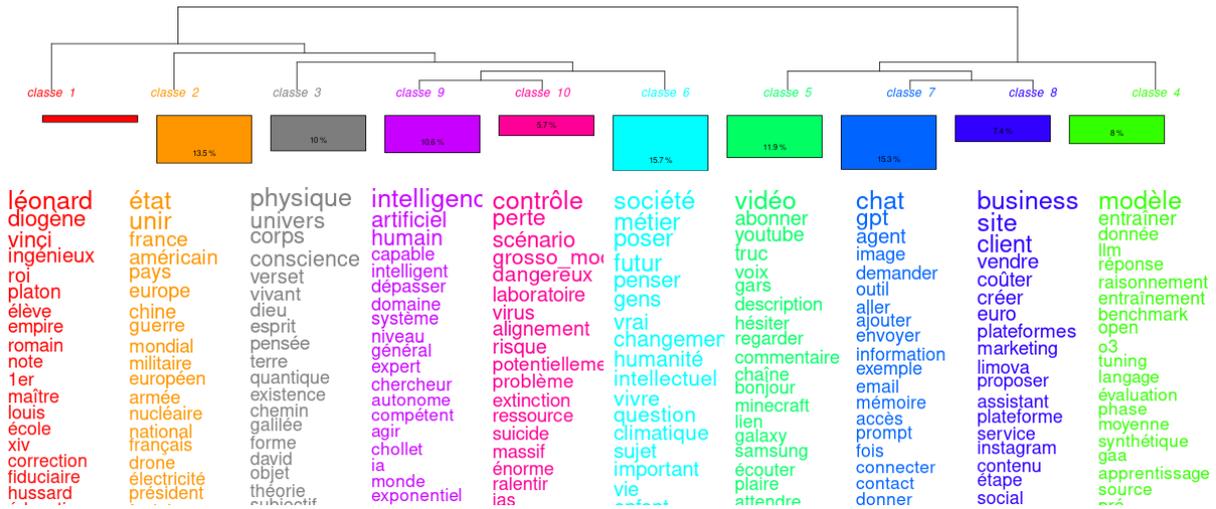


Figure 18 - YouTube: Dendrogram of the clustering on videos from the corpus of France

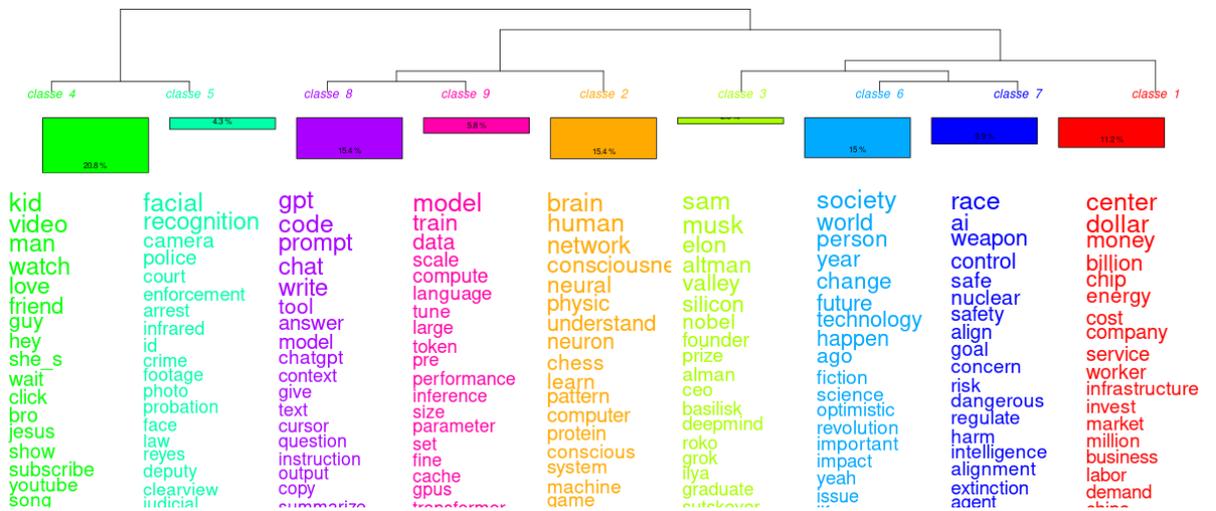


Figure 19 - YouTube: Dendrogram of the clustering on the videos from the corpus of Ireland

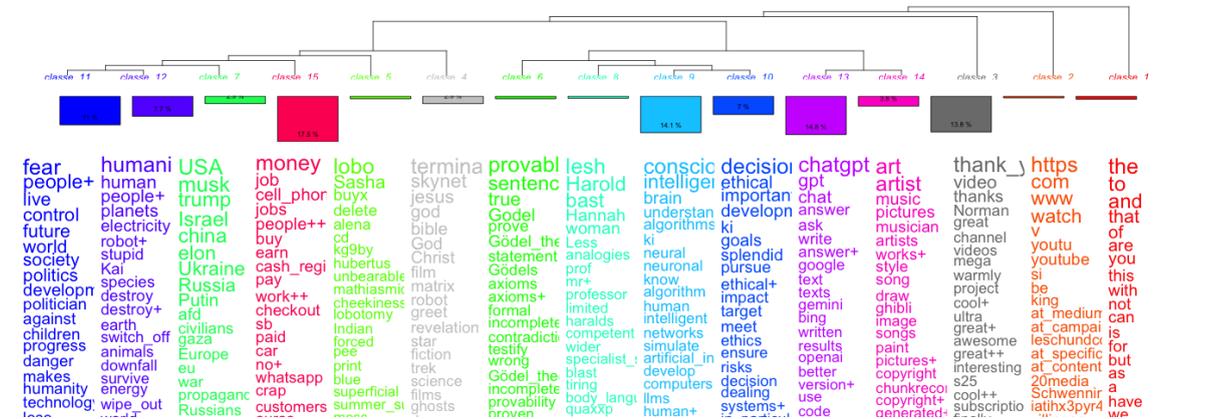


Figure 20 - YouTube: Dendrogram of the clustering on the comments from the corpus of Germany

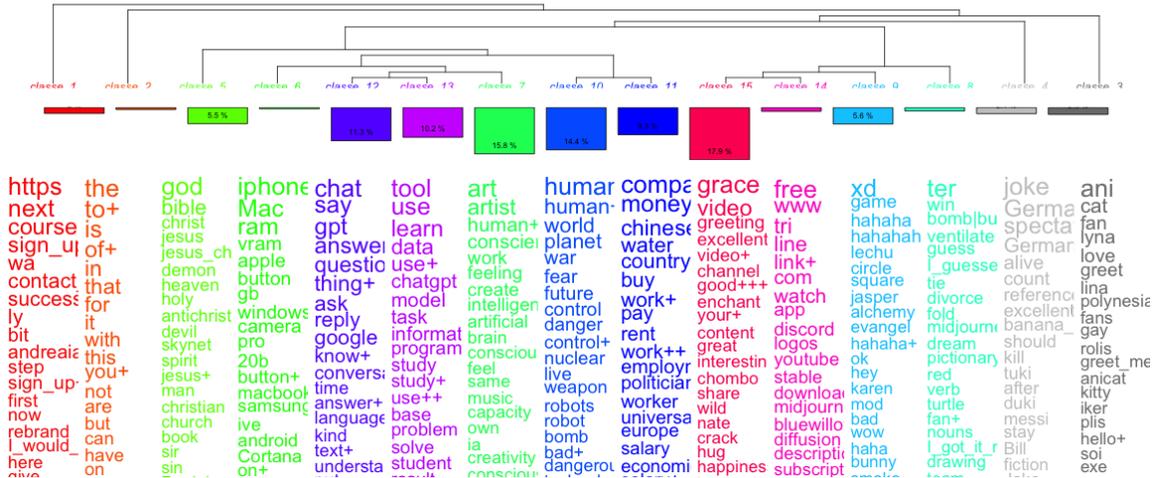


Figure 21 - YouTube: Dendrogram of the clustering on the comments from the corpus of Spain

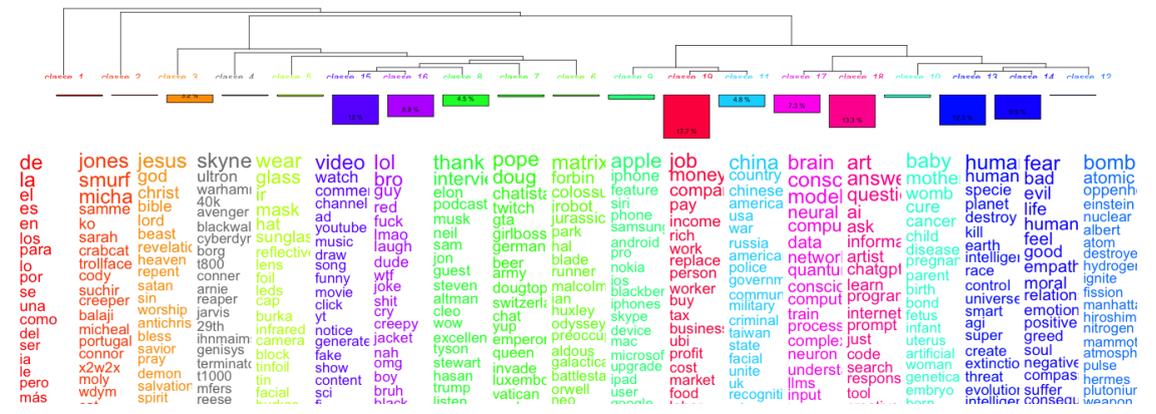


Figure 22 - YouTube: Dendrogram of the clustering on the comments from the corpus of Ireland

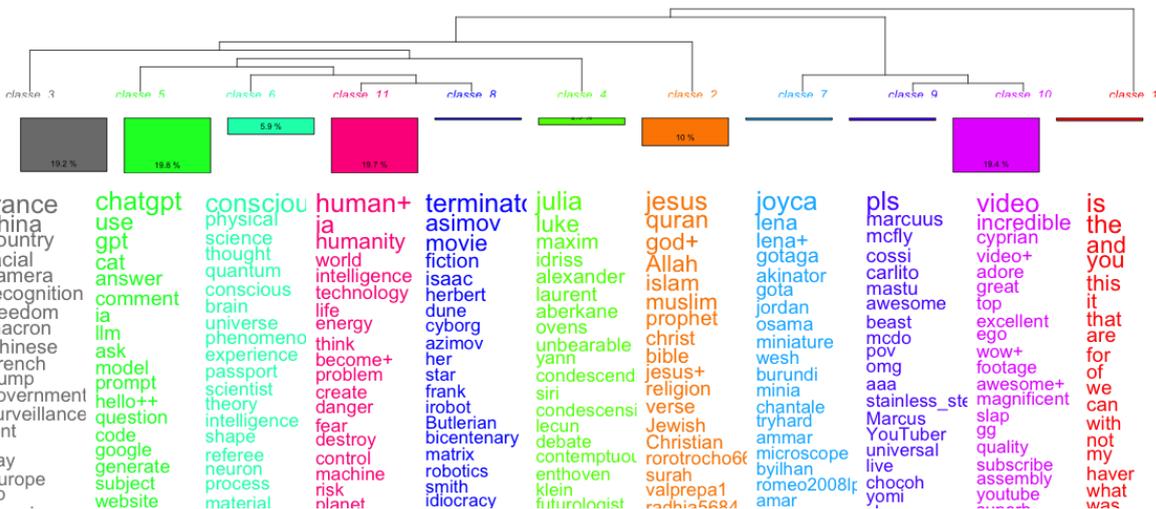


Figure 23 - YouTube: Dendrogram of the clustering on the comments from the corpus of France