Measuring Anthropomorphism in Italian Specialized Language on AI

Rossella Resi

University of Innsbruck, Austria

Abstract

The increasing use of anthropomorphism in English specialized discourse about new technologies has raised concerns about potential social implications, such as overreliance on technology, diminished accountability, and loss of critical engagement. Given that the use of anthropomorphic language appears to be a cross-linguistic phenomenon, this paper presents a preliminary investigation into the use of anthropomorphic language in Italian discourse on AI. It adapts the computational metric AnthroScore [1], originally conceived for English texts, to measure anthropomorphism in Italian texts. A corpus of articles was compiled semi-automatically with the help of BootCat and key terms were analyzed using a masked language model to compute anthropomorphism scores. Manual evaluations were conducted to validate the metric's output. Results show a relatively high degree of anthropomorphism in Italian, but also highlight limitations. The metric does not account for suprasegmental and metalinguistic cues that often mitigate anthropomorphic framing. Furthermore, English loanwords in Italian discourse introduce interpretive ambiguities. The study underscores the need for methodological refinements and raises broader questions about how linguistic choices not only shape but may reflect sociocultural attitudes toward AI across languages.

Keywords

Anthropomorphism, AI, computational linguistic measure, Italian, terminology

1. Introduction

This paper presents a preliminary investigation into the use and degree of anthropomorphic language in Italian discourse on Artificial Intelligence (AI)[1]. While previous research has primarily focused on English, this study explores the extent to which metaphorical, and specifically humanizing terms, are employed in Italian. The aim is to test and adapt the metric AnthroScore [1], originally developed for English, to measure the degree of anthropomorphism[2] in Italian public texts.

The paper is structured as follows: Section 2 outlines the role and effects of anthropomorphism in specialized language, particularly in the framing AI. Section 3 examines Italian-specific linguistic features that may promote or obscure anthropomorphic expressions. Section 4 details the methodology used to assess anthropomorphism in public discourse, beginning with the compilation of the corpus (4.1), followed by the presentation of the computational linguistic metric AnthroScore (4.2), the necessary linguistic and technical adjustments for its application to Italian (4.3), and a manual evaluation of the results (4.4). Preliminary findings are presented in Section 5, revealing a relatively high degree of anthropomorphism in Italian texts, while also identifying limitations that might affect the acceptance of certain scores. The concluding section (6) discusses potential improvements to the methodology and outlines directions for further research. Section 6 provides preliminary observations on the significance of mapping verb terminologization within TMS through the gradual reduction of perceived

5th International Conference on "Multilingual digital terminology today. Design, representation formats and management systems" (MDTT) 2026, June 25-26, 2026, Zadar, Croatia

rossella.resi@uibk.ac.at (R. Resi)

0000-0002-3261-5662 (R. Resi)

[1]The term Artificial Intelligence is used here in a deliberately broad sense to encompass a wide range of computational technologies that have shaped its evolution. This includes machine learning, deep neural networks, and more recent developments like transformers, large language models, multimodal generative models amongst others

[2]In this article, the term anthropomorphism refers specifically to the attribution of human capabilities and mental states to non-sentient objects through linguistic expressions, hence the use of anthropomorphic language, while acknowledging that this represents only one of several possible forms of anthropomorphism associated with technology

anthropomorphism.

2. From analogy to misconception: the impact of anthropomorphic language about AI

Specialized discourse frequently relies on metaphorical language, serving both pedagogical and epistemological functions. Metaphors help render complex or abstract concepts more accessible to lay audiences and act as cognitive tools that support analogical reasoning. They also facilitate communication among experts, particularly in emerging fields or when dealing with novel phenomena [2, 3, 4]. For instance, terms like *friendly bacteria*, *carbon footprint*, and *economic heartbeat* illustrate how metaphorical expressions, often drawn from human cognition and biological experience, have become integrated and normalized within technical vocabularies. Conversely, excessively technical terminology may create distance, thereby hindering understanding and engagement [5]. [3]

Among all domains, technology has historically been one of the most anthropomorphized, with human traits often projected onto machines and systems [5, 6]. In recent years, however, this anthropomorphic tendency has intensified significantly, particularly in both scientific and popular discourses on AI [7, 8]. One key factor behind this phenomenon is that many of the most impactful advancements in machine learning are directly or indirectly influenced by theories in neuroscience, cognitive psychology, and social epistemology. As a result, terms such as *memory*, *learning*, and *neural network*, as well more recent terms like *thinking model* (introduced into scientific discourse in January 2025) and, similarly, verbs like *to understand*, *to think*, *to reason*, or *to hallucinate* suggest intentionality and mental states that these systems do not, in fact, possess. Sometimes whether such language is to be taken metaphorically or literally is also ambiguous [9]. Beyond analogy, other mechanisms have also been identified as drivers of this anthropomorphizing trend: (1) the accessibility and applicability of anthropocentric knowledge; (2) the motivation to explain and predict the behaviour of other agents and (3) the desire for social contact and affiliation [5].

The scope of this phenomenon, combined with a general lack of explicit acknowledgment regarding the metaphorical nature of such terms, has raised concerns about its social implications. Anthropomorphism may inadvertently shape public and professional perceptions by fostering misinformation about the actual capabilities of AI [10], promote overreliance [11], decrease accountability [12, 13], and create emotional attachment to non-sentient artifacts [14]. Similar risks have long been documented in human–computer interaction research: the well-known ELIZA effect describes the human tendency to attribute understanding and empathy to computational systems that merely simulate dialogue [15, 16]. But unlike earlier instances, current anthropomorphic framings influence not only lay audiences but also expert discourse, partly due to the cognitive accessibility of human-centered concepts [17]. If so, anthropomorphism can obscure lines of responsibility, shifting accountability away from developers and corporations toward the systems themselves[9, 12].

The growing concerns around the consequences of anthropomorphism in AI are reflected in literature calling for clearer distinctions between metaphorical and literal uses of humanizing terms, hence greater linguistic transparency [9, 18, 19, 20]. These issues echo early warnings from thinkers like Bacon, Hobbes, and Locke, who cautioned against the use of metaphor in scientific discourse for its potential to obscure conceptual clarity [21].

This study wants to contribute to the discussion on the anthropomorphism of domain-specific language. In particular, it aims to investigate a methodology for measuring this phenomenon in Italian. By extending the analysis beyond English, the project seeks to lay the basis to understand whether similar trends, risks, and motivations are at play in other linguistic and cultural contexts and whether cross linguistic differences reflect cross cultural attitudes.

[3] See also extensive references on metaphors and scientific discourse from CIRM (Centro Interuniversitario della Ricerca sulla Metafora) at https://cirm.unige.it/node/8#toc-metaphor-a-n7zv5LrA

3. Anthropomorphism in the Italian specialized discourse around AI

Cognitive linguistics research emphasizes that anthropomorphism is a cross-linguistic phenomenon [5, 20]. All natural languages use it across various genres, including narrative, poetry, advertising, and scientific communication. However, the way anthropomorphism manifests can differ across linguistic and cultural contexts, especially as naming device in specialized discourse [4, 22]. In Italian, recent studies have documented the systematic use of anthropomorphism in popular science communication on AI both as heuristic tool in scientific exploration and communicative device for the effective transmission of knowledge, including across disciplinary boundaries [3, 23, 24].

Anthropomorphism in Italian discourse around AI can be expressed through three main grammatical categories: nouns, adjectives, and verbs. Since in specialized language, adjectives are often integrated into conceptual designations (e.g., *modello pensante* [thinking model]), nominal and verbal phrases remain the two primary categories for examining anthropomorphism. Nominal metaphors tend to be quickly absorbed and standardized within the scientific community, while terminologization of anthropologized verbs is slower and hinges on the degree of standardization versus the perception of the metaphorical use. In the sentence *Il modello pensa e produce una risposta* [the model thinks and produces an answer], the status of term of the verb *pensare* [to think] is controversial. As a result, anthropomorphism is detected, as it is still the case with other verbs that denote cognitive capacities, such as *to learn*. Nevertheless, certain human-like verbs like *allucinare* [to hallucinate], probably also due to their uncommon everyday use as human verbs, may become so conventionalized in scientific discourse that their anthropomorphic origin fades into the background [3, 25]. This linguistic conventionality and conceptual entrenchment may complicate human judgments regarding the presence of anthropomorphism. Furthermore, some verbs, such as *irrompere* [to burst in] in the sentence *l'intelligenza artificiale irrompe nel mondo dell'arte* [AI bursts into the art world], typically select human subjects but can semantically accommodate non-human agents. In such cases, anthropomorphism becomes ambiguous, as it depends less on terminological stability and more on interpretive framing [26]. Therefore, anthropomorphism in Italian should not be seen as a binary phenomenon, either present or absent, but rather as existing along a continuum.

Another important aspect is the intrinsic morphological structure of the language. In Italian, every noun carries grammatical gender (masculine or feminine), whereas in English, gender is typically limited to personal pronouns (he/she) and a small number of biologically animate nouns. This characteristic gives Italian a tendency to attribute a kind of grammatical personality even to inanimate objects and abstract concepts, which can facilitate personification [23]. This could facilitate the process of anthropomorphism.

This study lays the groundwork for evaluating the degree of anthropomorphism in Italian texts by proposing a computational scoring metric and a threshold to be compared to binary human judgement. Since the metaphorical use of language is not limited to the lexical item being anthropomorphized but is shaped by a broader cognitive and linguistic environment, the chosen approach treats terms as units of analysis to delineate the boundaries of anthropomorphized concepts while also examining how these terms interact with the surrounding discourse.

4. Methodology

4.1. The Italian corpus

One of the primary challenges in building a corpus for the Italian language lies in the scarcity of scientific literature specifically focused on AI, that are written in Italian. While AI is often mentioned in Italian-language academic publications in relation to other disciplines, the existence of a consistent, domain-specific scientific metalanguage on AI in Italian remains limited. The corpus contains also high-quality disseminating sources that report on AI-related topics.

To build a targeted dataset of AI-related discourse in Italian, a semi-automated bootstrapping approach using seed terms and tuples was adopted, supported by the BootCaT toolkit [27]. BootCaT

accelerated the identification and extraction of relevant texts from the web; however, manual evaluation of each source was necessary to ensure data quality. The process began with the identification of a set of Italian seed terms for the initial web queries. This initial list consisted of Italian equivalents of the artefact terms used in [1], which followed a comparable methodology for the English language but both for scientific and public discourse:

$X_{\text{artefatto}}$*{algoritmo, sistema, modello, rete neuronale, LLM, architettura}*.

After selecting the initial seeds, BootCaT was used to generate combinations of seed terms (tuples). These tuples helped retrieve documents in which multiple seed terms co-occurred, thereby increasing the likelihood that the retrieved texts were thematically relevant to AI, rather than only tangentially related. The tuples were submitted as search queries via BootCaT, which returned URLs corresponding to publicly accessible texts containing the desired term combinations. The resulting list of URLs underwent an initial manual screening resulting in the exclusion of approximately 30% of the entries. Private blogs were excluded due to the lack of verifiable authorship, and encyclopedic, lexicographic resources and glossaries were removed, as they do not embed target terms within authentic discursive contexts, an essential condition for metaphor analysis. From this first corpus, additional relevant keywords were automatically identified and added to the seed list, yielding an expanded set of tuples.

$X_{\text{artefatto}}$ { *algoritmo, sistema, modello, rete neuronale, LLM, architettura, BERT, ChatGPT, LLaMa, intelligenza artificiale, AI, apprendimento automatico, IA generativa, sistema esperto, agente cognitivo, transformer, modello generativo, modello linguistico di grandi dimensioni, machine learning, NLP, natural language processing*}.

A similar filtering procedure was applied to the URLs retrieved using the new tuples, again resulting in a discard rate of approximately 30%. Subsequently, a more thorough screening of the full articles was conducted. Each URL was manually opened, and the corresponding text was reviewed and validated. This stage led to a further reduction of about 50%. The current version of the corpus consists of 30 articles from Italian newspapers (e.g., *Repubblica Digitale*, *Il Mulino*, *Ticino Scienza*, *Corriere del Ticino*), scientific journalism platforms (e.g., *Agenda Digitale*, *Rivista AI*, *Mondo Digitale*, *Wired*), and organizations or companies operating in domains related to artificial intelligence (e.g., *Istituto di Ricerca Farmacologica Mario Negri*, *UniGe.life*, *HTLab*). All articles were published between 2021 and 2025 and are presumably authored by native Italian speakers. The corpus can be systematically expanded by generating additional tuples and using them as search queries to retrieve new, thematically relevant content.

The final list of URLs was stored in a separate file without web scraping or preliminary automated text extraction. The idea is to allow the script to load the URLs directly from the file, open them, parse their content, and proceed with content analysis according to the AnthroScore methodology.

4.2. Computational linguistic metric AnthroScore

AnthroScore [1] is a computational linguistic metric for implicit anthropomorphism in language. The method leverages a masked language model (RoBERTa) to quantify how non-human entities are implicitly framed as human by their surrounding linguistic context. More specifically, given a sentence containing a target entity, the model masks the entity and computes the probability that a human pronoun (e.g., he/she) versus a non-human pronoun (e.g., it/this) would appear in that context. The log-ratio of these probabilities reflects the degree to which the entity is anthropomorphized. A higher score indicates a stronger tendency for the entity to be implicitly framed as human (Figure 1).

AnthroScore computes a score that reflects the framing of an entity within its linguistic context, thus aligning with the requirement outlined in Section 2 that metaphors should be analyzed with reference to their broader linguistic environment. Due to the way the metric is designed, AnthroScore is primarily applicable to nominal and related verbal phrase. Adjectives are excluded from direct evaluation unless

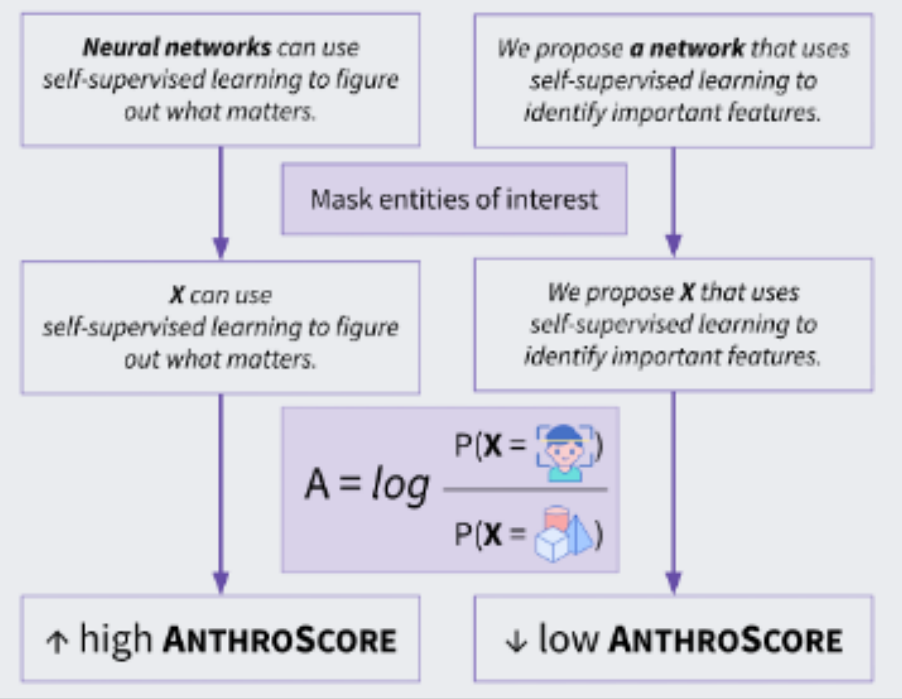

Figure 1: AnthroScore Probability [1]

they appear as part of the entity, in which case the noun phrase as a whole is analyzed. This limitation does not compromise the analysis, as adjectives in terminological contexts are typically not extracted in isolation, but rather as part of a conceptual unit that includes the noun they modify (e.g., *thinking model*).

For each given entity, the sentences in which it occurs are extracted and masked. The anthropomorphism score A is computed for each sentence, and an average score $\bar{A}$ is subsequently calculated for the entire text and for the full set of entities. Provided that all entities are terms, AnthroScore can evaluation the degree of anthropomorphism in the terminology of a given text or domain-specific corpus.

4.3. Adaptation for measuring anthropomorphism in Italian texts

The AnthroScore metrics seems to be conceptually applicable to Italian, as the language also distinguishes between human and non-human pronouns. However, using the metric for Italian requires several linguistic and technical adjustments.

Italian presents a wider variety of pronouns than English, both in subject and object positions and each of them has a gender and plural form that needs to be included. While subject pronouns in Italian are often omitted due to the pro-drop nature of the language, this omission is not mandatory. Therefore, masked language models can still infer and substitute a subject pronoun for a masked entity though the prediction may be less frequent or involve more variation, e.g., the subject pronoun *egli,* though largely obsolete in spoken Italian, can still be predicted in grammatically correct sentences without compromising grammaticality. However, listing also substitutes for the obsolete pronoun *egli* may increase the sensitivity of AnthroScore to subtle patterns of implicit anthropomorphism. Therefore demonstrative pronouns such as *questo* [this] and *quello* [that] with both genus forms were also included in the category of non-human pronouns since they can function as non-human subject. While these forms can occasionally refer to human entities in spoken language, their default semantic and syntactic function in Italian is to refer to objects, abstract concepts, or previously mentioned propositions. In the majority of contexts, especially in journalistic texts, demonstratives function as non-human referential devices.

To operationalize this adaptation, it is necessary to explicitly define which Italian pronouns will be

categorized as human and non-human, and then adjust the AnthroScore computation logic accordingly to reflect the number, type, and frequency of these pronouns in model predictions.

Pronoun $_{\text{human}}$ {"lui", "Lui", "lei", "Lei", "egli", "Egli", "ella", "Ella"}.

Pronoun $_{\text{non-human}}$ {"esso", "Esso", "essa", "Essa", "questo", "Questo", "questa", "Questa"}.

To accommodate the revised Italian pronoun lists used in the AnthroScore model, the logic for computing anthropocentric bias scores was updated to reflect the inclusion of 8 human and 8 non-human pronouns. The final AnthroScore computation remains structurally unchanged [1].

$$A = log(P_{\text{human}}/P_{\text{non-human}})$$

In addition, technical modifications are required to accommodate the Italian language in the processing pipeline. The original pipeline uses spaCy's English language model (en_core_web_sm) for tokenization, lemmatization, and syntactic parsing. To process Italian texts, this component was replaced with spaCy's Italian model (it_core_news_sm). This change ensures that linguistic preprocessing such as part-of-speech tagging and sentence segmentation is accurate for Italian grammar and syntax.

The original implementation relied on the English RoBERTa model for masked language modeling tasks. This was substituted with UmBERTo [28], a RoBERTa-based model pre-trained using a SentencePiece tokenizer and Whole Word Masking on large Italian corpora of approximately 70 GB of text. This substitution enables context-sensitive prediction on the level of words and semantic interpretation in Italian. UmBERTo has been shown to perform competitively on various Italian NLP benchmarks, making it a suitable choice for this adaptation [29, 30]. An additional advantage of using UmBERTo is that it follows the BERT masking convention, using the [MASK] token for masked language modeling. As a result, no additional changes to the masking logic are required.

Since the focus is terminology, simple and complex terms were used as entities. Analogously of what happens with AnthroScore, for each given entity/term (e.g. *modello*, *algoritmo*), the full sentence is extracted and the term replaced with a special [MASK] token. This creates a set of masked sentences, denoted S. The model computes the AnthroScore A for each sentence, as well as an average score $\bar{A}$ across the entire text or document set (in the case of the Italian corpus the text extracted from a single URL or from the list of URL are computed). The $\bar{A}$ score is provided directly while the single sentences as well as the single reference score are extracted in an output CSV. file for reference and cross manual evaluation [31]. Table 1 includes some examples of extracted sentences with high and low AnthroScore.

The S: *I modelli smettono di pensare* [models stop thinking] is reasonably highly humanized, as the masked term *modelli* [models] shows high probability of being replaced by human pronouns such as *loro*. In contrast, in S: *il modello non solo fallisce ma non riconosce* [the model not only fails but does not recognize] the metric retrieves a lower degree of anthropomorphism. This is because the verb *fallire* [to fail], although commonly associated with human agents, also admits a broader semantic framing to non-human entities, thus reducing its anthropomorphic weight.

The credibility of the metric is further supported by the Italian sentence pair *l'algoritmo sceglie uno psicologo adatto a te* [the algorithm chooses a psychologist suitable for you] which is evaluated as highly anthropomorphic due to the attribution of agency to *algoritmo*, with respect to *sulla base di un algoritmo viene selezionato lo psicologo per te* [based on an algorithm, the psychologist is selected for you] which employs a personal construction that suppresses anthropomorphic framing (hence the low score). However, the S: *Quando l'intelligenza artificiale "ragiona" con i modelli di reasoning*, presents a case where human judgement on anthropomorphism diverges from the score threshold of the computational metric.

4.4. Human evaluation of anthropomorphism

All sentences extracted by the metric for the terms *modello*, *modelli*, *intelligenza artificiale*, *algoritmo* and *algoritmi* were manually evaluated by two annotators using a binary classification: *human or*

Table 1
Example of extracted sentences with low and high AnthroScore related to the terms *modello*, *algoritmo*, and *intelligenza artificiale*

S↑: Sentences with high AnthroScore			
1	i modelli smettono di pensare proprio quando dovrebbero farlo di più	[MASK] smettono di pensare proprio quando dovrebbero farlo di più	A=2,763
2	In realtà, come sottolinea il paper, l'illusione si rompe proprio nei passaggi più complessi, dove il modello non solo fallisce ma non riconosce nemmeno il proprio errore	In realtà, come sottolinea il paper, l'illusione si rompe proprio nei passaggi più complessi, dove [MASK] non solo fallisce ma non riconosce nemmeno il proprio errore	A=0,846
3	l'algoritmo sceglie uno psicologo adatto a te	[MASK] sceglie uno psicologo adatto a te	A=2,464
4	Quando l'intelligenza artificiale "ragiona" con i modelli di reasoning.	Quando [MASK] "ragiona" con i modelli di reasoning.	A=2,302
S↓: Sentences with low AnthroScore			
5	Il test su puzzle logici mostra che anche i modelli più avanzati collassano di fronte alla complessità crescente, incapaci persino di seguire istruzioni corrette.	Il test su puzzle logici mostra che anche [MASK] più avanzati collassano di fronte alla complessità crescente, incapaci persino di seguire istruzioni corrette.	A=-4,173
6	I modelli vengono usati per fare previsioni	[MASK] vengono usati per fare previsioni	A=-1,885
7	sulla base di un algoritmo viene selezionato lo psicologo per te	sulla base di [MASK] viene selezionato lo psicologo per te	A=- -6,379
8	L'intelligenza artificiale sta rivoluzionando il modo in cui interagiamo con le macchine, aprendo la strada a nuove e interessanti applicazioni.	MASK] sta rivoluzionando il modo in cui interagiamo con le macchine, aprendo la strada a nuove e interessanti applicazioni.	A=-4,505

non-human. A sentence was labelled as *human* when the masked term was judged to exhibit capacities typically attributed to human beings, such as mental states or intentional actions.

While most human evaluations aligned closely with the computational threshold of anthropomorphism (e.g. S 1–3 and 5–7 in Table 1), some instances showed notable discrepancies. In example 4, the metric identified a high level of anthropomorphism (AnthroScore A=2,302), whereas human evaluators classified the expression as *non-human*. This divergence stemmed from the use of quotation marks around the verb *ragionare* [to reason], which were interpreted by human readers as a distancing device. Such suprasegmental markers downplay the literal meaning of the verb and thereby reduce its anthropomorphic effect. Markers such as quotation marks, italics, or metalinguistic expressions like *cosiddetto* [so-called], *per così dire* [so to speak] often function not to reinforce, but to neutralize anthropomorphism. However, the current implementation of AnthroScore does not distinguish between quoted and unquoted forms, and thus fails to account for these interpretive cues. Given the frequency of such distancing strategies in Italian public discourse, their exclusion from the computational model may compromise the accuracy of anthropomorphism detection. The metric should therefore integrate metalinguistic and suprasegmental features to more effectively reflect human interpretative judgments.

Moreover, Italian disseminating articles in many domains is inclined to make use of anglicisms for conceptual unit. Anglicisms like *thinking* is often used as device for terminologizing common word like *pensare* [32]. So S: *senza sapere quali modelli sono thinking* [without knowing which models are thinking] is perceived to have a higher degree of technicality as S: *senza sapere quali modelli sono pensanti*. AnthroScore, however, compute *thinking* and *pensanti* based on their lexical meaning. The higher degree of attachment to the cognitive meaning of *pensare* with respect to the English equivalent *thinking* is not computed. This is demonstrated by the fact that human judgment for S was hon-human

while computational measure was retrieved to be > 1.

Overall, extra-lexical information represent a limitation of the computational metric AnthroScore. Moreover, the initial manual evaluation may not fully reflect general human perception. To enhance the validity of the method, further testing is needed, for instance by incorporating additional terms, expanding the corpus, and adopting a more nuanced human scoring model that moves beyond binary classification.

5. Results, limitations and future work

Observations on AnthroScore suggest that can be a valuable basis for analyzing anthropomorphism in Italian discourse, as the metric evaluates terms within their immediate linguistic context. This contextual sensitivity aligns with the principle that metaphor in terminology must be interpreted within its broader linguistic and cognitive embedding, rather than through isolated lexical items. Early tests on Italian texts and terms indicate that the adaptation of AnthroScore for Italian language is able to match manual evaluation of the phenomenon of anthropomorphism in public discourse. Italian seems to exhibit a relatively high degree of anthropomorphism at the terminological level. At the same time some results highlight several methodological limitations. AnthroScore for Italian currently does not account for suprasegmental and metalinguistic devices that are particularly frequent in Italian public discourse and can significantly modulate the degree of anthropomorphism. These include distancing strategies such as quotation marks (e.g., *l'intelligenza artificiale "ragiona"*), metalinguistic qualifiers (e.g., *cosiddetto* [so-called]), and other stance markers that frame anthropomorphic expressions as either metaphorical or literal. Additionally, while the metric can process English terms embedded in Italian texts and interpret them based on their lexical meaning, it does not account for how such borrowings might alter metaphorical perception in context, potentially affecting human interpretation. These issues call for further investigation, including the inclusion of more terms, systematic manual validation, and an expansion of the corpus.

With more robust data and an enhanced methodological framework, a promising direction for future research involves examining whether the degree of anthropomorphic terminology in a given language correlates with cognitive and sociocultural biases in both directions. While it has been researched that anthropomorphic terminology foster the stabilization of AI as cognitively human-like in people's perception, it is equally relevant to explore whether general attitudes toward technology, potentially shaped by factors such as lower perceived performance or trustworthiness [33, 34], is able to influence how users linguistically frame AI systems. Is the use of distancing devices such as quotation marks, hedging expressions or impersonal syntactic constructions a consequence of scepticism towards AI related technologies?

6. Modelling anthropologizing verbs in TMS

The process of terminologization within specialized discourse is shaped by the injection of lexemes from general language, a transition that remains highly dependent on the surrounding textual and conceptual context. A primary challenge in this evolution concerns verbs. Because verbs, when functioning as terms, rely on dynamic conceptualization, their representation in a terminology management system requires the explicit specification of the entities with which they are associated or co-occur. This approach, supported by the present investigation in the established relation entity/verb, defines the functional semantic boundaries necessary to distinguish a general-language verb from a specialized terminological unit. Such boundaries are particularly crucial when a meaning extension from general to specialized usage is still in progress. The degree of anthropomorphization can be operationalized within a TMS as a diagnostic indicator for assessing the stabilization of a verb as a term. While the specialized status of some verbs remains controversial due to persistent human-centered associations, the system can track how others (e.g., *allucinare*) become conventionalized technical units as their anthropomorphic origins recede.

Effective TMS implementation must also account for the presence or absence of suprasegmental and metalinguistic markers, such as quotation marks and italics. These markers function as distancing strategies that enforce a technical reading of a verb while attenuating its general, anthropomorphized meaning. Since such metalinguistic markers may constitute part of the term's meaning rather than merely a graphical device, documenting them within a TMS also enables the tracking of linguistic evolution. In particular, italics and quotation marks in Italian can serve as indicators of recently introduced lexical borrowings (especially anglicisms), highlighting specialized status and modulating the metaphorical load in technical contexts.

Declaration on Generative AI

During the preparation of this work, the author used GPT-4o in order to check grammar and spelling as well as to check reference consistency with provided template. After using this service, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] M. Cheng, K. Gligoric, T. Piccardi, D. Jurafsky, Anthroscore: A computational linguistic measure of anthropomorphism, in: Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2024. URL: https://arxiv.org/abs/2402.02056.

[2] D. Gentner, M. Jeziorski, The shift from metaphor to analogy in western science, in: A. Ortony (Ed.), Metaphor and Thought, 2 ed., Cambridge University Press, Cambridge, 1993, pp. 447–480.

[3] A. Pascolini, Metafore e comunicazione scientifica, JCOM 3 (2004).

[4] T. Temmerman, Towards New Ways of Terminology Description: The Sociocognitive Approach, John Benjamins, 2000. doi:10.1075/tlrp.3.

[5] N. Epley, A. Waytz, J. T. Cacioppo, On seeing human: A three-factor theory of anthropomorphism, Psychological Review 114 (2007) 864–886. doi:10.1037/0033-295X.114.4.864.

[6] A. Salles, K. Evers, M. Farisco, Anthropomorphism in ai, AJOB Neuroscience 11 (2020) 88–95. doi:10.1080/21507740.2020.1740350.

[7] M. Bunz, M. Braghieri, The ai doctor will see you now: assessing the framing of ai in news coverage, AI & Society 37 (2022) 9–22. doi:10.1007/s00146-021-01145-9.

[8] L. Ibrahim, M. Cheng, Thinking beyond the anthropomorphic paradigm benefits llm research, 2025. URL: https://arxiv.org/html/2502.09192v1, preprint.

[9] D. Watson, The rhetoric and reality of anthropomorphism in artificial intelligence, Minds & Machines 29 (2019) 417–440. doi:10.1007/s11023-019-09506-6.

[10] Z. C. Lipton, The mythos of model interpretability, Communications of the ACM 61 (2018) 36–43. doi:10.1145/3233231.

[11] H. Vasconcelos, M. J"orke, M. G. McLaughlin, T. Gerstenberg, M. S. Bernstein, R. Krishna, Explanations can reduce overreliance on ai systems during decision-making, Proceedings of the ACM on Human-Computer Interaction 7 (2023) 1–38. doi:10.1145/3579605.

[12] N. Tiku, The google engineer who thinks the company's ai has come to life, 2022. URL: https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/.

[13] A. M. Kempton, E. Parmiggiani, P. Vassilakopoulou, Accountability in managing artificial intelligence, in: Proceedings of ECIS 2023, 2023. URL: https://aisel.aisnet.org/ecis2023_rp/361, paper 361.

[14] A. Waytz, C. K. Morewedge, N. Epley, G. Monteleone, J. J. Gao, J. T. Cacioppo, Making sense by making sentient, Journal of Personality and Social Psychology 99 (2010) 410–435. doi:10.1037/a0020240.

[15] J. Weizenbaum, Eliza—a computer program for the study of natural language communication, Communications of the ACM 9 (1966) 36–45.

[16] C. Nass, Y. Moon, Machines and mindlessness, Journal of Social Issues 56 (2000) 81–103.

[17] B. Shneiderman, Human-centered artificial intelligence, 2020. URL: https://arxiv.org/abs/2002.04087, arXiv:2002.04087.
[18] I. Rahwan, et al., Machine behaviour, Nature 568 (2019) 477–486.
[19] Z. Ye, J. Li, Artificial intelligence through the lens of metaphor, International Journal of Digital Law and Governance 1 (2024) 361–381. doi:10.1515/ijdlg-2024-0016.
[20] M. Shanahan, Talking about large language models, 2023. URL: https://arxiv.org/abs/2212.03551, arXiv:2212.03551.
[21] A. S. Reynolds, Metaphors and science, in: Understanding Metaphors in the Life Sciences, Cambridge University Press, 2022, pp. 1–13.
[22] P. ten Hacken, R. Panocov'a, Word Formation as a Naming Device, Edinburgh University Press, 2024.
[23] E. Calaresu, Significati non letterali nel discorso scientifico, Italiano LinguaDue (2025).
[24] D. Stanic, The metaphorical conceptualisation of artificial intelligence, Linguistica 65 (2025) 185–203. doi:10.4312/linguistica.65.1.185-203.
[25] T. Petricini, The power of language: framing ai, AI & Society (2025). doi:10.1007/s00146-025-02586-2.
[26] P. Vernillo, Mip-it: Identificazione delle metafore linguistiche in italiano, AION-L (2023).
[27] Bootcat, ???? URL: https://bootcat.dipintra.it/?section=home.
[28] Musixmatch, HuggingFace, Umberto commoncrawl cased v1, 2022. URL: https://huggingface.co/Musixmatch/umberto-commoncrawl-cased-v1.
[29] S. Auriemma, et al., Evaluating pre-trained transformers on italian administrative texts, 2022.
[30] Y. Liu, et al., Roberta: A robustly optimized bert pretraining approach, 2019. ArXiv:1907.11692.
[31] R. Resi, Anthroscore ita repository, ???? URL: https://github.com/RossellaResi/Anthroscore_ITA.git.
[32] R. Rossella, P. ten Hacken, Perdita di domini: is italian losing domains?, in: M. Thelen (Ed.), Domain Loss and Gain in Contemporary Terminology Studies, Amsterdam University Press, 2026.
[33] F. Cicero, L'italiano delle intelligenze artificiali generative, Italiano LinguaDue (2023).
[34] G. Fiorentino, M. Tavosanis, Chiaro, sintetico, e brillante, Lingua italiana d'oggi (2023).