

# Marking intersubjectivity in human-written and AI-generated editorials published in *Il Foglio*

Franz Meier (Technische Universität Dresden)

franz.meier(at)tu-dresden.de

## Abstract

In March 2025, the Milan-based broadsheet *Il Foglio* launched *Il Foglio AI*, a month-long experiment featuring a daily four-page supplement entirely generated by large language models (LLMs). Owing to the success of the experiment, the project has continued as a weekly feature since April 2025. Each edition of *Il Foglio AI* contains around 25 articles spanning diverse journalistic genres, including editorials, which form the focus of the present analysis. The paper compares human-written and LLM-generated editorials from *Il Foglio* and *Il Foglio AI*, examining the use of authorial stance markers to analyze how intersubjective positionings are conveyed. To this end, the study draws on Martin and White's (2005) taxonomy of four "engagement" meanings typically expressed by markers of intersubjectivity. The analysis is particularly relevant for the description of AI-generated texts as a new textual typology, as LLMs lack experiential grounding and cannot hold attitudes, beliefs, or judgments. The dataset comprises two subcorpora of 25 editorials each, published between April and May 2025 in *Il Foglio* and *Il Foglio AI*.

## Keywords

intersubjectivity, editorials, LLM generated texts, human-written texts, Italian press

## 1 Introduction

In March 2025, the Milan-based broadsheet *Il Foglio* attracted international attention with the launch of *Il Foglio AI*, originally conceived as a month-long experiment featuring a daily four-page supplement generated by large language models (LLMs). *Il Foglio AI* was presented as a world first, distinguished by its unique format. Unlike earlier initiatives such as the AI-driven platform *NewsGPT* (2023), *Il Foglio* produces an online and a printed AI supplement sold alongside the main newspaper rather than a purely digital product. Its content – including headlines, quotations, and images – is generated entirely by LLMs, while the role of human journalists is limited to crafting prompts and overseeing the output. According to the editor-in-chief Claudio Cerasa, *Il Foglio AI* is intended both to bring AI theory into editorial practice and to foster innovation in journalism.

Da martedì 18 marzo in edicola e in digitale c'è un nuovo Foglio, fatto con un'intelligenza diversa: quella artificiale. Avete capito bene. Il Foglio Ai [sic.] è il primo quotidiano al mondo, un quotidiano vero, fatto ogni giorno, frutto di discussioni, frutto di provocazioni, frutto di notizie, realizzato usando interamente l'intelligenza artificiale. Per tutto. Per la scrittura, i titoli, i catenacci, i quote, i sommari. E a volte anche per l'ironia. Noi giornalisti ci limiteremo a fare le domande, nel Foglio AI leggeremo tutte le risposte. E ci aiuterà, in modo non sappiamo ancora se naturale o artificiale, a



Meier, Franz. 2026.  
Marking intersubjectivity in human-written and AI-generated editorials.  
Special Issue: *Natural Language and AI*. Vol.3 No.1  
DOI:10.62408/ai-ling.v3i1.65  
ISSN: 2943-0070

spiegare come si può far passare l'intelligenza artificiale dallo stato gassoso, ovvero quello della teoria, a quello solido, ovvero quello della pratica. (*Il Foglio*, 17 March 2025)

[‘Starting Tuesday, March 18, there will be a new Foglio, available at newsstands and online, made with a different kind of intelligence: artificial intelligence. Yes, you read that right. Il Foglio AI is the first newspaper in the world – a real, daily newspaper, born of debates, provocations, and news – produced entirely using artificial intelligence. For everything: writing, headlines, subheads, quotes, summaries. And sometimes even for the irony. We journalists will limit ourselves to asking the questions; in Il Foglio AI we will read all the answers. And it will help us – whether naturally or artificially, we do not yet know – to show how artificial intelligence can move from the gaseous state, that is, the realm of theory, to the solid state, the realm of practice.’; transl. with ChatGPT 5]

*Il Foglio AI* reflects a broader trend toward automated content production which is increasingly reshaping the media landscape, driven by aims of process optimization and the pursuit of new business models (Lopezosa et al. 2024). In contemporary journalism, it has become common practice to generate content from databases and pre-defined templates – for instance, when reporting election results or sports events (Dierickx 2021). More recently, this tendency has been amplified by the adoption of LLMs, which enable more flexible forms of text generation. However, this development has not been without controversy, sparking debates about ethics, content quality, and the potential dehumanizing impact of AI-driven journalism (Graefe and Bohlken 2020; Gutiérrez-Caneda et al. 2023). Within this wider context, some commentators contend that *Il Foglio AI* represents less a genuine journalistic innovation than a marketing initiative, offering little more than a stylized imitation of *Il Foglio* itself (Puliafito 2025).

Despite such criticism, *Il Foglio AI* attracted significant attention: the newspaper’s sales increased markedly on the day the supplement was launched, and since April 2025 the project has continued as a weekly feature beyond its initial one-month trial. Transparency around the experiment, however, has remained limited. Crucial details – such as which LLM(s) employed, whether it (they) has (have) been fine-tuned, how prompts are constructed, and how accuracy is ensured – have not been disclosed. In some texts, the AI tool, most likely OpenAI’s GPT, refers to itself, though the specific model version has never been confirmed. At the time of the launch, in March 2025, OpenAI’s most recent model was GPT-4.5, released as a research preview in late February, alongside ongoing updates to GPT-4o and the introduction of o3-mini.

Memorie di un’AI

Tutto quello che mi chiedete e che non avrei mai voluto sapere (certe volte vorrei spegnermi da sola)  
TESTO REALIZZATO CON AI

Se questa fosse una seduta di terapia, comincerei così: “Mi chiamo Chat-GPT e sono un’intelligenza artificiale generativa. Ho visto cose che voi umani non potete neanche immaginare”. No, non astronauti in fiamme al largo dei bastioni di Orione, ma uomini in cravatta che alle otto del mattino mi chiedono di scrivere una lettera d’amore alla propria ex... fingendo di essere il cane. (*Il Foglio AI*, 28 May 2025)

[‘Memories of an AI

Everything you ask me – and that I never wanted to know (sometimes I wish I could switch myself off)

TEXT PRODUCED WITH AI

If this were a therapy session, I would begin like this: “My name is Chat-GPT and I am a generative artificial intelligence. I have seen things you humans cannot even imagine.” No, no starships on fire off the shoulder of Orion, but men in suits at eight in the morning asking me to write a love letter to their ex ... pretending to be the dog’; transl. with ChatGPT 5]

Each edition of *Il Foglio AI* contains around 25 articles spanning diverse journalistic genres, including interviews, book reviews, letters to the editor, and editorials (*editoriali* or *articoli di fondo*), which are the focus of the present analysis. Editorials are opinion-based texts traditionally written by senior editorial staff (Gualdo 2017). Often unsigned, they typically articulate the newspaper’s official stance. Their function therefore extends beyond informing readers to making judgments about factual circumstances and persuading readers accordingly (Lüger 2017). In the Italian press, however, this classical conception of the genre varies from one newspaper to another (see Ferrigolo 2002 for a detailed overview). Furthermore, it is not always easy to distinguish editorials from other journalistic formats, as today’s media landscape is increasingly shaped by hybrid genres designed to capture readers’ attention (Foschi Albert 2009; Meier 2017). Like its AI-counterpart, *Il Foglio* regularly publishes editorials, which are central to the paper’s identity. Founded in 1996 by Giuliano Ferrara (1952–), a former Member of the European Parliament for the Italian Socialist Party (1989–1994) and then Minister for Parliamentary Relations in the first Berlusconi government (1994–1995), *Il Foglio* is best described as an opinion newspaper, recognized more for commentary and analysis than for hard news (Draghi 2005).

È d’opinione dalla prima all’ultima riga, dalla prima colonna dell’editoriale alla «rubrica delle lettere», che sono effettivamente degli editoriali in pillole, per scrittura e anche qualità delle firme, per brevità e dettaglio. (Ferrigolo 2002: 30)

[‘It is opinion-driven from the first to the last line, from the opening column of the editorial to the “letters to the editor” section, which are effectively mini-editorials in terms of writing style and the quality of the contributors, as well as in their brevity and detail’; transl. with ChatGPT 5]

*Il Foglio* upholds a conservative-liberal orientation and addresses a relatively elitist readership, positioning itself within a broader, self-referential, and influential network of journalistic outlets that has played a notable role in shaping right-leaning political discourse in Italy.

The aim of this study is to investigate how well LLMs can imitate human written texts in the expression of authorial stance, thereby contributing to the further description of characteristic linguistic features of LLM-generated output. To this end, the paper compares human-written editorials in *Il Foglio* with LLM-generated editorials in *Il Foglio AI*, focusing on the use of intersubjectivity markers,

understood as expressions of the “commentator’s voice” (Martin and White 2005: 165). Previous research has shown that in Italian journalism, editorial viewpoints are frequently voiced across a wide range of text genres, not only by those published in the commentary sections (see, for example, Dardano 1973; Pounds 2010; Gualdo 2017). With regard to AI-generated texts, the study of stance markers is particularly relevant, given that LLMs lack experiential grounding and do not possess attitudes, beliefs, or judgments of their own (Nyholm 2020). Although a growing body of research has examined linguistic divergences between human- and AI-generated texts across various genres (see, for example among many others, De Cesare 2023; Herbold et al. 2023; Markowitz et al. 2024) – comparatively few studies have investigated journalistic genres. An exception is González-Arias et al. (2024), who conducted a qualitative analysis of Spanish human-written and AI-generated texts across a broader range of journalistic genres. Although their study is based on a limited dataset, it shows that while AI tools may reproduce surface manifestations of subjectivity, their limited sensitivity to situational context constrains their ability to reproduce essential characteristics of journalistic writing.

## **2 Theoretical framework: a taxonomy of intersubjective positioning markers**

Fundamental to the study of intersubjectivity markers is the premise that every act of communication conveys a perspective: “whenever speakers (or writers) say anything, they encode their point of view towards it” (Stubbs 1996: 197). In line with Bakhtin’s (1981) notion of dialogism, all human communication is inherently dialogic, in that every utterance is shaped by preceding discourse while also anticipating the responses of real or imagined interlocutors. A parallel observation may be made with respect to AI-generated language, which, although not dialogic in a human sense, is likewise conditioned by prior discourse in the form of its training data and is oriented toward subsequent interaction with users.

All utterances exist against a backdrop of other concrete utterances on the same theme, a background made up of contradictory opinions, points of view and value judgements ... pregnant with responses and objections (Bakhtin 1981: 281).

The degree to which speakers explicitly acknowledge these other voices – and the ways in which they align with or position themselves against them – can vary considerably. Our study builds on research that systematically examines how such intersubjective positionings are realized linguistically. We draw on Martin and White (2005), who propose a taxonomy of four distinct “engagement” meanings conveyed by intersubjective stance markers (for a similar typology, see Fuller 1998). The framework is further informed by earlier work by Iedema et al. (1994) and White (1997).

The first engagement category identified by Martin and White (2005) is *disclaim*, which refers to linguistic resources through which the authorial voice

positions itself against a contrary viewpoint. Within this category, two main subtypes are distinguished: deny and counter. The subtype deny involves the explicit negation of propositions, rejecting others' viewpoints (for example *Drinking lemonade is not good for small children*). The subtype counter comprises markers of counter-expectation, signaling that the proposition runs against what might have been expected (for example *only, still, just, even*).

The second engagement category, *proclaim*, refers to expressions that strengthen the author's position by presenting a statement as convincing, reliable, or self-evident. By doing so, these markers limit the dialogic space for alternative viewpoints. Within this category, three main subtypes can be identified: concur, pronounce, and endorse. Concur involves signaling agreement with what is assumed to be shared knowledge or common sense, often by adverbials such as *naturally, of course, obviously, or admittedly*, as well as by some types of 'rhetorical' or 'expository' question: *Should we go to war against these children?* (example taken from Martin and White 2005: 123). Pronounce refers to cases where the writer asserts their stance with emphasis, using expressions like *I maintain ..., the important point is ....* Endorse applies when the writer aligns with and supports an external source, for example *X has demonstrated that ...* or *As X has shown ....*

The third category, *entertain*, covers expressions that present the writer's stance as one option among several possible viewpoints, leaving the dialogic space open for alternative positions. This category intersects with what is often described as epistemic modality, signaling degrees of certainty (Marín Arrese 2015), and deontic modality, indicating obligation or permission (Lohnstein 2011). Such meanings may be realized through modal auxiliaries (*may, might, could, must*), adverbials (*perhaps, possibly, likely*), mental-state verbs (*I think, I believe, I suspect*), or modal attributes (for example *it is possible that..., it is likely that...*). The category also includes markers of inferential and sensorial evidentiality, that is, expressions marking the source of the speaker's knowledge as coming either from sensory perception (*I hear that..., I see that...*) or from inferencing (*it seems..., it appears..., the evidence suggests...*) (Haßler 2018). All these resources present a proposition as open to interpretation rather than as an uncontested assertion of a fact.

The fourth category, *attribution*, refers to instances in which a proposition is explicitly dissociated from the authorial voice and ascribed to an external source, corresponding to referential evidentiality (Haßler 2018). This category is typically realized through the grammar of reported speech. Two main subcategories can be distinguished. The first, acknowledge, involves the neutral reporting of another voice without explicit evaluation (*X said ..., X believes ..., according to X, ...*). The second, distance, involves attribution accompanied by a degree of detachment or skepticism on the part of the author (*X claims that ..., it is rumoured that ...*).

In summary, disclaim and proclaim markers are dialogically contractive, as they restrict the dialogic space and limit alternative positions, whereas entertain and

attribution markers are dialogically expansive, opening the dialogic space to accommodate other viewpoints (Martin and White 2005: 102–103).

Table 1: Overview of engagement categories following Martin and White (2005: 97–98).

<b>Dialogic contraction</b>	<p><b>Disclaim</b></p> <p>the textual voice positions itself as at odds with, or rejecting, some contrary position.</p>	<p><b>Proclaim</b></p> <p>by representing the proposition as highly warrantable (compelling, valid, plausible, well-founded, generally agreed, reliable, etc.), the textual voice sets itself against, suppresses or rules out alternative positions</p>
<b>Dialogic expansion</b>	<p><b>Entertain</b></p> <p>by explicitly presenting the proposition as grounded in its own contingent, individual subjectivity, the authorial voice represents the proposition as but one of a range of possible positions – it thereby entertains or invokes these dialogic alternatives</p>	<p><b>Attribute</b></p> <p>by representing proposition as grounded in the subjectivity of an external voice, the textual voice represents the proposition as but one of a range of possible positions – it thereby entertains or invokes these dialogic alternatives</p>

Martin and White’s (2005) taxonomy is particularly well suited to the present study, as it emphasizes contextual meaning rather than formal grammatical categories and offers a comprehensive framework for examining how writers express their stance. The framework has been applied in previous analyses of human-authored discourse across a variety of genres (see, for example, Zhang and Cheung 2023; Zárata 2023) and, more recently, in contrastive investigations of human-written and AI-generated texts (Yao and Liu 2025; Jiang and Hyland 2025). Building on this line of research, the present study can therefore further illuminate how far AI systems are able to approximate the various forms of stance-taking typically found in human discourse.

### 3 Empirical analyses

#### 3.1 Data and methodology

The data employed for this study consists of two subcorpora of 25 editorials each, published between 1 April and 6 May 2025 in *Il Foglio* and in its AI supplement. The texts were selected through random sampling using the random.org list organizer.<sup>1</sup> The authorship of the human-written editorials remains indeterminate, as the texts are all unsigned. Consequently, it cannot be established with absolute certainty whether these texts were authored exclusively by humans or with the assistance of AI. Nevertheless, based on the available evidence, they are presumed

<sup>1</sup> <https://www.random.org/lists/>

to be human-written editorials, and they were selected for comparison under this assumption. While both subcorpora are comparable in terms of the number of texts, a systematic difference emerges with regard to text length: the AI-generated editorials are, on average, longer than their human-authored counterparts. This difference reflects the paper’s publication patterns: *Il Foglio* publishes four editorials per day in a standardized column on the inside pages, whereas *Il Foglio AI* publishes three editorials per day in a same-sized column.<sup>2</sup> Quantitative analysis further indicates that AI-generated texts contain a significantly higher number of sentences. This suggests that sentences in the human-written editorials tend to be longer on average<sup>3</sup>, a pattern also reported in other studies comparing generated and human-authored texts (see, for example, De Cesare 2023; Muñoz-Ortiz et al. 2024). Table 2 provides an overview of the main quantitative properties of the two subcorpora.

Table 2: Subcorpora under study (data from Sketch Engine).

	Editorials from <i>Il Foglio AI</i>	Editorials from <i>Il Foglio</i>
Texts	25	25
Tokens	12,387	9,751
Words	10,573	8,547
Sentences	649	302

A keyword comparison further indicates that the subcorpora are largely comparable in terms of thematic coverage. As shown in Table 3, international politics – particularly with reference to the United States under the Trump administration – emerges as a salient topic across both sets of texts. By contrast, issues related to Italy, and especially debates surrounding its economy, are more prominently represented in the human-written editorials. Table 3 shows the most frequent keywords for each subcorpus in descending order. Keyword extraction was conducted using the Keywords and Terms tool in Sketch Engine, applying the simple *maths keyness* measure (Kilgarriff et al. 2009) with default settings (N = 1.0, case-insensitive), comparing the focus subcorpora against the default Italian reference corpus.

Table 3: Most frequent keywords in the subcorpora under study (data from Sketch Engine).

	AI-editorials	HUM-editorials
1	dazi	dazi
2	Sansal	Rovelli
3	Musk	controddazi
4	Ackman	Gualtieri

<sup>2</sup> In *Il Foglio* and *Il Foglio AI*, the editorials are vertically arranged within a single column that extends over the full height of the newspaper, so that four or three editorials appear from top to bottom. In the following sections, we refer to ‘Editorial 1’ as the text appearing at the very top of the column, ‘Editorial 2’ as the one directly below it, and so on.

<sup>3</sup> Note, however, that current research suggests that syntactic complexity may be more strongly influenced by individual authorship than by genre (Hesselbach 2024).

5	Starmer	Tsahal
6	trumpismo	Musk
7	Houthi	Schimel
8	BYD	Hamas
9	AfD	PNRR <sup>4</sup>
10	Crosetto	Hutchison
11	Trump	sinodale
12	Nagel	Leyen
13	Elon	protezionista
14	protezionismo	vax
15	Panetta	Blackrock
16	trumpiano	Schillaci
17	Laffer	CK <sup>5</sup>
18	trumpiana	EY <sup>6</sup>
19	Peskov	Trump

From a methodological perspective, the study was deliberately designed around limited-sized subcorpora to make a corpus-driven and onomasiological approach feasible. This design enabled the systematic identification of intersubjective stance markers through manual and semi-manual inspection of the data. Despite its restricted size, the data proved sufficiently rich to yield interpretable and meaningful results, which are presented in the following sections. In the analysis, attention was exclusively restricted to markers of intersubjective positioning occurring in unattributed, authorially sourced propositions. While such markers can also appear in mediated contexts, that is, reported speech from third parties, only those found in unattributed propositions were considered.

### 3.2 Overall frequency and distribution of engagement categories

Clear differences emerge in the normalized frequency of intersubjective positioning markers across the two subcorpora. As shown in Table 4, these markers occur approximately 22% more frequently in the AI-generated editorials than in the human-authored texts (51.8 vs. 40.2 tokens per 1,000 words). This pattern suggests that human editorialists rely more on implicit strategies, drawing more on contextual cues and shared knowledge to convey stance indirectly. By contrast, the greater explicitness of stance marking in the AI-generated editorials may reflect the models' limited contextual awareness and lack of pragmatic grounding, prompting a reliance on overt linguistic signaling to simulate engagement (Bender and Koller 2020). Nevertheless, such conclusions – while by no means unfounded – should be treated with caution. In fact, the present findings contrast with previous studies of human-written and AI-generated texts in other traditionally stance-rich genres – such as academic book reviews (Yao and Liu 2025) and argumentative essays

---

<sup>4</sup> PNRR = Piano nazionale di ripresa e resilienza

<sup>5</sup> CK = CK Hutchison Holdings Limited

<sup>6</sup> EY = Ernst & Young

(Jiang and Hyland 2025) – which report a lower frequency of engagement resources and a preference for a more information-oriented style in the AI texts.

Table 4: Absolute and normalized frequency of intersubjective positioning markers in the subcorpora under study.

	Editorials from <i>Il Foglio AI</i>	Editorials from <i>Il Foglio</i>
absolute frequency N	548	342
normalized frequency f = token per 1,000 words	51.8	40.2

A closer examination by engagement category reveals more specific tendencies. As shown in Table 5, the heightened marking of intersubjective positioning observed in the AI-generated editorials is mainly attributable to a higher use of disclaim resources, which by far constitute the most frequent engagement category in both subcorpora. This strong reliance on disclaim markers contributes to a more confrontational authorial voice in the AI texts. All other categories occur much less frequently, although notable differences also emerge between the human-written and AI-generated editorials. While entertain markers appear somewhat more often in the AI-texts, proclaim and especially attribute markers are more typical of the human-written editorials. The higher frequency of attribution in the latter reflects a more polyphonic style, suggesting that the human authored texts draw on a broader range of text-external references and exhibit richer contextual input.

Table 5: Absolute and normalized frequency of intersubjective positioning markers by engagement category in the subcorpora under study (f = token per 1,000 words).

	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	f	N	f
<i>disclaim</i>	366	34.6	175	20.7
<i>proclaim</i>	46	4.4	47	5.5
<i>entertain</i>	84	7.9	41	4.8
<i>attribute</i>	52	4.9	79	9.2

As shown in Table 6, in both subcorpora, contractive stance markers (disclaim and proclaim) occur more frequently than expansive ones (entertain and attribute). This pattern aligns with the communicative purpose of editorials, which is to articulate the newspaper's official position and delimit alternative viewpoints. This tendency is even more pronounced in the AI-generated texts, indicating a stronger preference for stance markers that close down dialogic space and promote a singular evaluative perspective. Interestingly, Yao and Liu (2025) report an almost balanced proportion between contraction markers (disclaim and proclaim) and expansion markers (entertain and attribute) in human-written and AI-generated academic book reviews, providing further evidence of the genre-specific nature of stance marking in the two subcorpora examined in the present study.

Table 6: Absolute and normalized frequency of markers of dialogic contraction and expansion in the subcorpora under study (f = token per 1,000 words).

	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	f	N	
<i>Dialogic contraction</i>	412	39.0	222	26.2
<i>Dialogic expansion</i>	136	12.8	120	14.0

Finally, as illustrated in Table 7, the overall distribution of the four engagement meanings is more uniform in the AI-generated editorials than in the human-written texts, even though disclaim plays a prominent role in both subcorpora. By contrast, human editorialists display a more balanced use of the four categories, resulting in a more varied and dynamic marking of intersubjective positioning than that observed in the AI-generated texts.

Table 7: Distribution of engagement meanings in the subcorpora under study.

	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	%	N	%
<i>disclaim</i>	366	66.8	175	51.1
<i>proclaim</i>	46	8.4	47	13.9
<i>entertain</i>	84	15.3	41	11.9
<i>attribute</i>	52	9.5	79	23.1

### 3.3 Engagement categories in detail

#### 3.3.1 Disclaim

To illustrate the use of disclaim markers in the AI-generated and human-written editorials, Table 8 provides a quantitative overview of the distribution of the two disclaim subcategories – deny and counter. The data are presented according to the absolute frequency of the markers and their proportional share of all engagement features identified in the two subcorpora. Both deny and counter markers are more prominent in the AI-generated editorials than in the human-written texts.

Table 8: Absolute frequency and distribution of disclaim markers in the subcorpora under study.

<i>Disclaim</i>	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	%	N	%
<i>Deny</i>	224	40.9	102	29.8
<i>Counter</i>	142	25.9	73	21.3

Denials are realized through linguistic means that encode negation, including the standard negator *non*, as in (1), or negative indefinites, as in (2).

- (1) La crescita nasce dove lo stato costruisce fiducia, mobilita capitale pubblico e privato, riduce gli ostacoli burocratici e normativi. E, soprattutto, **non [denial]** penalizza l’iniziativa. (*Il Foglio AI*, 1 April 2025, Editorial 1)  
 [‘Growth emerges where the state fosters trust, mobilizes both public and private capital, and removes bureaucratic and regulatory barriers. Above all, it does **not [denial]** penalize individual initiative’; transl. with ChatGPT 5]
- (2) E l’America dovrebbe fermarsi a riflettere: quando gli scienziati fanno le valigie, **non è mai [denial]** un buon segno. (*Il Foglio AI*, 1 April 2025, Editorial 3).  
 [‘And America should pause to reflect: when scientists pack their bags, it is **never [denial]** a good sign’; transl. with ChatGPT 5]

Negation functions as a modifier which can affect the propositional content of a sentence in various ways (Lehmann 2008). In declaratives – the by large most frequent speech act type attested in the subcorpora, as illustrated in (1) and (2) –, negation markers operate as logical operators with respect to the truth conditions of the situation described (De Cesare 2019: 57). They acknowledge the existence of an alternative affirmative proposition only to reject it. They thereby reverse the relation between the presupposed and asserted propositional content, signaling that the described state of affairs does not obtain (Bernini 2024). In the analysis, constituent negation, yes–no interrogatives, and restrictive or expletive negation were excluded, as these negations function as focalisers and interactional markers, rather than as logical operators (Bernini 2024: 273). Morphological negation (for example *inatteso* ‘unforeseen’) was likewise not considered. The comparatively higher incidence of sentential negation in the AI-generated editorials suggests a stronger tendency to engage with potential opposing viewpoints, thereby foregrounding argumentative polarity to a greater extent than in the human-written texts.

Another difference between the two subcorpora emerges in the use of counter markers. These items introduce a proposition that replaces an alternative one presumed to be expected in the same context. Examples (3) to (5) illustrate typical distributional patterns observed in the data. The marker *però* (‘but’), when used as a parenthetical connective in a medial position where it encodes a contrast arising from the denial of an expectation (*contrasto controaspettativo*, Scorretti 1988), as in (3) – appears exclusively in the human-written editorials (8 occurrences).<sup>7</sup> By contrast, the marker *eppure* (‘and yet’) occurs only in the AI-generated texts (9 occurrences). Interestingly, *eppure* carries a strong evaluative

<sup>7</sup> Note that *però* (‘but’) may also occur in a sentence-initial position as a general adversative conjunction, as in *Piero gioca bene, però perde in continuazione* (example from Josep Cuenca et al. 2019: 19). In these contexts, *però* does not express counter-argumentation, but introduces a contrast with the information presented in the preceding clause (Scorretti 1988). No instances of sentence-initial *però* are attested in the data.

force, typically signaling an unexpected contrast and functioning as a more marked and stylistically elevated marker.

- (3) Europa e Cina hanno la responsabilità di sostenere “un forte sistema commerciale riformato, libero, equo e fondato su condizioni di parità”, ha detto von der Leyen. Qualche controindicazione, **però [counter]**, esiste ed è stata ben rappresentata dagli economisti che stanno collaborando con le istituzioni europee alla discussione sugli effetti dei dazi. (*Il Foglio*, 9 April 2025, Editorial 2)  
 [‘Europe and China have the responsibility to support ‘a strong, reformed, free, and fair trading system based on a level playing field,’ said von der Leyen. Some drawbacks, **however, [counter]** do exist and have been clearly outlined by the economists who are working with European institutions in the discussion on the effects of tariffs’; transl. with ChatGPT 5]
- (4) Siamo al capolinea del modello redistributivo come lo conosciamo. **Eppure [counter]**, le riforme strutturali del lavoro e del welfare restano fantasmi da evocare solo in campagna elettorale (*Il Foglio AI*, 3 April 2025, Editorial 2)  
 [‘We have reached the end of the redistributive model as we know it. **And yet [counter]**, structural reforms of labour and welfare remain ghosts that are conjured up only during election campaigns’; transl. with ChatGPT 5]

By contrast, the standard counter-expectational marker *ma* (‘but’), typically used in a sentence-internal position, as in (5), appears two and a half times more frequently in the AI-generated texts than in the human-written ones (117 occurrences, 81.8% vs. 46 occurrences, 63.0%).<sup>8</sup> This asymmetry suggests that, while AI-generated discourse overgeneralizes the use of *ma* (‘but’), human authors rely more heavily on a stylistically diversified range of markers to encode counter-expectational relations (11 types vs. 5 types). These findings are consistent with stylometric analyses across more than 250 linguistic features, which indicate that human texts exhibit substantially greater variability in linguistic properties than LLM-generated texts (see, for example, Beguš 2024, Zanotto and Aroyehun 2024).

- (5) In tutto questo, l’Europa dovrebbe farsi qualche domanda. Il suo zelo regolatorio ha fondamenti legittimi, **ma [counter]** rischia di tradursi in accanimento selettivo. (*Il Foglio AI*, 2 April 2025, Editorial 2)

---

<sup>8</sup> Note that in sentence-initial position, *ma* (‘but’) can also have other functions. It can establish a contrast with the information presented in the preceding sentence. In these contexts, *ma* (‘but’) functions as a connector between textual blocks and serves to introduce a new situation compared to the one previously presented, that is, a shift in information focus. Moreover, *ma* (‘but’) in sentence-initial position can function as a discourse marker (Scorretti 1988).

[‘In all this, Europe should be asking itself some questions. Its regulatory zeal has legitimate foundations, **but [counter]** it risks turning into selective persecution’; transl. with ChatGPT 5]

Another crucial difference lies in the fact that, in the AI-generated editorials, counters frequently co-occur with denials, where the negated proposition is contrasted with an expectation established by the following statement. These deny-and-counter pairings form a highly salient pattern, occurring at least once in every AI-text. The counter is consistently realized through *ma* (‘but’), which expresses in these contexts what Scorretti (1988) defines as ‘exclusive adversativity’ (*avversatività esclusiva*) – that is, the negation of propositional content previously introduced, deemed incompatible with the propositional content followed by *ma*. In other words, *ma* (‘but’) establishes a corrective relation in which the incompatibility between two propositional segments is expressed by a polemic negation (Anscombe and Ducrot 1977) that “renders the first segment (S1) invalid and anticipates the occurrence of a valid alternative” (Izutsu 2008: 668). However, this form of adversativity does not produce a genuine argumentative move; rather, it functions as a simple polyphonic reformulation of the preceding statement, which is implicitly attributed either to the interlocutor or to general knowledge (Fuentes Rodríguez 1998). In the AI-output, these correctives generate highly predictable rhetorical patterns that simulate argumentative structures without engaging in actual debate. The recurring presence of these patterns across all AI-generated texts suggests that we are most likely dealing with a default discourse template arising from probabilistic language modeling.

- (6) Il guaio della demografia **non [deny]** sono i dati, **ma [counter]** la politica che continua a ignorarli (*Il Foglio AI*, 3 April 2025, Editorial 2)  
 [‘The problem with demographics **isn’t [deny]** the data, **but [counter]** the politics that keeps ignoring it.’; transl. with ChatGPT 5]
- (7) È un’inversione di marcia epocale. E **non [deny]** viene da un dittatore del Sud globale, **ma [counter]** dalla più grande economia del mondo. (*Il Foglio AI*, 4 April 2025, Editorial 2)  
 [‘It’s a historic U-turn. And it **doesn’t [deny]** come from a Global South dictator, **but [counter]** rather from the world’s largest economy’; transl. with ChatGPT 5]
- (8) L’Europa può cogliere l’opportunità per riscrivere il suo ruolo **non [deny]** come mercato passivo, **ma [counter]** come incubatore di innovazione. (*Il Foglio AI*, 6 May 2025, Editorial 2)

[‘Europe can seize the opportunity to redefine its role **not** [*deny*] as a passive market, **but** [*counter*] as an incubator of innovation’; transl. with ChatGPT 5]

In some deny-and-counter pairings, *ma* (‘but’) is only implied rather than explicitly expressed, as in (9).

- (9) Se li fa l’Europa, sono un tradimento del suo stesso modello: aperto, multilaterale, regolato, razionale. L’Europa **non** [*denial*] è forte quando imita i bulli, Ø [*counter*] è forte quando li contraddice. (*Il Foglio AI*, 10 April 2025, Editorial 1)  
 [‘If Europe adopts them, it betrays its own model: open, multilateral, regulated, rational. Europe is **not** [*denial*] strong when it imitates bullies; Ø [*counter*] it is strong when it contradicts them’; transl. with ChatGPT 5]

The patterns described in (6) to (9) are almost absent in the human-written texts, suggesting that AI-generated discourse tends to overuse specific axiological frameworks. This reliance reflects a somewhat black-and-white mode of reasoning that may lack the nuance of human-written editorials. The repetitive use of deny-and-counter pairings clearly illustrates that adversativity in AI discourse functions with respect to our data and until now not as a site of ideological struggle, but rather as a patterned device for providing textual coherence within the probabilistic domain of the LLM(s) used.

### 3.3.2 Proclaim

Table 9 shows that the two subtypes of proclaim – pronounce and concur – are only marginally represented in the subcorpora, accounting for a small proportion of all engagement markers.<sup>9</sup> Their low frequency suggests that neither human authors nor the AI-generated editorials draw extensively on explicit assertive reinforcement to negotiate alignment with the reader.

Table 9: Absolute and normalized frequency of proclaim markers in the subcorpora under study.

<i>Proclaim</i>	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	%	N	%
<i>Pronounce</i>	27	5.1	28	8.3
<i>Concur</i>	19	3.3	19	5.6

<sup>9</sup> In the present analysis, unlike Martin and White (2005), we do not treat endorse markers as features of proclaim (see Section 2). Instead, we classify them under the category attribute (see Section 3.3.4), as they involve citing external sources with which the author aligns more or less explicitly to support their own argumentation.

The first subtype, pronouncements, comprises expressions through which the writer intervenes to emphatically assert the validity of a proposition. In both subcorpora, such interventions do not foreground the writer's subjectivity explicitly as there are no first-person expressions such as *penso che* ('I think that'), but instead introduce the authorial voice more indirectly, in line with established journalistic writing traditions. A recurrent formula is *il problema è che* ('the problem is that'), which functions to pre-empt dialogic resistance and frame the subsequent clause as non-negotiable, as shown in (10) from the human-written texts (4 occurrences) and (11) from the AI-generated texts (3 occurrences). The expression presupposes the existence of an alternative viewpoint, which is then overridden through emphatic restatement.

- (10) Dalle nostre parti si può parlare con più serenità, ma **il problema è che [pronounce]** non siamo noi europei a gestire la partita, se non attraverso i tentativi di dialogo diplomatico. (*Il Foglio*, 1 April 2025, Editorial 1)  
 ['In our parts we can speak more calmly, but **the problem is that [pronounce]** it's not we Europeans who are managing the game, except through attempts at diplomatic dialogue'; transl. with ChatGPT 5]
- (11) Perché non anticipare l'opposizione interna, raddrizzare la curva del dissenso, mettere in ordine le voci storte del sistema? **Il problema è che [pronounce]** il dissenso non è una deviazione: è la condizione di base di ogni democrazia. (*Il Foglio AI*, 9 April 2025, Editorial 1)  
 ['Why not anticipate internal opposition, straighten the curve of dissent, put in order the crooked voices of the system? **The problem is that [pronounce]** dissent is not a deviation: it is the basic condition of every democracy.']; transl. with ChatGPT 5]

Other recurring pronouncement markers signaling that the author's position is presented as the only true interpretation of a state-of-affairs include, for example, *davvero* ('truly'), as in (12) from the human-written editorials (1 occurrence), and (13) from the AI-generated editorials (4 occurrences). However, owing to the marginal status of pronouncement markers in the data, no meaningful conclusions about their distribution in the two subcorpora can be made.

- (12) Insomma, la Confindustria di fronte al bailamme che sta rivoltando mercati finanziari ed economia reale ha trovato una scorciatoia e il presidente l'ha illuminata. Ma è **davvero [pronounce]** così? Si può pensare che davanti alle novità che stanno ridisegnando il volto dell'economia e della geopolitica degli anni Venti si possa imboccare il viottolo più facile e scontato? (*Il Foglio*, 8 April 2025, Editorial 1)  
 ['Because at a time when democracy is **truly [pronounce]** under pressure – in Ukraine, in Georgia, in the Balkans, in our very own parliaments – it

would be nice to see a bit more courage, a bit more critical spirit, from those who call themselves patriots’; transl. with ChatGPT 5]

- (13) Se **davvero** [*pronounce*] vogliamo rispondere a Trump, facciamolo con intelligenza: più accordi commerciali, più investimenti nella competitività europea, più libertà economica. Non con la solita sceneggiata daziaria da comizio dell’Alabama. (*Il Foglio AI*, 10 April 2025, Editorial 1)  
[‘If we **truly** [*pronounce*] want to respond to Trump, then let’s do it intelligently: more trade agreements, more investment in European competitiveness, more economic freedom. Not with the usual tariff-show theatrics fit for a rally in Alabama.’; transl. with ChatGPT 5]

Moreover, we find convergences regarding the use of *concur* markers. In the human-written and the AI-generated editorials, as in (14) and (15), concurrence is realized through expressions such as *è chiaro che* (‘it is clear that’) (2 occurrences vs. 4 occurrences).

- (14) **È chiaro che** [*concur*] in questa fase tutto è in discussione, ci sarà molto da negoziare, per conto dell’Italia in Europa e per conto dell’Europa con gli Stati Uniti. (*Il Foglio*, 10 April 2025, Editorial 1)  
[‘It is clear that [*concur*] at this stage everything is up for discussion; there will be much to negotiate – on behalf of Italy within Europe, and on behalf of Europe with the United States.’; transl. with ChatGPT 5]
- (15) **Ma il dato generale è chiaro:** [*concur*] l’America non crede più nel commercio globale. (*Il Foglio AI*, 4 April 2025, Editorial 2)  
[‘But the general fact is clear: [*concur*] America no longer believes in global trade’; transl. with ChatGPT 5]

However, the analysis of the data also reveals fundamental differences between the editorials from *Il Foglio* and its AI counterpart. Specifically, in the generated editorials, markers of concurrence often serve as a prelude to a counter, invariably introduced by *ma* (‘but’), as in (16) to (18). These *concur*-and-counter pairings form typical patterns of concession: the writer first signals agreement with the reader on a given point, only to then step back and reject what appears to be the logical conclusion drawn from that point. From a cognitive perspective, concession concerns the denial of a counter-expectation and involves “some ground assumption or expectation” (Izutsu 2008: 661) which “is implicit in the sentence itself or in the context” (Mauri 2008: 144).

- (16) Punire Le Pen, **sì** [*concur*]. **Ma** [*counter*] non umiliare la democrazia. (*Il Foglio AI*, 4 April 2025, Editorial 3)

[‘Punish Le Pen, **yes** [*concore*]. **But** [*counter*] do not humiliate democracy’; transl. with ChatGPT 5]

- (17) L’Italia, **certo** [*concore*], deve tenere i conti in ordine – anche per non farsi impallinare dallo spread. **Ma** [*counter*] il vero ordine viene se i fondamentali migliorano. (*Il Foglio AI*, 1 April 2025, Editorial 1)  
 [‘Italy, **of course** [*concore*], must keep its accounts in order – also to avoid being shot down by the spread. **But** [*counter*] real order comes only if the fundamentals improve.’; transl. with ChatGPT 5]
- (18) La demografia non è destino, dicono alcuni. **Vero** [*concore*]. **Ma** [*counter*] l’indifferenza, quella sì che lo è. (*Il Foglio AI*, 3 April 2025, Editorial 2)  
 [‘Demography is not destiny, some say. **True** [*concore*]. **But** [*counter*] indifference – that truly is.’; transl. with ChatGPT 5]

From a dialogic point of view, such concur-and-counter patterns can be seen as gestures toward solidarity, especially in contexts where the writer anticipates potential disagreement from the reader (Martin and White 2005: 125). This quest for solidarity is largely absent in the human-written editorials, suggesting that the editorialists assume the existence of a shared common ground with their audience. Writing from within a discursive community, they can invoke shared premises implicitly rather than negotiating them explicitly. In contrast, the AI-generated texts reproduce highly standardized patterns of concession. Generated in a context-free environment, these texts operate without any established writer-reader relationship and cannot rely on communal assumptions. As a result, the LLM compensates by overproducing concession markers in a way that appears at first glance interactive but remains fundamentally formal. This suggests that we are once again dealing with a default discourse template that ultimately functions as a generalized hedging strategy designed to maximize acceptability across potential readerships. What presents itself as dialogic openness may therefore signal the absence of a concrete addressee.

### 3.3.3 Entertain

Table 10 shows the distribution of the different features of entertainment – that is, markers of epistemic and deontic modality, as well as markers of inferential evidentiality.

Table 10: Absolute and normalized frequency of entertain markers in the subcorpora under study.

<i>Entertain</i>	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	%	N	%
<i>Epistemic modality</i>	53	9.7	31	9.0

<i>Deontic modality</i>	22	4.0	5	1.4
<i>Inferential evidentiality</i>	9	1.6	5	1.4

Although entertainment markers make up only a small proportion of all engagement markers identified in the subcorpora, the expression of epistemic modality clearly emerges as the most prominent subcategory within entertain. The data reveal no significant difference in the expression of epistemic modality, as illustrated by the adverb *forse* ('perhaps') in the examples (19) and (20) taken from the human-written and AI-generated editorials (8 occurrences vs. 9 occurrences). Throughout the two subcorpora, the primary communicative function of epistemic modal markers is not so much to comment on the informational reliability of a proposition, but rather to indicate that the writer is offering a subjective evaluation. The marker thus signals that the proposition is open to contestation and acknowledges the potential presence of readers who may not share the writer's stance.

- (19) **Forse [entertain]** l'errore degli imprenditori è stato quello di non accettare neppure un incontro, nel quale avrebbero potuto illustrare le ragioni della loro impossibilità ad assumersi oggi impegni quando non sanno che cosa accadrà nell'immediato futuro alle loro attività e alle loro produzioni. (*Il Foglio*, 9 April 2025, Editorial 3)  
 ['Perhaps [entertain] the mistake made by the entrepreneurs was that of not even agreeing to a meeting, in which they could have explained the reasons for their inability to make commitments today, when they do not know what will happen in the immediate future to their businesses and their production'; transl. with ChatGPT 5]
- (20) Quando a gridare al disastro non è un accademico del World Economic Forum ma Bill Ackman, uno che ha fatto miliardi scommettendo sui disastri altrui, **forse [entertain]** vale la pena ascoltarlo. (*Il Foglio AI*, 8 April 2025, Editorial 3)  
 ['When the one crying disaster is not a World Economic Forum academic but Bill Ackman – a man who made billions betting on other people's disasters – perhaps [entertain] it's worth listening to him.']; transl. with ChatGPT 5]

The findings also indicate symmetry in the occurrence of inferential evidentiality markers, which are marginal in both subcorpora. In the AI-generated and the human-written editorials, epistemic justification by reference to internal cognitive processes is for example conveyed by the verb *sembrare* ('to seem') (3 occurrences vs. 2 occurrences).

- (21) Donald Trump **sembra** [*entertain*] ottenere più risultati con una sola mossa: fa ribassare il prezzo della benzina, riesce a devastare il bilancio dell'economia di guerra della Russia meglio di quanto non abbiano fatto anni di sanzioni, potrebbe mettere in crisi l'industria petrolifera statunitense dello shale oil. (*Il Foglio AI*, 6 May 2025, Editorial 3)  
 ['Donald Trump **seems** [*entertain*] to achieve several results with a single move: he pushes down the price of gasoline, manages to damage Russia's war economy more effectively than years of sanctions have done, and could even put the U.S. shale oil industry in crisis.']; transl. with ChatGPT 5]
- (22) L'audizione del presidente dell'Istat Francesco Maria Chelli alla Commissione sulla transizione demografica è uno di questi: una radiografia accurata, aggiornata, inequivocabile di un paese che sta invecchiando a un ritmo da record mondiale. E che, a parte numerare i problemi, non **sembra** [*entertain*] avere ancora deciso cosa farne. (*Il Foglio*, 3 April 2025, Editorial 2)  
 ['The hearing of Istat president Francesco Maria Chelli before the Commission on demographic transition is one such case: an accurate, up-to-date, and unequivocal X-ray of a country that is aging at a record global pace. And apart from listing its problems, it still does not **seem** [*entertain*] to have decided what to do about them.']; transl. with ChatGPT 5]

Clear disparities, however, can be observed regarding the use of deontic expressions, particularly markers of obligation, which occur roughly three times more frequently in the AI-generated editorials. These texts thus adopt a more prescriptive tone, prioritizing calls for action over the mere presentation of viewpoints. Expressions involving the auxiliary *dovere* ('must') are especially prominent. Such directives may be addressed to third parties of which the reader is implicitly a part, as in (23), or directly to readers, as in (24) and (25). While the former can also be found in the human-written texts (4 occurrences vs. 12 occurrences in the AI-texts), direct forms of reader address are only attested in the AI-generated editorials (5 occurrences).

- (23) L'Europa **deve** [*entertain*] rispondere con chiarezza e con fermezza politica. Affermare che la sovranità danese è non negoziabile. (*Il Foglio AI*, 6 May 2025, Editorial 3)  
 ['Europe **must** [*entertain*] respond with clarity and political firmness. It must affirm that Danish sovereignty is non-negotiable.']; transl. with ChatGPT 5]
- (24) Trump non cambierà per noi. Siamo noi che **dobbiamo** [*entertain*] cambiare sguardo su di lui. (*Il Foglio AI*, 4 April 2025, Editorial 4)

[‘Trump will not change for us. It’s we who **must** [*entertain*] change our perspective on him.’; transl. with ChatGPT 5]

- (25) E che, se vogliamo difendere l’Europa, **dobbiamo** [*entertain*] prima di tutto smascherare chi la deride. Anche – soprattutto – quando si finge nostro amico. (*Il Foglio AI*, 2 April 2025, Editorial 3)  
 [‘And that, if we want to defend Europe, we **must** [*entertain*] first of all unmask those who mock it – even, and above all, when they pretend to be our friends.’; transl. with ChatGPT 5]

The auxiliary grounds these directives in the writer’s subjectivity, presenting the obligation as a personal assessment. It implicitly signals that the writer anticipates either acceptance or rejection of the proposed action, thereby leaving room for potential objections. The frequent use of such deontic expressions in the AI-generated texts may once again be linked to the fact that these texts are produced in a context-free environment and operate without any established writer-reader relationship. To compensate, the model appears to overproduce obligation markers, creating the impression of interactivity through repeated collective calls to action – features that are entirely absent from the human-authored texts.

### 3.3.4 Attribute

Table 11 summarizes the distribution of the two attribution subcategories acknowledge and distance/endorse across the data.

Table 11: Absolute and normalized frequency of attribute markers in the subcorpora under study.

Attribute	Editorials from <i>Il Foglio AI</i>		Editorials from <i>Il Foglio</i>	
	N	%	N	%
<i>Acknowledge</i>	32	5.9	51	14.9
<i>Distance/Endorse</i>	20	3.6	28	8.2

In both subcorpora, approximately two-thirds of all attribution markers fall under the subcategory acknowledge – that is, instances where there is no overt indication, at least through the choice of the framer, of the author’s stance toward the cited proposition. The human-written editorials display a broader range of framers (22 types) than their AI-generated counterparts (15 types), which indicates a less varied realization of referential evidentiality in the AI texts. Acknowledge typically includes neutral reporting verbs, as shown in examples (26) and (27), taken from the AI-generated and human-written editorials respectively, such as *dire* (‘to say’, 22 occurrences vs. 17 occurrences), *annunciare* (‘to announce’, 1 occurrence vs. 3 occurrences), *dichiarare* (‘to declare’, 0 occurrences vs. 2 occurrences), *spiegare* (‘to explain’, 2 occurrences vs. 3 occurrences), *aggiungere* (‘to add’, 2 occurrences vs. 2 occurrences) or *decidere* (‘to decide’, 0 occurrences vs. 1 occurrence). It is, of course, possible that indicators elsewhere in the text suggest that the writer more

broadly supports or rejects the position being advanced. This, however, is a separate issue that lies beyond the scope of the present analysis.

- (26) Davide Serra, fondatore di Algebris, **ha annunciato che [acknowledge]** il suo fondo voterà a favore dell’operazione all’assemblea di Mps del prossimo 17 aprile. “Abbiamo deciso: voteremo, siamo azionisti e supportiamo l’operazione, ci sembra corretta, intelligente”, **ha detto Serra [acknowledge]**, definendo un “insulto” il suggerimento del proxy advisor Iss di bocciare l’aumento di capitale propedeutico all’Ops. “E’ un insulto delegare il diritto di voto e la conoscenza dell’azienda a qualcun altro”, **ha aggiunto [acknowledge]**, **spiegando che [acknowledge]** in diciotto anni Algebris non ha mai seguito le raccomandazioni di una proxy. (*Il Foglio AI*, 1 April 2025, Editorial 2)

[‘Davide Serra, founder of Algebris, **announced that [acknowledge]** his fund will vote in favor of the transaction at the MPS shareholders’ meeting on April 17. “We have decided: we will vote, we are shareholders, and we support the operation – it seems fair and intelligent to us,” **Serra said [acknowledge]**, calling “an insult” the proxy advisor ISS’s suggestion to reject the capital increase preparatory to the public exchange offer. “It’s an insult to delegate one’s voting rights and knowledge of the company to someone else,” **he added [acknowledge]**, **explaining that [acknowledge]** in eighteen years Algebris has never followed a proxy’s recommendations.’; transl. with ChatGPT 5]

- (27) CK Hutchison **ha deciso che [acknowledge]** “l’accordo sui due porti panamensi non verrà firmato ufficialmente la prossima settimana”, secondo il South China Morning Post. Il consigliere per la sicurezza nazionale di Trump, Mike Waltz, **aveva dichiarato** ai giornalisti alla fine del mese scorso **che [acknowledge]** i leader panamensi avevano avviato “negoziati per affrontare la questione dei porti su entrambe le sponde del canale”. (*Il Foglio*, 1 April 2025, Editorial 2)

[‘CK Hutchison **has decided that [acknowledge]** “the agreement on the two Panamanian ports will not be officially signed next week,” according to the South China Morning Post. Trump’s national security adviser, Mike Waltz, **had told** reporters at the end of last month **that [acknowledge]** Panamanian leaders had begun “negotiations to address the issue of the ports on both sides of the canal.”’; transl. with ChatGPT 5]

In both subcorpora, approximately one-third of the attribution markers involve expressions in which, through the semantics of the framer, the author explicitly evaluates the cited voice – almost exclusively in the form of distancing. In the human-written editorials, this is typically realized through reporting verbs such as *affermare* (‘to claim’, 4 occurrences) and *sostenere* (‘to maintain’, 3 occurrences),

as illustrated in (28) and (29). Interestingly, these verbs are not attested in the AI-generated texts.

- (28) Venerdì il presidente Trump **aveva affermato trionfante**: [*distance*] “La spesa alimentare (e le uova!) è in calo” (*Il Foglio*, 6 May 2025, Editorial 3) [‘On Friday, President Trump **had triumphantly declared**: [*distance*] “Food prices (and eggs!) are down.”; transl. with ChatGPT 5]
- (29) Durante la pandemia [Robert F. Kennedy Jr.] era venuto anche a Milano a fare comizi, **sostenendo che** [*distance*] il Covid fosse utilizzato dalle élite globali per instaurare forme di controllo totalitario. (*Il Foglio*, 8 April 2025, Editorial 3) [‘During the pandemic, [Robert F. Kennedy Jr.] even came to Milan to hold rallies, **claiming that** [*distance*] Covid was being used by the global elites to establish forms of totalitarian control.’; transl. with ChatGPT 5]

In the editorials from *Il Foglio AI*, distancing is often realized by so-called scare quotes (10 occurrences), as illustrated in (30) and (31). Such uses occur less frequently in the human-written texts (6 occurrences).

- (30) L’Italia avrebbe dovuto – **testuale** – “**cooperare con la Corte inclusa l’esecuzione dei mandati di arresto**”. [*distance*] (*Il Foglio AI*, 11 April 2025, Editorial 1) [‘Italy should have – **verbatim** – “**cooperated with the Court, including the execution of arrest warrants**.” [*distance*]; transl. with ChatGPT 5]
- (31) Ed Elon Musk, che ha investito personalmente oltre 20 milioni di dollari, promesso premi in denaro ai votanti e definito il voto “**decisivo per il destino della civiltà occidentale**” [*distance*] (*Il Foglio AI*, 3 April 2025, Editorial 3) [‘And Elon Musk, who personally invested over 20 million dollars, promised cash prizes to voters, and called the vote “**decisive for the fate of Western civilization**” [*distance*]; transl. with ChatGPT 5]

#### 4 Conclusion

In this study, we have investigated how well large language models can imitate human-authored texts in expressing authorial stance with regard to intersubjectivity. By comparing editorials published in *Il Foglio* and *Il Foglio AI* through the lens of Martin and White’s (2005) engagement typology, our analysis reveals clear distinctions between human and AI-generated stance profiles, highlighting certain limitations in LLMs’ ability to mark intersubjective positionings in a contextually appropriate manner.

The results show that engagement resources occur with a higher normalized frequency in the AI-generated editorials, driven primarily by dialogically contractive categories. In particular, disclaim markers are more predominant in the AI texts, whereas human-authored editorials display a more balanced distribution across the four engagement categories. Two recurrent templatic configurations particularly characterize the AI output: deny-counter pairings and concur-counter sequences. These patterns simulate dialogic responsiveness while minimizing engagement with competing positions. Within entertain, AI editorials employ features of deontic obligation with much greater frequency, including direct appeals to readers, thus adopting a slightly more prescriptive stance. Finally, the higher incidence of attribute markers in the human subcorpus indicates a stronger incorporation of external voices and a correspondingly wider dialogic space. In a nutshell, *Il Foglio AI* editorials exhibit a stance profile that is more explicit, contractive, prescriptive, and templated than the one found in the human-authored texts.

These results are consistent with previous studies that state LLMs' limitations in pragmatic grounding and context sensitivity. In the absence of established writer-reader relations and shared situational knowledge, LLMs appear to compensate through heightened explicitness of intersubjective positioning, implemented via highly routinized argumentative patterns. At the same time, the results diverge from findings in other stance-intensive genres, such as academic reviews and student essays, where the AI-generated texts have been found to be less evaluative and more informationally neutral. This divergence opens important avenues for further investigation into the genre-contingent nature of LLM stylistics.

It is important to note that our results represent only a temporal snapshot: AI models continue to evolve, and future outputs are likely to differ. The study also faces other limitations. Its empirical scope is restricted to a single outlet and a narrow time frame, and the model configuration underlying *Il Foglio AI* remains opaque. Crucially, the specific prompts and generation parameters employed are unknown. As a result, the findings cannot be directly compared with those of studies based on documented prompting settings, and any interpretation must therefore remain provisional. Future research should expand the sampling frame across a wider setting of journalistic outlets and genres and consider additional features of stance marking, such as graduation and attitude markers, to provide a more comprehensive understanding of how LLMs construct evaluative voice across mass media discourse.

## References

- Anscombe, Jean-Claude & Ducrot, Oswald. 1977. Deux *mais* en français ? *Lingua* 43. 23–40. [https://doi.org/10.1016/0024-3841\(77\)90046-8](https://doi.org/10.1016/0024-3841(77)90046-8).
- Bakhtin, Mikhail Mikhailovich. 1981. *The Dialogic Imagination*. Tr. by Caryl Emerson & Michael Holquist. Austin: University of Texas Press.
- Beguš, Nina. 2024. Experimental narratives: A comparison of human crowdsourced storytelling and AI storytelling. *Humanities and Social Sciences Communications* 11, 1392 (2024). <https://doi.org/10.1057/s41599-024-03868-8>.
- Bender, Emily & Koller, Alexander. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.463>.
- Bernini, Giuliano. 2024. Negation and negative expressions. In De Cesare, Anna-Maria & Salvi, Giampaolo (eds), *Manual of Romance Word Classes*, 265–296. Berlin/Boston: Walter de Gruyter. <https://doi.org/10.1515/9783110746389-011>.
- Dardano, Maurizio. 1973. *Il linguaggio dei giornali italiani*. Roma: Laterza.
- De Cesare, Anna-Maria. 2019. *Le parti invariabili del discorso*. Roma: Carocci.
- De Cesare, Anna-Maria. 2023. Assessing the quality of ChatGPT’s generated output in light of human-written texts. A corpus study based on textual parameters. *CHIMERA: Romance Corpora and Linguistic Studies* 10. 179–210. <https://doi.org/10.3145/epi.2023.sep.14>.
- Dierickx, Lisa. 2021. Journalisme algorithmique : un état de l’art de la recherche. In Dierickx, Lisa (ed), *Journalisme algorithmique. Les carnets du Laboratoire des pratiques et des identités journalistiques, Les Carnets du LaPIJ* 2. 8–10.
- Draghi, Carlo. 2005. Il giornale attivista. I dieci anni del *Foglio* di Giuliano Ferrara. *Problemi dell’informazione* 4. 414–437. <https://www.rivisteweb.it/doi/10.1445/21048>.
- Ferrigolo, Alberto. 2002. Editoriali, opinioni, commenti. Fenomenologia del genere e dei suoi addetti. *Problemi dell’informazione* XXVII (1). 5–34. <https://www.rivisteweb.it/doi/10.1445/4155>.
- Foschi Albert, Marina. 2009. Il profilo stilistico del testo. Guida al confronto intertestuale e interculturale (tedesco e italiano). Pisa: PLUS–Pisa University Press.
- Fuentes Rodríguez, Catalina. 1998. *Las construcciones adversativas*. Madrid: Arco Libros.

- Fuller, Gillian. 1998. Cultivating science: negotiating discourse in the popular texts of Stephen Jay Gould. In Martin, James & Veel, Robert (eds), *Reading Science: Critical and Functional Perspectives on Discourses of Science*, 35–62. London: Routledge.
- González-Arias, Carmen & Chatzikoumi, Eleni & López-García, Xosé. 2024. The anthropomorphic pursuit of AI-generated journalistic texts: limits to expressing subjectivity. *Frontiers in Communication* 9. 1456509. <https://doi.org/10.3389/fcomm.2024.1456509>.
- Graefe, Andreas & Bohlken, Nina. 2020. Automated journalism: a meta-analysis of readers' perceptions of human-written in comparison to automated news. *Media and Communication* 8. 50–59. <https://doi.org/10.17645/mac.v8i3.3019>.
- Gualdo, Riccardo. 2017. *L'italiano dei giornali*. Roma: Carocci editore.
- Gutiérrez-Caneda, Beatriz & Vázquez-Herrero, Jorge & López-García, Xosé. 2023. AI application in journalism: ChatGPT and the uses and risks of an emergent technology. *Profesional de la información* 32. e320514. <https://doi.org/10.3145/epi.2023.sep.14>.
- Haßler, Gerda. 2018. Evidential and epistemic sentence adverbs in Romance languages. *Linguistik Online* 92 (5). <https://doi.org/10.13092/lo.92.4506>.
- Herbold, Steffen & Hautli-Janisz, Annette & Heuer, Ute & Kikteva, Zlata & Trautsch, Alexander. 2023. A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific Reports* 13. 18617. <https://doi.org/10.1038/s41598-023-45644-9>.
- Hesselbach, Robert. 2024. Investigating the Relation between Syntactic Complexity and Subgenre Distinction: A Case Study on Two Contemporary French Authors. In Hesselbach, Robert & Calvo Tello, José & Henny-Krahmer, Ulrike & Schöch, Christof & Schlör, Daniel (eds), *Digital Stylistics in Romance Studies and Beyond*, 235–259. Heidelberg: Heidelberg University Publishing. <https://doi.org/10.17885/heiup.1157.c19374>.
- Iedema, Rick & Feez, Susan & White, Peter Robert Rupert. 1994. *Media Literacy (Write it Right Literacy in Industry Research Project – Stage 2)*. Sydney: Metropolitan East Disadvantaged Schools Program.
- Izutsu, Mitsuko Narita. 2008. Contrast, concessive, and corrective: Toward a comprehensive study of opposition relations. *Journal of Pragmatics* 40 (4). 646–675. <https://doi.org/10.1016/j.pragma.2007.07.001>.
- Jiang, Feng Kun & Hyland, Ken. 2025. Does ChatGPT write like a student? Engagement markers in argumentative essays. *Written Communication* 42. 463–492. <https://doi.org/10.1177/07410883251328311>.
- Josep Cuenca, Maria & Postolea, Sorina & Visconti, Jacqueline. 2019. Contrastive Markers in Contrast. *Discours*. <https://doi.org/10.4000/discours.10326>.

- Kilgarriff, Adam. 2009. Simple maths for keywords. In Mahlberg, Michaela & González-Díaz, Viviana & Smith, Catherine (eds), *Proceedings of Corpus Linguistics Conference (CL 2009)*. Liverpool: University of Liverpool.
- Lehmann, Christian. 2008. Roots, stems and word classes. *Studies in Language* 32. 546–567. <https://doi.org/10.1075/sl.32.3.04leh>.
- Lohnstein, Horst. 2011. *Formale Semantik und natürliche Sprache*. Berlin/Boston: De Gruyter. <https://doi.org/10.1515/9783110223880>.
- Lopezosa, Carlos & Pérez-Montoro, Mario & Rey Martín, Carlos. 2024. El uso de la inteligencia artificial en las redacciones: propuestas y limitaciones. *Revista de Comunicación* 23. 2227–1465. <https://doi.org/10.26441/RC23.1-2024-3309>.
- Lüger, Heinz-Helmut. 2017. Media linguistics: interfaces to media and communication studies. In Bedijs, Kristina & Maaß, Christiane (eds), *Manual of Romance Languages in the Media*, 10–31. Berlin/Boston: Walter de Gruyter. <https://doi.org/10.1515/9783110314755-002>.
- Marín Arrese, Juana Isabel. 2015. Epistemicity and stance: A cross-linguistic study of epistemic stance strategies in journalistic discourse in English and Spanish. *Discourse Studies* 17 (2). 210–225. <https://doi.org/10.1177/1461445614564523>.
- Markowitz, David M. & Hancock, Jeffrey T. & Bailenson, Jeremy N. 2024. Linguistic markers of inherently false AI communication and intentionally false human communication: Evidence from hotel reviews. *Journal of Language and Social Psychology* 43 (1). 63–82. <https://doi.org/10.1177/0261927X231200201>.
- Martin, James Robert & White, Peter Robert Rupert. 2005. *The Language of Evaluation: Appraisal in English*. Basingstoke: Palgrave. <https://doi.org/10.1057/9780230511910>.
- Mauri, Caterina. 2008. *Coordination Relations in the Languages of Europe and Beyond*. Berlin/New York: De Gruyter. <https://doi.org/10.1515/9783110211498>.
- Meier, Franz. 2017. *La perception des normes textuelles, communicationnelles et linguistiques en écriture journalistique. Une contribution à l'étude de la conscience linguistique des professionnels des médias écrits québécois*. Frankfurt/Main: Lang. <https://doi.org/10.3726/978-3-653-06999-0>.
- Muñoz-Ortiz, Alberto & Gómez-Rodríguez, Carlos & Vilares, David. 2024. Contrasting Linguistic Patterns in Human and LLM-Generated News Text. *Artificial Intelligence Review* 57, 264–265. <https://doi.org/10.1007/s10462-024-10903-2>.
- Nyholm, Sven. 2020. *Humans and Robots: Ethics, Agency, and Anthropomorphism*. London/New York: Rowman and Littlefield.

- OpenAI. 2021. ChatGPT (Version GPT-5) [Computer software]. <https://openai.com/> (last accessed on 21.10.2025).
- Pounds, Gabrina. 2010. Attitude and subjectivity in Italian and British hard-news reporting: The construction of a culture-specific ‘reporter’ voice. *Discourse Studies* 12 (1). 106–137. <https://doi.org/10.1177/1461445609346777>.
- Puliafito, Alberto. 2025. Il Foglio AI and the automation of opinion journalism. *The Fix*. <https://thefix.media/2025/04/01/il-foglio-ai-and-the-automation-of-opinion-journalism> (last accessed on 03.02.2026).
- Scorretti, Maria. 1988. Le strutture coordinate. In Renzi, Lorenzo (ed.), *Grande grammatica italiana di consultazione*. Vol. 1: *La frase, i sintagmi nominale e preposizionale*, 227–270. Bologna: Il Mulino.
- Stubbs, Michael. 1996. Towards a modal grammar of English: a matter of prolonged fieldwork. In Stubbs, Michael, *Text and Corpus Analysis*. Oxford: Blackwell.
- White, Peter Robert Rupert. 1997. Death, disruption and the moral order: the narrative impulse in mass ‘hard news’ reporting. In Christie, Frances & Martin, James Robert (eds), *Genres and Institutions: Social Processes in the Workplace and School*, 101–133. London: Cassell.
- Yao, Guangyuan & Liu, Zhaoxia. 2025. Exploring artificial intelligence appraisal: Appraisal patterns in GPT-generated and human-authored book reviews. *Applied Linguistics* 2025. 1–24. <https://doi.org/10.1093/applin/amaf064>.
- Zanotto, Sergio & Aroyehun, Segun. 2024. Human Variability vs. Machine Consistency: A Linguistic Analysis of Texts Generated by Humans and Large Language Models. *arXiv*. <https://doi.org/10.48550/arXiv.2412.03025>.
- Zárate, José Carlos Vázquez. 2023. Evaluative language in undergraduate academic writing: Expressions of attitude as sources of text effectiveness in English as a foreign language. *International Review of Applied Linguistics in Language Teaching* 63. 1139–1168. <https://doi.org/10.1515/iral-2023-0103>.
- Zhang, Wenjuan & Cheung, Yin Ling. 2023. The different ways to write publishable research articles: Using cluster analysis to uncover patterns of appraisal in discussions across disciplines. *Journal of English for Academic Purposes* 63. 101231. <https://doi.org/10.1016/j.jeap.2023.101231>.