https://ese-journals.unisalento.it/index.php/paco/index

ISSN: 1972-7623 (print version)
ISSN: 2035-6609 (electronic version)
PACO, Vol. 19, No. 3, 2026: 1017-1034

https://doi.org/10.1285/i20356609v19n3p1017-1034

Published 15 November, 2026

OPEN SECTION / ARTICLE

Euroscepticism on TikTok: Validating Automated Stance Inference in Reactions to the European Parliament

Tănase Tasenţe
Faculty of Law and Administrative Sciences, Ovidius University of Constanţa, Romania

Cristina-Dana Popescu
Faculty of Letters, Ovidius University of Constanţa, Romania

ABSTRACT: Public responses to the European Parliament on TikTok offer insight into institutional communication, but automated text classifications require validation. We analyse 326 posts and 27,644 public comments (27 February 2024–21 November 2025), applying multilingual sentiment and zero-shot EU-stance models. A stratified benchmark of 500 comments was independently coded by two blinded researchers, with disagreements resolved by consensus; an additional 100 comments were examined for coverage. Pre-adjudication agreement reached 85.2% for sentiment (Cohen's kappa = .757) and 92.2% for EU stance (kappa = .763). Against the adjudicated reference, sentiment accuracy was .416 (macro-F1 = .403), and EU-stance accuracy was .330 (macro-F1 = .327). Human coding identified 81.8% neutral/weak, 10.0% anti-EU, and 8.2% pro-EU comments, whereas automated predictions assigned 24.4%, 21.8%, and 53.8%, respectively. Pro-EU precision was .130 and anti-EU precision .275, indicating substantial over-identification of explicit stance. Accordingly, temporal series, the attention-weighted anti-EU index, and stance-coloured co-commenting networks are reported as model-conditioned diagnostics, not validated measures of public opinion or Euroscepticism. The findings show that human benchmarking must precede substantive political inference from automated classifications.

KEYWORDS: Automated stance detection, European Parliament, Euroscepticism, TikTok, Validation

CORRESPONDING AUTHOR: tanase.tasente@365.univ-ovidius.ro

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

1. Introduction

Questions of European Union legitimacy have long been linked to a communication deficit: citizens often encounter a multilevel system in which responsibilities are difficult to attribute and institutional communication can appear distant or technocratic (Meyer 1999). This article treats that problem as contextual rather than as an outcome measure. Comment valence on a single platform cannot establish political legitimacy, which also involves trust, representation, procedural fairness, responsiveness, and institutional performance. The empirical object here is therefore narrower: expressed support for, opposition to, or ambivalence toward the EU within comments retrieved from the European Parliament's official TikTok account.

The European Parliament (EP) occupies a distinctive position in this communication ecology. It is simultaneously a directly elected political institution and a public institutional communicator, and its communicative role differs from that of executive bodies such as the European Commission. Research on EU communication has shown both the expansion of platform-based institutional outreach and the persistence of fragmentation, top-down information practices, and multilevel coordination problems (Barberio et al. 2020; Lovari and Belluati 2023). Treating 'EU communication' as homogeneous would therefore obscure the EP's specific representative and communicative functions.

These processes are also hybrid rather than purely digital. In a hybrid media system, platform activity interacts with legacy news, electoral campaigns, institutional events, other social media, and offline political action (Chadwick 2017). Accordingly, changes observed in TikTok comment activity are not interpreted as effects of TikTok alone; the platform is one component of a broader communication environment.

Most research on digital Euroscepticism has nevertheless concentrated on parties, movements, or campaign actors. Studies conducted on Twitter (now X) in Spain, Italy, France, and the United Kingdom show how populist actors combine crisis, corruption, sovereignty, and 'people versus elites' frames to contest EU authority (Alonso-Munoz and Casero-Ripolles 2020 2021; Berti and Loner 2020). Comparative work on radical-right communication similarly documents anti-integration, anti-multicultural, and re-nationalisation themes (Kasekamp et al. 2019; Vochocova and Rosenfeldova 2023).

Research on transnational contention shows that EU issues can also generate cross-border waves of mobilisation, as in anti-TTIP activism and election-related online debate (Bouza et al. 2019; Segesten and Bossetta 2019). These studies are important for explaining organised contestation, but they do not directly describe the reactions that accumulate beneath an EU institution's own posts.

Visual and user-generated studies of the 2014 and 2019 European elections likewise show that anti-EU and populist expression can be salient within particular issue publics and campaign moments (Marchal et al. 2021; Segesten and Bossetta 2019). The predominant unit of analysis, however, remains content produced by political or organised actors rather than the distribution of reactions to routine institutional communication.

The distinction matters because comment sections are neither representative surveys nor simple engagement counters. They are self-selected, platform-curated traces of participation around a specific account. Analysing them can reveal how explicit stances become visible, how they fluctuate over time, and how stance-bearing users are co-present under the same posts, but it cannot by itself identify the attitudes of European citizens at large.

TikTok makes this problem especially relevant. Its short-form audiovisual format, recommendation system, and youth-oriented cultural repertoires have made it increasingly important for European political communication. The 2024 European Parliament elections occurred amid heightened geopolitical tension, concerns about disinformation, intensified platform campaigning, and institutional efforts to mobilise younger voters. Recent work documents both cross-platform campaigning in the 2024 elections and substantial TikTok adoption by Members of the European Parliament, while research in this journal highlights the role of youth, information

disorder, and hybrid communication in the same electoral context (Daniel and Obholzer 2025; D'Ambrosi et al. 2026; Ferro-Santos and Sepulveda 2026; Pane 2026).

Conceptually, Euroscepticism should not be reduced to a single negative label. Classic work distinguishes hard from soft Euroscepticism and differentiates diffuse support for European integration from specific support for the EU as an institutional system (Kopecky and Mudde 2002; Taggart and Szczerbiak 2002). Criticism of a particular policy, distrust of elites, nationalist claims, and rejection of integration are therefore not equivalent. In the full-corpus pipeline, 'anti-EU' and 'pro-EU' are algorithmic candidate labels assigned to comments containing explicit EU-related terms. The validation benchmark applies a stricter conceptual rule: issue-specific criticism or approval is not coded as EU-wide stance unless the text clearly generalises to the Union as a political or institutional project. The automated labels are therefore treated as signals requiring validation rather than direct measurements of Eurosceptic ideology.

A second analytical issue concerns network structure. Co-commenting can identify users who appear beneath the same posts and can show whether stance-bearing users are concentrated in similar co-presence patterns. It does not demonstrate direct communication, mutual awareness, algorithmic exposure, deliberation, or the existence or absence of echo chambers. The network analysis below is therefore interpreted strictly as a projection of shared participation environments.

Against this background, the study examines the EP's official TikTok account from 27 February 2024 to 21 November 2025, a period that includes the June 2024 European elections and the subsequent parliamentary term. The corpus contains 326 posts and 27,644 public comments retrieved during data collection. The analysis combines full-corpus model-assisted sentiment and EU-stance classification with a stratified coding benchmark, daily temporal aggregation, an attention-weighted anti-EU salience index, and a co-commenting projection. The validation layer is analytically prior to substantive interpretation: it tests whether the automated labels are sufficiently reliable to support the temporal and network claims that follow.

Three research questions guide the analysis. RQ1: What distributions of sentiment and EU-related stance are produced by the automated models in comments retrieved from the European Parliament's TikTok account? RQ2: How accurately do those automated labels reproduce a blinded human reference coding in a stratified benchmark sample? RQ3: Given the observed measurement performance, what can and cannot be inferred from model-conditioned temporal patterns and stance-centred co-commenting structures?

2. Literature Review

2.1. Euroscepticism, Populism and the European Public Sphere

Euroscepticism is conceptually heterogeneous. Taggart and Szczerbiak (2002) distinguish 'hard' opposition to the integration project from 'soft' or qualified opposition to particular policies or trajectories, while Kopecky and Mudde (2002) separate diffuse support for European integration from specific support for the EU's institutional practice. These distinctions matter for computational research because a negative statement about a policy, a criticism of an EU institution, and rejection of European integration need not represent the same political orientation.

Digital political communication can nevertheless make diverse forms of EU contestation highly visible. Populist and radical-right actors routinely link EU criticism to sovereignty, corruption, migration, cultural threat, or democratic representation (Alonso-Munoz and Casero-Ripolles 2020 2021; Kasekamp et al. 2019; Vochocova and Rosenfeldova 2023). Such findings explain important discursive repertoires, but they primarily concern organised political communication and should not be projected automatically onto ordinary comment populations.

At the level of mass opinion, survey evidence also suggests that online information environments interact with education and other socio-demographic factors in shaping support for European integration. Fortunato and Pecoraro (2022), for example, report an association between frequent reliance on online political information and support for leaving the EU, especially among less educated respondents. Because these results derive from survey designs, they offer a useful contrast to digital trace data: observed comments reveal expression within a platform setting, not the prevalence of attitudes in a population.

That distinction is central to the present study. A TikTok comment can register dissatisfaction, irony, policy criticism, or opposition to the EU without allowing the analyst to infer a stable ideology, citizenship status, or broader institutional trust. The analytical claims are therefore confined to explicit textual stance as operationally coded within the retrieved corpus.

Information disorder adds a further layer. The weakening of conventional gatekeeping and the circulation of partisan, conspiratorial, and emotionally charged content can complicate how institutional claims are encountered and evaluated (Ksenzenko 2022). In the 2024 European-election context, research on communicative resilience similarly emphasises the interaction between disinformation, youth participation, and hybrid media practices (D'Ambrosi et al. 2026; Pane 2026).

Accordingly, the article uses 'Euroscepticism' as the broader theoretical field but reserves 'anti-EU stance' for the model-derived empirical category. The distinction prevents an automated label from being treated as a direct measurement of the full conceptual range of Euroscepticism.

2.2. Public Institutional Communication, the European Parliament and Platformised Campaigning

EU institutions' social media activity belongs not only to political communication but also to public institutional communication. Institutional communicators must explain policies, make decision processes visible, respond to information disorder, and maintain relationships with heterogeneous publics. Research has documented both strategic use of social platforms and recurrent limits associated with top-down communication, institutional fragmentation, and the multilevel character of EU governance (Barberio et al. 2020; Lovari and Belluati 2023).

The EP should be distinguished from other EU bodies within this ecosystem. As the Union's directly elected parliamentary institution, it combines representative politics, legislative communication, electoral mobilisation, and public-information functions. The 2024 elections sharpened this combination: cross-platform campaigning expanded, TikTok became more prominent among MEPs, and institutional communication unfolded amid geopolitical conflict, disinformation concerns, and attempts to reach younger audiences (Daniel and Obholzer 2025; Ferro-Santos and Sepulveda 2026).

This activity operates within a hybrid media system rather than an isolated platform sphere. Institutional posts can respond to parliamentary events, news agendas, campaign activity, civil-society mobilisation, and offline political developments, while platform reactions may in turn be recirculated elsewhere (Chadwick 2017). For this reason, temporal associations in TikTok comments are treated as contextual patterns, not as platform-exclusive causal effects.

The longstanding literature on the EU's communication deficit provides an additional rationale for studying these interfaces. Difficulties of visibility, attribution, and communicative distance can affect how citizens understand EU authority and responsibility (Meyer 1999). Digital outreach may reduce some forms of distance, but platform presence does not itself resolve questions of democratic legitimacy or institutional trust.

Recent work in Partecipazione e Conflitto is particularly relevant to the 2024 election context. D'Ambrosi et al. (2026) show how young activists used digital and hybrid communication to respond to disinformation, while Pane (2026) examines youth-led civil-society organisations as intermediaries within European public

institutional communication. Together, these studies reinforce the need to situate institutional platform activity within a wider ecology of participation, resilience, and information disorder.

Despite this growing literature, much empirical research still analyses messages produced by parties, candidates, movements, or institutions. Audience reactions often enter as aggregate engagement metrics rather than as a systematic object of analysis. That supply-side emphasis leaves open how explicit support, contestation, and ambivalence are distributed among users who choose to comment beneath institutional posts.

Comment sections provide one direct, but selective, feedback interface. They capture visible expressions from users who encountered and chose to engage with a post; they do not provide a representative cross-section of EU citizens, and their visibility is conditioned by platform ranking and moderation.

TikTok intensifies these methodological issues because its recommendation and interface logics differ from those of X, Facebook, or YouTube. Recent evidence from the 2024 campaign confirms that TikTok has become part of MEPs' communication repertoire, but systematic longitudinal work on reactions to the EP's institutional account remains limited (Ferro-Santos and Sepulveda 2026).

The present study therefore focuses narrowly on the observed reaction environment surrounding the EP's official account, rather than attempting to generalise about 'EU communication' as a homogeneous field or to infer the institution's underlying strategy from comment data alone.

2.3. From Categorical Stance to Measurement Validity, Dynamic Salience and Co-Presence

Most computational research on digital Euroscepticism depends on categorical coding of messages, actors, frames, or symbols. Such categorisation is useful only to the extent that the categories are measured reliably. Classification error can be amplified by downstream aggregation: a systematic tendency to over-predict explicit stance may create apparently coherent daily trends, cumulative imbalances, or stance-coloured network structures even when the underlying texts are more ambiguous.

Election and mobilisation research suggests that anti-EU and populist expression can intensify around contentious events (Marchal et al. 2021; Segesten and Bossetta 2019). A longitudinal design can therefore add information about temporal concentration, but only after the validity of the labels being aggregated has been assessed. High false-positive rates are especially consequential because they can turn model calibration into an apparent political event.

To examine that temporal dimension, the computational pipeline constructs an attention-weighted index from the share of model-classified anti-EU comments among EU-referencing comments and overall daily comment volume. Earlier versions of the workflow labelled this variable 'euroscepticism pressure'; in the present article it is interpreted more narrowly as a model-based anti-EU salience diagnostic. Its substantive status depends on the validation results and it is not a scale of Eurosceptic ideology, political legitimacy, persuasion, or electoral effect.

At the structural level, the co-commenting projection addresses a different question: how users assigned different stance labels are distributed across shared-post commenting environments. Because an edge records shared commenting on a post rather than a reply or verified exposure, network proximity is co-presence only. Moreover, if stance labels are noisy, stance-conditioned node and edge compositions inherit that measurement error and cannot be interpreted as evidence of ideological integration or segregation.

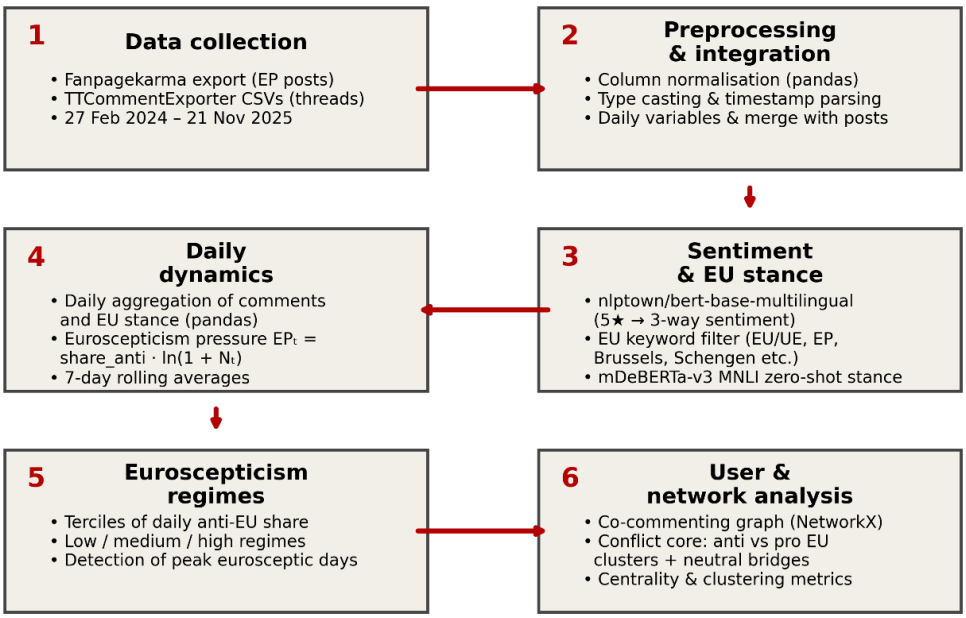
The article's contribution is therefore both substantive and methodological. It shifts attention from elite-produced content to reactions around an institutional account, but it also tests the measurement layer on which temporal and network interpretation depends. Rather than treating model outputs as self-validating, the analysis uses a blinded human-coded benchmark to determine which claims can survive measurement scrutiny and which outputs should remain exploratory diagnostics.

3. Methodology

The study employs a quantitative, computational, validation-oriented and longitudinal design centred on the European Parliament's official TikTok account between 27 February 2024 and 21 November 2025. The period covers 326 posts and 27,644 public comments retrieved for analysis and spans the June 2024 European elections as well as the subsequent post-election period. The design combines full-corpus automated classification with a reproducible stratified benchmark before interpreting daily temporal patterns or stance-conditioned user co-presence.

Figure 1 summarises the original computational workflow: (1) acquisition of publicly visible post and comment data; (2) preprocessing and integration of post- and comment-level metadata; (3) model-assisted sentiment and EU-stance inference; (4) construction of daily series and the model-based anti-EU salience index; (5) descriptive segmentation of model outputs; and (6) user-level co-commenting analysis. The present revision adds a blinded human-validation layer immediately after stage 3. The benchmark therefore determines the inferential status of stages 4-6 rather than being treated as a post hoc robustness note.

Figure 1. Overview of the analytical workflow



Source: Authors' elaboration.

3.1 Data Collection

Data were obtained from publicly visible interfaces associated with the European Parliament's official TikTok account. Post metadata were exported with Fanpagekarma, which supplied identifiers, timestamps, and engagement fields. Comments were retrieved with TTCommentExporter from the comments visible to the collection tool at the time of extraction. Each retrieved record included a comment identifier, user identifier, timestamp, engagement indicators, and comment text. Files were matched to posts through the TikTok video identifier. Because platform content can be deleted, moderated, reordered, or added after extraction, the

resulting corpus is a time-specific reconstruction of retrievable public comments rather than a claim to an immutable or independently auditable universe of every comment ever displayed.

All subsequent processing was conducted locally in Python 3.12.1.

3.2 Data Preprocessing and Integration

Preprocessing standardised metadata by harmonising column names, coercing data types, and parsing timestamps into a common datetime format. Comment files were loaded and matched to their corresponding posts using the TikTok video identifier extracted from each filename.

Comment text received light normalisation: whitespace was trimmed, empty entries were removed, and basic surface features such as word count and markers of expressive style (emojis, hashtags, and laughter tokens) were extracted. Numeric and logical fields, including likes, replies, and author-like flags, were converted to consistent formats. Comment-level records were then merged with post metadata to link each retrieved comment to its parent post, timestamp, and available engagement context.

3.3 Automated Sentiment Analysis and EU Stance Detection

Two transformer-based models from the HuggingFace ecosystem were used to generate scalable text labels. These outputs are treated as candidate measurements whose performance is evaluated against the benchmark described below, not as ground-truth annotations.

General sentiment was estimated with nlptown/bert-base-multilingual-uncased-sentiment, which produces a five-star polarity score. Scores were mapped to negative, neutral, and positive categories, and only comments containing at least two lexical items were processed to reduce noise. Inference was performed in mini-batches, retaining both the categorical label and model confidence. Because the model was not trained specifically on European political TikTok discourse, its labels are used as broad polarity indicators rather than fine-grained measurements of political affect.

The stance model assigns comments to pro-EU, anti-EU, or neutral/weak categories according to the operational definitions used in the computational pipeline.

3.4 Time-Series Construction and the Model-Based Anti-EU Salience Index

Daily aggregates include total comment volume, model-predicted sentiment shares, and model-predicted EU-stance shares, including the neutral/weak category.

For continuity with the original computational workflow, the script-generated index combines the daily model-predicted anti-EU share with total comment volume as follows:

$$EP_t = s_t\text{^-} \times ln(1 + N_t)$$

where s_t^- is the share of model-classified anti-EU comments among keyword-eligible EU-referencing comments on day t and N_t is the total number of retrieved comments. The logarithmic term dampens extreme volume. In light of the validation results, EP_t is interpreted only as a model-output diagnostic and not as a validated measure of Euroscepticism, political significance, public opinion, or institutional legitimacy.

Daily series were regularised and 7-day rolling averages were calculated for visualisation of medium-term patterns. Descriptive intensity bands were derived from the empirical distribution of valid anti-EU shares rather than from externally imposed substantive thresholds. Because daily proportions can be unstable when very few EU-referencing comments are present, extreme shares and regime assignments are interpreted jointly with their underlying counts and are not used as evidence of population-level polarisation.

3.5 Co-Commenting Network Analysis

To examine structural co-presence, a co-commenting network was constructed in NetworkX. Nodes represent distinct users, and an undirected edge links two users when both posted comments under the same EP TikTok post; edge weights record the number of posts on which they co-appeared. The relation is therefore a projected co-occurrence tie, not a reply, follow, mention, or verified exposure relation. Extremely large threads were excluded to limit mechanically generated clique density.

A filtered graph retained users above the specified degree threshold and the largest connected component. A stance-centred subgraph was then extracted containing users assigned anti-EU or pro-EU labels and users connected to them through shared-post co-commenting. Kamada-Kawai coordinates were used only for visual layout. Degree, betweenness, closeness, eigenvector centrality, and local clustering describe positions within this co-presence projection; they are not interpreted as direct interpersonal influence, message diffusion, or brokerage unless the underlying edge relation can support such claims.

3.6 Ethical Considerations

All analyses used comments that were publicly visible to the collection tools at the time of extraction. No private or restricted content was intentionally accessed, and the study did not attempt to identify, contact, or profile individual users beyond the pseudonymous identifiers required for aggregation and network construction. Identifiers were treated as technical keys and pseudonymised during processing. The analysis reports aggregate patterns and does not reproduce usernames or identifiable comment excerpts. Ethical interpretation follows principles of contextual integrity, data minimisation, and proportionality; the manuscript makes no blanket claim about ethics-review exemption, which depends on the relevant institutional and jurisdictional requirements.

3.7 Validation Design, Human Coding Protocol and Performance Metrics

Validation was conducted on a reproducible stratified sample drawn from the classification workflow. A primary pool of 4,485 comments with model outputs was sampled proportionally by model sentiment, model EU stance and empirical confidence terciles using random seed 20260907, yielding 500 benchmark comments. A separate 100-comment coverage audit was drawn proportionally from 23,159 comments that were sentiment-ineligible or lacked an ordinary stance label; those records were inspected for coverage but excluded from the performance metrics reported below.

Model labels were hidden from both coders during independent annotation. Each benchmark comment was assigned a sentiment label (negative, neutral, or positive) and an EU-stance label (anti-EU, pro-EU, or neutral/weak) according to a prespecified codebook. EU stance required explicit support for or rejection of the Union as a political or institutional project; criticism or approval of a single policy, politician, or European Parliament action was coded neutral/weak unless the text clearly generalised to the EU as a whole. Two members of the research team independently coded the full 500-comment benchmark. Percent agreement and Cohen's kappa were calculated from their independent labels before any reconciliation, separately for sentiment and EU stance.

After the independent coding stage, all disagreements were reviewed jointly by the two coders during an online consensus meeting. Each disputed case was reconsidered against the prespecified codebook and a single consensus label was assigned. The resulting consensus labels constituted the human-adjudicated reference standard used for model evaluation. Evaluation reports accuracy, balanced accuracy, class-specific precision, recall and F1, macro-F1, weighted-F1, confusion matrices, and bootstrap 95% confidence intervals for accuracy and macro-F1. Cohen's kappa therefore describes agreement between the two coders' independent pre-adjudication labels, not agreement between the model and the human reference labels.

3.8 Computational Transparency and Reproducibility Scope

Processing steps, model names, filters, aggregation rules, network definitions, benchmark strata and the fixed sampling seed are reported so that the analytical logic can be inspected. The workflow is deterministic conditional on the same raw exports, software environment, model versions and implementation code. However, the manuscript does not itself provide the raw comment corpus or an executable replication package; independent reproduction is therefore not claimed from the article alone. The appropriate claim is procedural transparency and re-runnability from identical inputs, subject to platform-access, data-sharing and ethical constraints.

4. Results

4.1. Descriptive Overview of the TikTok Corpus

The retrieved corpus contains 326 posts from the European Parliament's official TikTok account between February 2024 and November 2025 and 27,644 public comments available in the assembled files. Activity is strongly right-skewed: at the daily level, the mean is 45.9 comments (SD = 175.8), the median is 12, and the maximum is 3,144 (Table 1). These quantities describe the collected commenting environment and should not be interpreted as representative levels of engagement among EU citizens.

Table 1. Descriptive statistics for daily comment volume, EU references, sentiment, and EU stance

Variable	N Days	Mean	SD	Min	25%	Median	75%	Max
Total comments (n_comments)	593	45.93	175.78	1	5	12	31	3,144
EU-referencing comments (n_with_eu_reference)	593	7.56	30.24	0	0	2	5	497
Anti-EU comments (n_anti_EU)	593	1.54	4.80	0	0	0	1	66
Pro-EU comments (n_pro_EU)	593	4.26	20.51	0	0	1	3	346
Neutral/weak EU stance (n_neutral_EU)	593	1.76	5.75	0	0	0	1	85
Share anti-EU (share_anti_EU)	416	0.23	0.29	0.00	0.00	0.17	0.33	1.00
Share pro-EU (share_pro_EU)	416	0.54	0.35	0.00	0.28	0.50	0.82	1.00
Share neutral/weak (share_neutral_EU)	416	0.23	0.28	0.00	0.00	0.17	0.35	1.00
Model-based anti-EU salience index	416	0.75	0.83	0.00	0.00	0.58	1.22	3.93
Share of negative sentiment (share_negative)	585	0.39	0.22	0.00	0.25	0.38	0.50	1.00

Explicit EU references are unevenly distributed. Across the 593 observed days in Table 1, the mean number of EU-referencing comments is 7.56 (SD = 30.24) and the median is 2. The raw classifier outputs yield daily mean counts of 1.54 anti-EU, 4.26 pro-EU and 1.76 neutral/weak labels. These quantities describe the behaviour of the automated pipeline; Section 4.2 tests whether they can be interpreted as reliable stance measurements.

Among the 416 days with at least one EU-referencing comment, the raw model outputs have a mean anti-EU share of 0.23, a pro-EU share of 0.54 and a neutral/weak share of 0.23. At this stage these are prediction distributions, not audience distributions. The later benchmark shows that the apparent predominance of pro-EU labels is strongly affected by classification error, so the 0.54 value must not be read as evidence that pro-EU stance predominates among commenters.

General sentiment is distinct from EU stance. The raw pipeline produces a mean daily negative-sentiment share of 0.39 (SD = 0.22), with an interquartile range from 0.25 to 0.50. A negative comment may concern a policy, actor, event or communicative style rather than the EU itself, and the benchmark below additionally shows that the sentiment classifier substantially under-detects neutral comments. Sentiment is therefore neither a proxy for Euroscepticism nor a validated prevalence estimate in this corpus.

The script-generated attention-weighted anti-EU index has a mean of 0.75 (SD = 0.83) and a maximum of 3.93 among days on which it is defined. Mathematically, it identifies days on which model-predicted anti-EU share and comment volume are jointly higher. Because the validity of the anti-EU labels is evaluated in Section 4.2, the index is treated as a diagnostic of model output rather than a validated measure of political salience.

Overall, the corpus exhibits highly variable commenting volume and the automated pipeline produces substantial temporal variation in sentiment and stance labels. No substantive conclusion about the prevalence of pro-EU or anti-EU attitudes is drawn from these raw predictions before the benchmark results are considered.

4.2. Validation of Automated Sentiment and EU-Stance Labels

The 500-comment primary benchmark shows substantial agreement between the two human coders before adjudication. Sentiment agreement was 85.2% with Cohen's kappa = .757, while EU-stance agreement was 92.2% with kappa = .763. These values indicate that the prespecified codebook produced comparatively consistent human decisions before reconciliation, especially for EU stance. The agreement statistics concern the coders' independent annotations prior to the online consensus meeting; they do not measure model performance.

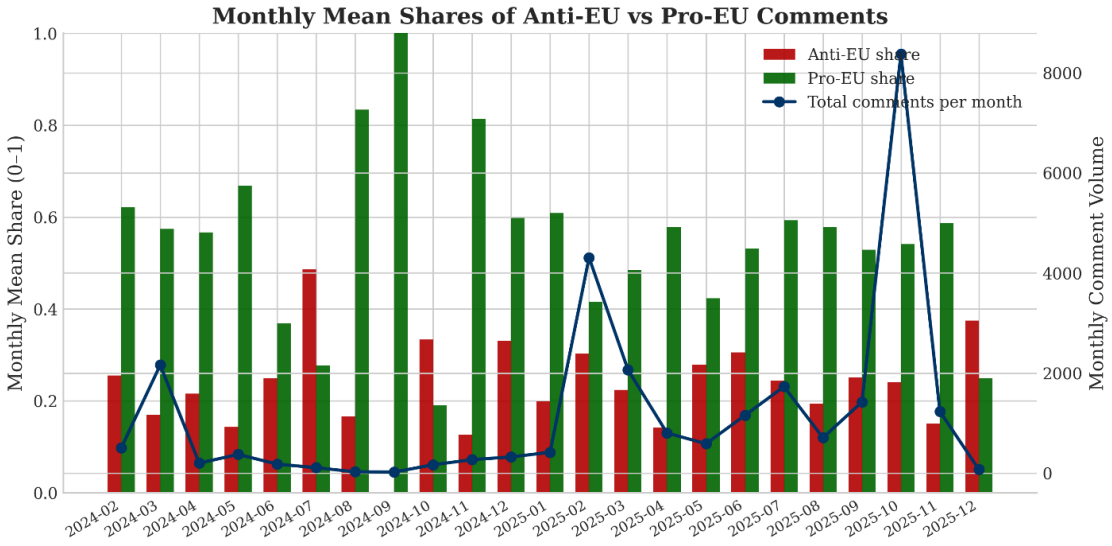
Against the adjudicated reference labels, sentiment performance was limited: accuracy = .416 (95% bootstrap CI [.374, .460]), balanced accuracy = .491, macro-F1 = .403 (95% CI [.360, .443]) and weighted-F1 = .358. Class-specific precision/recall/F1 were .392/.724/.509 for negative, .519/.118/.192 for neutral, and .424/.632/.508 for positive. The confusion matrix shows that 148 of 238 reference-neutral comments were predicted as negative and 62 as positive, indicating a strong tendency to force neutral material into valenced categories.

EU-stance performance was weaker still: accuracy = .330 (95% bootstrap CI [.290, .372]), balanced accuracy = .566, macro-F1 = .327 (95% CI [.282, .369]) and weighted-F1 = .364. For anti-EU, precision/recall/F1 were .275/.600/.377; for pro-EU, .130/.854/.226; and for neutral/weak, .820/.244/.377. Within the 500-comment benchmark, final reference coding was 81.8% neutral/weak (409 comments), 10.0% anti-EU (50) and 8.2% pro-EU (41), whereas the model predicted 24.4% neutral/weak, 21.8% anti-EU and 53.8% pro-EU. The particularly low precision of the pro-EU label shows that high recall was achieved by over-predicting explicit support. Consequently, the full-corpus stance distributions cannot be treated as validated estimates of audience orientation.

4.3. Temporal Outputs as Model-Conditioned Diagnostics

Comment volume itself remains directly observed and varies substantially across the period: monthly totals range from 23 comments in September 2024 to 8,378 in October 2025, with the highest-volume months occurring mainly after the June 2024 European elections. The stance proportions superimposed on those volume patterns, however, are model predictions whose interpretation is constrained by the validation results above. Figures 2-5 are therefore retained as diagnostics of the original computational pipeline rather than as validated estimates of changing political orientation.

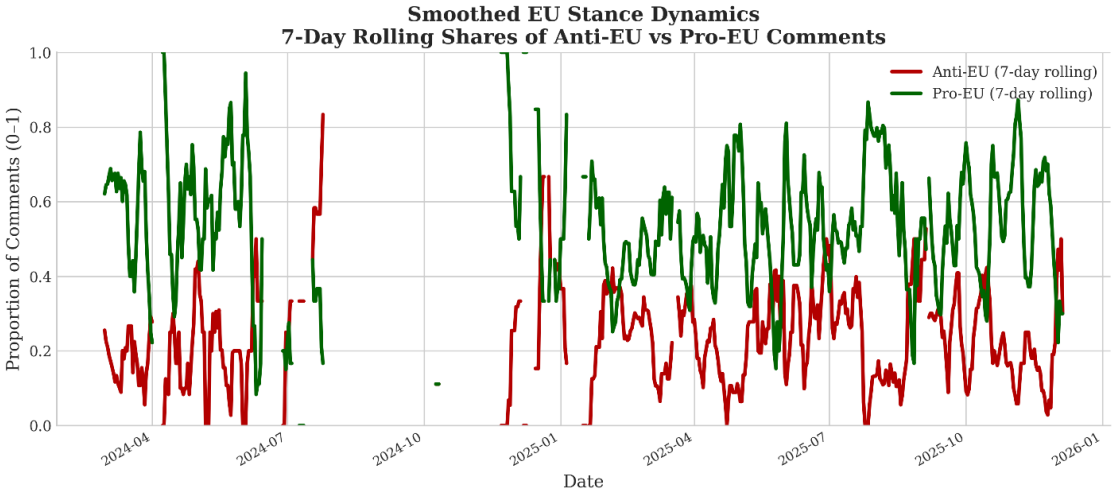
Figure 2. Monthly distribution of comment volume and stance proportions



Source: Authors' elaboration.

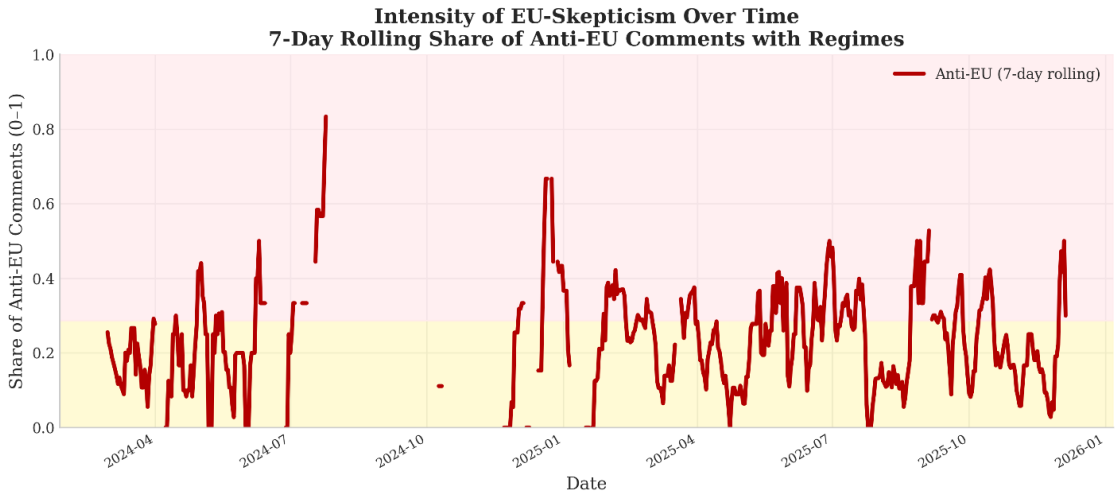
At the daily level, Figures 3 and 4 visualise the classifier's 7-day stance shares. Some days reach 100% anti-EU or 100% pro-EU among keyword-eligible comments, but the benchmark shows that these proportions combine denominator instability with substantial classification error. In particular, the model's apparent predominance of pro-EU labels cannot be treated as a substantive finding because pro-EU precision is only .130. The rolling series therefore displays temporal variation in automated assignments, not validated shifts in political orientation.

Figure 3. Seven-day rolling shares of anti-EU and pro-EU comments



Source: Authors' elaboration.

Figure 4. Model-based anti-EU salience regimes

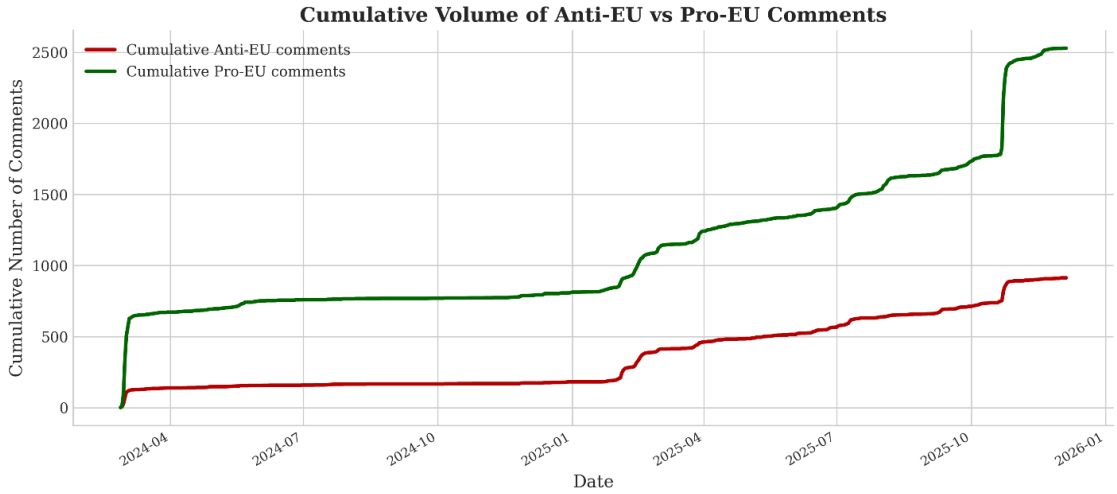


Source: Authors' elaboration.

Figure 4 should be read as a visualisation of model-assigned categories rather than evidence of stable political regimes. The anti-EU share is defined only when keyword-eligible comments are present, small denominators can generate extreme proportions, and benchmark precision for the anti-EU label is only .275. Accordingly, dates such as 30 January 2025, 22 May 2025, 7 June 2025 and 20 September 2025 are useful as candidates for targeted qualitative inspection, not as validated episodes of heightened Euroscepticism. The term 'Euroscepticism intensity' in the original figure is therefore an operational script label, not a validated construct.

Cumulative counts in Figure 5 show that the model assigns pro-EU labels more frequently than anti-EU labels across the collected period. The benchmark changes the status of this pattern: pro-EU precision is .130 and the model labels 53.8% of benchmark comments as pro-EU although only 8.2% receive that final reference label. The cumulative pro-EU dominance must therefore be interpreted as a property of classifier output, not as evidence of a supportive commenting public.

Figure 5. Cumulative anti-EU vs. pro-EU comment volume



Source: Authors' elaboration.

In sum, the temporal series remains useful for locating periods in which the automated system produces unusually high volumes or unusual combinations of stance labels, but the benchmark does not support reading those outputs as direct measurements of changing EU support or Euroscepticism. Their defensible use is diagnostic and hypothesis-generating: they can identify intervals for subsequent manual or qualitative analysis after measurement quality has been established.

4.4. Co-Commenting Structure under Measurement Uncertainty

The network analysis examines shared participation rather than interpersonal interaction, but its stance-centred filtering also depends on the automated labels evaluated above. Users assigned anti-EU or pro-EU labels are retained together with users connected to them through co-commenting on the same posts. Because stance precision is low, especially for pro-EU predictions, membership in the stance-centred subgraph is itself measurement-conditioned. The graph is therefore retained as a diagnostic projection of the pipeline rather than a validated map of ideological positions.

The resulting model-conditioned projection contains 2,227 nodes and 46,801 edges. Density is 0.0189 and mean degree is approximately 42, while the diameter is 5 and the average clustering coefficient is 0.0466. These statistics describe overlap in posting under the same EP videos within the selected subgraph. They do not establish diffusion, exposure or openness, and because node inclusion is partly determined by noisy stance labels, even the composition of this subgraph should not be compared politically without reclassification or sensitivity analysis.

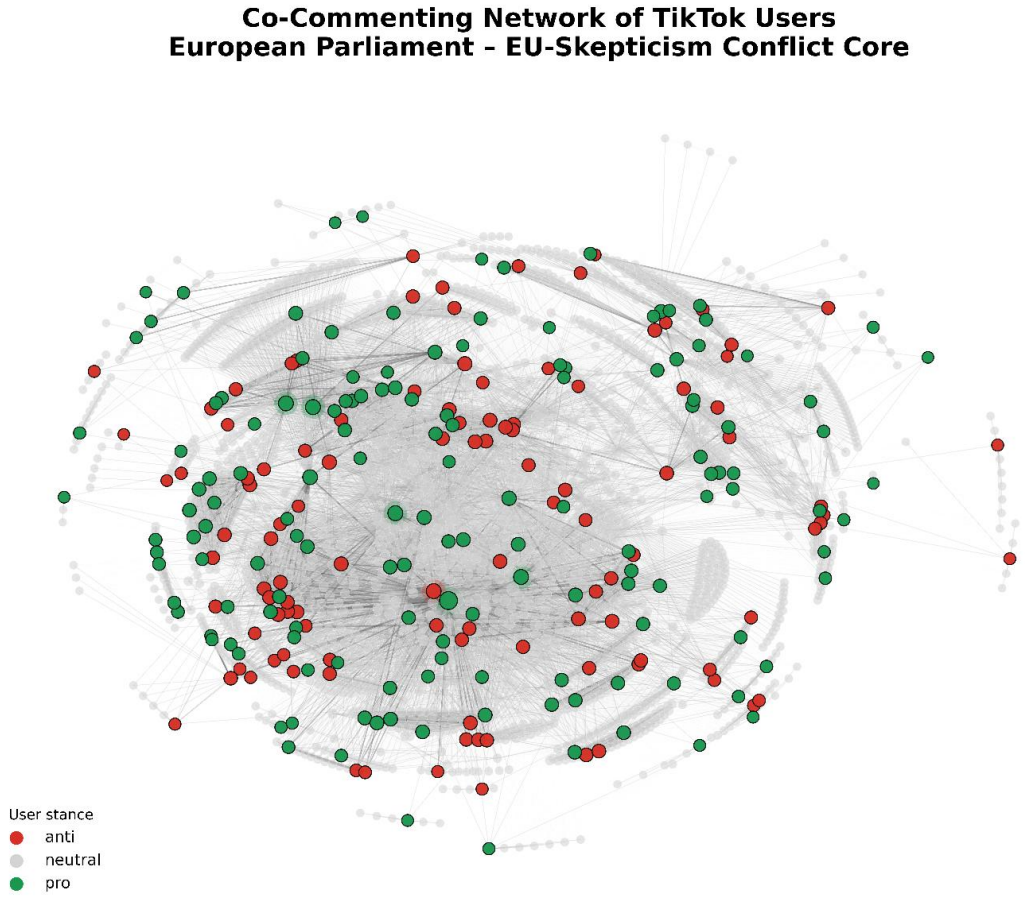
Within the stance-centred subgraph, the automated pipeline assigns 88.5% of nodes as neutral, 6.6% as pro-EU and 4.9% as anti-EU. These percentages are reported only to document how the original graph was constructed. The benchmark shows that comment-level stance assignments are insufficiently reliable for these node colours to be treated as verified political positions, and user-level aggregation does not by itself remove that error.

Edge composition in the selected graph is 81.4% neutral-neutral, 0.48% anti-pro, 0.25% anti-anti, 0.31% pro-pro, 7.44% anti-neutral and 10.08% pro-neutral under the automated user labels. Given the stance benchmark, these are model-conditioned edge categories rather than verified cross-ideological relations. They cannot support claims about same-stance cohesion, cross-stance mixing or mutual exposure without first establishing reliable user-level stance assignment.

Degree and betweenness remain descriptive of position within the selected co-presence projection. The highest-degree node co-occurs with 993 others, but the association of high-degree or high-betweenness nodes with particular stance colours is not interpreted politically because those colours inherit classification error. Even apart from measurement error, the underlying ties do not record communication, persuasion or information flow.

Taken together, the network outputs do not permit a defensible inference about ideological segregation under the observed classifier performance. Figure 6 should be read as a map of co-commenting coloured by model assignments, not as evidence that anti-EU, pro-EU and neutral users actually occupy the depicted political positions. A substantive test of segregation would require reliable user-level labels and, ideally, reply, follow or exposure data rather than shared-post co-presence alone.

Figure 6. Model-conditioned co-commenting subgraph showing automated anti-, pro-, and neutral-EU user labels



Source: Authors' elaboration.

5. Discussion

The most consequential finding is methodological rather than attitudinal: the automated classifiers do not reproduce the reference coding with sufficient accuracy to support direct political interpretation of the full-corpus stance and sentiment outputs. The apparent pro-EU predominance, anti-EU temporal spikes and stance-coloured network patterns are visually coherent model outputs, but the benchmark shows that coherence can coexist with substantial misclassification. This result directly addresses the measurement problem raised by computational analyses of short, multilingual and context-dependent platform comments.

The stance errors are theoretically informative. The benchmark codebook distinguishes EU-wide support or rejection from issue-specific praise, policy criticism, criticism of politicians and other ambiguous reactions. This distinction mirrors the classic separation between broad opposition to integration and more specific or qualified criticism (Kopecky and Mudde 2002; Taggart and Szczerbiak 2002). The model's low pro-EU precision (.130) and anti-EU precision (.275), together with the predominance of neutral/weak reference labels, indicate that a zero-shot stance prompt can collapse these conceptually distinct forms of expression into apparently explicit political positions.

The temporal analysis illustrates why validation must precede aggregation. A classifier with low precision can generate sharp daily shares and cumulative trends that appear politically meaningful when plotted over

time. The attention-weighted index still has value as a screening device for locating intervals that deserve close reading, but the present benchmark does not justify interpreting its peaks as verified increases in Euroscepticism. Observed comment-volume spikes remain real platform events; their political meaning requires qualitative event analysis and more reliable text labels within the broader hybrid-media context of the 2024 election, disinformation and geopolitical conflict (Daniel and Obholzer 2025; D'Ambrosi et al. 2026; Ferro-Santos and Sepulveda 2026; Pane 2026).

The network analysis adds a second form of error propagation. Co-commenting already measures only shared participation under posts, not replies, exposure or deliberation; when nodes are additionally coloured and selected through a low-precision stance classifier, claims about cross-stance mixing or same-stance cohesion become doubly indirect. The appropriate conclusion is therefore not that the observed public is open or non-polarised, but that stance-conditioned network inference should be postponed until user-level labels are validated or recalibrated.

These findings contribute to three literatures without collapsing their levels of inference. For Euroscepticism research, they show that the conceptual distinction between EU-wide opposition and issue-specific criticism is not merely theoretical but a concrete measurement problem for automated stance detection. For EU public institutional communication, they demonstrate that reaction data can illuminate audience expression only when communicative approval, policy criticism and institutional legitimacy are kept separate (Lovari and Belluati 2023). For platform and computational communication research, the case demonstrates that benchmark validation is not a decorative robustness exercise: it can materially reverse the interpretation of downstream temporal and network outputs.

6. Conclusion

This article analysed 27,644 reactions retrieved from the European Parliament's official TikTok account between February 2024 and November 2025 and, crucially, evaluated the automated sentiment and EU-stance labels on a stratified 500-comment benchmark before accepting their political interpretation. The validation results materially change the conclusions that can be drawn from the full-corpus temporal and co-commenting outputs.

Three conclusions follow. First, the automated labels are not sufficiently accurate for direct prevalence claims: sentiment accuracy is .416 with macro-F1 .403, while stance accuracy is .330 with macro-F1 .327. Second, the benchmark reveals systematic over-assignment of explicit stance, especially pro-EU stance: 81.8% of final reference labels are neutral/weak, compared with only 24.4% of model predictions. Third, temporal stance curves, the attention-weighted anti-EU index and the stance-coloured co-commenting graph should therefore be treated as model-conditioned screening diagnostics rather than validated evidence of support, Eurosceptic mobilisation or ideological segregation.

For institutional communication, automated classifiers may still be useful for triage. High-volume intervals or comments receiving strong model scores can identify material for closer qualitative review, but the present evidence does not justify using the model to estimate the prevalence of support or opposition. Monitoring systems should therefore combine automated screening with validated coding rather than translating classifier outputs directly into claims about public attitudes.

The study has several limitations. It examines one institutional account and a self-selected population of commenters; the corpus is a time-specific reconstruction of comments retrievable through third-party tools; TikTok ranking, moderation, deletion and recommendation processes are only partially observable; and the analysis is primarily textual, excluding audiovisual context and non-commenting audiences. The 500-comment benchmark is drawn from the model-eligible pool rather than every retrieved comment, while the separate 100-comment coverage audit is not used in performance estimates. Although the benchmark was independently

coded by two members of the research team who were blinded to model labels and disagreements were subsequently resolved by human consensus, it remains a finite, corpus-specific reference sample. Most importantly, the low classification performance means that the existing full-corpus stance-derived temporal and network analyses require reclassification before they can support stronger substantive claims.

The next analytical step is not to add stronger interpretation to the present model outputs but to improve the measurement layer. The present human-adjudicated reference standard provides a validated basis for diagnosing classifier error; future model development should use additional human-coded data for training or calibration while preserving an independent holdout set for final evaluation. After model improvement, the complete 27,644-comment corpus should be reclassified and all temporal and stance-conditioned network analyses rerun. Robustness should be tested across alternative multilingual models, keyword gates and confidence thresholds. Comparative designs across EU institutions and platforms, together with multimodal analysis and reply or exposure data where ethically and technically available, could then assess whether any validated temporal or structural patterns generalise.

Within these boundaries, the study offers a transparent demonstration of why measurement validation matters for computational political communication. A large platform corpus and sophisticated downstream analyses can produce persuasive patterns even when the classification layer is weak. By reporting the benchmark results rather than concealing them, the study converts a potential validity failure into an empirically documented methodological finding: automated stance labels must be validated before they are used to infer Euroscepticism, institutional support or ideological network structure.

References

Alonso-Munoz L., A. Casero-Ripolles (2020), “Populism against europe in social media: The eurosceptic discourse on twitter in spain, italy, france, and united kingdom during the campaign of the 2019 european parliament election”, Frontiers in Communication, 5: 54. https://doi.org/10.3389/fcomm.2020.00054

Alonso-Munoz L., A. Casero-Ripolles (2021), “Looking for the guilty? The discursive strategy on twitter of the european populist political actors in times of crisis”, Cultura, Lenguaje y Representación, 26: 29–45. https://doi.org/10.6035/clr.5827

Barberio V., I. Kuric, E. Mollona, and L. Pareschi (2020), “The use of social media in EU policy communication and implications for the emergence of a European public sphere”, Investigaciones Regionales – Journal of Regional Research, 46: 111–129.

Berti C., E. Loner (2020), “The 2019 european elections on twitter between populism, euroscepticism and nationalism: The case of italy”, SocietàMutamentoPolitica, 11(22): 275–288. https://doi.org/10.13128/smp-12654

Bouza L., A. Oleart, and J. Tunon (2019), “Framing EU trade policy online: The case of @NoAlTTIP on Twitter”, Communication & Society, 32(4): 257–273. https://doi.org/10.15581/003.32.4.257-273

Ceron A., L. Curini, and S.M. Iacus (2016), “First- and second-level agenda setting in the Twittersphere: An application to the Italian political debate”, Journal of Information Technology & Politics, 13(2): 159–174. https://doi.org/10.1080/19331681.2016.1160266

Chadwick A. (2017), The Hybrid Media System: Politics and Power (2nd ed.), Oxford University Press. https://doi.org/10.1093/oso/9780190696726.001.0001

D'Ambrosi L., M. Parito, and F. Pasciuto (2026), “Young activists' communicative practices of resilience: The challenge of disinformation toward the 2024 European elections”, PArtecipazione e COnflitto, 19(1): 77–93. https://doi.org/10.1285/i20356609v19i1p77

Daniel W.T., L. Obholzer (2025), "Social media campaigning across multiple platforms: Evidence from the 2024 European elections", West European Politics, 49(1): 293–319. https://doi.org/10.1080/01402382.2025.2501464

Elshehawy A., K. Gavras, N. Marinov, F. Nanni, and H. Schoen (2022), "Illiberal communication and election intervention during the refugee crisis in germany", Perspectives on Politics, 20(3): 860–878. https://doi.org/10.1017/S1537592721003108

Fenoll V. (2022), "Third-order election. Spanish political parties' communication on Facebook during the 2019 European Parliament election campaign", Profesional de la Información, 31(3): e310306. https://doi.org/10.3145/epi.2022.may.06

Ferro-Santos S., R. Sepulveda (2026), "From Brussels to the 'for you' page: Members of the European Parliament TikTok adoption and campaign use", Revista de Comunicación, 25(1): 97–120. https://doi.org/10.26441/RC25.1-2026-4303

Fortunato P., M. Pecoraro (2022), "Social media, education, and the rise of populist Euroscepticism", Humanities & Social Sciences Communications, 9(1): 301. https://doi.org/10.1057/s41599-022-01317-y

Kasekamp A., M.-L. Madisson, and L. Wierenga (2019), "Discursive opportunities for the estonian populist radical right in a digital society", Problems of Post-Communism, 66(1): 47–58. https://doi.org/10.1080/10758216.2018.1445973

Kopecky P., C. Mudde (2002), "The two sides of Euroscepticism: Party positions on European integration in East Central Europe", European Union Politics, 3(3): 297–326. https://doi.org/10.1177/1465116502003003002

Ksenzenko O.A. (2022), "Post-truth, post-press, post-europe. Euroscepticism and the crisis of political communication", MGIMO Review of International Relations, 15(3): 279–292. https://doi.org/10.24833/2071-8160-2022-3-84-279-292

Lovari A., M. Belluati (2023), "We are all Europeans. EU institutions facing the Covid-19 pandemic and information crisis", in G. La Rocca, M.-E. Carignan, and G. Boccia Artieri (eds.), Infodemic Disorder: Covid-19 Coping Strategies in Europe, Canada and Mexico, Palgrave Macmillan, pp. 65–96. https://doi.org/10.1007/978-3-031-13698-6_4

Marchal N., L.-M. Neudert, B. Kollanyi, and P.N. Howard (2021), "Investigating visual content shared over twitter during the 2019 EU parliamentary election campaign", Media and Communication, 9(1): 158–170. https://doi.org/10.17645/mac.v9i1.3421

Meyer C. (1999), "Political legitimacy and the invisibility of politics: Exploring the European Union's communication deficit", JCMS: Journal of Common Market Studies, 37(4): 617–639. https://doi.org/10.1111/1468-5965.00199

Pane S. (2026), "Resilience in post-truth and post-digital European Union: Youth-led civil society organisations in the 2024 European elections", PArtecipazione e COnflitto, 19(1): 94–113. https://doi.org/10.1285/i20356609v19i1p94

Rivas-de-Roca R., M. Garcia-Gordillo (2022), "Brussels will pay for everything. The shaping of the European public sphere before NextGenerationEU", Profesional de la Información, 31(3): e310301. https://doi.org/10.3145/epi.2022.may.01

Segesten A.D., M. Bossetta (2019), "The Eurosceptic Europeanization of public spheres: Print and social media reactions to the 2014 European Parliament elections", Comparative European Politics, 17(3): 361–379. https://doi.org/10.1057/s41295-017-0099-5

Silva B.C., L. Schuermann, and S.-O. Proksch (2024), "Modulation of democracy: Partisan communication during and after election campaigns", British Journal of Political Science, 54(2): 339–354. https://doi.org/10.1017/S0007123423000169

Taggart P., A. Szczerbiak (2002), "Europeanisation, Euroscepticism and party systems: Party-based Euroscepticism in the candidate states of Central and Eastern Europe", Perspectives on European Politics and Society, 3(1): 23–41. https://doi.org/10.1080/15705850208438827

Vochocova L., J. Rosenfeldova (2023), "The EU as a weak and authoritative traitor: Signs of post-socialist ressentiment and populist rhetoric in online civic anti-European discourses", Journal of Contemporary European Studies, 31(3): 916–928. https://doi.org/10.1080/14782804.2022.2092460

Notes on Contributors

Tănase Tasențe. Lecturer at the Faculty of Law and Administrative Sciences, Ovidius University of Constanța, Romania. His research focuses on institutional communication, social media analytics, and computational political communication. E-mail: tanase.tasente@365.univ-ovidius.ro. ORCID: https://orcid.org/0000-0002-3164-5894.

Cristina-Dana Popescu. Lecturer at the Faculty of Letters, Ovidius University of Constanța, Romania. Her research covers linguistics, interdisciplinary communication, education, and personal development. E-mail: cristina.popescu@365.univ-ovidius.ro. ORCID: https://orcid.org/0009-0009-5280-6608.