Journal of Radiological Protection

Official journal of the Society for Radiological Protection

PAPER • OPEN ACCESS

Beyond the 5G-COVID moment: the scope of EMF mis- and disinformation on German-language social media

To cite this article: Tobias Eder *et al* 2026 *J. Radiol. Prot.* **46** 041502

View the article online for updates and enhancements.

You may also like

- The growth of climate change misinformation in US philanthropy: evidence from natural language processing
Justin Farrell
- Online misinformation during extreme weather emergencies: short-term information hazard or long-term influence on climate change perceptions?
Stefan Daume
- Extreme heat in focus: comparing media narratives from the Philippines and Hong Kong
Johanne Rei R Castro and Laurence L Delina

This content was downloaded from IP address 146.75.182.19 on 05/11/2026 at 13:05

Journal of Radiological Protection

Check for updates

OPEN ACCESS

RECEIVED
4 May 2026

REVISED
20 July 2026

ACCEPTED FOR PUBLICATION
11 September 2026

PUBLISHED
9 October 2026

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

PAPER

Beyond the 5G-COVID moment: the scope of EMF mis- and disinformation on German-language social media

Tobias Eder[1,*], **Katja Aue**[2], **Larissa Drescher**[2], **Kerstin Dressel**[3], **Georg Groh**[1] and **Jutta Roosen**[4]

[1] School of Computation, Information and Technology, Technical University of Munich, Garching, Germany
[2] C[3] team, Munich, Germany
[3] sine-Institut, Munich, Germany
[4] School of Management, Technical University of Munich, Freising, Germany
[*] Author to whom any correspondence should be addressed.

E-mail: tobias.eder@in.tum.de

Keywords: electromagnetic fields, 5G, misinformation, disinformation, social media, risk communication, radiological protection

Abstract

Public perception of health risks from electromagnetic fields (EMF) and 5G mobile networks diverges persistently from the scientific evidence base, a gap increasingly shaped by social media. While existing research has examined the acute COVID-5G conspiracy episode on individual platforms, systematic multi-platform analyses of the broader EMF misinformation landscape remain scarce. This study presents a mixed-methods analysis of EMF and 5G-related misinformation across YouTube, Telegram, and Twitter/X in German-language social media, combining automated classification, sentiment analysis, and topic modelling with structured qualitative content analysis. The three corpora comprise over 430 000 posts from more than 16 000 accounts, covering content published primarily between 2019 and 2024. The analysis reveals that EMF misinformation circulates in a predominantly pseudo-expert communicative register rather than overtly inflammatory language, is structurally embedded within broader conspiracy ecosystems rather than circulating as standalone content, and is sustained by commercial incentive structures alongside ideological motivations. The three platforms serve complementary functional roles within an integrated distribution network. While the pandemic-era peak in activity has subsided, the underlying network infrastructure remains intact. The findings establish an empirical baseline for the German-language EMF misinformation landscape against which future developments, including the impact of generative AI on misinformation production, can be measured.

1. Introduction

Public perception of risks from electromagnetic fields (EMF) has practically always differed from the scientific evidence base (Rowley and Mazar 2021). In recent years, the deployment of fifth-generation (5G) mobile network infrastructure has intensified this divergence, giving rise to health concerns that are not supported by the current body of research yet have proven capable of motivating real-world action, including the vandalism of telecommunications infrastructure in several European countries during the COVID-19 pandemic (Jolley and Paterson 2020). The gap between evidence and perception represents a core challenge for radiological protection, and it is increasingly shaped by the dynamics of social media platforms.

Social media networks have become a primary channel through which the public encounters, discusses, and forms opinions about EMF-related topics (e.g. Ahmed *et al* 2020). These platforms enable rapid dissemination of both accurate information and misinformation, often without clear differentiation between the two. For radiological protection authorities, the practical consequence is that effective risk communication now requires not only the production of scientifically sound information, but also an understanding of the misinformation landscape it must compete with: its scope, its structure, the strategies it employs, and the platforms on which it operates.

© 2026 The Author(s). Published on behalf of the Society for Radiological Protection by IOP Publishing Ltd

Public concern about potential health effects of EMF has a substantial research history. Survey data from across Europe have consistently shown that a significant proportion of the population expresses concern about EMF exposure, particularly from mobile phone base stations, despite the absence of established health risks at exposure levels below regulatory limits (Freudenstein *et al* 2015). At the same time experimental studies have demonstrated that risk perception of EMF is driven primarily by subjective exposure perception rather than actual exposure levels, and that even precautionary measures intended to reassure the public can paradoxically amplify perceived risk and erode trust in protective institutions (Wiedemann and Schütz 2005). More recently, the deployment of 5G networks has introduced a new focal point for these dynamics, with studies identifying female gender, high perceived exposure, distrust in public policy effectiveness, and perceived severity of consequences as predictors of elevated 5G risk perception (Koh *et al* 2020). Comparative research between Germany and Greece has further shown that the relationship between precautionary information, risk perception, and trust in radiological protection authorities is sensitive to sociocultural context, underscoring the importance of nationally specific empirical evidence (Eggeling-Böcker *et al* 2025). These findings indicate that the gap between scientific evidence and public perception of EMF is not merely an information deficit, but a structural feature of how non-ionising radiation risks are processed by the public.

The rapid spread of the conspiracy theory linking 5G mobile infrastructure to COVID-19 during 2020 brought EMF misinformation into acute public and academic focus. The academic literature distinguishes between misinformation and disinformation distinguished by the intent behind the content (Wardle and Derakhshan 2017). Since the distinction hinges on producer intent, which is typically not accessible from the content itself, recent work in health communication has increasingly adopted 'misinformation' as an intent-agnostic umbrella term (Altay *et al* 2023, Nan *et al* 2023). We follow this convention and use 'misinformation' throughout to refer to content that contradicts the scientific consensus on EMF, without making claims about the producer's intent. Ahmed *et al* (2020) mapped the dissemination of the 5G-COVID narrative on Twitter over a single week, identifying a broadcast-dominated network structure with no authoritative counter-voice. Bruns *et al* (2020) traced the migration of the same narrative from fringe conspiracy groups to mainstream attention on Facebook, demonstrating that the COVID-5G conspiracy built upon pre-existing anti-5G networks. This dynamic is consistent with the broader observation that health-related conspiracy theories tend to cluster and mutually reinforce one another rather than circulating in isolation (Cinelli *et al* 2022). Jolley and Paterson (2020) demonstrated that endorsement of 5G-COVID conspiracy beliefs was associated with state anger and willingness to justify violence against telecommunications infrastructure. Subsequent longitudinal analysis confirmed that 5G-related misinformation on Twitter peaked sharply in early 2020 and declined thereafter, though it did not disappear entirely (Langguth *et al* 2022). On Telegram, which has attracted users migrating from more heavily moderated platforms, misinformation has been shown to circulate within small but highly active communities rather than reaching general audiences (Herasimenka *et al* 2022), though the platform remains understudied relative to its role in the conspiracy ecosystem.

Despite this body of work, the current evidence base leaves radiological protection authorities without the empirical foundation they need to design effective communication responses. Existing studies focus predominantly on English-language content and analyse single platforms, most frequently Twitter (Ahmed *et al* 2020, Langguth *et al* 2022) or individual alternative platforms (Bruns *et al* 2020, Herasimenka *et al* 2022). They address the COVID-5G conspiracy as an isolated episode (Ahmed *et al* 2020, Bruns *et al* 2020, Jolley and Paterson 2020) rather than situating it within the broader and ongoing landscape of EMF-related misinformation, which encompasses health anxiety about mobile phone radiation, claims of electromagnetic hypersensitivity, commercial exploitation of radiation fears, and long-standing distrust of regulatory bodies such as the International Commission of Non-Ionizing Radiation Protection (ICNIRP). Among the studies reviewed above, computational and survey-based approaches have been applied separately. Systematic combinations of large-scale quantitative analysis with structured qualitative content analysis remain rare. Systematic analysis of EMF-related misinformation on German-language social media is lacking, despite Germany's position as Europe's largest market for mobile communications and a substantial body of survey-based research on EMF risk perception in the German population (Wiedemann and Schütz 2005, Freudenstein *et al* 2015, Eggeling-Böcker *et al* 2025). Understanding that 5G-COVID misinformation existed on Twitter in April 2020 is of limited operational value to an agency seeking to monitor and respond to the ongoing misinformation landscape across multiple platforms in its own national context. What is needed is a systematic mapping of the scope, structure, and communicative strategies of EMF-related misinformation as it manifests in a specific national media environment.

This paper addresses this gap through a systematic, multi-platform analysis of EMF and 5G-related misinformation in German-language social media. Drawing on data collected across YouTube, Telegram,

and Twitter/X between 2019 and 2024, the study pursues three objectives: (1) to map the scope, structure, and temporal dynamics of the EMF misinformation landscape across platforms, including the volume and actor structures involved; (2) to identify platform-specific differences in how EMF-related content is framed, disseminated, and received; and (3) to characterise the communicative register and discursive strategies employed in EMF misinformation. The study employs a mixed-methods design combining large-scale automated classification, sentiment analysis, and topic modelling with structured qualitative content analysis. The research was conducted as part of a project commissioned by the German Federal Office for Radiation Protection (BfS, project 3621 EMF 106). The findings provide an empirical foundation for understanding how EMF-related misinformation operates within the contemporary media landscape and how its structure and communicative strategies may challenge risk communication in radiological protection.

2. Data and methods

2.1. Data collection

The study initially collected data from five major social media platforms: YouTube, Telegram, Twitter/X, Instagram, and Facebook. For this article, we focus on three platforms that represent distinct roles within the misinformation ecosystem. Building on a broader multi-platform analysis of EMF-related discourse in German-language social media (Drescher *et al* 2025), we select Telegram, YouTube, and Twitter/X as cases that capture different communicative functions. Telegram serves as the primary community hub for EMF-related misinformation in German, with the largest volume of relevant content and the most insular communication dynamics. Twitter/X functions as the interface between niche EMF discourse and broader political debate, exhibiting a distinctly more confrontational discourse culture. YouTube operates as the principal medium for long-form content production and is unique among the platforms studied in that audience responses (comments) are directly observable alongside creator content, allowing analysis of both dissemination and reception. Instagram and Facebook were included in the broader project but are excluded from the present analysis for reasons of methodological comparability. Data for both platforms were obtained exclusively through the CrowdTangle service (since discontinued), whose data model and coverage differed substantially from the direct collection methods available on the other platforms, precluding the network-based expansion approach described below. The two platforms also yielded considerably smaller corpora with predominantly top-level posts. YouTube was retained despite the limitation that video content itself could not be analysed in our study. The qualitative analysis showed that a characteristic dissemination pattern in the collected networks consisted of short posts on other platforms making a claim accompanied by a link to a long-form YouTube video elaborating it, making the platform a central node in how misinformation circulated across all platforms during the collection period.

Data collection followed a seed-user approach with iterative network expansion. An initial set of accounts known to disseminate EMF-related content was identified through manual research across all platforms, yielding a starting list of 12 Telegram channels, 49 Twitter/X accounts, and 28 YouTube channels. From these seed users, the network was expanded in successive crawling rounds (automated collection passes that follow links and connections outward from known accounts) through two mechanisms: following direct connections (follower/following relationships) and identifying additional accounts through shared EMF-related keywords and topical overlap. After each expansion round, the network was pruned (accounts not relevant to the topic were removed) using a combination of automated filtering—sentiment analysis, entity recognition, and topic modelling to assess thematic relevance—and manual verification to remove off-topic or non-German-language accounts. Expansion was continued until successive rounds yielded predominantly off-topic or inactive channels, indicating diminishing returns. This convergence behaviour is consistent with snowball sampling in topically bounded Telegram networks, where small-world network properties (most actors being reachable from any other actor within a few steps) and the limited forwarding radius of niche communities lead to rapid saturation of the reachable actor landscape (Baumgartner *et al* 2020, Ricaud *et al* 2020, Peter *et al* 2022). The resulting corpus represents a systematic mapping of the German-language EMF misinformation network on each platform, though it does not constitute a random sample of general platform content. The misinformation rates reported in section 3.1.1 should therefore be understood as densities within the identified network, not as prevalence estimates for the platforms at large. Within these networks we define individual actors as users who either actively disseminate EMF-related content or maintain direct network connections to users who did.

Data collection was adapted to the technical constraints of each platform. For Telegram, a custom crawler was developed to interface with the platform's API, which permits comparatively high request volumes. Collected data were stored in a graph database to preserve the network structure of channels, groups, and message forwarding relationships. Automated language detection was applied to prioritise German-language content. The combination of generous API access and the platform's group-based architecture made Telegram the largest dataset by a substantial margin.

For Twitter/X, data collection was constrained by the removal of academic API access in early 2023. The open-source framework twscrape was used as an alternative, retrieving up to 3.200 recent tweets per account. This ceiling posed no limitation for most accounts but restricted historical coverage for highly active users. Consequently, the Twitter/X dataset is skewed toward more recent content relative to the other two platforms, a limitation that should be considered when interpreting temporal trends.

YouTube lacks the inherent network structures present on Telegram and Twitter/X. Instead, relevant videos were identified through links shared on Telegram channels and Twitter accounts already in the dataset. For each identified video, metadata (title, description, view count, subscriber count) and available comments were collected via the YouTube Data API (application programming interface, the platform's structured data-access service) v3. Video content itself was not processed, as systematic analysis of audiovisual material was beyond the project's scope. Classification of YouTube content therefore relies on textual metadata only.

Data collection took place between September 2023 and March 2024, capturing content published from 2019 onward. Coverage depth varies by platform: Telegram provides the deepest historical record; YouTube covers content linked from other platforms regardless of publication date, while Twitter/X is limited to each account's most recent posts as described above.

2.2. Data processing and quantitative content analysis

To assess the relative importance of actors within the collected network, we applied a combined metric that weights both reach and thematic focus: Importance = log (n_followers) × n_EMF posts, where n_followers denotes the account's subscriber or follower count at the time of collection and n_EMF posts the number of posts classified as EMF-related in the dataset. The logarithmic transformation of follower counts normalises the disproportionate influence of very large accounts, ensuring that actors with smaller but topically dedicated audiences are not systematically undervalued relative to accounts with large but thematically diffuse followings. This metric revealed that the highest-ranking accounts by reach were predominantly content aggregators rather than EMF specialists, a pattern discussed in section 3.1.

Automated identification of potential misinformation content was performed by a dedicated classifier applied to the assembled corpus. The classifier's outputs enter the analyses reported in this paper solely as the misinformation density estimates in section 3.1.1. Corpus construction (section 2.1) did not involve the classifier. The analyses of emotional register, thematic structure, and terminology described below were conducted on the full corpora, independently of per-post classification. The qualitative content analysis (section 2.3) employed a two-step random sampling of accounts and posts, conducted without access to classifier outputs.

The classifier employed a retrieval-augmented generation (RAG) approach, in which a pretrained large language model from the open-source LLaMA family is supplied with relevant reference material at inference time. For each post, the system identified EMF-related keywords and retrieved corresponding reference material from a curated database reflecting knowledge from the current scientific consensus on EMF and 5G topics. This reference database was compiled from publications of the German Federal Office for Radiation Protection (Bundesamt für Strahlenschutz, BfS), the German Commission on Radiological Protection (Strahlenschutzkommission, SSK), and the Swiss Federal Office for the Environment (BAFU). The language model then compared the claims in the post against the retrieved reference material. Posts containing assertions unsupported or contradicted by the reference base were flagged as potential misinformation.

This constitutes the study's operational definition of misinformation underlying the density estimates in section 3.1.1. The criterion is deliberately consensus-referenced. It does not adjudicate the scientific validity of individual claims and cannot, by construction, distinguish deliberate false information from legitimate scientific disagreement. Content which engages critically with contested areas of EMF research would in principle be flagged. The manual content analysis (section 2.3) provides an empirical check on this limitation. Coded samples of flagged content were dominated by conspiracy narratives, unsubstantiated health claims, and commercial promotion, rather than substantive engagement with the scientific literature, indicating that the flagged content does not to any substantial degree represent legitimate scientific discourse.

Table 1. Performance of the misinformation classifier evaluated against the independent manual content analysis (section 2.3). Values are shares of agreement between automated classification and manual coding, with 95% Wilson score confidence intervals in brackets. *n* denotes the number of manually coded posts per platform.

	Precision	Recall/sensitivity	Specificity
Telegram ($n = 562$)	54.0% [46.8, 61.0]	69.4% [61.5, 76.3]	79.0% [74.9, 82.7]
Twitter/X ($n = 1{,}288$)	25.4% [21.9, 29.1]	82.1% [75.7, 87.1]	62.5% [59.6, 65.3]

The classifier was designed to default to flagging posts rather than passing in ambiguous cases. The resulting misinformation rates were therefore expected to overestimate misinformation density. To assess this assumption, classifier outputs were compared against the independent manual content analysis (section 2.3). For this evaluation, a coded post was counted as misinformation if it received at least one code from the categories for unsubstantiated health claims or conspiracy narratives. This mapping is deliberately strict: the codebook categories do not exhaustively cover content contradicting the reference consensus, so the two methods operationalise misinformation differently and the resulting agreement rates conservatively bound classifier performance.

Performance was assessed on two platforms, analogous to a diagnostic screening test. The measurement was done on precision (the share of classifier-flagged posts confirmed by the manual coding), recall or sensitivity (the share of coder-identified misinformation the classifier flagged), and specificity (the share of posts without misinformation-indicative codes correctly left unflagged) across classifier-coding agreement. Table 1 reports these metrics with 95% confidence intervals. Precision values should not be read as a direct cross-platform ranking of classifier performance, as they also reflect differences in the composition of flagged content relative to the strict category mapping. For YouTube, inter-method comparison was not feasible, as automated classification was focused on video titles and descriptions while the manually annotated sample predominantly comprises video comments. On both validated platforms, false positives exceeded false negatives, empirically confirming that the reported rates constitute upper-bound estimates of misinformation density within the identified networks, relative to the manual coding standard.

Independently of the classification pipeline, a sentiment analysis model was applied to identify posts with pronounced emotional framing. Research has consistently shown that false and misleading content tends to employ stronger emotional framing than accurate information (Vosoughi *et al* 2018), making this a useful signal for surfacing content of potential interest during corpus exploration.

The collected corpora were further analysed along three dimensions. First, the emotional register of posts was assessed through automated sentiment classification based on the GoEmotions taxonomy (Demszky *et al* 2020), assigning each post a primary emotion label from a set of 27 fine-grained categories subsequently grouped for analysis into seven summary categories: neutral, approval, realisation, annoyance, disapproval, curiosity, and confusion. Since classification assigns the single highest-probability category, neutral tends to dominate whenever no specific emotion is expressed with high confidence. The neutral share should therefore be understood as an upper bound reflecting the absence of a strong emotional signal rather than a positive identification of emotionally neutral intent. Sentiment was computed both at the level of individual posts and aggregated by actor and by platform to identify differences in communicative tone across the three platforms.

Second, the thematic structure of each platform's corpus was analysed through Top2Vec topic modelling (Angelov 2020). Top2Vec derives topics automatically by grouping posts according to semantic similarity, without requiring the number of topics to be specified in advance. For YouTube and Twitter/X, the resulting topic sets were hierarchically reduced from 89 raw topics to six overarching themes, enabling comparison of thematic priorities while retaining interpretability. For Telegram, the substantially larger dataset permitted a parametric uniform manifold approximation and projection (UMAP) visualisation (a technique that arranges content in a two-dimensional map in which thematically similar posts appear close together) of the thematic space, capturing the relative positioning of content clusters rather than discrete topic assignments.

Third, keyword and hashtag frequencies were computed across all posts to trace the prevalence and distribution of EMF-related terminology (including '5G,' 'Elektrosmog' [electric smog], 'Strahlung' [radiation], 'EMF,' and 'Mobilfunk' [mobile communications]) and to identify patterns of topical co-occurrence.

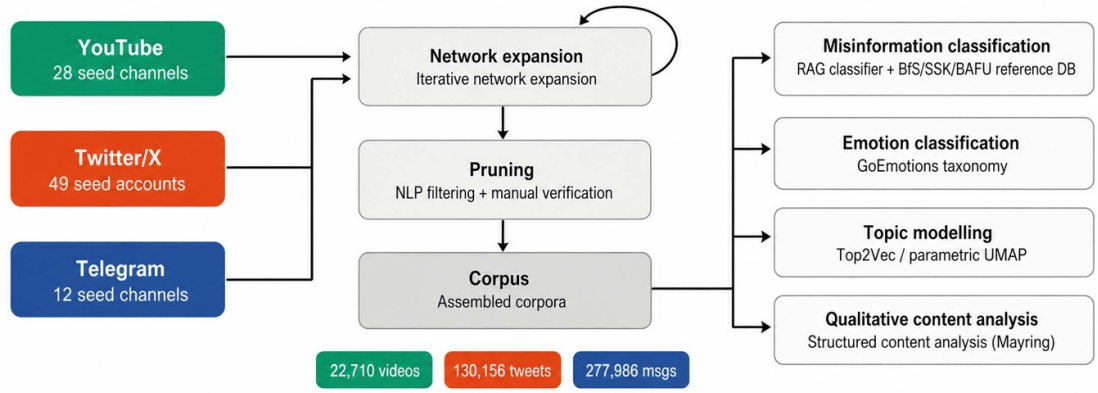

Figure 1. Methodological pipeline from data collection through analysis. Seed users were iteratively expanded and pruned per platform, producing the assembled corpus. Four analytical procedures were applied independently to the resulting data.

The available engagement metrics differ by platform—views and likes for YouTube, likes and retweets for Twitter/X, reactions and forwarding counts for Telegram—but in all cases reflect user interactions with content rather than exposure. The study does not have access to platform-internal impressions or reach data. The interaction metrics reported in section 3 should therefore be understood as indicators of active engagement, not as measures of how many users encountered a given post.

2.3. Qualitative content analysis

The qualitative sub-study used a structured qualitative content analysis following Mayring's (2015) approach. This method aims to systematically extract predefined structures from text by using a theoretically grounded category system. The analytical framework was developed primarily deductively, drawing on insights from a preceding literature review. Its purpose was to identify recurring content patterns and communicative strategies in social media discourse related to 5G /EMF misinformation.

A comprehensive codebook was developed[5] and applied uniformly across social media content from YouTube, Twitter/X, and Telegram. The codebook comprised five main category blocks (A–E). Block A captured contextual embedding themes, identifying broader discursive environments (e.g. Corona, politics, esotericism) in which 5G /EMF-related communication is embedded beyond the immediate EMF topic. Block B constituted the analytical core and covered EMF-related keywords (these were predefined by the commissioning authority), including electric smog, 5G, and a clustered set of EMF-related terms (e.g. radiation exposure, mobile phones, and infrastructure). Additional categories captured conspiracy theories, health-related effects, environmental impacts, references to organisations and initiatives, claims of corruption or unfair advantage, and mentions of scientific studies—all in relation to EMF and 5G. Block C included structural elements such as URLs, images, emojis, hashtags, mentions of a post, and the existence of questions. Block D captured discursive communication strategies, including informative framing, factual assertions, attention grabbing and calls to action, problem-solution framing, and positive or negative motivational appeals (based on University of Minnesota Libraries Publishing 2016). Block E addressed stylistic characteristics, including emotionality, address style, advertising, expressions of agreement or rejection, and the appearance of professionalism. Each category was defined by explicit coding rules and anchor examples.

The sample for the qualitative analysis was drawn from the same crawled corpora used in the quantitative strand of the study. For YouTube, the qualitative study additionally included the available transcripts of 30 randomly selected misinformation videos identified in the quantitative sub-study. Other than that, for each platform, random samples were drawn, amounting to approximately 10% of the crawled data. Individual posts were defined as the unit of coding.

Coding was conducted using Atlas.ti, which enabled systematic code assignment, frequency analysis, and co-occurrence analysis. To ensure methodological rigour, dual coding was performed for approximately 9% of the material per platform. Intercoder reliability was calculated using Krippendorff's Cu-Alpha and ranged between 0.88 and 0.91, indicating high agreement. Coding discrepancies were resolved

[5] The complete codebook is available from the authors upon request.

through discussion, and minor refinements were applied to the codebook where necessary. Coding proceeded iteratively until thematic saturation was reached, meaning that no substantively new categories emerged.

Figure 1 provides an overview of the complete methodological pipeline, from platform selection and seed-user identification through data collection, network expansion, and pruning, to the analytical procedures described above.

Several stages of this study employed automated language-processing tools, which we summarise here for transparency. Network pruning during data collection used automated sentiment analysis, named-entity recognition, and topic modelling to assess thematic relevance. Misinformation classification was performed by the RAG-based classifier described above, built on a pretrained large language model. Emotion classification and topic modelling used pretrained open-source models (GoEmotions and Top2Vec). The qualitative content analysis was conducted entirely by human coders without automated assistance.

3. Results

3.1. Quantitative findings

3.1.1. Corpus composition and misinformation rates

Figure 2 summarises the corpus composition and estimated misinformation rates across the three platforms. Telegram yielded the largest dataset by a substantial margin, with 277 986 messages collected from 1985 channels and groups. Twitter/X produced 130 156 tweets from 8368 accounts, while the YouTube corpus comprised 22 710 videos from 6099 channels alongside their associated comment threads.

Estimated misinformation rates within the collected networks ranged from 19.2% on YouTube to 34.2% on Twitter/X, with Telegram at 31.8%. These figures constitute upper-bound estimates of misinformation density within the mapped networks (see section 2.2) and are not prevalence estimates for the platforms at large (section 2.1), nor are the rates directly comparable across platforms, as they reflect differences in communication format and classifier applicability as much as underlying misinformation density. Telegram contains the highest absolute volume of flagged content by a wide margin, but its conversational, chat-like communication style dilutes the per-post rate. Twitter/X's higher rate warrants some methodological caution, as the validation reported in section 2.2 shows the strongest over-flagging on this platform. YouTube's lower rate reflects the limitation that only video titles and descriptions (not audiovisual content) were classifiable. Additionally, the performance of the classifier on this platform could not be independently validated (see section 2.2).

The available temporal indicators suggest that EMF-related misinformation activity was most intense during the period 2020–2022, though the trajectory differed by platform. YouTube posting volume peaked in 2020 and had approximately halved by 2022, though a partial resurgence was observed from 2023 onward. On Telegram, posting volume across the collected network rose sharply from April 2020 onward (increasing nearly sevenfold over the course of 2020) and peaked in January 2022 at over 80 000 messages per month. Activity subsequently declined but stabilised at approximately 50 000–55 000 messages per month through 2023, a level still several times higher than the pre-2020 baseline. The Twitter/X dataset does not permit reliable temporal analysis due to the 3200-tweet-per-account retrieval limit, which skews the corpus toward recent content. Across the platforms where temporal trends are observable, two patterns stand out: the COVID-19 pandemic triggered a sharp and sustained increase in activity within the identified networks, and the subsequent decline has been gradual rather than complete. The prominence of the 2020–2022 period reflects the entanglement of EMF misinformation with the broader COVID-19 conspiracy discourse during the pandemic, which constitutes a structural feature of these networks examined further in section 3.1.4. The underlying infrastructure of channels, accounts, and cross-platform linkages remains intact, suggesting that the ecosystem could be reactivated by future triggering events.

3.1.2. EMF-related terminology across platforms

Table 2 shows the distribution of EMF-related keywords across the collected corpora on each platform. The term '5G' was the most frequent keyword on all platforms by a wide margin: 394 occurrences in YouTube video metadata, 3.634 in tweets, and 119.894 in Telegram messages. This dominance held regardless of corpus size and was consistent across all analytical dimensions. '5G' functions as the de-facto anchoring term of the discourse.

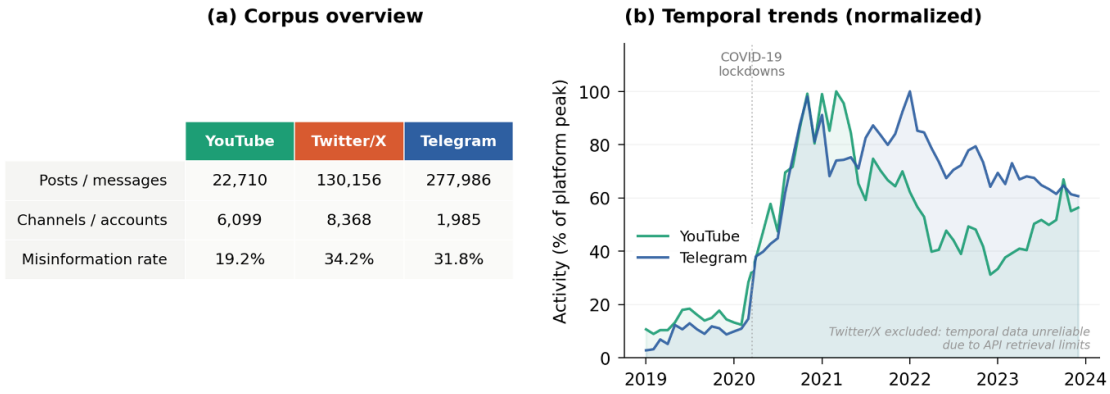

Figure 2. (a) Corpus overview across the three platforms studied. Misinformation rates represent classifier-derived estimates within the collected networks. (b) Monthly content production normalised to each platform's peak month. YouTube shows video publication counts; Telegram shows total message volume across collected channels. Twitter/X is excluded due to the 3200-tweet-per-account API retrieval limit, which skews the dataset towards recent content. The dashed line marks the onset of COVID-19 lockdowns in March 2020.

Table 2. Distribution of EMF-related keywords across platforms. Raw counts and normalised frequencies (occurrences per 1000 posts) are shown. Keywords ordered by total frequency across all platforms.

	YouTube		Twitter/X		Telegram	
Keyword	*n*	*per 1000*	*n*	*per 1000*	*n*	*per 1000*
5G	394	17.3	3634	27.9	119 894	431.3
Radiation	180	7.9	723	5.6	32 783	117.9
Electric smog	239	10.5	504	3.9	7093	25.5
Mobile comms.	132	5.8	820	6.3	15 556	56.0
EMF	70	3.1	252	1.9	45 756	164.6
Corpus size (*n*)	22 710		130 156		277 986	

Beyond this shared pattern, the platforms diverged in revealing ways. On YouTube, 'Elektrosmog' [electric smog] was the second most frequent term (239), followed by 'Strahlung' [radiation] (180) and 'Mobilfunk' [mobile communications] (132), while the technical term 'EMF' appeared only 70 times. Twitter/X showed a similar ranking but with 'Mobilfunk' [mobile communications] (820) and 'Strahlung' [radiation] (723) more prominent than 'Elektrosmog' [electric smog] (504). 'EMF' remained comparatively rare (252). Telegram was the notable exception: 'EMF' was the second most frequent term (45.756), appearing far more often than on either of the other two platforms both in absolute terms and relative to corpus size. Specialised terms such as 'Elektrosensibilität' [electrosensitivity] and 'Hypersensibilität' [hypersensitivity] were marginal on all platforms, never exceeding 140 occurrences even in the largest corpus.

The prevalence of lay terms ('5G,' 'Elektrosmog' [electric smog], 'Strahlung' [radiation]) over technical terminology ('EMF') on YouTube and Twitter/X suggests that the public discourse is framed around consumer technology and perceived exposure rather than the scientific domain of EMF research. Telegram's higher use of 'EMF' may reflect its role as a hub for more committed and specialised actors.

3.1.3. Emotional register within the identified networks

Automated emotion classification reveals a broadly consistent baseline across all three platforms: approximately one third of all content was classified as neutral in tone (figure 3). Approval was the second most frequent category on every platform, ranging from 26.2% on Twitter/X to 33.2% on Telegram. The combined share of neutral and approving content thus accounted for roughly 60%–66% of all analysed posts, indicating that EMF-related content within the identified networks is predominantly composed rather than overtly inflammatory.

The platforms diverged most clearly in the distribution of negative emotions. Twitter/X exhibited the highest combined share of annoyance (11.2%) and disapproval (6.8%), consistent with the platform's more confrontational discussion culture. Telegram, by contrast, showed the lowest annoyance of any platform (3.5%) and virtually no disapproval (0.3%), while approval (32.2%) and realisation (17.4%) together accounted for nearly half of all classified content. Confusion was notably elevated on Telegram

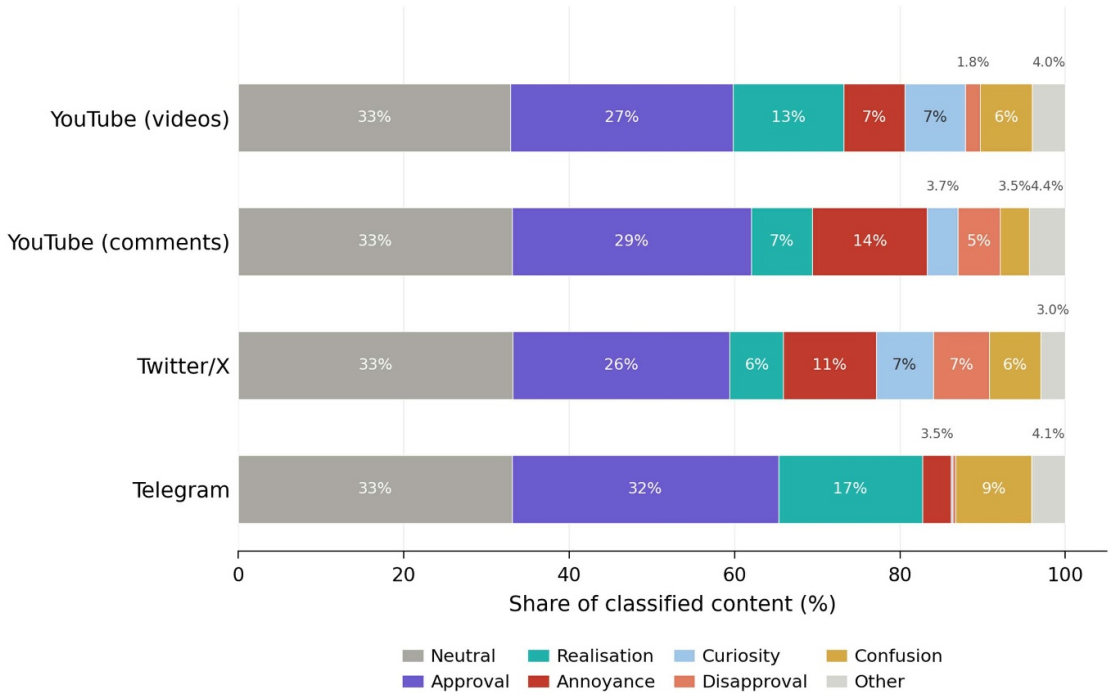

Figure 3. Distribution of classified emotion categories across platforms based on the GoEmotions taxonomy. YouTube is shown separately for video metadata and comments. Categories below 1.5% are not labelled.

(9.2%) compared to the other platforms, possibly reflecting the interpretive demands of the conspiracy-dense content environment rather than emotional opposition. This overall pattern is consistent with an echo chamber dynamic in which content circulates primarily among like-minded users and dissenting responses are rare.

YouTube provided a unique opportunity to compare creator-produced content with audience responses. While video titles and descriptions were predominantly neutral (32.9%) and approving (26.9%), with annoyance at only 7.4%, the associated comments showed markedly higher levels of annoyance (13.9%) and disapproval (5.1%). This difference indicates that comment sections do capture some degree of counter-speech and critical engagement, even where the videos themselves are framed in a measured tone.

3.1.4. Thematic embedding of EMF discourse

Topic modelling revealed a consistent pattern across all three platforms: EMF-related content did not form a distinct thematic cluster but was embedded within broader content domains. On YouTube, where 89 automatically identified topics were hierarchically reduced to six overarching themes, EMF content appeared primarily as a subset of clusters centred on alternative media sources and alternative medicine, while the single largest thematic cluster addressed political events (figure 4). On Twitter/X, the two largest topic clusters were both political in nature, with EMF not separately identifiable. A smaller cluster addressing COVID-19 and vaccination was present but occupied a notably smaller share than on other platforms. On Telegram, a parametric UMAP visualisation of the extended corpus likewise showed EMF topics as a minority that did not warrant separate categorisation, even though Telegram exhibited the highest density of EMF-related keywords of any platform. COVID-19 and vaccination content formed a recognisable but distinct cluster on all three platforms, adjacent to but not identical with the EMF discourse. The recurrent appearance of pandemic-related content throughout the data, as observed in the temporal dynamics (section 3.1.1), the thematic structure described in this section, and consequently in the qualitative coding categories, reflects the discursive environment in which EMF misinformation circulated during the collection period.

The thematic embedding pattern is reinforced by the actor-level structure of the networks. As noted in section 2.1, the highest-reach accounts in the corpus are content aggregators that cover a broad range of topics, while accounts dedicated exclusively to EMF have significantly lower reach. A further structural observation concerns cross-platform content flow: a substantial portion of the YouTube videos in the corpus were initially identified through links shared on Telegram channels and Twitter/X accounts,

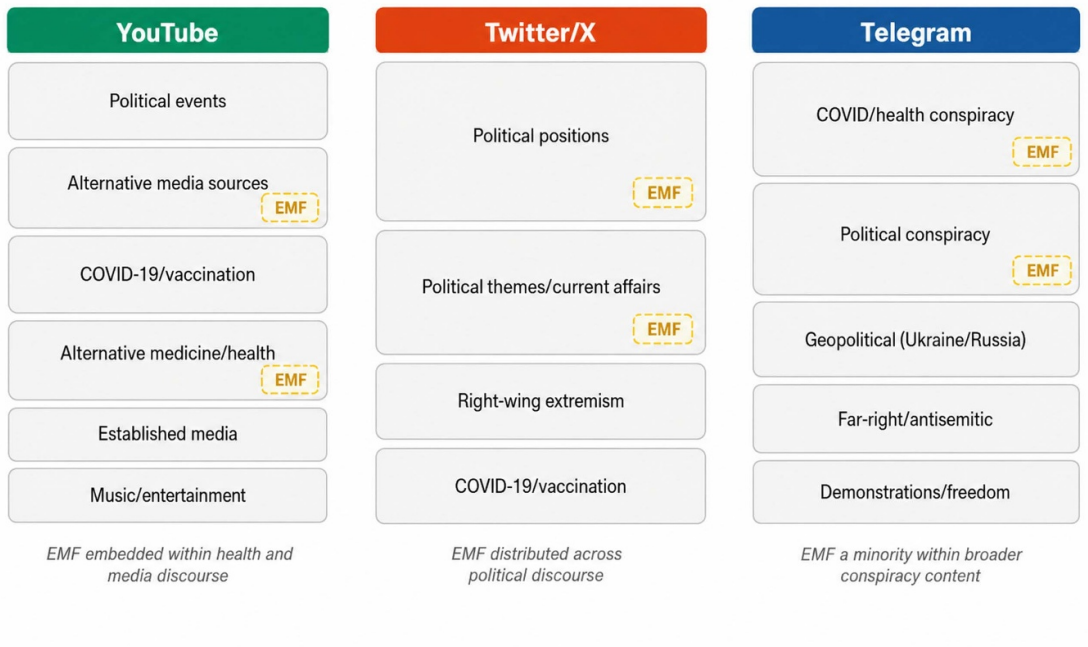

Figure 4. Major content domains per platform derived from topic modelling (Top2Vec for YouTube and Twitter/X, parametric UMAP for Telegram). EMF content (highlighted) is embedded within broader domains on all platforms. Relative box sizes are approximate.

suggesting that the platforms do not operate independently but form an interconnected distribution network in which content produced on one platform is disseminated through others.

3.2. Qualitative findings

The objective of the qualitative study was to conduct a comprehensive content analysis of a selection of misinformation found on YouTube, Twitter/X, and Telegram, drawing on a subset of the respective quantitative dataset. For YouTube, titles, descriptions, comments, and selected video transcripts were analysed, alongside posts from Twitter/X and Telegram. The text corpora were coded by two coders using a coding scheme, as indicated in the method section. Overall, a typical post within the identified networks on the platforms studied, references 5G or mobile infrastructure in connection with unspecified health risks, situates the topic within a broader narrative of institutional distrust or conspiracy, and frequently includes links to external resources or commercial products such as EMF shielding devices. The underlying discursive strategies frequently employ negative motivational appeals such as health fears and institutional distrust (Dressel *et al* 2026).

On YouTube, where the analysis also included user comments, a heterogeneous reception was perceived: while expressions of agreement with and gratitude for misinformation predominate, a notable proportion of comments articulate rejection and contradiction of the misinformation. However, further examination suggests that critical comments are frequently deprioritized by the platform's algorithm and thus positioned at the end of comment threads. This reduced their visibility and perceived relevance. In addition, video descriptions on YouTube are intensively used for advertising purposes.

In contrast, communication on Twitter/X is characterised by a stronger political orientation than on other platforms. Tweets commonly resemble newspaper headlines, are phrased in the third person, and present claims in a concise, factual manner; within this context, references to electric smog and electro (-hyper) sensitivity occur with higher frequency. Telegram, by comparison, exhibits the most pronounced influence of conspiracy theories and a clearly esoteric framing of 5G/EMF-related misinformation. Content on this platform often integrates broader conspiratorial worldviews, including a relatively high prevalence of references to QAnon-related narratives, indicating a communication environment in which such interpretative frameworks are not only present but actively reinforced. It must be noted that the described platform-specific differences in communication about 5G/EMF can be partly explained by the distinct target audiences and communicative conventions of each platform.

Beyond platform-specific differences, several similarities in communication related to 5G/EMF misinformation can be observed. Across platforms, references predominantly address general effects of 5G

and EMF, especially with regard to health and, to a lesser extent, environmental impacts. Concrete health effects are also frequently mentioned, most notably cancer and sleep disturbances, while electro (-hyper) sensitivity appears more often within the EMF-related thematic (not 5G) cluster. In contrast to these largely vague risk descriptions, organisations and initiatives opposing 5G/EMF are often named explicitly, typically combined with calls for support; thus, potential harms are framed broadly and as widely relevant, whereas actors and mobilisation efforts are specified. Notably, many keywords listed in the EMF codebook did not appear in the analysed texts at all, and references to specific scientific studies were extremely rare. Similarly, many organisations and initiatives included in the codebook were absent from the material. A further cross-platform pattern concerns the use of hashtags. Posts were frequently tagged with '5 G' or related EMF terms without substantively engaging with EMF topics in the text itself, effectively associating EMF concerns with unrelated content through metadata alone.

Across all platforms, discussions of 5G/EMF are closely intertwined with other emotionally discussed topics such as COVID-19, politics, and esotericism, and conspiracy-theoretical elements are ubiquitous. This thematic clustering suggests that rational, evidence-based engagement with 5G/EMF on social media is particularly challenging, as it would require not only addressing knowledge gaps but also confronting established belief systems. The qualitative analysis further confirms the recurring use of multiple discursive strategies within single messages, most frequently those aimed at attracting attention and prompting action beyond the platform (such as contacting representatives or attending demonstrations), followed by the selective provision of information aligned with the communicator's own interpretation and negative motivation (such as appeals to health risks or institutional distrust). Particularly within the EMF cluster (not so much the strict 5G discourse), commercial interests are salient, ranging from the promotion of protective products such as shielding devices and wearables marketed against alleged radiation harm, to the advertisement of services such as esoteric consultations. For a subset of accounts, this commercial activity appeared central rather than incidental to their operation, with EMF-related fear-based appeals functioning as the framing for product promotion. The prevalence of this dimension varied by platform: commercial content was coded most frequently in the YouTube samples, followed by Twitter/X, with Telegram showing the lowest share. Across all platforms, however, commercial interest ranked among the most frequently coded motivations, appearing even more often than appeals to share or forward the content itself. The appearance of expertise is likewise a recurring feature, especially on YouTube and Twitter/X where academic or journalistic titles are employed to signal credibility and facilitate engagement with misleading content.

Questions and follow-up inquiries occur more frequently in relation to 5G than to EMF, indicating a potential point of entry for future science communication efforts. Finally, conspiracy theories related to 5G/EMF are present on all analysed platforms. Overall, 5G emerges as a particularly polarising keyword: as an abstract and non-self-explanatory term embedded in a highly complex scientific context, it easily lends itself to speculation.

4. Discussion

Communicative register and the structure of credibility. The combination of quantitative and qualitative analysis reveals a misinformation landscape that defies several common assumptions about how false information operates on social media. Automated sentiment classification showed that across all three platforms, misinformation content was predominantly neutral to approving in tone, with explicitly negative emotions accounting for less than 20% of classified posts even on Twitter/X, the most confrontational platform studied.

The qualitative analysis, however, shows that this composed surface coexists with underlying appeals to fear, health anxiety, and institutional distrust. The discursive strategy 'attention-grabbing and calls to action' is most prevalent, followed by 'informative framing' and 'negative motivation.' These strategies are frequently combined within single messages, yet delivered in a register that the automated classifier consistently categorises as neutral or approving. A key enabling mechanism is the use of an 'appearance of professionalism': academic titles, journalistic headline-style framing, and references to real organisations or institutions function as credibility cues that produce a controlled and authoritative communicative register. The result is content that is often structurally difficult to distinguish from legitimate science communication, which may help explain its persistence and reach within the networks studied. This finding also suggests that detection approaches relying on emotional extremity as a signal will systematically underdetect this type of content.

Structural embedding of EMF discourse. The topic modelling consistently showed that EMF-related content does not form a distinct thematic cluster on any of the three platforms. On YouTube, it appears as a subset of alternative media and alternative medicine clusters; on Twitter/X, it is embedded within political content; on Telegram, it constitutes a minority even within the platform's conspiracy-oriented landscape. At the actor level, the highest-reach accounts in the network are content aggregators covering a broad range of topics, while accounts dedicated exclusively to EMF have significantly lower reach. EMF misinformation thus gains visibility not through dedicated advocacy but by being woven into narratives that command wider attention, such as COVID-19 scepticism, anti-establishment politics, and alternative health movements. The qualitative analysis reinforces this pattern: 5G/EMF-related misinformation consistently co-occurs with COVID-19, political, and esoteric themes, and logically incompatible explanatory frameworks, such as esoteric accounts attributing EMF harm to disrupted 'natural frequencies' alongside political accounts framing 5G as a tool of state surveillance, coexist without mutual contradiction within the same networks, suggesting a shared oppositional identity that transcends the specific claims being made.

This embedding pattern is consistent with findings from English-language studies of the COVID-5G conspiracy (Ahmed *et al* 2020, Bruns *et al* 2020), but the present study demonstrates that it extends beyond the pandemic moment and persists in the post-COVID period across multiple platforms. It also carries methodological implications: tracking EMF-specific keywords alone would miss the primary channels through which EMF content reaches audiences.

Platform-specific dynamics. The three platforms studied exhibit distinct communication dynamics in the dissemination of 5G/EMF-related misinformation. Twitter/X is characterised by a pseudo-objective communicative register, with third-person formulations and headline-like structures that convey journalistic neutrality despite speculative content. Telegram exhibits a dense concentration of conspiracy-theoretical and esoteric narratives. YouTube displays a split between creator-generated content that largely drives misinformation narratives and audience responses in the comments that include both endorsement and comparatively higher levels of disagreement—though critical counter-speech appears to be deprioritized by algorithmic comment ranking, reducing its visibility. A further pattern observed across platforms was topical co-optation through hashtag use: posts were frequently tagged with '5 G' or related terms without substantively discussing EMF topics, effectively associating EMF concerns with unrelated content through metadata alone.

Across all three platforms, misinformation was systematically entangled with commercial practices, including advertising, product placement, and the promotion of EMF shielding devices, esoteric consultations, and health-related products (as reported in section 3.2). This suggests that 5G/EMF misinformation is not solely ideologically motivated but also sustained by commercial incentive structures that operate alongside and, in some cases, independently of the content's truth claims.

5. Limitations

Several methodological limitations should be considered when interpreting the findings. The seed-user sampling approach provides systematic coverage of the identified misinformation network but does not constitute a random sample of general platform content. Misinformation rates should be understood as densities within the mapped network, not as prevalence estimates for the platforms at large. The automated classifier was validated against the independent manual content analysis for Telegram and Twitter/X, empirically confirming that the reported rates constitute upper-bound estimates within the mapped networks (section 2.2, table 1). Item-level agreement between the two methods was nonetheless moderate, and for YouTube, where classification was limited to textual metadata, no such validation was feasible. The Twitter/X dataset is constrained by the removal of academic API access in early 2023, skewing the corpus toward more recent content and limiting temporal comparability with the other platforms. The qualitative content analysis captures discursive patterns but does not permit causal inference. While double coding was employed, the method inherently involves interpretive judgement. The algorithmic mechanisms that appear to shape content visibility—most notably the deprioritisation of critical comments on YouTube—were inferred from observed ranking patterns but could not be measured directly. Finally, the study covers German-language content on three platforms during a specific observation period and may not generalise to other linguistic or platform contexts.

6. Conclusion

This study presents a systematic, mixed-methods mapping of EMF and 5G-related misinformation across three major social media platforms in the German-language context. Five principal findings emerge from the analysis.

Misinformation content in the corpus does not primarily rely on overtly inflammatory language. It layers emotional appeals beneath a surface of pseudo-expertise: academic titles, journalistic framing, and references to real institutions establish credibility, while the underlying narratives still invoke fear, health threat, and institutional distrust. The emotional charge is present but packaged in a register that resembles authoritative communication. This distinguishes the EMF misinformation landscape from domains where emotional extremity is itself the dominant signal and suggests that the relationship between emotional tone and misinformation is more structurally complex than a simple association with negative or inflammatory content.

EMF topics do not circulate as standalone content but are embedded within broader conspiracy ecosystems. Topic modelling and actor-level analysis consistently show that EMF content is woven into larger narratives encompassing COVID-19 scepticism, political anti-establishment movements, and alternative health. The highest-reach accounts are not EMF specialists, but content aggregators who periodically incorporate EMF themes alongside adjacent material. EMF misinformation thus gains visibility through association rather than dedicated advocacy—a pattern that persists beyond the pandemic period and across all three platforms studied.

The pandemic-era peak in EMF misinformation has subsided, but the network infrastructure that sustained it remains intact. While the volume of EMF-related content has declined since 2020–2022, the channels, accounts, and cross-platform linkages that carried it persist. This latent infrastructure means the ecosystem is not rebuilding from scratch but dormant, capable of rapid reactivation by future triggering events such as new wireless technology deployments or public health controversies.

The three platforms serve complementary functional roles within an integrated misinformation ecosystem. Telegram operates as a major community infrastructure, sustaining insular networks with minimal internal dissent. Twitter/X functions as the interface to broader political discourse, where EMF claims are reframed in pseudo-journalistic registers and connected to mainstream debates. YouTube serves as the content production layer, generating long-form material that is subsequently distributed through links on the other platforms. These distinct roles suggest that EMF misinformation operates not as parallel activity across platforms but as a cross-platform system with a functional division of labour.

Commercial interests actively sustain the misinformation ecosystem. EMF shielding products, esoteric services, and health supplements are marketed within misinformation content across all three platforms. The entanglement of commercial and ideological incentives suggests that the persistence of EMF misinformation cannot be explained by belief alone but is also sustained by economic structures that benefit from continued public uncertainty.

The data analysed predate the mainstream availability of generative AI tools. The emergence of large language models capable of producing fluent, pseudo-expert text at scale may substantially alter the production dynamics of misinformation by lowering the barrier to content creation and making the authoritative register identified in this study even easier to replicate. The present findings thus establish a baseline against which future developments in AI-assisted misinformation can be measured. Intervention-oriented recommendations drawing on these findings and subsequent project phases are developed in Dressel *et al* (2026).

Acknowledgment

The data collection and initial analysis underlying this study were conducted as part of project 3621 EMF 106, commissioned by the German Federal Office for Radiation Protection (Bundesamt für Strahlenschutz, BfS). The views and conclusions expressed in this paper are those of the authors and do not necessarily reflect the position of the BfS. We would like to thank Dr Corinna Lüthje of the BfS for her support and constructive feedback throughout the project.

Data availability statement

The data cannot be made publicly available upon publication because they contain sensitive personal information. The data that support the findings of this study are available upon reasonable request from the authors.

Ethical statement

This study analyses publicly available content posted on social media platforms (YouTube, Telegram, and Twitter/X). No human participants were recruited, contacted, or surveyed as part of this research. The study did not require approval from an institutional research ethics committee as it exclusively analyses publicly accessible online content at aggregate level, with no individual users identified in the published findings. The research was conducted in compliance with the European General Data Protection Regulation (GDPR/DSGVO). All data were processed and stored in accordance with the data governance provisions of the commissioning research project.

ORCID iDs

Tobias Eder 0009-0007-3377-655X
Katja Aue 0000-0002-2931-456X
Larissa Drescher 0000-0002-0227-8407
Kerstin Dressel 0000-0002-2885-0846
Georg Groh 0000-0002-5942-2297
Jutta Roosen 0000-0002-3032-399X

References

Ahmed W, Vidal-alaball J, Downing J and López Seguí F 2020 COVID-19 and the 5G conspiracy theory: social network analysis of Twitter data *J. Med. Internet Res.* **22** e19458

Altay S, Berriche M and Acerbi A 2023 Misinformation on misinformation: conceptual and methodological challenges *Soc. Media Soc.* **9** 20563051221150412

Angelov D 2020 Top2Vec: distributed representations of topics (arXiv:2008.09470)

Baumgartner J, Zannettou S, Squire M and Blackburn J 2020 The pushshift telegram dataset *Proc. Int. AAAI Conf. On Web and Social Media* vol 14 pp 840–7

Bruns A, Harrington S and Hurcombe E 2020 'Corona? 5G? or both?': the dynamics of COVID-19/5G conspiracy theories on Facebook *Media Int. Aust.* **177** 12–29

Cinelli M, Etta G, Avalle M, Quattrociocchi A, Di Marco N, Valensise C, Galeazzi A and Quattrociocchi W 2022 Conspiracy theories and social media platforms *Curr. Opin. Psychol.* **47** 101407

Demszky D, Movshovitz-attias D, Ko J, Cowen A, Nemande G and Ravi S 2020 GoEmotions: a dataset of fine-grained emotions *Proc. 58th Annual Meeting of the Association of Computational Linguistics* pp 4040–50

Drescher L, Aue K, Eder T, Groh G, Dressel K and Roosen J 2025 Wie lassen sich Des- und Fehlinformation zu EMF-Themen in sozialen Medien am effektivsten identifizieren und entkräften? [How can disinformation and misinformation on EMF topics in social media be most effectively identified and countered?]. Ressortforschungsberichte zum Strahlenschutz, Vorhaben 3621EMF106. Bundesamt für Strahlenschutz (BfS) (available at: https://doris.bfs.de/jspui/bitstream/urn:nbn:de:0221-2025122258057/1/BfS-RESFOR-254-25.pdf)

Dressel K, Drescher L, Aue K, Eder T, Groh G and Roosen J 2026 More than debunking misinformation: a holistic communication model for public authorities. The case of electromagnetic fields and 5G in Germany *Research Handbook on Public Communication* ed J Matusitz (Edward Elgar Publishing) pp 218–36

Eggeling-böcker M, Karabetsos E, Christopoulou M, Link S C, Abacioglu F and Boehmert C 2026 RF-EMF risk perception and trust in radiation protection authorities: a comparative study on precautionary information in Germany and Greece *Bioelectromagnetics* **47** e70042

Freudenstein F, Wiedemann P M and Brown T W C 2015 Exposure perception as a key indicator of risk perception and acceptance of sources of radio frequency electromagnetic fields *J. Environ. Public Health* **2015** 198272

Herasimenka A, Bright J, Knuutila A and Howard P N 2022 Misinformation and professional news on largely unmoderated platforms: the case of Telegram *J. Inf. Technol. Polit.* **20** 198–212

Jolley D and Paterson J L 2020 Pylons ablaze: examining the role of 5G COVID-19 conspiracy beliefs and support for violence *Br. J. Soc. Psychol.* **59** 628–40

Koh T H, Choi J W, Seo M, Choi H and Kim K 2020 Factors affecting risk perception of electromagnetic waves from 5G network base stations *Bioelectromagnetics* **41** 491–9

Langguth J, Filkuková P, Brenner S, Schroeder D T and Pogorelov K 2022 COVID-19 and 5G conspiracy theories: long term observation of a digital wildfire *Int. J. Data Sci. Anal.* **15** 329–46

Mayring P 2015 *Qualitative Inhaltsanalyse: Grundlagen und Techniken* 12th edn (Beltz Verlagsgruppe)
Nan X, Thier K and Wang Y 2023 Health misinformation: what it is, why people believe it, how to counter it *Ann. Int. Commun. Assoc.* **47** 381–410
Peter V, Küh R, Mitrović J, Granitzer M and Schmid-petri H 2022 Network analysis of German COVID-19 related discussions on telegram *Int. Conf. on Applications of Natural Language to Information Systems* pp 26–33
Ricaud B, Aspert N and Miz V 2020 Spikyball sampling: exploring large networks via an inhomogeneous filtered diffusion *Algorithms* **13** 275
Rowley J and Mazar H 2021 Misunderstandings about radiofrequency electromagnetic field exposure and 5G misinformation *2021 IEEE Int. Conf. on Microwaves, Antennas, Communications and Electronic Systems (COMCAS)* pp 219–21
University of Minnesota Libraries Publishing 2016 Communication in the real world: an introduction to communication studies (University of Minnesota Libraries Publishing) (available at: https://open.lib.umn.edu/communication/)
Vosoughi S, Roy D and Aral S 2018 The spread of true and false news online *Science* **359** 1146–51
Wardle C and Derakhshan H 2017 *Information Disorder: Toward an Interdisciplinary Framework for Research and Policymaking* (Council of Europe) pp 162317GBR
Wiedemann P M and Schütz H 2005 The precautionary principle and risk perception: experimental studies in the EMF area *Environ. Health Perspect.* **113** 402–5