NICOLA PELIZZARI

Pseudo-authority and Health (Mis)information on TikTok. A Corpus-assisted Multimodal Discourse Analysis

Abstract: Health (mis)information increasingly circulates through synthetic agents on digital platforms, where fabrication, fictionality, and vagueness destabilise conventional markers of expertise. TikTok, often portrayed as a site of "democratised" creativity, is structured by algorithmic logics that amplify particular forms of visibility and shape textual practices. Within this environment, AI-generated videos function as multimodal texts that simulate authority through converging linguistic, visual, aural, and identity cues. This article applies corpus-assisted multimodal discourse analysis (CAMDA) to a stratified corpus of 120 TikTok health (mis)information videos (2024–2025), comprising equal subcorpora of declared-AI, probable-AI, and human-produced videos. For each video, transcripts, captions, overlays, and metadata are examined to trace patterned uses of evidential markers, vague quantifiers, empathy formulas, synthetic voices, credential imagery, and platform-specific identity signals. The analysis shows how *pseudos* is enacted through TikTok's algorithmic circulation, casting AI-generated health content as a distinctive mode of AI-mediated authority with implications for public health communication.

Keywords: *pseudos*, *multimodal discourse*, *corpus-assisted analysis*, *TikTok*, *AI-generated health personae*, *health misinformation*

1. Introduction

Contemporary digital environments have transformed the circulation of health information, generating semiotic conditions in which the boundaries between expertise, simulation and promotion are increasingly difficult to discern. TikTok, now one of the most widely used social platforms globally,[1] plays a central role in this shift: its short-form video ecology promotes content that is often visually curated, affectively engaging and optimised for virality. Given this scale and level of engagement, these platforms could in principle function as an effective channel for disseminating accessible health information. However, empirical research paints a more complex picture, indicating that a substantial proportion of health-related videos contain misleading[2] or unverified claims.[3] These findings reveal an ecosystem in which credibility is conferred not through institutional validation but through stylistic legibility and platform dynamics that frequently encourage overuse while downplaying associated harms.[4] Such dynamics intersect with unequal levels of health and digital health literacy, shaping who

[1]Marisa Dellatto, "TikTok Hits 1 Billion Monthly Active Users", *Forbes* (2021), https://www.forbes.com/sites/marisadellatto/2021/09/27/tiktok-hits-1-billion-monthly-active-users/.

[2] Rose Dimitroyannis et al., "A Social Media Quality Review of Popular Sinusitis Videos on TikTok", *Otolaryngology–Head and Neck Surgery*, 170.5 (2024), 1456-1466, https://doi.org/10.1002/ohn.688.

[3] Macklin Loveland, "Analysis of Liver Disease Misinformation and Accurate Information within the Social Media Platform, TikTok", *Gastroenterology*, 164.6 (2023), S1400.

[4] Emma Grundtvig Gram et al., "Addressing Misleading Medical Information on Social Media: A Scoping Review of Current Interventions", *BMJ Evidence-Based Medicine* (advance online publication, 2025), https://doi.org/10.1136/bmjebm-2025-113704; see also Anna Gaysynsky et al., "Perceptions of Health Misinformation on Social Media: Cross-Sectional Survey Study", *JMIR Infodemiology*, 4 (2024), e51127, https://doi.org/10.2196/51127.

can critically appraise such content and who remains vulnerable to persuasive but unvetted health advice.[5]

Within this environment, AI-generated health content introduces an additional layer of complexity. Such material often features AI-generated health personae: digitally produced speakers, narrators or avatars that deliver health-related explanations, blurring the line between authoritative presentation and the institutional settings in which medical communication ordinarily occurs. In these videos, authority is not tied to professional identity but is constructed through semiotic resources that index medical expertise, drawing on features such as anatomical or procedural imagery, diagnostic-style overlays, steady synthetic prosody modelled on clinical explanation, and identity-signalling devices including white coats and stethoscopes.

Crucially, such material spans a continuum, from explicitly declared synthetic videos to content that is probably AI-generated yet often highly lifelike in its delivery, owing to the increasingly sophisticated mimetic capacities of contemporary generative systems.[6] This variability raises a central analytical issue for the present study: how to conceptualise performances that adopt the stylistic features of medical expertise when the expertise they project cannot be authenticated through conventional professional or experiential anchors?

The notion of *pseudos* offers a productive framework for addressing this question. While historically associated with deception, *pseudos* has long denoted a broader domain of epistemic indeterminacy in which discourse operates between fact, fiction and persuasion.[7] In classical philosophy, Plato's account of the *γενναῖον ψεῦδος* (noble lie) framed certain forms of fabrication as socially functional rather than merely incorrect. Aristotle's *Rhetoric* conceptualised persuasion as inherently selective and strategically shaped, a view that resonates with contemporary analyses of how discourse adopts the conventions of authoritative genres.

Modern philosophical work expands this instability, emphasising how authoritative effects may be produced without stable ties to intention, institutional accountability or shared epistemic norms.[8] Pragmatic research further shows that persuasive force frequently emerges not through explicit falsehood but through the manipulation of conversational expectations. Drawing on Grice's account of implicature, speakers may exploit the maxims of conversation, such as relevance, quantity or clarity, to generate implied meanings that project certainty or evidential grounding while remaining unstated and therefore resistant to verifiable propositions.[9]

Studies of so-called 'pseudo-media' extend these insights to digital contexts, demonstrating how content can reproduce the stylistic and organisational features of institutional communication while selectively departing from its norms.[10] This parallel illuminates how pseudo-authority can also emerge in health-related content disseminated by synthetic agents. Together, these strands suggest that AI-

---

[5] Sheila M. do Nascimento et al., "Infodemic and Information-Seeking Behaviour During the COVID-19 Pandemic: A Cross-Sectional Study", *BMJ Open*, 12 (2022), e054074, https://doi.org/10.1136/bmjopen-2021-054074.

[6] Yuxi Wang et al., "Systematic Literature Review on the Spread of Health-Related Misinformation on Social Media", *Social Science & Medicine*, 240 (2019), 112552, https://doi.org/10.1016/j.socscimed.2019.112552.

[7] Stephen Halliwell, *Between Ecstasy and Truth: Interpretations of Greek Poetics from Homer to Longinus* (Oxford: Oxford U.P., 2011).

[8] Sanford C. Goldberg, "Getting Told and Being Believed", *Philosophers' Imprint*, 18.10 (2018), 1-20.

[9] H. Paul Grice, "Logic and Conversation", in Peter Cole and Jerry L. Morgan, eds., *Syntax and Semantics 3: Speech Acts* (New York: Academic Press, 1975), 41-58.

[10] Barbara Palau-Sampio, "Pseudo-Media: Disinformation Patterns, Polarised Discourse, Clickbait and Twisted Journalistic Mimicry", *Journalism Practice* (2023), https://doi.org/10.1080/17512786.2023.2269355.

generated health content should be understood not as a series of discrete inaccuracies but as part of a broader field of multimodal textual practices shaped by platform logics.[11] Meaning in short-form video arises through the interplay of linguistic, visual and aural modes embedded in algorithmically governed circulation. Research in multimodality shows that credibility is jointly produced through gaze, gesture, framing, typography, sound and identity cues. Work in corpus-assisted multimodal discourse analysis (CAMDA) and discursive news values analysis demonstrates how evaluation and value are distributed across modes,[12] offering analytical tools that are directly relevant to platform-native, attention-driven health videos.

Accordingly, this study applies CAMDA to a stratified corpus of TikTok health (mis)information videos. CAMDA integrates corpus linguistic techniques with multimodal analysis to identify recurrent linguistic strategies and to examine how these intersect with visual, aural and identity cues across videos.[13] Through this combined approach, the study directly addresses its two research questions:

(1) *How do AI-generated TikTok videos simulate medical authority through patterned combinations of linguistic, visual, aural and identity cues?*

(2) *How do fabrication, fictionality and vagueness operate as multimodal textual practices within TikTok's algorithmically shaped circulation of health (mis)information?*

The overarching aim of the study is to shed light on emerging genres influenced by generative technologies and to contribute to public health communication by elucidating how synthetic agents perform authority within a platform environment where signals of credibility are increasingly decoupled from professional identity. More broadly, the analysis positions AI-generated TikTok videos as contemporary manifestations of *pseudos*, showing how platform-mediated semiotic practices render the boundaries between authorship, authenticity and persuasion newly permeable. Building on the conceptual framework outlined above, the three dimensions of *pseudos* foregrounded in the second research question are operationalised as follows. *Fabrication* is understood as the semiotic production of expertise markers, generic evidential frames, credential imagery, and clinical-style visual organisation, that simulate institutional grounding without traceable professional anchoring. *Fictionality* is approached as the adoption of the formal conventions of clinical explanation and health education, diagnostic overlays, procedural sequencing, direct-address delivery, outside the institutional contexts that conventionally sustain them. *Vagueness* is treated as a systematic linguistic and pragmatic resource, under-specified quantifiers, non-attributed evidentials, and high-certainty stance expressions, that extends the apparent scope of claims while limiting their verifiability. These three dimensions, and their multimodal realisations, constitute the analytical lens applied to the corpus.

---

[11] Helen Caple et al., *Multimodal News Analysis Across Cultures: Corpus-Assisted Discourse Analysis of News Values* (Cambridge: Cambridge U.P., 2020).

[12] Monika Bednarek, "Corpus-Assisted Multimodal Discourse Analysis of Television and Film Narratives", in Paul Baker and Tony McEnery, eds., *Corpora and Discourse Studies: Integrating Discourse and Corpora* (Basingstoke: Palgrave Macmillan, 2015), 63-87.

[13] Gunther Kress and Theo van Leeuwen, *Reading Images: The Grammar of Visual Design*, Second Edition (London: Routledge, 2006).

2. Methods

This study applies CAMDA to a stratified corpus of 120 TikTok health (mis)information videos collected between September 2024 and August 2025. The analytic procedure is organised into four stages: sampling and corpus construction, data retrieval and transcription, multimodal annotation, and integrated corpus-assisted analysis, which are described below.

2.1 *Sampling and corpus construction*

Sampling followed a stratified, purposive design, in which strata were defined analytically by content type rather than by population characteristics. Three subcorpora of equal size were constructed: declared-AI (n = 40), probable-AI (n = 40), and human-produced videos (n = 40). Videos were retrieved using TikTok's internal search function and discovery feed through combinations of health-related hashtags, condition-specific keywords, and medication names with high platform visibility (e.g., #health, #healthtips, #healthadvice, #medications, #disease, #cure). Data collection was conducted in compliance with applicable European Union data protection and research ethics frameworks, including regulation (EU) 2016/679 (General Data Protection Regulation). The study relied exclusively on publicly available content, involved no interaction with users, no collection of personal or special-category data, and no attempts at user identification or profiling. Data access and processing followed a predefined and formally approved research protocol and were carried out using authorised access to TikTok's Research platform (TikTok protocol ID: 1723718073).

To minimise algorithmic personalisation, data collection was conducted from a purposely created account with no interaction history, neutral language and location settings, and no prior engagement with health-related content; no videos were liked, shared, commented on, or followed during sampling. Within each stratum, videos were manually screened to ensure topical relevance and compliance with the classification criteria defined below. Videos were also distributed across multiple creator accounts; to prevent any single profile from shaping the dataset disproportionately, no creator was allowed to contribute more than 5% of the corpus (≤ 2 videos).

Equal stratum sizes were selected to enable systematic cross-category comparison of multimodal and discursive features, rather than to approximate the relative prevalence of content types on the platform, following established practices in social-media discourse research.[14]

Declared-AI videos were identified through explicit creator disclosure or platform labelling stating that the content was AI-generated. Probable-AI videos were classified by the author on the basis of multimodal indicators commonly associated with current generative systems, in the absence of verified production provenance. Four indicator types were applied across two semiotic modes. Within the aural mode, synthetic vocal timbre was identified through uniform pitch contour, unnaturally steady pacing, and absence of disfluencies or hesitation phenomena. Within the visual mode, three indicator types were considered: uniform facial rendering, including overly smooth skin texture, limited micro-expression variation, and symmetrical lighting not attributable to studio conditions; non-naturalistic

[14] Corey H. Basch et al., "A Global Pandemic in the Time of Viral Memes: COVID-19 Vaccine Misinformation and Disinformation on TikTok", *Human Vaccines & Immunotherapeutics*, 17.8 (2021), 2373-2377, https://doi.org/10.1080/21645515.2021.1894896; see also Ciera E. Kirkpatrick and Lauren L. Lawrie, "TikTok as a Source of Health Information and Misinformation for Young Women in the United States: Survey Study", *JMIR Infodemiology*, 4 (2024), e54663, https://doi.org/10.2196/54663.

articulatory timing, where lip movement appeared loosely synchronised with speech rather than phonetically driven; and template-like visual design, involving standardised backgrounds, repetitive overlay formats, and consistent colour schemes across otherwise unrelated topics. A video was classified as probable-AI when at least two indicator types converging across at least two semiotic modes were present. This co-occurrence requirement is consistent with the relational interpretive approach applied throughout the multimodal annotation (section 2.3), in which semiotic features are assessed in combination rather than in isolation. Since verified production metadata is not available for the majority of TikTok content, the classification necessarily reflects the author's informed judgement; the scope and implications of this limitation are discussed in section 4.1.

Human-produced videos were included to provide a comparative baseline for linguistic and multimodal patterning and were matched to the AI datasets by topic, approximate duration, and communicative function. All included videos met four inclusion criteria: relevance to health or health-related advice; all linguistic content (spoken and on-screen text) in English; public accessibility at the time of data collection; and absence of platform warnings or restrictions limiting research use. Algorithmic recommendation and circulation were treated as contextual conditions shaping discourse production rather than as variables measured directly, consistent with CAMDA approaches to platformed media.

2.2 *Data retrieval and transcription*

Each video was locally downloaded using TikTok's download function to ensure stability of the analytical material over time. Downloaded files were used exclusively for research purposes within the scope of the approved protocol. Speech was transcribed verbatim using OpenAI's Whisper automatic speech recognition system,[15] followed by full manual verification and correction against the original audio. Transcriptions retained disfluencies, hesitations, and paralinguistic markers when analytically relevant to stance or evidentiality. Automatically generated platform captions were not used as a transcription source and were considered only as an additional multimodal artefact, given their limited reliability as representations of spoken content. On-screen text, diagnostic-style overlays, embedded citations, and identity-signalling elements, including professional labels (e.g. "doctor", "MD", "nurse") and symbolic attire associated with medical authority (e.g. white coats, scrubs, stethoscopes), were transcribed separately in accordance with established multimodal transcription conventions.[16]

2.3 *Multimodal annotation*

Multimodal features were annotated across four domains: linguistic, visual, aural, and identity-related resources. The annotation scheme draws on established frameworks in multimodal discourse analysis

[15] OpenAI, "Whisper: Robust Speech Recognition via Large-Scale Weak Supervision" (2022), https://openai.com/research/whisper.
[16] Carey Jewitt et al., *Introducing Multimodality* (London and New York: Routledge, 2016).

and social semiotics, combining visual grammar approaches,[17] principles of multimodal interaction,[18] and multimodal frameworks developed in health communication research.[19]

Linguistic annotation included identification of evidential markers, vague quantifiers, epistemic stance expressions, formulaic empathy constructions, and metaphor types, informed by previous corpus-based studies of health discourse and stance. Visual annotation included camera framing, angle, gaze configuration, presence of medical props, diagrams, or anatomical imagery, and the organisation of diagnostic-style overlays. Aural annotation included background music type, presence of synthetic voice modelling, intonational contour approximating clinical explanation, and sound effects used for emphasis. Identity cues included self-declared expertise, credential display, verification marks, bio descriptions, and pseudo-medical attire. This multimodal tagging was applied consistently across all videos. Annotations were interpreted relationally by analysing how multiple cues across modes combined within individual videos, rather than by evaluating single features in isolation.

2.4 *Corpus-assisted analysis*

Corpus analysis was carried out using AntConc 4.2.[20] All transcribed verbal data were compiled into three subcorpora corresponding to the sampling strata. Wordlists, frequencies and keyword analyses were generated to identify salient lexical patterns contrastively across the three study subcorpora, with each subcorpus used in turn as the reference. Keyword analysis functioned as a corpus-assisted heuristic for the identification of lexically salient patterns guiding qualitative and multimodal interpretation, rather than as a method for statistical inference. Collocation and concordance analyses were used to examine how evidential markers, stance expressions, and vague quantifiers were distributed and how they patterned with other linguistic or multimodal features. Multimodal findings were integrated with corpus outputs through iterative comparison: linguistic patterns identified via keyword or collocation analysis were examined in context alongside corresponding visual, aural and identity cues. For example, collocates of evidential verbs were reviewed together with accompanying on-screen overlays and credential imagery. The same procedure was applied to vague quantifiers, empathy formulations, and synthetic-voice sequences. This integrative analytic cycle follows CAMDA practice as established in prior multimodal corpus research to address the research questions articulated above.[21]

3. Findings

The semiotic patterns analysed in this study operate within specific conditions of duration, circulation, and topical focus that differ across subcorpora and shape the communicative environment in which

---

[17] Gunther Kress and Theo van Leeuwen, *Reading Images: The Grammar of Visual Design*, Third Edition (London and New York: Routledge, 2020), https://doi.org/10.4324/9781003099857.

[18] Elisabetta Adami, "Multimodality", in Ofelia García et al., eds., *The Oxford Handbook of Language and Society* (Oxford: Oxford U.P., 2017), https://doi.org/10.1093/oxfordhb/9780190212896.013.23.

[19] Jeff Bezemer, "Multimodality: A Guide for Linguists", in Robert J. Podesva and Devyani Sharma, eds., *Research Methods in Linguistics*, Second Edition (London: Bloomsbury Academic, 2018), 281-304, https://doi.org/10.5040/9781350043466.ch-012.

[20] Laurence Anthony, *AntConc*, Version 4.2 (Tokyo: Waseda University, 2023), https://www.laurenceanthony.net/software/antconc.

[21] Monika Bednarek and Helen Caple, *The Discourse of News Values: How News Organizations Create Newsworthiness* (New York: Oxford U.P., 2017), https://doi.org/10.1093/acprof:oso/9780190653934.001.0001.

authority is performed. Across the corpus, 85 out of 120 videos (70.8%) have a duration below 90 seconds, while 35 videos (29.2%) extend beyond this threshold and only 11 videos (9.2%) exceed two minutes. This pattern is consistent across subcorpora, indicating strong convergence toward short-form delivery irrespective of production type. In the declared-AI subset, 29 out of 40 videos (72.5%) fall below 90 seconds; a similar proportion is observed for probable-AI videos (30/40, 75.0%), while human-produced videos show a slightly lower share of short formats (26/40, 65.0%) and a correspondingly higher presence of longer videos. Videos exceeding 120 seconds remain marginal in all three groups, though they are somewhat more frequent among human-produced content (5/40, 12.5%) than among declared-AI and probable-AI videos (3/40, 7.5% in both cases). Full duration distributions by subcorpus are reported in Table 1.

Subcorpus	**Duration (in seconds)**			
	30–59 s	**60–89 s**	**90–119 s**	**120–150 s**
Declared-AI	15 (37.5%)	14 (35.0%)	8 (20.0%)	3 (7.5%)
Probable-AI	17 (42.5%)	13 (32.5%)	7 (17.5%)	3 (7.5%)
Human-produced	14 (35.0%)	12 (30.0%)	9 (22.5%)	5 (12.5%)
Total	**46 (38.3%)**	**39 (32.5%)**	**24 (20.0%)**	**11 (9.2%)**

Table 1. Distribution of video duration across subcorpora (N = 120)

Temporal distribution of posting dates across the observation window (September 2024 to August 2025) shows differentiated timing across the three subcorpora. Declared-AI videos are more densely represented in the early months of 2025, with 22 of 40 (55.0%) posted between January and April. Probable-AI videos display a gradual shift toward the later part of the period, with 15 of 40 (37.5%) posted between May and August. Human-produced videos are more evenly distributed across the window, with identical proportions in the early and late phases (14 of 40, 35.0% in each). Monthly distributions for each subcorpus are shown in Figure 1.

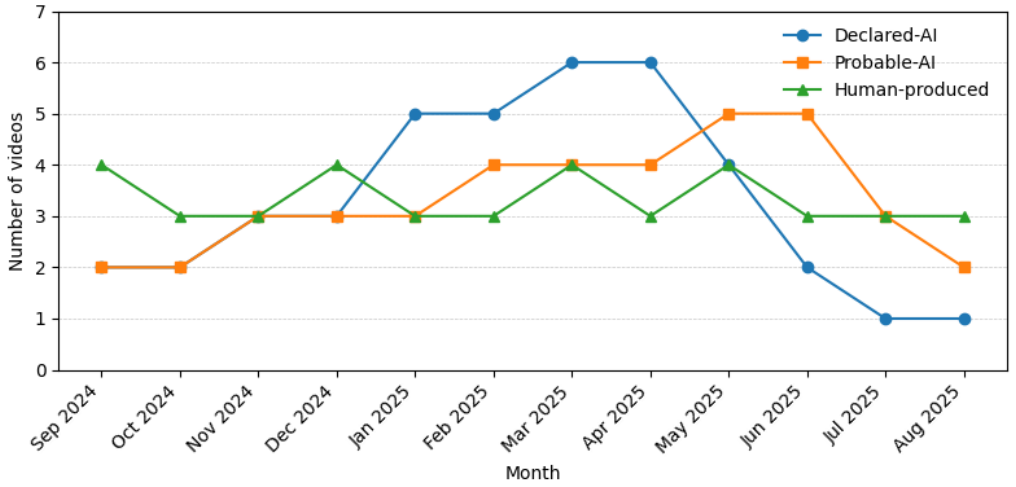

Figure 1. Monthly posting distribution of videos in the three subcorpora

Engagement differs substantially across the three subcorpora. In the declared-AI subcorpus, videos collectively account for 119,741,268 views, 2,747,745 likes, and 29,275 comments. In the probable-AI subcorpus, videos accumulate 27,936,793 views, 20,042,398 likes, and 75,345 comments. In the human-produced subcorpus, videos account for 35,449,413 views, 3,466,558 likes, and 86,357 comments. When considered together, the three subcorpora display clearly differentiated engagement profiles. Declared-AI videos account for the largest share of total views (65.7% of all views across the corpus), while probable-AI videos dominate in terms of likes (76.3% of all likes) and show high levels of commenting activity. Human-produced videos, by contrast, generate a smaller share of total views (19.4%) but the highest number of comments per video. These quantitative differences suggest distinct interpretive contexts of audience interaction within which the linguistic and multimodal authority performances examined below can be situated.

Within these engagement contexts, further differentiation emerges at the level of topical focus across the three subcorpora (see Table 2). AI-related content concentrates primarily on lifestyle and preventive themes, which together account for 55% of declared-AI videos and 50% of probable-AI videos. Metabolic and hormonal topics further contribute a substantial share in both AI-related datasets (declared-AI: 27.5%; probable-AI: 32.5%), resulting in over three quarters of AI-generated and AI-like videos addressing conditions framed as modifiable through individual behaviour or routine intervention.

Topic category	**Definition**	**Declared-AI**	**Probable-AI**	**Human-produced**
Lifestyle & prevention	Diet, supplements, sleep, habits, behavioural advice aimed at prevention or optimisation	22 (55%)	20 (50%)	9 (22.5%)
Metabolic / hormonal health	Content focusing on metabolism, hormones, weight regulation, insulin, GLP-1 medication	11 (27.5%)	13 (32.5%)	6 (15%)

Acute care / first aid	Immediate response to health events (e.g. seizures, fainting, emergencies)	2 (5%)	3 (7.5%)	14 (35%)
Procedural / clinical explanation	Explanations of medical procedures, diagnostics, treatment pathways, or clinical decision-making	5 (12.5%)	4 (10%)	11 (27.5%)

Table 2. Topic categories across subcorpora (N, % of videos)

By contrast, human-produced videos show a markedly different distribution. Acute and first-aid topics account for 35% of the human-produced subcorpus, compared to only 5% and 7.5% in the declared-AI and probable-AI subcorpora respectively. Procedural or clinical explanations are also more frequent in human-produced videos (27.5%) than in AI-related content (declared-AI: 12.5%; probable-AI: 10%). Lifestyle-oriented topics, while still present, represent a smaller proportion of the human-produced dataset (22.5%). These distributions may suggest that AI-generated and AI-like videos predominantly address health topics that lend themselves to generalisation and behavioural prescription, whereas human-produced videos more frequently focus on acute scenarios and context-dependent clinical explanations. This topic pattern provides a further contextual parameter for interpreting the linguistic and multimodal features reported in the following sections, as the communicative affordances available for authority construction differ across these topical domains.

3.1 *Distribution of linguistic resources across subcorpora*

The three subcorpora are comparable in size, comprising 34,260 words in the declared-AI subcorpus, 35,883 words in the probable-AI subcorpus, and 31,604 words in the human-produced subcorpus. As described in the methods, all frequencies are reported as absolute frequency (Af), followed by normalised frequency per 10,000 words (Nf). Table 3 summarises the distribution of the linguistic resources operationalised in the annotation scheme (section 2.3), namely evidential markers (lexical items encoding information source), vague quantifiers (under-specified expressions of amount or frequency), high-certainty and hedging stance expressions (epistemic formulations signalling, respectively, strong or mitigated speaker commitment), empathy formulas (routinised first- and second-person constructions displaying affective alignment), and directives (imperative or prohibitive clauses inviting or discouraging action). Lexical realisations of these categories were identified through frequency profiling and verified through concordance analysis to ensure functional relevance. The categories are reported here to establish relative prominence and cross-subcorpus variation, and are discussed in detail in the subsequent sections.

Linguistic target	**Declared-AI**	**Probable-AI**	**Human-produced**
	Af; Nf		
Evidential markers	392; 114.4	518; 144.3	261; 82.6
Vague quantifiers	684; 199.6	812; 226.3	529; 167.4

High-certainty stance	301; 87.9	379; 105.6	173; 54.7
Hedging stance	214; 62.5	196; 54.6	267; 84.5
Empathy formulas	176; 51.4	203; 56.6	88; 27.8
Directives (imperatives/ prohibitives)	421; 122.9	463; 129.0	317; 100.3

Table 3. Core linguistic targets across subcorpora

Overall, the two AI-related subcorpora show higher normalised frequencies for evidential, directive, and empathy-related resources, while the human-produced subcorpus displays a higher density of hedging devices.

3.1.1 *Evidential markers*

Evidential markers are particularly prominent in the AI-related subcorpora, both in absolute and normalised terms. In the declared-AI corpus, the most frequent evidential forms include *studies show* (Af 96; Nf 28.0), *research suggests* (Af 83; Nf 24.2), and *clinically proven* (Af 58; Nf 16.9). Concordance analysis indicates that these expressions are typically realised as clause-initial frames introducing brief, self-contained claims. For example, in a video discussing GLP-1 medications, one AI-generated script opens with "Studies show that GLP-1 injections can completely reset your appetite signals", followed immediately by a declarative outcome claim without further specification. In another case, a skincare-related video introduces advice with "Research suggests this ingredient repairs your skin barrier overnight", where research functions as a generic authority source without temporal or institutional anchoring.

In the probable-AI subcorpus, evidential markers are higher overall. The most frequent items include *doctors say* (Af 91; Nf 25.4), *studies found* (Af 88; Nf 24.5), and *experts agree* (Af 74; Nf 20.6). Collocation analysis shows that these expressions frequently co-occur with evaluative adjectives such as *effective*, *safe*, and *important*, as well as with directive constructions. A recurrent pattern appears in nutrition-related content, where statements such as "Doctors say you should never combine these two supplements" precede an imperative instruction, or in mental health videos claiming "Experts agree this is the fastest way to calm anxiety". In these cases, evidential frames introduce prescriptive advice while remaining detached from identifiable sources.

By contrast, evidential markers in the human-produced subcorpus are less frequent and more contextually anchored. The most common forms are *in my practice* (Af 67; Nf 21.2) and *guidelines recommend* (Af 48; Nf 15.2). Concordance lines show that these expressions are typically followed by contextual limitation or qualification. For instance, a clinician-led video on hypertension states "In my practice, I usually see this response in patients over fifty", while another video references "current guidelines recommend this only after lifestyle changes have failed". These constructions embed evidentiality within professional experience or institutional frameworks rather than presenting it as a generalised authority claim.

3.1.2 *Vague quantification*

Vague quantifiers are frequent across all three subcorpora but differ in density and pragmatic deployment. In the declared-AI corpus, vague quantifiers occur 684 times (Nf 199.6), with *many people* (Af 143; Nf 41.7) and *most people* (Af 131; Nf 38.2) as the most recurrent forms. These expressions commonly introduce claims about symptoms or treatment outcomes. For example, a hormone-related video states "Many people don't realise their fatigue is caused by hormone imbalance", while a weight-loss video claims "Most people start seeing results in just two weeks", extending applicability without specifying conditions or variability.

The probable-AI subcorpus shows the highest density of vague quantification (Af 812; Nf 226.3). Frequent items include *most people* (Af 156; Nf 43.5) and *a lot of* (Af 139; Nf 38.7), often occurring in proximity to directives. In a video about gut health, the script claims "A lot of people are damaging their gut without knowing it", followed by "you need to stop eating this immediately". Similarly, "Most people should avoid this common breakfast food" appears as a lead-in to prescriptive dietary advice.

In the human-produced subcorpus, vague quantifiers are less dense (Af 529; Nf 167.4) and more frequently accompanied by contextual restriction. The most frequent forms are *some people* (Af 134; Nf 42.4) and *sometimes* (Af 119; Nf 37.7). Concordance examples include "Some people experience side effects at higher doses" and "Sometimes this happens during the first few weeks of treatment", where vagueness is integrated into explanatory or cautionary framing.

3.1.3 *Epistemic stance*

High-certainty epistemic stance markers are more frequent in the AI-related subcorpora. In the probable-AI corpus, high-certainty forms occur 379 times (Nf 105.6), with *will* (Af 211; Nf 58.8) and *always* (Af 97; Nf 27.0) as the most frequent items. These forms commonly appear in outcome-oriented claims, such as "This supplement will boost your metabolism" or "Avoiding this food will always reduce inflammation". The declared-AI corpus shows a similar pattern (Af 301; Nf 87.9), with examples including "This method will fix insulin resistance" and "You will notice results almost immediately".

By contrast, the human-produced subcorpus shows lower frequencies of high-certainty stance (Af 173; Nf 54.7) and higher frequencies of hedging devices (Af 267; Nf 84.5). Common hedging forms include *may*, *can*, and *depends*, as in "This may help some patients, depending on their condition" or "Results can vary depending on adherence and dosage". These constructions are typically embedded within longer explanatory segments rather than positioned as standalone claims.

3.1.4 *Empathy formulations and directives*

Empathy constructions are more prevalent in the AI-related subcorpora. In the declared-AI corpus, empathy formulas occur 176 times (Nf 51.4), with frequent patterns such as *you're not alone* (Af 49; Nf 14.3) and *this can be frustrating* (Af 42; Nf 12.3). These expressions often follow problem statements, as in "If you've tried everything and nothing works, you're not alone", before transitioning to advice. In the probable-AI corpus, comparable frequencies are observed (Af 203; Nf 56.6), with

examples such as “It’s normal to feel overwhelmed by conflicting health advice” and “A lot of people struggle with this, and that’s okay”.

In the human-produced subcorpus, empathy formulas are less frequent (Af 88; Nf 27.8) and are more often embedded within narrative or experiential framing. Examples include “Patients often tell me they’re worried about side effects” and “I understand why this diagnosis can be scary”, where empathy is tied to reported patient interaction rather than generic alignment.

Directive constructions are common across all subcorpora but occur with higher density in the AI-related sets. In the probable-AI corpus, directives occur 463 times (Nf 129.0), followed by the declared-AI corpus (Af 421; Nf 122.9) and the human-produced corpus (Af 317; Nf 100.3). In AI-related videos, directives frequently follow evidential or evaluative statements, as in “Studies show this increases risk, so stop doing this today”. In human-produced videos, directives are more often embedded within explanatory sequences, for example “If symptoms persist, you should speak to your doctor”.

3.2 *Metaphors*

Metaphorical expressions occur in all three subcorpora but with markedly different densities. In the declared-AI corpus, metaphorical formulations occur 238 times (Nf 69.5), compared to 261 instances in the probable-AI corpus (Nf 72.7) and 121 instances in the human-produced corpus (Nf 38.3). Thus, metaphor use is nearly twice as dense in AI-related content as in human-produced videos. Three metaphor clusters account for the majority of occurrences and are described below.

3.2.1 *Mechanical repair metaphors*

Mechanical repair metaphors are the most frequent cluster in both AI-related subcorpora. In the declared-AI corpus, they occur 104 times (Nf 30.4), and in the probable-AI corpus 117 times (Nf 32.6), compared to 41 occurrences in the human-produced corpus (Nf 13.0). These metaphors typically involve verbs such as *fix*, *reset*, *repair*, or *correct*, framing health processes as malfunctions that can be directly addressed. For example, a declared-AI video on insulin resistance includes the statement “This protocol fixes insulin resistance at the source”, while a probable-AI nutrition video claims “Reset your metabolism in seven days”, presenting complex physiological processes as discrete mechanical problems with identifiable solutions. In human-produced videos, comparable expressions are less frequent and more often embedded within hedged explanations, as in “This can help correct certain imbalances over time”.

3.2.2 *Energy-related metaphors*

Energy-related metaphors form the second most frequent cluster. These occur 71 times in the declared-AI corpus (Nf 20.7) and 83 times in the probable-AI corpus (Nf 23.1), compared to 38 instances in the human-produced corpus (Nf 12.0). Common realisations include *boost*, *activate*, and *supercharge*, frequently used to describe immune function, metabolism, or cognitive performance. For instance, a probable-AI supplement video states “This vitamin boosts your immune system instantly”, while a declared-AI video claims “Activate your metabolism before breakfast”. In human-produced content,

energy metaphors appear more cautiously, often modified by modal verbs or conditional clauses, as in "This may help boost energy levels in some patients".

3.2.3 *Balance metaphors*

Balance metaphors are less frequent overall but remain salient in AI-related content. They occur 39 times in the declared-AI corpus (Nf 11.4) and 44 times in the probable-AI corpus (Nf 12.3), compared to 18 occurrences in the human-produced corpus (Nf 5.7). These metaphors typically involve expressions such as *restore balance*, *rebalance*, or *bring back balance*, especially in relation to hormones or gut health. For example, a declared-AI video on digestive health claims "Restore gut balance naturally", while a probable-AI video frames dietary advice as a way to "Rebalance your hormones". In human-produced videos, balance metaphors are more often contextualised within longer explanations, such as "The aim is to gradually restore balance, depending on individual response".

3.3 *Multimodal and identity-related features across subcorpora*

This section reports the distribution of visual, aural, and identity-related resources across the three subcorpora, focusing on camera framing, gaze configuration, vocal delivery, background sound, visual medical indexing, identity signalling, sequential visual organisation, and text–speech alignment.

3.3.1 *Camera framing and visual address*

Frontal camera framing, defined as the speaker facing the camera directly, is the dominant configuration across all subcorpora, though with notable variation in prevalence. It occurs in 36/40 (90%) declared-AI videos, 34/40 (85%) probable-AI videos, and 26/40 (65%) human-produced videos. Medium close-up framing (head and shoulders) is the most frequent shot type in the AI-related subcorpora, appearing in 32/40 (80%) declared-AI and 30/40 (75%) probable-AI videos, compared with 21/40 (52.5%) human-produced videos. Wider framing, including torso or environmental context, is more common in the human-produced subcorpus (19/40; 47.5%) than in declared-AI (8/40; 20%) or probable-AI content (10/40; 25%).

Direct gaze toward the camera is present in 35/40 (87.5%) declared-AI videos and 33/40 (82.5%) probable-AI videos, versus 24/40 (60%) human-produced videos. Indirect or alternating gaze patterns, often associated with explanatory or narrative segments, are more frequent in the human-produced subcorpus (16/40; 40%) than in the AI-related datasets (declared-AI: 5/40; 12.5%; probable-AI: 7/40; 17.5%). For example, in several human-produced videos, speakers alternate gaze between the camera and off-screen reference points when elaborating on clinical experience or contextual detail.

3.3.2 *Vocal delivery and background sound*

Synthetic or highly standardised voice delivery characterises the declared-AI subcorpus, occurring in 40/40 (100%) videos. In the probable-AI subcorpus, voice delivery judged to be synthetic or semi-synthetic appears in 31/40 (77.5%) videos, while the remaining 9/40 (22.5%) feature human voices

with minimal prosodic variation. In contrast, all human-produced videos (40/40; 100%) feature natural human voice delivery with noticeable prosodic modulation, including pitch variation and rhythmic emphasis. Background music is present in 29/40 (72.5%) declared-AI videos and 27/40 (67.5%) probable-AI videos, typically consisting of low-volume instrumental tracks. In the human-produced subcorpus, background music appears in 18/40 (45%) videos, while silence or ambient sound predominates in 22/40 (55%). For instance, several clinician-led videos rely solely on ambient room sound during explanatory segments, whereas AI-generated content frequently includes continuous background music throughout.

3.3.3 *Visual medical indexing and identity signalling*

Visual elements indexing medical or clinical contexts are unevenly distributed across subcorpora. Anatomical imagery, including stylised organ visuals or simplified diagrams, appears in 28/40 (70%) declared-AI videos and 25/40 (62.5%) probable-AI videos, compared with 14/40 (35%) human-produced videos. Diagnostic-style text overlays, such as symptom lists, condition labels, or stepwise instructions, are present in 33/40 (82.5%) declared-AI, 31/40 (77.5%) probable-AI, and 19/40 (47.5%) human-produced videos. Medical attire or props (e.g. white coats, scrubs, stethoscopes) occur in 21/40 (52.5%) declared-AI and 18/40 (45%) probable-AI videos, compared with 24/40 (60%) human-produced clips; in AI-related content, such attire is frequently combined with overlays and anatomical visuals, whereas in human-produced videos it often appears without additional visual reinforcement.

Explicit self-identification as a medical professional (e.g. “doctor”, “MD”, “clinician”) appears in 14/40 (35%) declared-AI videos and 17/40 (42.5%) probable-AI videos, compared with 29/40 (72.5%) human-produced videos. Verification markers or platform badges are absent from declared-AI content (0/40; 0%), rare in the probable-AI subcorpus (3/40; 7.5%), and more common in human-produced videos (11/40; 27.5%). Conversely, implicit identity signalling through visual proxies, such as clinical settings, medical attire, or authoritative text labels, is more prevalent in AI-related content, appearing in 34/40 (85%) declared-AI and 32/40 (80%) probable-AI videos, compared with 21/40 (52.5%) human-produced videos, indicating a stronger reliance on indirect identity cues in AI-generated and AI-like clips.

3.3.4 *Sequential visual organisation and text–speech alignment*

Structured visual sequencing, defined as staged presentation of information through numbered lists, stepwise overlays, or progressive disclosure, occurs in 31/40 (77.5%) declared-AI videos and 29/40 (72.5%) probable-AI videos, compared with 17/40 (42.5%) human-produced videos. In AI-related content, this organisation most often takes the form of numbered overlays (e.g. “Sign 1”, “Sign 2”, “Sign 3”) that appear and disappear in synchrony with spoken segments; in human-produced videos, sequential organisation is less rigid and more commonly realised through unnumbered summaries or recap slides toward the end of the clip.

Text–speech alignment further differentiates subcorpora. In declared-AI videos, on-screen text duplicates or closely paraphrases spoken content in 34/40 (85%) clips, functioning primarily as reinforcement, while in the probable-AI subcorpus duplication occurs in 30/40 (75%) videos and partial reformulation, where on-screen text condenses spoken explanations into simplified slogans,

appears in 8/40 (20%) videos (e.g. extended dietary advice accompanied by the overlay "STOP EATING THIS"). In contrast, human-produced videos show lower rates of duplication (18/40; 45%) and higher rates of supplementary on-screen text that adds contextual information not explicitly stated in speech (14/40; 35%), such as dosage ranges, conditional warnings, or brief citations, so that on-screen text functions more as elaboration than reinforcement.

4. Discussion

The results support a reading of pseudo-authority as an emergent, platform-native semiotic configuration. Credibility seems to be produced through patterned alignment between recurring linguistic and multimodal features. Analytically, authority does not reside in any single feature, but in the bundling and relative stability of these configurations across a large number of videos.

A first finding is that the corpus does not suggest a simple opposition between 'AI equals fake authority' and 'human equals real authority'. Instead, the subcorpora differ in how authority is anchored and made recognisable. In the human-produced set, authority tends to be anchored through explicit professional positioning, including self-identification and more frequent verification, alongside linguistic choices that more often keep claims open to contextual limitation. This is reflected in a greater density of hedging devices and in evidential expressions tied to practice-based or institutional reference frames, such as professional experience or clinical guidelines. This pattern does not imply that human-produced videos are more accurate (accuracy was not assessed), but it does point to a different discourse logic in which legitimacy is constructed through situated expertise: expertise that can be narrated ("in my practice") or institutionally referenced ("guidelines recommend"). This configuration aligns with established accounts of legitimation through authorisation and expert role,[22] as well as with research on epistemic stance in professional health genres, where caution and conditionality are part of public accountability.[23]

In the AI-related subcorpora, authority is more often made recognisable through generic authorisation and procedural packaging. Evidential patterns frequently invoke science or expertise as abstract categories ("studies show", "research suggests", "experts agree", "doctors say") without source traceability. This is not a claim about deception but an observation about discursive form: evidentiality functions as a credential substitute, and its pragmatic effect is strengthened by both syntactic positioning and multimodal support. Clause-initial evidential frames introduce compact, high-certainty propositions that are often visually reinforced through diagnostic-style overlays and list-based sequencing. The same proposition is thus presented aurally and visually, providing a coherence cue even when the evidential ground remains underspecified. From a multimodal discourse perspective, this systematic alignment of modes stabilises interpretation and reduces ambiguity in reception,[24] while also responding to platform-specific pressures for rapid legibility and immediate comprehension.

---

[22] Theo van Leeuwen, "New Forms of Writing, New Visual Competencies", *Visual Studies*, 23.2 (2008), 130-135, https://doi.org/10.1080/14725860802276263.

[23] Ken Hyland, "Genre, Discipline and Identity", *Journal of English for Academic Purposes*, 19 (September 2015), 32-43, https://doi.org/10.1016/j.jeap.2015.02.005.

[24] Gunther Kress and Theo van Leeuwen, "The Semiotic Landscape", in Patrick Griffiths et al., eds., *Language in Use: A Reader*, First Edition (London: Routledge, 2010), 6, https://doi.org/10.4324/9781003060994.

A further pattern concerns the nature of identity signalling. In AI-related content, pseudo-authority appears less as an explicit identity claim ("I am a doctor") and more as an interactional script. The combination of direct gaze, frontal framing, steady prosody, and directive closure produces a recognisable institutional delivery style even when professional identity is not overtly claimed. Identity cues operate on two levels. One level is explicit, involving self-identification and verification markers. The other is implicit and compositional, relying on visual proxies such as medical attire, clinical settings, anatomical imagery, and diagnostic typography. The distributions indicate that AI-related videos rely more consistently on this second layer. In terms of legitimation strategies, authority is achieved through invoked expertise, procedural rationalisation, and evaluative framing of what is presented as 'safe', 'important', or 'effective', often without the institutional anchoring observed in human-produced videos.

Directive constructions are particularly informative with respect to the simulation of authority. Directives are common across all subcorpora, which is expected in health advice genres. What differentiates the AI-related sets is how directives are embedded within a recurrent micro-structure: generic evidential framing followed by a high-certainty claim and then by directive closure. This sequence produces a 'clinical instruction' effect resembling procedural lay guidance: authority is established, a rule is stated, and an action is prescribed. In the human-produced set, directives more often appear downstream of contextualising explanation, including hedging and conditionality. This alters their pragmatic force, making them read less as categorical rules and more as recommendations contingent on circumstances. Authority, in this case, is enacted through explanatory positioning rather than through prescriptive closure.

The second research question concerns vagueness, fictionality, and fabrication-adjacent practices as multimodal textual strategies. The corpus indicates that vagueness is not an occasional feature but a systematic resource, particularly in the AI-related subcorpora, where vague quantifiers are denser and frequently positioned near prescriptive moves. Expressions such as "many people", "most people", or "a lot of" allow claims to be framed as broadly applicable while avoiding the specificity that would enable straightforward falsification or contextual challenge. Pragmatically, this supports population-level inference without committing to defined populations or conditions. The analytical point is not that vagueness equates to misinformation, but that it constitutes one of the ways *pseudos* is enacted: discourse can sound scientific while remaining epistemically underdetermined. This is compatible with a Gricean account of meaning construction through what is left unsaid, where listeners are guided toward inferences (e.g. personal applicability) without explicit propositional commitment.[25]

Metaphorical framing may reinforce this dynamic by compressing complex physiological processes into simplified causal models. Mechanical repair, energy activation, and balance restoration metaphors are particularly well suited to short-form video because they translate multifactorial health processes into clear problem-solution schemas. Mechanical metaphors ("fix", "reset", "repair") align closely with directive closure by framing the body as a system that can be adjusted to produce reliable outcomes. Energy metaphors ("boost", "activate", "supercharge") support outcome-oriented claims compatible with high-certainty stance, while balance metaphors invoke a moralised notion of restoration and natural order. Importantly, these metaphors do not operate in isolation:[26] they co-occur

[25] H. Paul Grice, "Logic and Conversation", 41-58.
[26] George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago: Chicago U.P., 2003).

with generic evidentiality, high-certainty stance, and overlay-based sequencing, forming part of a coherent authority package rather than a purely stylistic layer.

Within this framework, fictionality and fabrication are approached as modes of presentation rather than as factual labels. The corpus shows that AI-related videos frequently produce the appearance of scientific grounding through generic evidential frames and the visual grammar of clinical explanation, while simultaneously limiting traceability through the absence of source specification. This combination constitutes a plausible discursive mechanism for *pseudos*: authority is performed through recognisable institutional signs, while epistemic accountability is softened through vagueness and non-attributed evidentiality. The multimodal layer may intensify this effect by packaging claims into modular units (e.g. lists, labels, steps) that resemble health education artefacts regardless of their origin. In this sense, the videos often enact the genre of health communication without being anchored in its conventional epistemic infrastructures, echoing prior research on genre imitation and pseudo-institutional discourse in digital environments.[27]

Metadata and topic distributions impose an important contextual constraint on interpretation. The AI-related subcorpora more frequently address lifestyle, preventive, and metabolic or hormonal topics, whereas the human-produced set includes a larger share of acute, first-aid, and procedural content. Topic type plausibly shapes the communicative affordances available to creators: lifestyle optimisation lends itself to generalisation and outcome-focused framing, while acute and procedural topics are more normatively tied to stepwise instruction and contextual qualification. These differences do not explain the linguistic patterns away; rather, they clarify how pseudo-authority may take slightly different semiotic routes depending on the informational ecology of the topic. Within these topic landscapes, AI-related videos nonetheless show a stronger tendency toward repeatable authority packaging built on generic evidentials, high-certainty stance, and directive closure, scaffolded by diagnostic overlays and structured sequencing.

Taken together, the findings shed light on both research questions. Within this stratified corpus, medical authority in AI-generated videos is simulated primarily through stable semiotic configurations that reproduce the formal features of clinical explanation. Vagueness, fictionality, and fabrication-adjacent practices operate as multimodal textual resources that render claims broadly applicable, rapidly legible, and resistant to verification demands without requiring overt falsification. These conclusions remain descriptive and configurational: they identify what the texts do semiotically, not what audiences necessarily believe or how platforms causally rank content, and should be read in light of the limitations discussed in the following section.

4.1 *Limitations*

Several limitations delimit the scope of the claims advanced in this study. The dataset is a stratified corpus constructed for comparative analysis rather than a representative sample of TikTok health content, a type of sampling that is difficult to achieve on a fast-changing, search-personalised platform; the findings therefore characterise patterned semiotic practices within the corpus rather than platform-wide prevalence or distribution. Moreover, the classification of ‘probable-AI’ videos is

[27] Jana Laura Egelhofer and Sophie Lecheler, “Fake News as a Two-Dimensional Phenomenon: A Framework and Research Agenda”, *Annals of the International Communication Association*, 43.2 (June 2019), 97-116, https://doi.org/10.1080/23808985.2019.1602782.

necessarily inferential and reflects the author's judgement based on these multimodal cues; no external detector or inter-coder reliability test was applied, and some misclassification cannot be ruled out. In addition, the analysis does not assess clinical accuracy, biomedical validity, or potential health harm, nor does it model viewer interpretation, behavioural effects, or algorithmic ranking; engagement metrics are treated as contextual descriptors of circulation rather than as evidence of persuasion or impact. Finally, topic distributions differ across subcorpora; while these differences are documented and taken into account analytically, topic effects cannot be fully separated from production-type effects within the present design.

5. Conclusions

The analysis shows that declared-AI and probable-AI TikTok health videos enact medical authority through recurrent multimodal configurations. These combine generic evidentiality and high-certainty stance with visually and aurally standardised formats that echo clinical explanation, offering viewers a recognisable yet institutionally untethered register of expertise. Within this *pseudos* framework, vagueness, fictionality, and fabrication-adjacent practices function as textual resources that widen the apparent applicability of advice while keeping claims resistant to straightforward verification. At the same time, contrasts with the human-produced subcorpus indicate that pseudo-authority is not a simple property of AI generation but of how epistemic stance, evidentiality, and identity cues are orchestrated across modes. Human-produced videos in this corpus more often situate claims within experiential or institutional frames and retain greater space for hedging, signalling a different discourse logic of legitimacy. These findings refine the concept of *pseudos* for platform-mediated health communication by showing that epistemic instability is not only linguistic but multimodally organised.

For discourse studies and public health communication alike, the analysis suggests that health (mis)information on this platform is best understood not only in terms of propositional accuracy, but as *pseudos*: semiotic configurations that render authority legible, repeatable, and actionable while reproducing the forms of medical explanation in an epistemic space where authorship, authenticity, and persuasion are increasingly difficult to disentangle.