



Daria Evangelista

University of Amsterdam

d.evangelista@uva.nl

TRANSLATED WITH AI: HOW ARE SPECIALIZED, YET HIGHLY RHETORICAL TEXTS TRANSLATED BY DEEPL? AN ANALYSIS OF PREAMBLES IN NORMATIVE TEXTS

<https://doi.org/10.17234/9789533793337.18>

Abstract

DeepL, a widely used Neural Machine Translation (NMT) system, can translate entire paragraphs holistically (Belinkov et al. 2020). The program has been found to improve fluency compared to traditional machine translation methods (Volkart, Bouillon and Girletti 2018). However, to date, no research has explored how AI systems like DeepL handle the translation of rhetorical phenomena and genre-sensitive traits in specialized texts. This paper examines how DeepL translates rhetorical devices in the preambles of normative texts, which typically contain a multitude of rhetorical phenomena (Evangelista 2025, 54–64). Using a qualitative, corpus-based approach informed by translation theory, rhetorical analysis, and text linguistics (cf. e.g. Bassnett 2014; Prandi 2015; Ferrari, Lala and Lampese 2021), the study investigates whether DeepL captures key rhetorical devices at various textual levels. The analysis reveals that DeepL frequently fails to translate critical rhetorical features essential for meaning and stylistic accuracy in preambles. These findings underscore the limitations of current AI translation systems in contexts requiring rhetorical precision, and highlight the need for human oversight and expert post-editing.

Keywords: translation, rhetoric, figures of speech, preambles, normative texts, DeepL



1. Introduction

In an era where our thoughts, and consequently our words, are influenced by the massive use of artificial intelligence, the task of linguists and rhetoricians alike is often to evaluate the impact of these new instruments on communication and language use. A tough task indeed, which is further complicated by the limited understanding, both among the general public and within academic circles, of how such tools actually function (Belinkov et al. 2020). A simple search on scholarly databases reveals a gap, as (apart from very specialized and extensive works, such as Koehn 2020), there are few articles dedicated solely to explaining aspects concerning the inner workings of these programs. Instead, existing studies that attempt to evaluate their performance, studies just like the present one, as well as the ones consulted for this paper, typically spend only a few paragraphs in order to explain their general functioning mechanisms.

The degree to which we are not aware of what happens when we, for example, generate a text by using an AI-powered tool, or translate one with the help of an AI-powered translation software, is so high that in order to cope with this lack of awareness, we end up humanizing the program: we conceive it as if it was a rational human being. Metaphors in this sense are developing at a rapid pace (Mitchell 2024), fuelled by the original framing process that gave birth to the entire conceptual field, a process that found its fortune in what constitutes the head of the syntactical compound “artificial intelligence”, that is, *intelligence*. Undoubtedly, the adjective *artificial* is placed before the noun, reminding us that we are not dealing with *natural intelligence*, that is, the intelligence of a living being. But what started as a creative metaphor, artificial intelligence, with all the connotative latitude such a term allows, is now latently solidifying into a conventional way of thinking – and into lexicalized language. This shift inevitably carries the risk of disorienting us, leading to uncritical acceptance of these metaphorical framing processes. Needless to say, that, as a consequence, when confronted with AI-generated content, it is necessary to redouble our efforts in critical thinking – and linguistic discernment.

2. Research questions and methods

With this in mind, the aim of this article is to conduct a linguistic and rhetorical analysis of a particular kind of material which can be produced by artificial intelligence, namely, translated content. The intent of this study is not to focus on just any kind of text, but more specifically on a highly specialized genre that typically contains plenty of rhetorical phenomena, namely preambles at the beginning of normative texts. The choice of focusing on this text genre is justified by a specific aim: the ultimate objective is to assess how an AI-powered translation program such as DeepL handles the various degrees of human creativity conveyed through rhetorical phenomena. Such phenomena are also crucial for the text genre of normative preambles itself, which is stylistically prone to containing many rhetorical traits (Evangelista 2023; 2024). This analysis thus seeks to address two main questions: i) how are rhetorical phenomena rendered by the translation program? And ii) are the specific rhetorical features that are codified into the style of the text genre kept intact in the translation process? Additionally, if the features are retained (or not), what are the potential consequences?

2.1. Translating rhetoric: theoretical approaches for the analysis

The primary theoretical framework adopted in this paper is that of rhetorical analysis. Drawing on linguistic theories such as, most importantly, the theorization of Michele Prandi (2023), while still keeping as reference points the studies of Fahnestock (e.g. 2011), rhetorical analysis is understood



as a broad discipline encompassing a wide range of conceptual and linguistic components. Particular attention is given to figures of speech, and especially to metaphors, which are nowadays increasingly examined through the lens of cognitive science. But conspicuous space in such theorization is also taken by theoretical tools from other disciplines, e.g. pragmatics and text linguistics, that have now been made explicitly part of an overarching view of rhetoric (Evangelista 2025, 19–21; Prandi 2023, 29–35).

In addition to the broad theoretical framework rooted in rhetorical analysis, this study also engages with key insights from the field of traductology. The interplay between rhetorical analysis, on the one side, and key theories in traductology, on the other, is important for two reasons: first, because until now the two disciplines have “regrettably been one of those reciprocally ignorant neighbors” (Dávila-Montes 2017, 3). Second, because translation has increasingly prioritized meaning: as we know, a good translation is not one that keeps the exact same words of the original, but rather one that, despite changing words and their order, succeeds in conveying the communicative content contained in the source text. In its hermeneutical and dialectic process, translation is, in Dávila-Montes (2017) words, “a rhetoric of meaning”, which pays particular attention also to those levels that are oftentimes the most neglected, namely those that fall within the scope of rhetorical analysis. Moreover, having as the object of analysis translated rhetorical texts produced by AI systems, this article contributes to a still relatively small but emerging body of rhetorical scholarship (cf. e.g. Majdik and Graham 2024) that takes up AI-generated content as an object of critical rhetorical analysis.

Lastly, in analyzing AI-produced translations from the rhetorical point of view, particularly useful are theories drawn from traductology, and especially the concept of translation universals, which has played a central role in identifying linguistic patterns observed in translation corpora that tend to emerge in translated texts regardless of the languages involved (cf. Scarpa 2008, 108–113; Bassnett 2014). Universals offer valuable tools for evaluating how linguistic and rhetorical features are transformed in translation. Furthermore, specific attention in the study is given to theoretical perspectives on metaphor translation, with a focus on how different types of metaphor (ranging from conventional to highly creative) are treated in the translation process. In this regard, Prandi’s (2015) typology of metaphor proves particularly useful, as it allows for a more nuanced analysis of the challenges that AI-powered translation systems face when confronted with metaphorical language of varying complexity and originality.

2.2. The corpus

The analysis draws on a qualitative and corpus-based methodology and is done manually, by systematically observing the rhetorical traits of preambles in their original versions, and then comparing this original version with the translation in Italian made by the AI-powered translation program DeepL, which uses a neural machine translation (NMT) system. The corpus of texts that was analyzed for this paper is divided into two parts. The first part consists of a group of 24 preambles¹ coming from Swiss constitutional texts (that of the Federal Constitution and 23 preambles of the cantonal Constitutions²), which original languages are mostly French and German; the second part consists of 24 preambles of international treaties in their original versions, which are English or

¹ The number corresponds to the number of cantonal preambles, 23, to which the federal preamble is added.

² Swiss Cantons are 26, yet the Constitutions of the Cantons of Aargau, Appenzell Innerrhoden and Zug have no preamble.



French. Both parts of the corpus containing original texts are available on the Swiss publication platform for federal law “Fedlex”.³

The idea of considering preambles of international treaties along with preambles of Swiss Constitutions is justified by literature on the Swiss Federal Constitution, which states that the structure of the preamble of the Swiss Constitution is based on the structure of the preambles of international treaties (Confederazione Svizzera 1996, 114). Moreover, it is also justified by the intention to broaden the pool of analysis compared to the 24 Swiss preambles alone, and thus to be able to count on a broader and more balanced repertoire of translation examples. The selection of international preambles was made according to a criterion of importance in the hierarchy of treaty law and according to the function of the text in which they were found, preferring texts with the constitutive function of bringing into being an international body or organ (e.g. the Constitution of the United Nations Educational, Scientific and Cultural Organization, UNESCO). This is to allow a closer comparison with the preambles of the Swiss constitutions, considering their similar pragmatic intent.

2.3. Preambles of laws as a text genre: rhetorical traits

The object of analysis of the present research is AI-produced translations of preambles of normative texts, and as this text genre so far has not been the focus of broad linguistic studies, it is useful to introduce here its textual characteristics. Despite not being a well-known text genre, preambles constitute the optimal testing ground both for observing the translation of various rhetorical phenomena – as they contain many of them – and for discovering if and how text genre is taken into account in AI-produced translations – as they are a highly specialized genre that largely relies on rhetoric for meaning construction.

That said, it is true that when thinking about a plausible prototype of a creative text, what comes to mind might not be the preamble of a normative text. This does not come as a surprise: the language of normative texts is considered by many scholars (cf. e.g. Sabatini 1999; Irti 2016) to be denotative, explicit, clear, and interpretively binding. Despite this common conception of the language of law texts, it is crucial to note that recent research, on the contrary, has shown many examples of rhetorical phenomena found in the law that are, instead, connotative, implicit, vague, leave space for interpretation, and serve different pragmatic functions (Evangelista 2025). These phenomena are contained throughout the normative text, but above all, they are condensed precisely in the introductory section of the preamble, which has the important function of offering a key for guiding the interpretation of the whole text by mentioning contextual information, values, and principles in an introductory partition which is not legally binding.

From the linguistic point of view, preambles are characterized by a macrostructure based on the rhetorical figure of the anaphora. Moreover, their sentences contain a multitude of rhetorical figures. Together, the incipit position, macrostructure, and rhetorical figures within the sentences add up to form a textual section which stylistically may seem more like a literary composition rather than a normative text. As a representative example, we can observe the preamble of the Swiss Constitution, here in its official English translation:

³ https://www.fedlex.admin.ch/en/treaty?news_period=last_year&news_pageNb=2&news_order=desc&news_itemsPerPage=10



(1) Preamble

In the name of Almighty God!

The Swiss People and the Cantons,

mindful of their responsibility towards creation,

resolved to renew their alliance so as to strengthen liberty, democracy, independence and peace in a spirit of solidarity and openness towards the world,

determined to live together with mutual consideration and respect for their diversity,

conscious of their common achievements and their responsibility towards future generations,

and in the knowledge that only those who use their freedom remain free, and that the strength of a people is measured by the well-being of its weakest members,

adopt the following Constitution:

(Federal Constitution of the Swiss Confederation)

Set aside some small details, the structure of the preamble is identical to that of preambles of international treaties, with a very present and rhythmical syntactic anaphora (*mindful, resolved, determined, etc.*) and the figures of speech within the sentence (e.g. the antithesis: “the *strength* of a people is measured by the well-being of its *weakest* members”). We can note the similarity by observing the preamble of a treaty, for example, the Rome Statute of the International Criminal Court, which is also built through anaphora and contains figures of speech within its sentences (e.g. the metaphor of the *mosaic* and the *accumulatio* of nouns: *children, women and men, etc.*):

(2) Preamble

The States Parties to this Statute,

Conscious that all peoples are united by common bonds, their cultures pieced together in a shared heritage, and concerned that this delicate mosaic may be shattered at any time,

mindful that during this century millions of children, women and men have been victims of unimaginable atrocities that deeply shock the conscience of humanity,

recognizing that such grave crimes threaten the peace, security and well-being of the world,

affirming that the most serious crimes of concern to the international community as a whole must not go unpunished and that their effective prosecution must be ensured by taking measures at the national level and by enhancing international cooperation,

determined to put an end to impunity for the perpetrators of these crimes and thus to contribute to the prevention of such crimes,

[...]

resolved to guarantee lasting respect for and the enforcement of international justice,

have agreed as follows:

(Rome Statute of the International Criminal Court)

3. Results: a general overview

Observing the comparison between preambles in their original language and their translations in Italian produced by DeepL, the first striking remark is that rhetorical phenomena in the translated versions are frequently not preserved, or are only partially rendered.

In order to have an overview of how rhetoric is treated by DeepL, we can observe below the German original version of the preamble to the Swiss Constitution, on the left, and its Italian translation made by DeepL on the right. In this comparison, the present rhetorical phenomena are



marked in bold, while those that have not been accurately translated are in italics in the target text in Italian. This overview already makes it possible to see that a good half of the rhetorical phenomena present in the text are left untranslated:

(3) Präambel Im Namen Gottes des Allmächtigen! Das Schweizervolk und die Kantone, in der Verantwortung gegenüber der Schöpfung, im Bestreben, den Bund zu erneuern, um Freiheit und Demokratie, Unabhängigkeit und Frieden in Solidarität und Offenheit gegenüber der Welt zu stärken, im Willen, in gegenseitiger Rücksichtnahme und Achtung ihre Vielfalt in der Einheit zu leben, im Bewusstsein der gemeinsamen Errungenschaften und der Verantwortung gegenüber den künftigen Generationen, gewiss, dass frei nur ist, wer seine Freiheit gebraucht, und dass die Stärke des Volkes sich misst am Wohl der Schwachen, <i>geben sich folgende Verfassung:</i>	Preambolo Nel nome di Dio onnipotente! Il popolo svizzero e i Cantoni, <i>nella responsabilità</i> verso il creato, <i>si sforzano</i> di rinnovare la Confederazione per rafforzare la libertà e la democrazia, l'indipendenza e la pace nella solidarietà e nell'apertura verso il mondo, <i>nella volontà</i> di vivere la propria diversità nell'unità, nella considerazione e nel rispetto reciproci, consapevoli delle nostre conquiste comuni e della nostra responsabilità nei confronti delle generazioni future, <i>riconoscendo che solo chi usa la propria libertà è libero e che la forza del popolo si misura dal benessere dei deboli</i>, <i>adottano la seguente Costituzione:</i>
---	--

(Federal Constitution of the Swiss Confederation)

As we can see in the comparison, the solemn anaphora that builds the structure of the German text is missing in the translation. This anaphora in preambles is usually translated with a past participle or gerund by Italian-speaking human translators (Evangelista 2024), whereas in this case, the program's output consists of a prepositional syntagma at the beginning of each sentence. This output clearly follows the syntactic structure of German (*in der Verantwortung* becomes *nella responsabilità*, *im Willen* is translated with *nella volontà*). Therefore, an evident interference with the structure of the German sentence is to be observed.

The mistranslation by DeepL of the structural anaphora is data that cannot be overlooked. The anaphora represents the core structure of preambles, a structure that is encoded in this textual genre and, by being so, forms one of its most characteristic rhetorical traits. More specifically, the macrostructural anaphora in preambles has two functions: to make the whole text portion more salient within the normative text, compensating for its lack of normativity, and to formulate the principles contained therein as presuppositions (Pecorari 2022) through the use of participial or gerundive constructions. This rhetorical strategy positions such principles as foundational to the legal provisions that follow. If we think of the preamble as a text genre, it is therefore significant that DeepL does not holistically recognize one of its most characteristic traits, and does not reproduce it in a coherent structure in the target language, as these are exactly the sorts of patterns one expects ML systems to classify and predict.

If we focus on figures of structure, we can observe that the anastrophe *frei nur ist, wer seine Freiheit gebraucht*, is also not translated. The respective sentence contained in the Italian AI translation, *solo chi usa la propria libertà è libero*, is not built through a word inversion, but instead follows the canonical word order of the Italian grammar. In this case, we are confronted with a case of normalization (a tendency to conform to the syntactic norms of the target language), as the content of the message is rendered in unmarked, non-rhetorical, denotative terms.

Figures that intertwine content and structure, but in which the word order is unmarked from the syntactical point of view, such as the *enumeratio*, are well-reproduced in the target language. This



is evidenced, for instance, in the accurate translation of the excerpt *um Freiheit und Demokratie, Unabhängigkeit und Frieden in Solidarität und Offenheit gegenüber der Welt zu stärken*, as well as in the rendering of the antithesis between the concepts of *strong* and *weak*, in the sentence *gewiss (..) dass die Stärke des Volkes sich misst am Wohl der Schwachen*. This data suggests that the translation program tends to favor a word-for-word approach in such cases. When the syntactic structure of the source language can be directly mapped onto the target language without significant adaptation, the output remains both structurally sound and semantically faithful, minimizing the likelihood of error.

Amid this general overview, a minor yet curious detail deserves brief mention, although it may ultimately be attributed to coincidence and thus does not warrant deeper investigation. In the final line of the German preamble, the text features an alliteration on the letter *s*, which is accurately, although partially, reproduced in the Italian AI translation: *die Stärke des Volkes sich misst am Wohl der Schwachen*. This alliteration, which was originally conceived by the Swiss German writer Adolf Muschg (Egger and Evangelista 2022, 385), is kept almost intact in the Italian AI translation, which echoes *la forza del popolo si misura dal benessere dei deboli*, thus compensating the missing “*s*” sound in *deboli* and *popolo* by adding it earlier on, in *benessere*.

3.1. How the rhetorical macrostructure is translated

As we have seen, rhetorical phenomena in preambles can be found at three different levels: at the level of the macrostructure, which invests the whole text; at the level of the utterance and of the sentence, where we find figures of construction and content; and finally at the level of syntagms, where we find tropes which do not invest the syntactical construction of the sentence. We can now focus on the translation of each of these levels, starting from the macrostructure.

As already observed, the macrostructure of preambles is based on a very solemn and rhythmic anaphora. In this regard, it is worth noticing that the program is inclined to reproduce the same morphosyntactic structures contained in the original version. As a result, the more the original version uses regular forms of *repetitio*, that are, in this case, forms that repeat identical syntactic structures at the beginning of each sentence, the more the program translates them into regular anaphoras, thus reproducing the identical repeating structure.

One of the representative examples for this case is found in the preamble of the UN Charter, which we can observe below, along with its Italian translation made by DeepL. The comparison shows that every -ing form opening a sentence in the English original version corresponds to a verb in the gerund form in the Italian DeepL translation, so that the repetition is marked by the same syntactic element.

(4)	The States Parties to the present Convention, Considering the fundamental role of treaties in the history of international relations, Recognizing the ever-increasing importance of treaties as a source of international law [...] Noting that the principles of free consent and of good faith and the pacta sunt servanda rule are universally recognized, [...]	Gli Stati contraenti della presente Convenzione, considerando il ruolo fondamentale dei trattati nella storia delle relazioni internazionali, riconoscendo l'importanza sempre crescente dei trattati come fonte di diritto internazionale [...] constatando che i principi del libero consenso e della buona fede e la regola pacta sunt servanda sono universalmente riconosciuti, [...]
-----	---	--

(United Nations Charter)



It is important to highlight that the performance of AI-based translation systems varies significantly depending on the language pair involved. In general, translations from English into Italian tend to exhibit relatively high levels of fluency and accuracy, whereas translations from French or German into Italian are often less consistent and prone to errors. This disparity is largely attributable to the size and quality of the training corpora available for different language combinations. Given the dominance of English in digital communication and published content, NMT systems such as DeepL are typically trained on far larger and more diverse datasets for English than for other source languages (cf. Toral and Sánchez-Cartagena 2017; Koehn 2020). As a result, the quality of output can vary considerably depending on the linguistic input. This discrepancy is clearly observable in the translation of the preamble to the Constitution of the Canton of Vaud, where a comparison between the French original and its Italian rendering by DeepL reveals notable divergences:

(5) Pour favoriser l'épanouissement de chacun dans une société harmonieuse qui respecte la Création comme berceau des générations à venir, soit ouverte au monde et s'y sente unie, mesure sa force au soin qu'elle prend du plus faible de ses membres, et conçoive l'État comme l'expression de sa volonté, <i>le peuple du Canton de Vaud se donne la Constitution suivante:</i>	Per promuovere lo sviluppo di ogni individuo in una società armoniosa che rispetti il Creato come culla delle generazioni future, sia aperta al mondo e si senta unita ad esso, <i>misura</i> la sua forza dalla cura che ha dei suoi membri più deboli, <i>e vede</i> lo Stato come espressione della sua volontà, <i>il popolo del Cantone di Vaud si dà la seguente Costituzione:</i>
--	---

(Constitution of the Canton of Vaud)

The Italian translation produced by the program contains several errors that undermine the rhetorical structure of the original text. Notably, the French version constructs a morphosyntactic anaphora through the consistent use of the subjunctive mood, which plays a crucial role in creating rhythmic and syntactic coherence. In the Italian translation, however, only the first two clauses retain the subjunctive correctly, while the final two are rendered using the indicative. This grammatical error breaks the pattern and cancels the anaphoric repetition. Consequently, the passage loses much of its rhetorical impact and would require significant editing by a human translator to properly reproduce its structure.

3.2. How figures mixing content and structure are translated

As previously noted, figures at the utterance level that intertwine content and structure are generally well rendered in AI-generated translations. An example can be found in the well-known chiasmus from the preamble of the Constitution of the United Nations Educational, Scientific and Cultural Organization. In this instance, the Italian translation produced by the AI program faithfully reproduces the original structure, preserving the chiasmic form and maintaining the word order of the source text intact:

(6) The Governments of the States Parties to this Constitution on behalf of their peoples declare: That since wars begin in the minds of men, it is in the minds of men that the defences of peace must be constructed; [...]	I Governi degli Stati parti di questa Costituzione, a nome dei loro popoli, dichiarano: che, poiché le guerre iniziano nella mente degli uomini, è nella mente degli uomini che devono essere costruite le difese della pace; [...]
--	--

(Constitution of the United Nations Educational, Scientific and Cultural Organisation, UNESCO)



A similar observation applies to the translation of the rhetorical figure of *accumulatio* found in the International Convention on the Elimination of All Forms of Racial Discrimination. In this case, the AI-generated Italian version preserves the original word order, effectively maintaining the cumulative rhythm and semantic progression intended by the source text:

(7)	Convinced that any doctrine of superiority based on racial differentiation is scientifically false, morally condemnable, socially unjust and dangerous , and that there is no justification for racial discrimination, in theory or in practice, anywhere [...]	Convinti che qualsiasi dottrina di superiorità basata sulla differenziazione razziale sia scientificamente falsa, moralmente condannabile, socialmente ingiusta e pericolosa , e che non vi sia alcuna giustificazione per la discriminazione razziale, in teoria o in pratica, in alcun luogo
-----	--	---

(International Convention on the Elimination of All Forms of Racial Discrimination)

3.3. How metaphors and similes are translated

Focusing on the sentence level, it is worth noting that tropes, especially metaphors, tend to be only partially preserved in AI-generated translations. More precisely, these figures tend to be faithfully rendered when they are conventional and isolated in the sentence, *i.e.*, when they do not engender a semantic web of creative collocations that last through several sentences. An example of this case is present in the Constitution of the Canton of Vaud (example 5), where the isolated simile connotating the Creation *as the cradle of future generations* was kept intact in the translation process.

While the source domain of a metaphor is often accurately rendered at the lexical level, the full conceptual mapping – including the intricate network of semantic associations and pragmatic inferences – frequently becomes attenuated or fragmented in the target text when it lasts throughout several sentences. In this case, we observe a process of normalisation, in which stylistically salient or contextually rich expressions in the source language are replaced by conventional and unmarked equivalents in the target language. In NMT, this tendency can be due to the probabilistic nature of the models, which are trained to optimize for the most statistically likely output, often at the expense of figurative complexity.

An illustrative case is found in a metaphor-rich passage from the Rome Statute of the International Criminal Court, in which part of the metaphoric web is flattened in the Italian version:

(8)	PREAMBLE The States Parties to this Statute, Conscious that all peoples are united by common bonds, their cultures pieced together in a shared heritage, and concerned that this delicate mosaic may be shattered at any time , [...]	PREAMBOLO Gli Stati parti del presente Statuto, consapevoli che tutti i popoli sono uniti da legami comuni, le loro culture sono <i>unite</i> in un patrimonio condiviso, e preoccupati che questo delicato mosaico possa andare in frantumi in qualsiasi momento , [...]
-----	---	--

(Rome Statute of the International Criminal Court)

As we can see from the comparison, the metaphor in this passage, namely that all cultures are pieced together to form a shared heritage that is a delicate mosaic, a balance that may be subverted at any time, is not completely translated. The information that cultures are *pieced together*, like the tiles of a mosaic, is simply not present in the AI version, in which the meaning is normalised by using the word *unite* (united). Such an outcome highlights a limitation of current NMT systems, namely their difficulty in handling metaphoric webs that extend throughout several utterances



– a limitation that weighs heavily in genres such as preambles, where rhetorical force is integral to meaning construction.

4. Conclusions

The results of this study make it evident that significant rhetorical phenomena present in the source texts of the preambles remain untranslated or are only partially rendered in the AI-generated versions. In particular, figures that rely on syntactic structure and word order at the macro-structural level and that are typical of the text genre, such as the anaphora, can get lost in translation, as DeepL seems to privilege single sentence fluency over macrostructural organization.

By contrast, figures that combine syntactic and semantic features at the sentence level, such as chiasmic structures and *enumerations*, are generally well preserved. These constructions are often translatable without significant alteration to the sentence structure of the source text, which makes them more accessible to the model's pattern recognition capabilities. Similarly, figures of content, including metaphorical expressions, are often retained at the lexical level, especially if they are conventional and isolated. However, when metaphors are creatively woven into several sentences, the AI system can normalize them. This results in the loss of semantic nuance and implicit content.

A further consideration concerns the language pair under analysis, which significantly affects the quality of the translation. In this study, the German-Italian combination showed notable weaknesses, partially due to the syntactic distance between the two languages. German's rather fixed word order poses substantial challenges to translation systems trained predominantly on English-centric corpora. In comparison, French-Italian and English-Italian pairs tend to yield more accurate translations, and this reflects both greater corpus availability and closer typological proximity.

The analysis also highlighted that NMT errors can be qualitatively different from those made by human translators. As we have seen, NMT systems may produce grammatically plausible but contextually inappropriate outputs, such as a sudden change in grammatical mood that breaks the anaphoric rhythm. This type of change disrupts not only the grammatical coherence of the whole text but also undermines its rhetorical function.

This observation leads to a final, and perhaps the most critical, reflection. As we have seen, the preamble as a genre is characterized by a high degree of rhetorical density, marked by stylistic devices such as anaphora, metaphorical mappings, and a formal register. These features are not ornamental but central to the communicative purpose of the genre, which aims not only to inform but to symbolically frame foundational principles in order to keep them salient throughout decades. In this regard, the consistent inability of AI-powered systems to identify and reproduce such genre-specific features exposes a fundamental limitation, namely that in this case current NMT models remain ill-equipped for genre-sensitive translation. The reason underlying these observations is to be found in the architecture of NMT systems, which is designed to process language in probabilistic terms, generating output that is statistically likely rather than rhetorically or contextually optimal. As a result, key elements of the genre – its persuasive, symbolic, and performative dimensions – are attenuated or entirely lost in translation.

In light of these findings, ethical and deontological issues come to the fore, particularly in domains such as legal discourse, where rhetorical precision crucially influences the interpretative process. If AI-made translations of such texts are relied upon without the benefit of rigorous post-editing by expert human translators, the risks are manifold: the intended message may be misrepresented, the rhetorical impact diminished, and ambiguities introduced, potentially with



consequences on legal interpretation. For all these reasons, the growing reliance on AI-based translation tools must be tempered by a critical awareness of their serious limitations. Their use in high-stakes, rhetorically complex contexts demands responsible oversight, as these systems are still far from substituting the interpretive judgement that human translators bring, particularly in genres where language functions not merely as a vehicle of information, but as a specialized carrier of legal principles and values. If translation is to be understood as a rhetoric of meaning, as previous studies have suggested (Dávila-Montes 2017), it follows that this art continues to be a distinctly human prerogative.

References

- Bassnett, Susan. 2014. *Translation Studies*. London: Routledge.
- Belinkov, Yonatan, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2020. "On the Linguistic Representational Power of Neural Machine Translation Models." *Computational Linguistics* 46 (1): 1–52. https://doi.org/10.1162/coli_a_00367.
- Confederazione Svizzera. 1996. *Messaggio del 20 novembre 1996 concernente la revisione della Costituzione federale, FF 1997 I 1* (96.091).
- Dávila-Montes, Jose M. 2017. "Translation as a Rhetoric of Meaning." *Poroi* 13 (1).
- Egger, Jean-Luc, and Daria Evangelista. 2022. "La Costituzione scomparsa: Alcuni cambiamenti concettuali e lessicali nella riforma del 1999." [The missing constitution: Some conceptual and lexical changes in the 1999 reform]. In *L'italiano dei testi costituzionali. Indagini linguistiche e testuali tra Svizzera e Italia* [The Italian of constitutional texts. Linguistic and textual investigations between Switzerland and Italy], edited by Angela Ferrari, Letizia Lala, and Filippo Pecorari, 377–396. Alessandria: Edizioni dell'Orso.
- Evangelista, Daria. 2023. "Potissima pars principium est: Osservazioni retorico-pragmatiche sui preamboli delle costituzioni cantonali svizzere." [*Potissima pars principium est: Rhetorical and pragmatic observations on the preambles of the Swiss cantonal constitutions*]. *LTO – Lingue e testi di oggi* 1 (2022): 9–24.
- Evangelista, Daria. 2024. "Tra retorica e denotazione: Un'analisi contrastiva dei preamboli delle Costituzioni svizzere e dei trattati internazionali." [Between rhetoric and denotation: A contrastive analysis of the preambles of Swiss constitutions and international treaties]. *LeGes – Legislazione & Valutazione* 35 (2).
- Evangelista, Daria. 2025. *La legge ornata. Un'analisi dei testi normativi tra retorica e linguistica del testo*. [The ornate law: An analysis of legislative texts between rhetoric and textual linguistics]. Alessandria: Dell'Orso.
- Fahnestock, Jeanne. 2011. *Rhetorical Style: The Uses of Language in Persuasion*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199764129.001.0001>.
- Ferrari, Angela, Letizia Lala, and Luciano Zampese. 2021. *Le strutture del testo scritto. Teoria e esercizi*. [Written text structures: Theory and exercises]. Rome: Carocci.
- Irti, Natalino. 2016. *Un diritto incalcolabile*. [An incalculable right]. Turin: Giappichelli.
- Koehn, Philipp. 2020. *Neural Machine Translation*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108608480>.
- Majdik, Zoltan P., and Scott S. Graham. 2024. "Rhetoric of/with AI: An Introduction." *Rhetoric Society Quarterly* 54 (3): 222–231. <https://doi.org/10.1080/02773945.2024.2343264>.
- Mitchell, Melanie. 2024. "The Metaphors of Artificial Intelligence." *Science* 386 (6723). <https://doi.org/10.1126/science.adt6140>.
- Pecorari, Filippo. 2022. "Osservazioni linguistiche e testuali sul preambolo della Costituzione federale svizzera." [Linguistic and textual observations on the preamble to the Swiss Federal Constitution]. In *L'italiano dei testi costituzionali. Indagini linguistiche e testuali tra Svizzera e Italia* [The Italian of constitutional texts. Linguistic and textual investigations between Switzerland and Italy], edited by Angela Ferrari, Letizia Lala, and Filippo Pecorari, 435–445. Alessandria: Edizioni dell'Orso.
- Prandi, Michele. 2015. "Translating Metaphors." In *Language across Languages: New Perspectives on Translation*, edited by Emanuele Miola and Paolo Ramat, 77–89. Newcastle upon Tyne: Cambridge Scholars Publishing.
- Prandi, Michele. 2023. *Retorica. Una disciplina da rifondare*. [Rhetoric: A discipline to be re-established]. Bologna: Il Mulino.
- Sabatini, Francesco. 1999. "'Rigidità-esplicitzza' vs 'elasticità-implicitzza': Possibili parametri massimi per una tipologia dei testi." ['Rigidity-explicitness' vs 'elasticity-implicitness': Possible maximum parameters for a typology of texts].



In *Linguistica testuale comparativa: In memoriam Maria-Elisabeth Conte* [Comparative Textual Linguistics: In memoriam Maria-Elisabeth Conte], edited by Gunver Skytte and Francesco Sabatini, 141–172. Copenhagen: Museum Tusculanum Press.

Scarpa, Federica. 2008. *La traduzione specializzata. Un approccio didattico professionale*. [Specialized translation: A professional teaching approach]. Milan: Hoepli.

Toral, Antonio, and Victor M. Sánchez-Cartagena. 2017. "A Multifaceted Evaluation of Neural Machine Translation for the English-Spanish Language Pair." In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 206–217.

Volkart, Lise, Pierrette Bouillon, and Sabrina Girletti. 2018. "Statistical vs. Neural Machine Translation: A Comparison of MTH and DeepL at Swiss Post's Language Service." In *Proceedings of the 40th Conference Translating and the Computer*, 145–150.