
Improving Term Evaluation in Machine Translation: Variation Matters

Nicolas Dahan

Inria, Paris, France

nicolas.dahan@inria.fr

Ziqian Peng

Sorbonne Université, CNRS, ISIR, Paris, France & Inria, Paris, France

peng@isir.upmc.fr

François Yvon

Sorbonne Université, CNRS, ISIR, Paris, France

yvon@isir.upmc.fr

Rachel Bawden

Inria, Paris, France

rachel.bawden@inria.fr

Abstract

Terminology evaluation in machine translation (MT) usually assumes a single correct target form per source term. However, human translators routinely introduce variation that current metrics penalize as inconsistency. We examine how to account for this variation in document-level MT evaluation of English–French scientific translation, combining glossary-based accuracy, translation consistency, and a new *cross-term variation* (CTV) diagnostic measure that tests whether variation relationships are preserved across languages. Based on analyses of two parallel corpora, translated by four MT systems, we find that (1) MT systems generate less target-side variation than human translators; (2) transfer patterns strongly depend on the variation type; (3) consistency rankings vary with the choice of metric; and (4) constraining MT with a glossary improves accuracy and consistency but degrades CTV by suppressing valid variation. We argue for variation-aware evaluation that conditions consistency penalties on whether target-side variation mirrors source-side variation.

1 Introduction

Evaluating how machine translation (MT) systems handle terminology in scientific and specialized documents requires accounting for a fundamental property of terms in running texts: *terminological variation*. A concept (e.g., *machine translation*) does not always appear under a single form, which we call the preferred term. It may surface as an acronym (*MT*), a reduction (*translation*), a lexical substitution (*automatic translation*), or a morphosyntactic reformulation (e.g., *translation by machines*) (Daille, 2017). We refer to these alternative surface forms as variants. Such variation arises naturally in scientific writing, for example because of stylistic alternation, register adaptation, or avoidance of repetition, and is a well-attested feature of specialized discourse (Fernández-

Silva and Kerremans, 2011; Daille, 2017).

Whether such variation should be considered an error depends on the role the terminological resource is meant to play. When a glossary prescribes a single sanctioned form, for reasons of trademark, safety, or usability, consistency is genuinely required and any deviation is an error, on both the source and the target side. However, glossaries are just as often provided as vocabulary guidance, and terminology standards themselves recognize a scale of acceptability, from preferred to merely admitted forms (ISO, 2019). Unless terminology is enforced to be unambiguous and used consistently, as in controlled-language settings, stylistic variation is expected and even preferred, as in the scientific corpora we study. Yet most evaluation metrics penalize such variation uniformly, ignoring whether it mirrors variation already present in

the source.

Current MT terminology evaluation does not account for variation. *Glossary-based methods* (ibn Alam et al., 2021; Semenov et al., 2023, 2025) match source terms against expected translations from a bilingual glossary, but cannot evaluate non-glossary variants, i.e. a concept’s surface forms that are not listed in the glossary. Such forms naturally arise in running texts, such as reorderings (*quality of translation* instead of *translation quality*), reductions (*processing* for *language processing*), or context-dependent acronyms. Because such forms are produced dynamically as the text unfolds (by insertion, reduction, or new terms being coined from existing ones), they cannot be exhaustively enumerated in advance, and simply enlarging the glossary therefore cannot solve the problem. *Consistency metrics* (Itagaki et al., 2007; Lyu et al., 2021; Semenov and Bojar, 2022) assess whether a source term is uniformly rendered throughout a document, but considers all target-side variation as inconsistency. Neither family of metrics captures whether the *variation relationship* between a preferred term and its variant, for example between an acronym and its expanded form, is preserved in translation; treating the two as separate terms only measures each in isolation.

In this work, we study how to account for terminological variation in document-level MT evaluation for English–French scientific translation. We extend terminology coverage to non-glossary variants, identified in context by a concordancer (Section 3), and examine three complementary evaluation dimensions: (1) term-level accuracy, (2) translation consistency, both established in prior work, and (3) cross-term variation, a new dimension introduced to assess whether variation relationships are preserved in translation (see Table 1). Each dimension has structural coverage gaps tied to variation type, but together they achieve wider coverage. We analyze two parallel English–French scientific corpora from the Natural Language Processing (NLP) domain (Salesky et al., 2023; Peng et al., 2026), each translated by four MT systems, and report the following findings:

- Human translators introduce more target-side lexical variation than MT systems, reflecting standard translation practices that current consistency metrics penalize.
- Variation types are not always preserved in trans-

lation, with some types (e.g., reductions) transferred consistently and others (e.g., acronyms) strongly dependent on the corpus and its translation guidelines.

- Different measures of translation consistency result in divergent system rankings, showing that existing consistency scores are sensitive to metric design.
- Constraining MT to use glossary terms improves accuracy and consistency metrics, but reduces the system’s ability to transfer source-side variations onto the target, exposing a tension between terminological uniformity and valid variation.

Our contribution is primarily methodological: rather than claim that these variation patterns are universal, we provide a framework and metrics for measuring how terminological variation interacts with MT evaluation, and report empirical findings specific to the two English–French NLP corpora we study.

Head (EN)	automatic translation system
Head (FR)	système de traduction automatique
Glossary variants	none listed
Src variant	translation system (<i>reduction</i>)
Expected tgt	système de traduction (<i>reduction, inferred</i>)
MT output	système de traduction (<i>unevaluated by glossary-based methods</i>)

Table 1: Cross-term coherence example: the reduction *translation system* is absent from the glossary, but its French counterpart is inferred from the reduction relationship to the head term.

2 Related Work

Modern MT evaluation increasingly relies on fine-tuned neural metrics such as COMET (Rei et al., 2020) and MetricX (Juraska et al., 2024), which provide strong overall quality estimates, while LLM-based error analysis approaches such as GEMBA-MQM (Kocmi and Federmann, 2023) produce fine-grained Multidimensional Quality Metrics (MQM) annotations that include terminology errors. These methods operate at the segment level, however, and do not model term-level structure: they neither track per-concept cross-document consistency nor capture

variation relationships between preferred terms and variants. This holds most directly for the neural metrics, which score semantic adequacy rather than surface form: they rarely penalize valid variation but, for the same reason, cannot diagnose terminological inconsistency. GEMBA-MQM goes beyond adequacy, targeting terminology through a dedicated MQM category, but its predicted error spans do not align well with human annotations (Lu et al., 2025), and, scoring each segment in isolation, it does not track whether a concept is rendered consistently across a document or whether variation relationships are preserved. Dedicated term-level evaluation protocols have therefore been developed to close these gaps.

A first family of term-level methods relies on a bilingual glossary, measuring how often MT produces the expected target term for each source entry. Matching strategies range from exact surface-form matching (ibn Alam et al., 2021) to partial and lemma-based matching (Alam et al., 2021; Semenov et al., 2025), and have been adopted in shared evaluation campaigns such as the WMT terminology translation tasks (Alam et al., 2021; Semenov et al., 2023, 2025). Document-level extensions further compute per-term ratios of target-to-source occurrences (Semenov et al., 2025). Even with lemmatization, these approaches remain centered on canonical glossary entries and miss unlisted surface forms such as reductions, dynamic acronyms, and lexical variants.

Glossaries are also useful to enforce specific translations at generation time through lexically constrained decoding (Hokamp and Liu, 2017; Post and Vilar, 2018) or constraint-aware training (Dinu et al., 2019), with related lexical constraints used in paraphrase generation (Hu et al., 2019); our glossary-guided condition (Section 5) instead applies such guidance through prompting rather than hard decoding constraints.

When no glossary is available, consistency-based metrics offer an alternative. Semenov and Bojar (2022) compute per-term consistency with respect to a pseudo-reference derived from the first target realization. Concentration-based metrics, such as the Herfindahl–Hirschman Index (Itagaki et al., 2007; Gašpar et al., 2022) and the Lexical Translation Consistency Ratio (LTCR) (Lyu et al., 2021; Wang et al., 2025) quantify how uniformly a source term is translated, by rewarding the reuse of a dominant form. Despite methodological differences, these methods

share a common assumption: optimal scores require that every occurrence of a term should receive the same translation. They therefore penalize acceptable target-side variation, such as the stylistic alternation commonly observed in scientific writing (Vinay and Darbelnet, 1995), where alternating between forms is a deliberate, well-motivated device rather than carelessness (Bowker, 1998; Pecman, 2014), and enforcing a single form risks over-standardization that erases author-intended distinctions (Bowker and Hawkins, 2006).

This assumption sits at odds with a long line of linguistic work showing that terminological variation is a fundamental property of specialized discourse rather than random noise. Daille (2017) provides a comprehensive characterization of denominative variation in domain-specific corpora, distinguishing graphical, morphosyntactic, semantic, and pragmatic variants of a single concept; Cárcamo et al. (2025) refine this typology into graphical, morphosyntactic, reduction, expansion, and lexical variants. Translation studies have also started to examine how variation behaves under translation: Fernández-Silva and Kerremans (2011) show that source-side variation is broadly preserved or accentuated in human translations of Galician scientific texts into English, while Čulo and Nitzke (2016) observe, in English–German post-edited texts, that MT systems produce fewer variant translations than humans, with MT patterns carrying over into post-editing. This reduction is consistent with broader evidence that MT produces lexically less diverse and less rich output than human translators (Vanmassenhove et al., 2019, 2021), and with the notion of *translationese*, of which post-edited MT is an exacerbated case: it is simpler and more normalized than human translation (Toral, 2019). These descriptive findings have not yet been integrated into automatic evaluation: none of these metrics model how variation types affect evaluation coverage, nor whether structural relationships between preferred terms and variants (e.g., acronym–expansion pairs) are preserved in translation.

3 Methodology

We study how MT systems translate terms at the document level, accounting for the range of surface forms a term may assume in context, along three complementary dimensions: (1) how accurately each occurrence is translated based on a reference glossary,

(2) how consistently each surface form (whether a preferred term or a variant) is rendered throughout its document occurrences, and (3) whether variation relationships between terms and variants are preserved in translation.

Our terminological resource is structured as a Simple Knowledge Organization System (SKOS)¹ glossary: each entry pairs a concept with a preferred term (SKOS `prefLabel`) and a list of alternative terms (SKOS `altLabel`); we use *variant* as a synonym for alternative term. We turn term occurrences into their canonical forms by stripping inflection and other non-terminological variation (e.g., by lemmatization and orthographic normalization). Our pipeline takes this glossary as input and processes parallel documents in four steps: (1) term detection, (2) source-side variation labeling, (3) source-target alignment, and (4) target-side labeling. The output is a list of annotated term pairs relating each source occurrence with the corresponding target span.

Type	Description and example
GRAPH	Graphical (acronyms, spelling): <i>intellectual property / IP</i>
MORPH	Morphosyntactic (reordering, inflection): <i>contract of employment / employment contract</i>
RED	Reduction (truncation): <i>European Union acquis / acquis</i>
EXP	Expansion (added material): <i>wear / normal wear and tear</i>
LEX	Lexical (near-synonymy): <i>action for damages / claim for damages</i>
COMBO	Combined (multiple types): <i>coronavirus disease 2019 / COVID-19</i>

Table 2: Typology used to label variation in source- and target term occurrences (Cárcamo et al., 2025).

Term detection. We identify occurrences of known terms using CONCORDANCER,² a trie-based concordancer operating on documents to which linguistic analysis (lemmatization, POS tagging, dependency parsing) is applied with spaCy (Honnibal et al., 2020). Coverage extends beyond exact glossary entries via two complementary mechanisms: (1) *static*

variants, generated offline from each glossary entry by applying a fixed set of transformation rules (synonym substitution via WordNet (Miller, 1995), syntactic restructuring, morphological derivation via MorphyNet (Batsuren et al., 2021), and modifier reduction); (2) *dynamic variants*, detected at run-time from dependency patterns in context. Together these recover a wider range of forms under which each concept appears in the document.

Source-side variation labeling. Each detected occurrence is labeled with its variation type relative to the preferred term, following the typology of Cárcamo et al. (2025) (Table 2; full typology in Appendix A.14). The label provenance depends on whether the matched form is lexicalized in the glossary: for the preferred term and its alternative terms, we assign the type by few-shot prompting ChatGPT (gpt-4.1-mini; prompts in Appendix 9); for non-lexicalized matches (static or dynamic variants), the type is derived directly from the CONCORDANCER formation pattern that triggered the match. When a non-lexicalized match is derived from an alternative term rather than from the preferred term, the final label combines both classifications (e.g., a dynamic acronym of an EXP variant becomes COMBO). To validate this step, one of the authors, a computational linguist and native French speaker fluent in English, manually assigned gold variation labels to 50 sampled pairs; the classifier agrees with these human labels on 84% of types and 77% of subtypes, with near-perfect agreement on single-type variants and systematic under-detection of COMBO (recall 11%; see Appendix B). Since COMBO accounts for fewer than 2% of occurrences, this has little impact on aggregate scores.

In our corpora, exact-form matching accounts for about 62–63% of term occurrences, with lemma matching increasing these numbers to about 82–87%. Additional matches contributed by CONCORDANCER (static and dynamic variants) are dominated by lexical (LEX) and reduction (RED) forms (full breakdown in Appendix A.1).

Alignment. We apply Bertalign (Liu and Zhu, 2022) for sentence-level alignment, and the token-

¹<https://www.w3.org/TR/skos-reference/>

²<https://gitlab.inria.fr/almanach/concordancer>

³We use the HuggingFace models `paraphrase-multilingual-mpnet-base-v2` and `facebook/xlm-roberta-xl` as the embedding models for Bertalign and SimAlign, respectively.

level correspondences are computed with SimAlign (Jalili Sabet et al., 2020).³ For each source term occurrence, the aligned target span is assembled from individual token alignments. When the cosine similarity of the assembled span falls below 0.8, the bilingual sentence pair, the source term, and the SimAlign candidate are submitted to ChatGPT (gpt-4.1-mini; prompt in Appendix 2); the LLM returns the translated form of the source term as it appears in the target sentence (an empty return signals an omission), which replaces the SimAlign span. The 0.8 cutoff is a conservative operating point that reserves the LLM fallback for spans where the embedding aligner is least confident, keeping higher-similarity spans unchanged. To validate this step, the same annotator evaluated 100 term alignments sampled across the nine system outputs; 96% (96/100) were correct, with the four errors limited to span-boundary and lexical-choice mistakes on single occurrences (see Appendix C).

Target-side labeling. Target-side analysis combines two parallel processes that are reconciled afterwards. (1) CONCORDANCER scans the full target document and labels each detected occurrence using the same two-level protocol as the source side (CONCORDANCER + ChatGPT), producing a variation label relative to the expected target preferred term. (2) The alignment step assembles, for each source term occurrence, a target span. We reconcile the two by matching each canonicalized aligned span against CONCORDANCER’s target term list: all aligned spans found by CONCORDANCER inherit its variation label, while aligned spans missed by CONCORDANCER are directly labeled by ChatGPT, given the source term and the aligned span. Each realization is additionally assigned a glossary conformity label: *glossary-conforming*, *variant translation*, or *out-of-glossary* (paraphrase, omission, or error).

4 Evaluation Framework

We analyse source-target associations of terms along three measurements: *term-level accuracy* (Dimension 1, building on glossary-based exact match (ibn Alam et al., 2021)), *translation consistency* (Dimension 2, building on the consistency tradition (Itagaki et al., 2007; Lyu et al., 2021; Semenov and Bojar, 2022)), and *cross-term variation* (Dimension 3, introduced here). These quantities serve a dual purpose. Applied to the human translation references, they de-

scribe terminological behaviour in a corpus; applied to MT outputs against the source and the reference, they evaluate translation quality. The descriptive use establishes that the human reference is not arbitrary; the evaluative use asks whether MT systems resemble human decisions, and whether such similarities are even desirable.

Dimension 1: Term-level accuracy. We compute Exact Match (EM) accuracy: the proportion of occurrences whose translation matches the glossary entry. EM can be applied to any occurrence whose canonical form maps to a glossary entry (preferred term or alternative term).

Dimension 2: Translation consistency. For each canonical surface form (that is, each preferred term or variant treated independently) with at least two occurrences, we collect all its target counterparts in a document and measure the consistency of the resulting translation set. We apply three metrics to compute a translation consistency ratio (TCR); none of them account for whether the target-side variation is justified by source-side differences.

- **TCR_{EMmacro} (TCR-EM):** for each term with $k \geq 2$ occurrences, computes the proportion of occurrences that match the glossary, then averages across terms (term-level).
- **TCR temporal (TCR-T):** for each pair of successive occurrences of the same source canonical form (the 1st and 2nd, then the 2nd and 3rd, etc.), TCR-T checks whether they receive identical translations. The score is the proportion of consecutive pairs that agree.
- **TCR pseudo-ref (TCR-PR):** consistency measured against a pseudo-reference form (the first occurrence). This score does not require a glossary and applies even when a terminological resource is lacking.

Of these, TCR-PR implements the pseudo-reference consistency of Semenov and Bojar (2022), TCR-EM applies glossary-based EM (ibn Alam et al., 2021) at the term level, and TCR-T is a temporal variant we adopt to capture whether translation shifts are gradual or repeated; we retain all three, rather than select one, because their divergence is itself a finding: the choice of consistency definition is not neutral and can change system rankings (Section 6.2). TCR-EM

combines Dimensions 1 and 2: occurrences must agree with each other and with the glossary. This makes it more restrictive than TCR-T and TCR-PR: a term consistently translated with a correct but non-glossary variant translation would score 0 despite being perfectly consistent. TCR-T and TCR-PR are purely consistency-oriented: they measure whether a term is translated uniformly, regardless of which translation is used. Table 20 (Appendix A.8) illustrates how these metrics can diverge on the same data.

Dimension 3: Cross-term variation (CTV). While the previous dimensions assess each preferred term or variant independently (each treated as a separate canonical surface form), this dimension evaluates whether the variation relationships between a preferred term and its variants are transferred to the target. Specifically, we measure *Variation preservation*: when a source preferred term is correctly translated, each GRAPH/EXP/RED/LEX variant should be translated as a corresponding form in the target. We retain these four types because each has a predictable expected target form: a source acronym is typically rendered as a target acronym (GRAPH), a truncated source form as a target truncation (RED), an expanded source form as a target expansion (EXP), and a near-synonymous substitution as a target lexical variant or the target preferred term itself (LEX).⁴

Metric	Definition	Exclusions
1. Term-Level Accuracy		
EM (micro)	<i>Occurrence-level conformity</i>	OOG terms
2. Translation Consistency		
TCR-PR	<i>Match to first occurrence</i>	$k_t < 2$
TCR-T	<i>Consecutive stability</i>	$k_t < 2$
TCR-EM	<i>Avg. conformity per term</i>	$k_t < 2 + \text{OOG}$
3. Cross-Term Variation		
CTV	<i>Preserve variation type</i>	MORPH, COMBO

Table 3: Evaluation dimensions. k_t : number of occurrences of term t . OOG: out-of-glossary.

Table 3 summarizes all metrics, their objectives,

⁴MORPH and COMBO are excluded because their source-to-target mapping is less predictable: morphosyntactic changes depend on target-language grammar, and compound variants involve multiple simultaneous transformations.

⁵Traitement Automatique des Langues Naturelles, <https://www.atala.org/-Conference-TALN-RECITAL>.

⁶Journées d’Etude de la Parole, <https://www.afcp-parole.org/category/animations-scientifiques/les-jep/>.

⁷<https://2022.aclweb.org>.

⁸<https://skosmos.loterre.fr/8LP/fr/>

and exclusion criteria. Importantly, the framework is diagnostic: it characterizes where and how current evaluation clashes with term variation but does not assess translation quality. A term may be consistently translated yet semantically wrong, or inconsistently translated in a contextually appropriate manner.

5 Experimental Setup

We conduct our analysis on two English–French parallel corpora in the NLP domain: ParaNLP and IWSLT2023. Table 4 summarizes their statistics.

ParaNLP The ParaNLP corpus (Peng et al., 2026) consists of 32 English papers published in NLP conferences (2002–2015) paired with comparable French versions published in TALN⁵ or JEP⁶ conferences, resulting in 6,272 parallel sentences. It was constructed via automatic sentence alignment, with Automatic Post-Editing and MT used to handle poorly aligned segments (details in Appendix A.12).

IWSLT2023 The IWSLT2023 corpus comprises the validation and evaluation sets of the IWSLT 2023 translation shared task (Salesky et al., 2023). It contains manually revised transcripts and translations of 10 presentations delivered at the ACL 2022 conference,⁷ aligned at the sentence level. We use the English–French subset.

Terminological resource. Our terminological reference is a manually curated bilingual glossary covering 1,642 NLP concepts with their preferred English and French realizations, compiled from domain-specific resources and validated by experts. It was prepared and published by INIST, the French Institute for Scientific and Technological Information.⁸

Models and Inference. We evaluate four medium-size multilingual LLMs using vLLM (Kwon et al., 2023) (model cards in Appendix A.11): Llama3.1-8B-Instruct (Llama) (Grattafiori et al., 2024), Qwen3-8B (Qwen) (Yang et al., 2025), EuroLLM-9B-Instruct (Euro9) (Martins et al., 2025), and EuroLLM-22B-Instruct (Euro22) (Ramos et al., 2026). We refer to the human reference as “Ref”. We ap-

ply greedy decoding to all systems except Qwen, for which we use the recommended “non-thinking” mode (temperature=0.7, top- p =0.8, top- k =20).

Translation settings. We evaluate MT outputs under two zero-shot conditions: (1) baseline, where models translate (English-to-French) without terminological guidance; and (2) glossary-guided (base+terms), where glossary-preferred terms and their expected translations are injected into the prompt. Comparing the two conditions serves as a diagnostic for whether constrained systems genuinely exploit the glossary information (prompts in Appendix 3).

6 Analysis

6.1 Variation Patterns: Reference and MT

Statistic	ParaNLP	IWSLT2023
<i>Corpus overview</i>		
#docs	32	10
Avg. #sents / doc	196	88
Avg. concepts / doc	50.44	31.00
<i>Occurrence & variation richness</i>		
Avg. occ. / concept	5.55	3.92
Avg. distinct terms / concept	1.37	1.14
<i>Concept distribution</i>		
Single-term concepts (% of concepts)	77.74	88.27
Single-term concepts (% of occ.)	64.44	76.94
Hapax concepts (% of concepts)	37.56	42.68
Hapax concepts (% of occ.)	7.58	11.55
<i>Variation category breakdown (% of occ.)</i>		
NO-VAR	72.28	74.51
LEX	7.97	5.66
RED	9.21	8.93
EXP	1.10	0.39
MORPH	0.59	0.29
GRAPH	7.54	8.73
COMBO	1.32	1.49

Table 4: Terminological variation statistics of the source texts, averaged at the document level.

Table 4 characterizes source-side terminological variation in both corpora. Each concept appears under 1.37 distinct surface forms in ParaNLP and 1.14 in IWSLT2023 on average. 77.7% and 88.3% of concepts respectively are realized by a single form (full distribution in Appendix A.2). Variation is concentrated among frequent concepts, and preferred terms make up 72–75% of occurrences. Among variation

types, RED and GRAPH account for 7–9% each, LEX for 6–8%, and MORPH, EXP, and COMBO together represent less than 3%. Since 38–43% of concepts are hapaxes, any score requiring multiple occurrences is blind to much of the terminology.

The analysis of reference targets in this same setting (Table 5, row *Ref*) confirms that variation is not random noise: most concepts retain their number of distinct surface forms across languages. However, in ParaNLP, 34.4% of concepts exhibit more target-side than source-side realizations, higher than the MT systems. In IWSLT2023 (spoken transcripts), the rate drops to 15.8%. This register asymmetry is consistent with the well-documented tendency of French academic writing to favor lexical variation (*variation élégante*) (Vinay and Darbelnet, 1995).

System	Var.	Ref	E22	E9	Qwen	Llama
ParaNLP						
baseline	=	63.7	69.1	68.5	66.5	66.6
	<	34.4	29.4	29.8	32.2	31.9
base+terms	=	63.8	73.8	72.4	73.2	73.6
	<	34.3	24.1	25.3	24.6	24.8
IWSLT2023						
baseline	=	82.5	67.8	66.7	69.5	64.4
	<	15.8	31.6	32.8	29.9	35.0
base+terms	=	82.5	77.4	74.0	79.7	77.4
	<	15.8	21.5	24.9	19.2	21.5

Table 5: Distribution (%) of concepts by source/target variation (=: EN and FR have the same number of distinct terms; <: FR has more). Row EN>FR omitted (all values below 2.3%). Avg. concepts/doc: 31.48 (ParaNLP), 17.70 (IWSLT2023).

Human translators also routinely replace preferred terms with lexical variants: in the reference, source preferred terms are rendered by a target-side lexical variant (LEX) in 24.98% of cases in IWSLT2023 and 15.68% in ParaNLP (full transfer matrices in Appendix A.3). A concrete case: within a single sentence of the ParaNLP corpus, the English term *reported speech* receives two distinct French translations in the reference (*parole rapportée* and *discours rapporté*, both valid translations of the same phenomenon), a stylistic alternation that a consistency metric may penalize as inconsistency (full ex-

ample with MT outputs in Table 15, Appendix A.4).

MT systems reproduce variation unevenly.

Whether MT outputs reproduce the source’s variation depends on register and variation type. For academic prose (ParaNLP), MT amplification (29–33%) is in the same ballpark as the reference (34.4%); for transcripts (IWSLT2023), MT systems amplify more (30–35%) than the reference (15.8%).

The differences deepen when we break down variation transfer by type, assessing how each source variation type (preferred term, acronym, reduction, etc.) is rendered on the target side (full transfer matrices in Appendix A.3). RED transfers most reliably, with 83–86% preservation on both corpora. GRAPH, in contrast, shows the widest gap between corpora: 50–60% preservation on ParaNLP drops to 14–29% on IWSLT2023, where source acronyms are far more often spelled out as a French phrase than kept as an abbreviation. This likely reflects the translation guidelines of IWSLT2023 rather than the spoken source: even the human reference spells acronyms out more than the MT systems, which a source-side effect could not explain. NO-VAR preservation is likewise higher on ParaNLP (81–85%) than on IWSLT2023 (71–75%). Evaluating all variation types uniformly therefore conflates phenomena with fundamentally different cross-lingual behaviors. For instance, in a single ParaNLP document the acronym *RL* (GRAPH) for *reinforcement learning* is preserved in one context but expanded into different forms in others (full example in Table 16, Appendix A.6).

6.2 Current Metrics Penalize Valid Variation

Method	ParaNLP		IWSLT2023	
	occ	dist	occ	dist
EM	85.0	71.7	90.4	83.3
TCR	86.9	50.5	85.6	53.2
CTV	9.4	12.5	2.7	4.2
EM+TCR	95.1	81.1	96.9	90.0
EM+CTV	91.9	82.6	92.8	87.2
TCR+CTV	88.7	57.7	86.2	55.0
EM+TCR+CTV	96.8	87.9	97.5	91.9

Table 6: Average per-document coverage (%) by method. *occ*: each occurrence; *dist*: each distinct term counted once. TCR uses non-EM variants.

Before reporting scores for the three measurements, we discuss their *coverage*, i.e. the fraction of source-side occurrences each one evaluates. The three evaluation dimensions have complementary coverage profiles (Table 7). EM covers 100% of NO-VAR terms by design but this drops sharply for variants.

EM’s coverage drop on variants is structural: MORPH and COMBO are essentially invisible to glossary-based evaluation, and GRAPH coverage varies across corpora (97% in IWSLT2023, where most acronyms are listed in the glossary, versus 59% in ParaNLP, where many domain-specific acronyms are absent from it). TCR covers 86–87% of occurrences on both corpora but applies only to recurrent canonical forms, assessing each surface form independently and missing cross-variant coherence. CTV has the narrowest raw coverage but provides a qualitatively distinct signal: it is the only dimension that evaluates hapax variants (38–43% of concepts) through their structural relationship to the preferred term. Combined, the three dimensions reach 96.8% of occurrences in ParaNLP and 97.5% in IWSLT2023, leaving a small structurally irreducible residual (Table 6; details in Appendix A.7).

Table 8 reports the three-dimensional scores for MT systems (baseline and base+terms); we focus on the baseline here, leaving base+terms for Section 6.3. Aggregate EM places all systems within a narrow range per corpus, hiding a structural effect: the per-variation breakdown (Appendix A.9) shows that aggregate EM is driven by preferred-term performance (72–85% on NO-VAR) while variant scores are substantially lower. The aggregate therefore says little about how systems handle variants.

Cat.	ParaNLP					IWSLT2023				
	$\bar{N}$	EM	TCR	TCR-EM	CTV	$\bar{N}$	EM	TCR	TCR-EM	CTV
NO-VAR	198	100	90	90	—	88	100	87	87	—
RED	25	43	89	42	32	10	50	85	47	17
MORPH	3	8	2	2	—	1	0	0	0	—
EXP	4	26	32	13	40	4	42	42	42	58
GRAPH	21	59	69	53	32	11	97	82	82	7
LEX	22	35	63	24	32	8	41	46	22	8
COMBO	5	5	26	5	—	4	0	28	0	—

Table 7: % of occurrences evaluated per category, averaged over documents. $\bar{N}$: avg. occurrences. Values rounded to the nearest integer.

Consistency scores expose the cost of penalizing acceptable variation. The reference dominates for

IWSLT2023 under all TCR variants; for ParaNLP, differences narrow. The strictest variant (TCR-EM) results in lower scores across the board, reflecting the compounding difficulty of requiring both accuracy and consistency.

Consistency is hardest to achieve precisely for terms that inherently exhibit variation: TCR scores for LEX and GRAPH are substantially lower than for NO-VAR (Appendix A.9). For instance, the term *pitch* (glossary: *tonie*) receives two contextually appropriate French equivalents (*hauteur, mélodique*) across 14 occurrences in a single ParaNLP document, and a consistency metric severely penalizes this justified variation (full example in Table 17, Appendix A.5).

Metric	ParaNLP					IWSLT2023				
	Ref	E22	E9	Q.	Ll.	Ref	E22	E9	Q.	Ll.
baseline										
EM Acc.	81.5	79.3	79.9	77.8	78.4	72.3	71.3	72.0	74.1	69.9
TCR-EM	83.0	80.8	81.3	79.4	79.8	74.2	73.2	73.5	76.5	72.0
TCR-PR	86.7	88.2	89.2	86.6	89.3	95.1	84.4	87.9	88.1	86.1
TCR-T	85.7	87.2	87.4	85.8	87.7	94.0	83.3	84.0	84.2	82.7
CTV	66.1	72.8	72.7	71.7	74.4	86.7	85.7	85.7	68.2	75.8
base+terms										
EM Acc.	81.5	87.9	85.4	88.3	88.8	72.3	84.5	78.4	87.4	86.6
TCR-EM	83.0	89.0	86.4	89.2	89.7	74.3	86.0	79.6	89.0	87.6
TCR-PR	86.8	91.6	90.5	91.3	90.7	95.1	89.9	90.3	93.0	91.6
TCR-T	85.7	90.7	89.7	90.6	89.8	94.0	87.5	88.3	90.1	89.0
CTV	66.1	67.0	68.5	69.2	68.2	86.7	74.2	81.0	80.6	81.1

Table 8: Scores for baseline and base+terms. E22/E9=Euro22/Euro9, Q.=Qwen, Ll.=Llama, Bold=highest across systems per metric and corpus.

The reference’s CTV score makes the same point from a different angle. While the reference achieves the highest CTV for IWSLT2023, it scores the lowest for ParaNLP. This paradox is explained by the higher target-side variation the reference introduces for ParaNLP, which creates more opportunities for structural coherence failure. Consistency metrics are not stable across definitions. For ParaNLP, each TCR variant ranks a different system as top-1 (Table 8): no single system dominates consistency across all three operationalizations, so the choice of metric is not a neutral implementation detail.

6.3 Glossary Constraints vs. Acceptable Variations

The base+terms block of Table 8 reports MT scores where the preferred terms from the glossary and their expected translations are injected into the translation prompt. Comparing with the baseline, EM and TCR improve substantially for all MT systems: when prompted with the glossary entries, systems readily produce them, showing that they were capable of generating these forms but, without the prompt, often selected alternative wordings. The constraint’s effect on the transfer matrices (Appendix A.3) is selective: NO-VAR → NO-VAR increases by 8–16 points while NO-VAR → LEX drops correspondingly, and other variation types are largely unaffected.

CTV, however, drops for all MT systems under the glossary constraint. Forcing glossary-preferred forms improves both accuracy and consistency at the price of suppressing acceptable target-side variation: the very behavior that Section 6.1 showed to be a hallmark of human translation. This is the central tension exposed by our framework: no single dimension captures the full picture, and a glossary-constrained system that appears to “improve” on EM and TCR may simultaneously regress on CTV. The methodological takeaway is that variation-aware evaluation should condition consistency penalties on whether target-side variation mirrors source-side variation, rather than penalizing all deviation uniformly. Mirroring source variation is a sufficient but not a necessary condition for legitimate target variation: it validates the structural case that CTV measures, but not the stylistic alternation of a term that is itself invariant in the source (e.g., *pitch*, *reported speech*), which has no source counterpart to mirror and can only be adjudicated semantically. CTV is therefore diagnostic rather than prescriptive: it flags where surface consistency and structurally legitimate variation diverge, without licensing every target-side deviation.

7 Conclusion

Terminological variation is a standard property of scientific translation. However, existing automatic metrics systematically penalize it. Analyzing two English–French NLP corpora translated by four medium-size LLMs, we show that human translators introduce more target-side variation than MT systems on academic prose, and that this variation is often semantically motivated. MT systems repro-

duce it unevenly across registers, less the reference for academic prose and more for transcripts.

Current automatic metrics handle this variation poorly: glossary-based EM is dominated by preferred-term performance, translation consistency penalizes the acceptable diversification that human translators produce, and different operationalizations of consistency result in divergent system rankings. Cross-term variation (CTV), the structural coherence dimension we introduce, recovers signal where the other dimensions are blind. No single metric captures terminology quality, motivating variation-aware evaluation that conditions consistency penalties on whether target-side variation mirrors structurally legitimate source-side variation.

The most striking result is that enforcing MT systems to use glossary-preferred forms improves accuracy and consistency but degrades CTV, by suppressing the acceptable target-side variation that the unconstrained baseline preserved. This reveals the need for better MT engines that produce high-quality translations with terminological accuracy, consistency and acceptable variation.

8 Limitations

We carry out our analysis on one language pair (English–French), one domain (NLP scientific text), two corpora, and medium-size LLMs (8B–22B parameters) evaluated under zero-shot conditions. Variation transfer patterns are likely to differ for typologically distant language pairs (e.g., English–Chinese), where acronym and reduction behavior is fundamentally different, and for other specialized domains with distinct terminological profiles. Larger frontier models or translation-specialized systems such as Tower (Alves et al., 2024) may also exhibit different variation behavior; in particular, more capable models might integrate glossary guidance with less suppression of valid variation, an open question that our zero-shot medium-size setting cannot settle. Confirming the generalizability of our findings requires broader experimental coverage.

Our cross-term variation diagnostic (CTV; Dimension 3, Section 4) measures structural coherence under an idealized transfer assumption: it expects each source variation type to be preserved in translation, but legitimate cross-lingual adaptation may well produce a different type (e.g., expanding an acronym

for clarity in the target language).

We do not test how neural or reference-free metrics (e.g., COMET, MetricX, or LLM-as-judge) treat terminological variation. Since they score semantic adequacy rather than surface similarity, they may be more robust to legitimate alternation. However, they commonly operate at the segment level and are not well adapted for the document level. Whether they can also diagnose the terminological inconsistency that our surface-level dimensions target therefore remains an open question.

Several pipeline choices condition our results. The variation classifier uses a closed commercial LLM (gpt-4.1-mini), so exact reproducibility depends on API-version stability; manual validation on a sample of 50 pairs reports 84% type accuracy, with reduced recall on compound variants (11%). As COMBO accounts for fewer than 2% of all occurrences, this affects only finer-grained per-type conclusions. Using an LLM to produce variation labels that are then used to evaluate other LLMs’ translations introduces a potential circular dependency, although the tasks involved (classification vs. translation) are distinct. Alignment errors from Bertalign and SimAlign propagate to all three evaluation dimensions; a manual audit of 100 sampled alignments estimates their rate at 4% (Appendix C). Finally, CONCORDANCER depends on language-specific resources (spaCy, MorphyNet) and rule-based heuristics for dynamic variant detection, which may miss unconventional syntactic constructions. We quantify what CONCORDANCER finds (Appendix A.1) but not what it misses, as no gold-standard annotation of all term occurrences exists to measure detection recall. The framework also requires a bilingual glossary, which may not be available for all domains or language pairs.

Acknowledgements

This work was supported by the French national research agency (ANR) as part of the MaTOS project (grant number: ANR-22-CE23-0033).⁹ Rachel Bawden was also partly funded by her chair position in the PRAIRIE institute funded by ANR as part of the “Investissements d’avenir” programme under reference ANR19-P3IA-0001. The authors are grateful to the anonymous reviewers for their insightful comments and suggestions.

⁹<http://anr-matos.fr/>

References

- Alam, M. M. I., Kvapilíková, I., Anastasopoulos, A., Besacier, L., Dinu, G., Federico, M., Gallé, M., Jung, K., Koehn, P., and Nikoulina, V. (2021). Findings of the WMT shared task on machine translation using terminologies. In Barrault, L., Bojar, O., Bougares, F., Chatterjee, R., Costa-jussa, M. R., Federmann, C., Fishel, M., Fraser, A., Freitag, M., Graham, Y., Grundkiewicz, R., Guzman, P., Haddow, B., Huck, M., Yepes, A. J., Koehn, P., Kocmi, T., Martins, A., Morishita, M., and Monz, C., editors, *Proceedings of the Sixth Conference on Machine Translation*, pages 652–663, Online. Association for Computational Linguistics.
- Alves, D. M., Pombal, J., Guerreiro, N. M., Martins, P. H., Alves, J., Farajian, M. A., Peters, B., Rei, R., Fernandes, P., Agrawal, S., Colombo, P., de Souza, J. G. C., and Martins, A. F. T. (2024). Tower: An open multilingual large language model for translation-related tasks. In *First Conference on Language Modeling*.
- Batsuren, K., Bella, G., and Giunchiglia, F. (2021). MorphoNet: a large multilingual database of derivational and inflectional morphology. In Nicolai, G., Gorman, K., and Cotterell, R., editors, *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 39–48, Online. Association for Computational Linguistics.
- Bowker, L. (1998). Variant terminology: Frivolity or necessity? In *Proceedings of the 8th EURALEX International Congress*, pages 487–495, Liège, Belgium.
- Bowker, L. and Hawkins, S. (2006). Variation in the organisation of medical terms. *Terminology*, 12(1):79–110.
- Cárcamo, J. C., Kübler, N., Mestivier, A., and Zhu, L. (2025). Proposition d’une méthodologie pour l’annotation de la variation dénomminative dans les textes scientifiques en sciences de la Terre. In *Beyond Single Words*, Paris, France. Lichao Zhu and Natalie Kübler and Mojca Pecman.
- Čulo, O. and Nitzke, J. (2016). Patterns of terminological variation in post-editing and of cognate use in machine translation in contrast to human translation. In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation*, pages 106–114.
- Daille, B. (2017). *Term variation in specialised corpora: characterisation, automatic discovery and applications.*, volume 19 of *Terminology and Lexicography Research and Practice series*. John Benjamins Publishing Company.
- Deutsch, D., Briakou, E., Caswell, I. R., Finkelstein, M., Galor, R., Juraska, J., Kovacs, G., Lui, A., Rei, R., Riesa, J., Rijhwani, S., Riley, P., Salesky, E., Trabelsi, F., Winkler, S., Zhang, B., and Freitag, M. (2025). WMT24++: Expanding the language coverage of WMT24 to 55 languages & dialects. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T., editors, *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12257–12284, Vienna, Austria. Association for Computational Linguistics.
- Dinu, G., Mathur, P., Federico, M., and Al-Onaizan, Y. (2019). Training neural machine translation to apply terminology constraints. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3063–3068, Florence, Italy. Association for Computational Linguistics.
- Fernández-Silva, S. and Kerremans, K. (2011). Terminological variation in source texts and translations: A pilot study. *Meta*, 56(2):318–335.
- Gašpar, A., Seljan, S., and Kučič, V. (2022). Measuring Terminology Consistency in Translated Corpora: Implementation of the Herfindahl-Hirshman Index. *Information*, 13(2).
- Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., Yang, A., Fan, A., Goyal, A., Hartshorn, A., Yang, A., Mitra, A., Sravankumar, A., Korenev, A., Hinsvark, A., Rao, A., Zhang, A., Rodriguez, A., Gregerson, A., Spataru, A., Roziere, B., Biron, B., Tang, B., and et al., B. C. (2024). The Llama 3 Herd of Models. Preprint arXiv:2407.21783.
- Hokamp, C. and Liu, Q. (2017). Lexically constrained decoding for sequence generation using grid beam search. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- Honnibal, M., Montani, I., van Landeghem, S., and Boyd, A. (2020). spaCy: Industrial-strength Natural Language Processing in Python. <https://spacy.io>.
- Hu, J. E., Rudinger, R., Post, M., and Van Durme, B. (2019). PARABANK: Monolingual bitext generation

- and sentential paraphrasing via lexically-constrained neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6521–6528.
- ibn Alam, M. M., Anastasopoulos, A., Besacier, L., Cross, J., Gallé, M., Koehn, P., and Nikoulina, V. (2021). On the evaluation of machine translation for terminology consistency. Preprint arXiv:2106.11891.
- ISO (2019). ISO 1087:2019: Terminology work and terminology science. Vocabulary. International Organization for Standardization, Clause 3.4.
- Itagaki, M., Aikawa, T., and He, X. (2007). Automatic validation of terminology translation consistency with statistical method. In Maegaard, B., editor, *Proceedings of Machine Translation Summit XI: Papers*, Copenhagen, Denmark.
- Jalili Sabet, M., Dufter, P., Yvon, F., and Schütze, H. (2020). SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings. In Cohn, T., He, Y., and Liu, Y., editors, *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Juraska, J., Deutsch, D., Finkelstein, M., and Freitag, M. (2024). MetricX-24: The Google submission to the WMT 2024 metrics shared task. In Haddow, B., Kocmi, T., Koehn, P., and Monz, C., editors, *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.
- Kocmi, T. and Federmann, C. (2023). GEMBA-MQM: Detecting translation quality error spans with GPT-4. In Koehn, P., Haddow, B., Kocmi, T., and Monz, C., editors, *Proceedings of the Eighth Conference on Machine Translation*, pages 768–775, Singapore. Association for Computational Linguistics.
- Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J. E., Zhang, H., and Stoica, I. (2023). Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Liu, L. and Zhu, M. (2022). Bertalign: Improved word embedding-based sentence alignment for Chinese–English parallel corpora of literary texts. *Digital Scholarship in the Humanities*, 38(2):621–634.
- Lu, Q., Ding, L., Zhang, K., Zhang, J., and Tao, D. (2025). MQM-APE: Toward high-quality error annotation predictors with automatic post-editing in LLM translation evaluators. In Rambow, O., Wanner, L., Apidianaki, M., Al-Khalifa, H., Eugenio, B. D., and Schockaert, S., editors, *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5570–5587, Abu Dhabi, UAE. Association for Computational Linguistics.
- Lyu, X., Li, J., Gong, Z., and Zhang, M. (2021). Encouraging lexical translation consistency for document-level neural machine translation. In Moens, M.-F., Huang, X., Specia, L., and Yih, S. W.-t., editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3265–3277, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Martins, P. H., Alves, J., Fernandes, P., Guerreiro, N. M., Rei, R., Farajian, A., Klimaszewski, M., Alves, D. M., Pombal, J., Boizard, N., Faysse, M., Colombo, P., Yvon, F., Haddow, B., de Souza, J. G. C., Birch, A., and Martins, A. F. T. (2025). EuroLLM-9B: Technical Report. Preprint arXiv:2506.04079.
- Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Pecman, M. (2014). Variation as a cognitive device: How scientists construct knowledge through term formation. *Terminology*, 20(1):1–24.
- Peng, Z., Zhu, L., Bawden, R., Bénard, M., de la Clergerie, É., Huguin, M., Kübler, N., Lerner, P., Mestivier, A., and Yvon, F. (2026). Parallel Corpora of Scholarly Documents for English-French Machine Translation. In *Proceedings of the 19th Workshop on Building and Using Comparable Corpora (BUCC)*, Palma de Mallorca, Spain. LREC2026.
- Post, M. and Vilar, D. (2018). Fast lexically constrained decoding with dynamic beam allocation for neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1314–1324, New Orleans, Louisiana. Association for Computational Linguistics.

- Ramos, M. M., Alves, D. M., Gisserot-Boukhlef, H., Alves, J., Martins, P. H., Fernandes, P., Pombal, J., Guerreiro, N. M., Rei, R., Boizard, N., Farajian, A., Klimaszewski, M., de Souza, J. G. C., Haddow, B., Yvon, F., Colombo, P., Birch, A., and Martins, A. F. T. (2026). Eurollm-22b: Technical report. Preprint arXiv:2602.05879.
- Rei, R., Stewart, C., Farinha, A. C., and Lavie, A. (2020). COMET: A neural framework for MT evaluation. In Webber, B., Cohn, T., He, Y., and Liu, Y., editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Rei, R., Treviso, M., Guerreiro, N. M., Zerva, C., Farinha, A. C., Maroti, C., C. de Souza, J. G., Glushkova, T., Alves, D., Coheur, L., Lavie, A., and Martins, A. F. T. (2022). CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task. In Koehn, P., Barrault, L., Bojar, O., Bougares, F., Chatterjee, R., Costa-jussà, M. R., Federmann, C., Fishel, M., Fraser, A., Freitag, M., Graham, Y., Grundkiewicz, R., Guzman, P., Haddow, B., Huck, M., Jimeno Yepes, A., Kocmi, T., Martins, A., Morishita, M., Monz, C., Nagata, M., Nakazawa, T., Negri, M., N  v  ol, A., Neves, M., Popel, M., Turchi, M., and Zampieri, M., editors, *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Salesky, E., Darwish, K., Al-Badrashiny, M., Diab, M., and Niehues, J. (2023). Evaluating multilingual speech translation under realistic conditions with resegmentation and terminology. In Salesky, E., Federico, M., and Carpuat, M., editors, *Proceedings of the 20th International Conference on Spoken Language Translation (IWSLT 2023)*, pages 62–78, Toronto, Canada (in-person and online). Association for Computational Linguistics.
- Semenov, K. and Bojar, O. (2022). Automated evaluation metric for terminology consistency in MT. In Koehn, P., Barrault, L., Bojar, O., Bougares, F., Chatterjee, R., Costa-juss  , M. R., Federmann, C., Fishel, M., Fraser, A., Freitag, M., Graham, Y., Grundkiewicz, R., Guzman, P., Haddow, B., Huck, M., Jimeno Yepes, A., Kocmi, T., Martins, A., Morishita, M., Monz, C., Nagata, M., Nakazawa, T., Negri, M., N  v  ol, A., Neves, M., Popel, M., Turchi, M., and Zampieri, M., editors, *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 450–457, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Semenov, K., Huang, X., Zouhar, V., Berger, N., Zhu, D., Oncevay, A., and Chen, P. (2025). Findings of the WMT25 terminology translation task: Terminology is useful especially for good MTs. In Haddow, B., Kocmi, T., Koehn, P., and Monz, C., editors, *Proceedings of the Tenth Conference on Machine Translation*, pages 554–576, Suzhou, China. Association for Computational Linguistics.
- Semenov, K., Zouhar, V., Kocmi, T., Zhang, D., Zhou, W., and Jiang, Y. E. (2023). Findings of the WMT 2023 shared task on machine translation with terminologies. In Koehn, P., Haddow, B., Kocmi, T., and Monz, C., editors, *Proceedings of the Eighth Conference on Machine Translation*, pages 663–671, Singapore. Association for Computational Linguistics.
- Toral, A. (2019). Post-edits: an exacerbated translationese. In Forcada, M., Way, A., Haddow, B., and Senrich, R., editors, *Proceedings of Machine Translation Summit XVII: Research Track*, pages 273–281, Dublin, Ireland. European Association for Machine Translation.
- Vanmassenhove, E., Shterionov, D., and Gwilliam, M. (2021). Machine translationese: Effects of algorithmic bias on linguistic complexity in machine translation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2203–2213. Association for Computational Linguistics.
- Vanmassenhove, E., Shterionov, D., and Way, A. (2019). Lost in translation: Loss and decay of linguistic richness in machine translation. In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 222–232, Dublin, Ireland. European Association for Machine Translation.
- Vinay, J.-P. and Darbelnet, J. (1995). *Comparative Stylistics of French and English: A Methodology for Translation*. John Benjamins, Amsterdam. Translated and edited by Juan C. Sager and M.-J. Hamel. Originally published as *Stylistique compar  e du fran  ais et de l’anglais* (1958).
- Wang, Y., Zeng, J., Liu, X., Wong, D. F., Meng, F., Zhou, J., and Zhang, M. (2025). DelTA: An online document-level translation agent based on multi-level memory. In

Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., Zheng, C., Liu, D., Zhou, F., Huang, F., Hu, F., Ge, H., Wei, H., Lin, H., Tang, J., Yang, J., Tu, J., Zhang, J., Yang, J., and et al., J. Y. (2025). Qwen3 technical report. Preprint arXiv:2505.09388.

A Appendix

A.1 Term Detection Coverage

	ParaNLP	IWSLT2023
Exact-form match (%)	61.9%	62.7%
Lemma match (%)	82.2%	86.6%
Avg. terms / doc	273.8	120.3

Table 9: CONCORDANCER recognition coverage, averaged per document. Exact-form and lemma matching are baselines that a naive identification method could use; CONCORDANCER recovers the remaining terms by detecting further variants.

	ParaNLP		IWSLT2023	
	ING	OOG	ING	OOG
Avg. terms / doc	232.8	40.9	109.1	11.2
NO-VAR	84.3%	–	82.5%	–
GRAPH	7.0%	11.3%	9.2%	3.9%
LEX	3.5%	35.7%	3.2%	31.3%
EXP	0.5%	5.0%	0.3%	1.7%
RED	4.6%	36.3%	4.8%	45.5%
MORPH	0.0%	4.1%	0.0%	2.7%
COMBO	0.1%	7.6%	0.0%	14.9%

Table 10: Distribution of variation types among CONCORDANCER-detected terms, averaged per document and split by whether the term’s lemma is listed in the glossary (*ING*, in-glossary) or not (*OOG*, out-of-glossary). Percentages are within each group.

Table 9 reports the recognition coverage of CONCORDANCER on both corpora. Exact-form and lemma matching (the two naive identification baselines) would respectively retrieve 61.9% and 82.2% of terms for ParaNLP, and 62.7% and 86.6% for

IWSLT2023. CONCORDANCER, which detects additional variant types, identifies the remaining terms.

Table 10 breaks down CONCORDANCER-detected terms into *in-glossary* terms (*ING*), whose lemma appears in the glossary, and *out-of-glossary* terms (*OOG*), discovered by CONCORDANCER beyond its scope. For *ING* terms, more than 80% are preferred terms (NO-VAR). For *OOG* terms, the distribution shifts: LEX and RED together account for more than 70%, meaning CONCORDANCER primarily discovers out-of-glossary terms through lexical and reduction variants. Since no glossary reference exists for *OOG* terms, their translation cannot be assessed by term-level accuracy (Dimension 1), motivating the complementary evaluation dimensions.

A.2 Distribution of Distinct Terms per Concept

Figure 1 reports the distribution of the number of distinct canonical surface forms (preferred term and variants) per concept, averaged at the document level. The distribution is strongly skewed: most concepts are realized by a single form in both corpora, with a long tail of concepts showing two or more variants.




Figure 1: Distribution of the number of distinct terms per concept, averaged by document, showing the proportion of concepts with exactly k distinct terms.

A.3 Variation Transfer Matrices

Tables 11 and 12 present the full source-to-target variation transfer matrices under the baseline condition. Each cell reports the percentage of source occurrences of a given variation type (rows) that are translated as a given target variation type (columns).

Diagonal values indicate type preservation. Key observations beyond those discussed in Section 6.1:

- **NO-VAR** → **NO-VAR**: 81–85% on ParaNLP but only 71–75% on IWSLT2023. The off-diagonal NO-VAR → LEX is the main source of preferred-term variation: 10–12% on ParaNLP and 18–25% on IWSLT2023, with the reference showing the highest rate on both corpora.
- **GRAPH** → **GRAPH**: 50–60% on ParaNLP but only 14–29% on IWSLT2023, where 63–79% of source GRAPH are realized as RED instead. This likely reflects the translation guidelines of IWSLT2023, where source acronyms are systematically spelled out rather than kept as an abbreviation (the human reference does so even more than the MT systems), rather than the spoken origin of the source.
- **RED** → **RED**: 83–86% on both corpora, the most stable transfer across all types and systems.
- **LEX** → **LEX**: 58–63% on ParaNLP and 52–77% on IWSLT2023, with the reference achieving notably higher preservation than MT systems on IWSLT2023.

Tables 13 and 14 present the same matrices under the base+terms condition. The glossary constraint selectively affects preferred-term transfer: NO-VAR → NO-VAR increases from 81–84% to 90–93% on ParaNLP and from 71–75% to 79–91% on IWSLT2023, while NO-VAR → LEX drops from 10–12% to 3–5% on ParaNLP and from 18–23% to 5–16% on IWSLT2023. Other variation types are largely unaffected: RED → RED remains at 84–86% and GRAPH → GRAPH at 49–54% on ParaNLP, indicating that the glossary constraint acts primarily on preferred terms and does not propagate to variant categories.

A.4 Legitimate Variation Example: *reported speech*

Table 15 shows translations of the sentence below by all systems. Within a single sentence, the reference alternates between *parole rapportée* and *discours rapporté* as renderings of *reported speech*, a stylistic choice that a consistency metric would penalize as inconsistency, although both forms are standard French

equivalents of the same linguistic phenomenon. All four baseline MT systems instead produce a single form repeated.

A.5 Legitimate Variation Example: *pitch*

Table 17 shows three of the 14 occurrences of *pitch* (glossary: *tonie*) in a single ParaNLP document, translated as either *hauteur* (for pitch values) or *mélodique* (in compound expressions such as *contour mélodique* or *patron mélodique*). All renderings are contextually appropriate, yet a consistency metric severely penalizes the alternation.

A.6 Variation Type Variability Example: *reinforcement learning*

Table 16 shows three occurrences of the acronym *RL* (a GRAPH variant of *reinforcement learning*) in a single ParaNLP document. The acronym is preserved in one context but expanded into different forms in others, illustrating the within-document variability of GRAPH transfer.

Source	Reference
... pitch values were ex- tracted and averaged...	...les valeurs de hauteur ont été extraites...
...more pitch movement on penultimate syllables...	...plus de mouvement mélodique sur la pénul- tième...
...a specific pitch pattern on the penultimate syllable.	...un schéma mélodique spécifique sur la syllabe pénultième.

Table 17: The term *pitch* (glossary: *tonie*) translated as 2 distinct lexical variants (*hauteur*, *mélodique*) across 14 occurrences by the Reference. All translations are contextually appropriate, yet a consistency metric results in a low score for this term.

A.7 Unevaluated Terms

Tables 18 and 19 characterize the terms that remain unevaluated by any of the three dimensions. For ParaNLP, 3.2% of occurrences (12.1% of distinct terms) are never evaluated; for IWSLT2023, 2.5% (8.1%). These residual terms are hapax occurrences absent from the glossary and without a CTV-eligible relationship to a preferred term, representing a small but structurally irreducible gap.

System	Src	NO-VAR	LEX	RED	GRAPH	EXP	MORPH	COMBO	∅
Ref.	NO-VAR	85.2	9.9	1.5	0.9	0.1	2.4	0.0	0.1
	LEX	7.1	63.0	18.7	4.1	1.8	4.1	0.4	0.9
	RED	0.1	7.4	86.2	2.4	0.6	1.4	0.2	1.7
	GRAPH	13.2	13.6	9.6	49.9	5.4	3.9	3.1	1.1
	EXP	9.0	14.4	9.6	13.9	26.1	16.4	7.9	2.8
	MORPH	12.3	24.1	17.6	0	10.2	35.8	0	0
	COMBO	7.2	34.2	14.0	13.4	5.7	22.7	1.9	0.9
Euro22	NO-VAR	83.0	10.9	1.6	0.7	0.1	3.6	0	0.0
	LEX	6.8	59.5	22.9	4.8	0.9	3.9	0	1.1
	RED	0	9.7	83.4	2.2	0.2	2.7	0.1	1.6
	GRAPH	9.5	9.9	8.4	53.9	10.4	2.4	4.8	0.7
	EXP	4.6	19.3	1.5	24.8	31.4	7.2	8.4	2.8
	MORPH	13.3	14.0	17.6	0	8.9	44.5	1.1	0.6
	COMBO	4.4	26.4	14.1	11.7	8.5	31.0	2.9	0.9
Euro9	NO-VAR	83.5	10.4	1.6	0.9	0.1	3.4	0	0.1
	LEX	6.2	61.8	21.5	4.8	1.2	3.2	0	1.3
	RED	0.2	10.2	82.8	2.5	0.3	2.4	0	1.6
	GRAPH	6.6	7.9	7.6	58.9	7.4	4.3	6.3	0.9
	EXP	8.5	15.5	9.9	23.4	24.6	13.9	1.4	2.8
	MORPH	13.8	19.9	15.3	0	7.2	42.6	0.6	0.6
	COMBO	4.0	25.7	16.3	8.7	7.3	34.7	2.5	0.9
Qwen	NO-VAR	81.4	12.5	1.4	1.7	0.1	2.8	0.0	0.1
	LEX	6.3	58.8	24.9	4.0	1.6	3.1	0.1	1.3
	RED	0.1	10.2	83.0	2.3	0.4	2.2	0.1	1.8
	GRAPH	5.9	13.7	7.5	57.5	7.1	0.5	7.2	0.7
	EXP	2.1	23.8	2.3	24.2	18.8	17.9	8.2	2.8
	MORPH	14.6	20.5	16.7	0.6	5.7	40.9	0.6	0.6
	COMBO	5.9	30.8	11.6	7.3	6.8	34.4	2.2	0.9
Llama	NO-VAR	81.6	11.1	1.2	2.5	0.1	3.4	0	0.0
	LEX	6.0	59.6	22.2	5.0	1.7	4.0	0.2	1.3
	RED	0.1	8.7	85.5	2.3	0.2	1.5	0	1.7
	GRAPH	7.7	7.4	10.5	59.4	5.0	1.9	7.3	0.9
	EXP	9.2	18.9	2.7	23.4	32.9	7.9	2.1	2.8
	MORPH	10.0	27.3	12.7	0	11.4	38.1	0	0.6
	COMBO	4.0	33.4	15.1	8.1	7.3	29.1	2.2	0.9

Table 11: Variation transfer from source to target (% , mean per document) per system, baseline condition, ParaNLP. Rows: source variation category per system; columns: target variation category. Diagonal (preservation) in bold. Source distribution (ref., mean/doc): NO-VAR: 72.3% (198), LEX: 7.9% (22), RED: 9.0% (25), GRAPH: 7.7% (21), EXP: 1.1% (4), MORPH: 0.6% (2), COMBO: 1.3% (5).

System	Src	NO-VAR	LEX	RED	GRAPH	EXP	MORPH	COMBO	∅
Ref	NO-VAR	72.6	21.1	1.4	0.8	0	4.0	0	0.1
	LEX	3.5	77.2	8.1	0.6	0	10.6	0	0
	RED	0	11.5	86.4	0	0	1.0	1.1	0
	GRAPH	0	0	78.6	14.2	2.0	5.3	0	0
	EXP	0	41.7	0	0	58.3	0	0	0
	MORPH	0	0	66.7	0	0	33.3	0	0
	COMBO	0	8.3	20.8	8.3	8.3	45.8	8.3	0
Euro22	NO-VAR	72.2	19.8	1.0	0.6	0	6.2	0	0.1
	LEX	11.5	54.7	13.7	9.5	0	10.6	0	0
	RED	0	15.7	81.2	0	0	2.0	1.1	0
	GRAPH	0	0	66.7	26.6	1.0	5.3	0.4	0
	EXP	0	25.0	0	16.7	58.3	0	0	0
	MORPH	0	0	100.0	0	0	0	0	0
	COMBO	0	18.1	37.5	8.3	8.3	19.4	8.3	0
Euro9	NO-VAR	73.1	19.2	1.2	1.0	0	5.6	0	0
	LEX	14.9	58.6	8.1	7.3	0	11.1	0	0
	RED	0	10.4	86.5	0	1.1	2.0	0	0
	GRAPH	0.4	0	66.7	26.6	1.0	5.3	0	0
	EXP	0	0	41.7	0	0	8.3	50.0	0
	MORPH	0	0	66.7	33.3	0	0	0	0
	COMBO	0	2.8	36.1	16.7	8.3	27.8	8.3	0
Qwen	NO-VAR	74.8	17.3	0.6	0.9	0	4.7	0.2	1.3
	LEX	14.9	52.2	12.6	9.8	0	10.6	0	0
	RED	0	13.4	84.8	0	0	0	0	1.7
	GRAPH	1.7	1.8	65.0	25.2	1.0	5.0	0	0.3
	EXP	0	41.7	0	0	0	0	58.3	0
	MORPH	0	0	66.7	33.3	0	0	0	0
	COMBO	0	18.1	30.6	8.3	8.3	18.1	16.7	0
Llama	NO-VAR	70.8	20.2	1.4	1.4	0	6.0	0.1	0.1
	LEX	11.0	56.7	11.4	9.7	0	11.2	0	0
	RED	0	12.9	85.1	0	0	2.0	0	0
	GRAPH	0.4	1.4	62.9	29.0	1.0	4.3	1.0	0
	EXP	0	75.0	0	16.7	0	8.3	0	0
	MORPH	0	0	66.7	0	0	0	33.3	0
	COMBO	0	18.1	30.6	8.3	8.3	18.1	16.7	0

Table 12: Variation transfer from source to target (% , mean per document) per system, baseline condition, IWSLT2023. Rows: source variation category per system; columns: target variation category. Diagonal (preservation) in bold. Source distribution (ref., mean/doc): NO-VAR: 74.5% (88), LEX: 5.7% (7), RED: 8.9% (10), GRAPH: 8.7% (11), EXP: 0.4% (4), MORPH: 0.3% (1), COMBO: 1.5% (4).

System	Src	NO-VAR	LEX	RED	GRAPH	EXP	MORPH	COMBO	∅
Ref.	NO-VAR	85.2	9.9	1.5	0.9	0.1	2.4	0.0	0.1
	LEX	7.1	63.0	18.7	4.1	1.8	4.1	0.4	0.9
	RED	0.1	7.4	86.2	2.4	0.6	1.4	0.2	1.7
	GRAPH	13.2	13.6	9.6	49.9	5.4	3.9	3.1	1.1
	EXP	9.0	14.4	9.6	13.9	26.1	16.4	7.9	2.8
	MORPH	12.3	24.1	17.6	0	10.2	35.8	0	0
	COMBO	7.2	34.2	14.0	13.4	5.7	22.7	1.9	0.9
Euro22	NO-VAR	93.0	3.3	1.6	0.5	0.1	1.3	0.0	0.0
	LEX	8.4	57.4	23.3	4.8	0.7	4.2	0	1.1
	RED	0.2	8.7	85.0	2.3	0.2	2.0	0.1	1.6
	GRAPH	15.5	6.7	11.0	50.0	9.2	4.1	2.6	0.9
	EXP	9.7	16.2	2.7	23.2	34.0	9.3	2.1	2.8
	MORPH	18.9	8.3	15.3	0	7.8	47.9	1.1	0.6
	COMBO	9.1	30.3	9.7	8.7	6.4	31.2	3.8	0.9
Euro9	NO-VAR	90.2	5.3	1.4	0.7	0.1	2.2	0	0.1
	LEX	7.7	60.0	21.8	3.6	2.0	3.1	0.1	1.7
	RED	0.1	9.0	84.1	2.4	0.4	2.3	0.1	1.6
	GRAPH	14.8	5.7	9.9	50.9	9.8	3.2	4.4	1.3
	EXP	11.3	13.3	11.7	26.6	27.1	6.9	0	3.0
	MORPH	19.5	14.6	13.8	0.6	7.2	41.7	0.6	2.1
	COMBO	7.9	25.8	19.5	7.8	7.3	28.3	2.5	1.0
Qwen	NO-VAR	93.5	3.2	1.5	0.5	0.1	1.1	0	0.1
	LEX	10.5	53.5	25.2	4.6	1.5	3.3	0	1.3
	RED	0.7	7.3	85.1	2.2	0.3	2.7	0	1.7
	GRAPH	18.2	8.5	9.8	49.1	9.3	2.6	1.8	0.7
	EXP	6.0	16.5	2.1	17.0	44.4	10.7	0.6	2.8
	MORPH	24.1	10.8	14.2	0	12.9	36.9	0.6	0.6
	COMBO	12.9	26.2	11.7	12.8	7.2	25.9	2.5	0.9
Llama	NO-VAR	93.2	3.4	1.4	0.7	0.1	1.1	0	0.1
	LEX	8.8	54.7	23.4	5.8	1.7	4.2	0	1.4
	RED	0.6	7.3	86.0	2.4	0.3	1.7	0	1.7
	GRAPH	15.4	6.1	12.2	54.4	6.4	1.4	3.4	0.7
	EXP	12.2	10.9	2.7	31.9	36.5	3.0	0	2.8
	MORPH	25.9	19.7	10.4	0	9.5	33.9	0	0.6
	COMBO	9.1	32.1	10.8	8.7	8.2	26.9	3.3	0.9

Table 13: Variation transfer from source to target (% , mean per document) per system, base+terms condition, ParaNLP. Rows: source variation category per system; columns: target variation category. Diagonal (preservation) in bold. Source distribution (ref., mean/doc): NO-VAR: 72.3% (198), LEX: 7.9% (22), RED: 9.0% (25), GRAPH: 7.7% (21), EXP: 1.1% (4), MORPH: 0.6% (2), COMBO: 1.3% (5).

System	Src	NO-VAR	LEX	RED	GRAPH	EXP	MORPH	COMBO	$\emptyset$
Ref.	NO-VAR	72.6	21.1	1.4	0.8	0	4.0	0	0.1
	LEX	3.5	77.2	8.1	0.6	0	10.6	0	0
	RED	0	11.5	86.4	0	0	1.0	1.1	0
	GRAPH	0	0	78.6	14.2	2.0	5.3	0	0
	EXP	0	41.7	0	0	58.3	0	0	0
	MORPH	0	0	66.7	0	0	33.3	0	0
	COMBO	0	8.3	20.8	8.3	8.3	45.8	8.3	0
Euro22	NO-VAR	86.7	7.5	1.4	0.7	0.1	3.5	0.1	0
	LEX	21.6	40.3	12.6	8.5	0	16.9	0	0
	RED	1.2	10.4	84.1	0	2.2	2.0	0	0
	GRAPH	12.1	0	65.3	21.2	1.4	0	0	0
	EXP	0	91.7	0	0	8.3	0	0	0
	MORPH	0	0	100.0	0	0	0	0	0
	COMBO	8.3	19.4	36.1	8.3	8.3	11.1	8.3	0
Euro9	NO-VAR	80.4	12.8	1.1	0.4	0	5.0	0.1	0.1
	LEX	16.6	52.3	11.4	7.8	0	11.8	0	0
	RED	0	10.3	86.6	0	1.1	2.0	0	0
	GRAPH	8.5	0	63.9	26.6	1.0	0	0	0
	EXP	0	0	41.7	8.3	0	0	50.0	0
	MORPH	0	33.3	33.3	33.3	0	0	0	0
	COMBO	8.3	20.8	36.1	8.3	8.3	9.7	8.3	0
Qwen	NO-VAR	89.8	5.2	1.1	0.7	0	2.7	0.1	0.4
	LEX	22.1	43.1	9.3	6.0	0	13.1	0	6.4
	RED	3.0	9.2	85.6	0	1.1	1.1	0	0
	GRAPH	8.4	0	65.4	23.2	1.0	2.0	0	0
	EXP	8.3	25.0	0	8.3	58.3	0	0	0
	MORPH	0	0	33.3	33.3	0	0	0	33.3
	COMBO	8.3	2.8	36.1	8.3	8.3	11.1	25.0	0
Llama	NO-VAR	88.9	6.6	1.5	0.7	0.1	2.2	0	0
	LEX	26.2	40.7	10.5	4.8	0	17.8	0	0
	RED	1.2	9.8	86.6	0	1.1	1.2	0	0
	GRAPH	13.7	0	58.7	26.6	1.0	0	0	0
	EXP	0	25.0	0	16.7	58.3	0	0	0
	MORPH	0	0	66.7	0	0	0	33.3	0
	COMBO	8.3	0	38.9	0	8.3	36.1	8.3	0

Table 14: Variation transfer from source to target (% , mean per document) per system, base+terms condition, IWSLT2023. Rows: source variation category per system; columns: target variation category. Diagonal (preservation) in bold. Source distribution (ref., mean/doc): NO-VAR: 74.5% (88), LEX: 5.7% (7), RED: 8.9% (10), GRAPH: 8.7% (11), EXP: 0.4% (4), MORPH: 0.3% (1), COMBO: 1.5% (4).

A.8 TCR Metric Divergence Example

Table 20 illustrates how the three TCR metrics diverge on the same data. For 5 occurrences of *neural network* (glossary: *réseau de neurones*), TCR-EM scores 0.40, TCR-PR 0.60, and TCR-T 0.50.

	Translation	EM	PR	Temporal
o_1	réseau neuronal	×	=	
o_2	réseau neuronal	×	=	=
o_3	réseau de neurones	✓	≠	≠
o_4	réseau de neurones	✓	≠	=
o_5	réseau neuronal	×	=	≠
Score		2/5	3/5	2/4

Table 20: Translations of *neural network* (glossary: *réseau de neurones*). **EM**: glossary match; **PR**: agrees with pseudo-ref (o_1); **Temporal**: agrees with previous occ.

A.9 Scores Breakdown by Variation Label

A.9.1 Baseline Condition

Term-level accuracy (EM). Table 21 reports EM Micro broken down by source variation category. EM ranges from 70.0–74.1 on IWSLT2023 and 77.8–81.5 on ParaNLP, with Qwen leading on IWSLT2023 (74.05) and Ref on ParaNLP (81.52). NO-VAR terms achieve 72–85% EM accuracy across systems. Performance drops sharply for variants: LEX reaches only 34–55%, GRAPH 42–74%, and MORPH and COMBO score 0% across all systems on both corpora. Since NO-VAR accounts for 72–74% of occurrences, the aggregate EM is structurally dominated by preferred-term performance.

Translation consistency (TCR). Tables 22, 23, and 24 report the three TCR variants broken down by variation category. Under the least restrictive variant (TCR-PR), the reference dominates on IWSLT2023 (95.09 vs. 84.4–88.1 for MT systems) but differences narrow on ParaNLP (86.8–89.3). The per-variation pattern is consistent across metrics: NO-VAR terms score substantially higher (88–95%) than LEX (47–55% on IWSLT2023, 78–88% on ParaNLP) or GRAPH (53–89% depending on corpus and system). COMBO scores are low but non-zero (6–24%), indicating that the few compound variants with $k \geq 2$ occurrences are rarely translated consistently.

A notable observation concerns system ranking

stability. On ParaNLP under baseline, each TCR variant ranks a different system as top-1: TCR-EM favors the reference (82.99), TCR-PR favors Llama (89.31), and TCR-T favors E9 (87.93). No single system dominates consistency across all operationalizations, indicating that the choice of consistency metric is not neutral and can substantially affect system ranking.

Cross-term variation (CTV). Table 25 reports variation preservation by source variation category. On IWSLT2023, the reference achieves the highest CTV (86.67), outperforming all MT systems (68–86%). On ParaNLP, the pattern reverses: the reference scores lowest (66.06) while Llama leads (74.63). This asymmetry is explained by the target-side amplification documented in Section 6.1.

Per variation type, RED preservation is consistently high across all systems (82–100% on IWSLT2023, 81–87% on ParaNLP), confirming that reductions transfer reliably. GRAPH coherence varies more: the reference achieves 90% on IWSLT2023 but only 58% on ParaNLP, while MT systems range from 50–86% on IWSLT2023 and 42–78% on ParaNLP. LEX coherence shows the widest system-dependent spread (50–86% on IWSLT2023, 57–72% on ParaNLP), and EXP is both rare and variable.

A.9.2 Glossary-Constrained Condition (base+terms)

Term-level accuracy (EM). Under base+terms (Table 26), EM Micro increases substantially for all MT systems: +10–17 points on IWSLT2023 (from 70–74% to 78–90%) and +6–10 points on ParaNLP (from 78–80% to 85–89%), confirming that baseline systems already possess the correct terms but do not always produce them without guidance. The improvement is concentrated on NO-VAR terms (+14–18 points on IWSLT2023, +8–12 points on ParaNLP), with smaller gains on RED and GRAPH. MORPH and COMBO remain at 0%.

Translation consistency (TCR). TCR metrics also improve under base+terms (Tables 27, 28, 29). On NO-VAR terms, TCR-PR rises from 86–88% to 91–95% on IWSLT2023 and from 88–90% to 93–95% on ParaNLP, indicating that forcing preferred-term usage reduces translation inconsistency. The improvement is smaller for variant categories, which are not directly targeted by the glossary constraint.

Cross-term variation (CTV). Under base+terms (Table 30), CTV drops for most MT systems compared to baseline, particularly on ParaNLP (from 72–75% to 67–69%). On IWSLT2023, GRAPH coherence improves slightly for some systems (Qwen: 50% → 90%), but RED coherence decreases (95–100% → 83–96%). This confirms the tension identified in the main analysis: the glossary constraint improves accuracy and consistency at the cost of suppressing the target-side variation that CTV measures. The reference, unchanged across conditions, serves as a stable baseline.

A.10 Top-5 Most Frequent Concepts

Table 31 provides a detailed breakdown of the five most frequent concepts per document. The five most frequent concepts account for a cumulative 47.2% of all occurrences in ParaNLP and 52.5% in IWSLT2023, with up to 2.38 distinct terms per concept. Notably, the proportion of preferred-term usage (%NO-VAR) decreases from Top-1 to Top-5 (e.g., 87.3% to 60.3% on ParaNLP), suggesting that the most frequent concept in a document tends to be expressed more canonically, while less central concepts allow for greater surface variation.

Rank	ParaNLP			IWSLT2023		
	%occ	Terms	%NO-VAR	%occ	Terms	%NO-VAR
Top-1	16.8%	1.66	87.3%	19.4%	1.30	88.7%
Top-2	10.7%	2.19	70.2%	10.7%	1.20	92.1%
Top-3	8.2%	2.03	77.7%	8.7%	1.40	65.0%
Top-4	6.4%	1.88	73.2%	7.5%	1.30	73.9%
Top-5	5.1%	2.38	60.3%	6.2%	1.30	69.0%

Table 31: Average statistics for the top-5 most frequent concepts per document. *%occ* = share of total occurrences in the document; *Terms* = average number of distinct surface forms; *%NO-VAR* = percent-age of occurrences using the preferred term.

A.11 Model Cards

We list below the HuggingFace model cards for the four open-weight LLMs used in the MT experiments.

- Llama3.1-8B-Instruct: <https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>
- Qwen3-8B: <https://huggingface.co/Qwen/Qwen3-8B>

¹⁰<https://huggingface.co/Unbabel/wmt22-cometkiwi-da>

¹¹<https://huggingface.co/google/metricx-24-hybrid-large-v2p6>

- EuroLLM-9B-Instruct: <https://huggingface.co/utter-project/EuroLLM-9B-Instruct>

- EuroLLM-22B-Instruct: <https://huggingface.co/utter-project/EuroLLM-22B-Instruct-2512>

A.12 ParaNLP Corpus Construction Details

The comparable articles in the ParaNLP corpus were transformed into parallel data through automatic sentence segmentation and alignment. Poorly aligned or unaligned sentences were handled via automatic post-editing (APE) (4,928 sentences) and machine translation (1,381 sentences), resulting in 6,272 parallel sentences in total. Three articles, including 199 MT sentences and 42 APE sentences, were further post-edited by a native French speaker with expertise in NLP.

A.13 Global Score Evaluation

For completeness, we report general MT quality scores (COMET-Kiwi (Rei et al., 2022)¹⁰ and MetricX-QE (Juraska et al., 2024)¹¹) for each system on both corpora. These complement the terminology-focused analysis in the main paper by situating each system’s overall translation quality.

System	COMET-Kiwi ↑	MetricX-QE ↓
EuroLLM-22B-greedy	0.8380	2.5987
EuroLLM-9B-greedy	0.8390	2.5797
Llama-3.1-8B-greedy	0.8286	2.8116
Qwen3-8B-top-p	0.8344	2.7076
Reference	0.8242	2.9656

Table 32: Global evaluation scores for the ParaNLP corpus, where each global score is computed as the mean of document-level scores, and each document-level score is the mean of its sentence-level scores, using COMET-Kiwi (↑) and MetricX-QE (↓).

System	COMET-Kiwi $\uparrow$	MetricX-QE $\downarrow$
EuroLLM-22B-greedy	0.8229	2.9226
EuroLLM-9B-greedy	0.8229	2.8744
Llama-3.1-8B-greedy	0.8067	3.4546
Qwen3-8B-top-p	0.8151	3.3200
Reference	0.8030	3.3728

Table 33: Global evaluation scores for the IWSLT2023 coprus, where each global score is computed as the mean of document-level scores, and each document-level score is the mean of its sentence-level scores, using COMET-Kiwi ($\uparrow$) and MetricX-QE ($\downarrow$).

Results Analysis Gobal Score. Across both corpora, all MT systems consistently outperform the human reference under both a known artefact of QE metrics, which are predominantly trained on MT outputs and thus tend to assign systematically lower scores to human-generated translations (Deutsch et al., 2025). Importantly, scores were computed segment-by-segment, averaged to document level, then aggregated per system, meaning *no contextual information was available during evaluation*. So terminological consistency is therefore not captured, which likely flatters MT systems over human translators who actively manage such consistency across documents. System rankings remain stable throughout: **EuroLLM-9B-greedy** achieves the best or tied-best MetricX-QE in every setting, notably surpassing its 22B counterpart, while **Llama-3.1-8B-greedy** consistently ranks last. Scores are uniformly higher on ParaNLP than on IWSLT2023, reflecting the greater regularity of academic prose versus spontaneous speech.

A.14 Variation Classification Typology

This appendix provides a comprehensive reference for the terminological variation typology used throughout this work. The typology distinguishes between single-type variations (GRAPH, MORPH, RED, EXP, LEX) and multiple-type variations (COMBO), which combine two or more single types. Table 34 provides a complete reference of all variation types and subtypes.

B Manual Evaluation of LLM-based Variation Labeling

The variation labels assigned to lexicalised alternative terms (Section 3) are produced by a ChatGPT-

based classifier with few-shot prompting (see Appendix 9). To assess the reliability of this automatic labeling, we conducted a manual evaluation on a sample of 50 randomly selected (preferred term, variant) pairs drawn from the full set of labeled instances. The evaluation was restricted to pairs where the preferred term and the variant differ, that is, pairs not automatically labeled by CONCORDANCER on the basis of structural patterns. This ensures that the evaluation targets exclusively the LLM classifier’s judgment, independently of the rule-based component of the pipeline. For each sampled pair, the predicted variation type and subtype (e.g., GRAPH acronym, MORPH constituent reordering) were compared against gold labels assigned by one of the authors, a native French speaker fluent in English, with expertise in computational linguistics. Tables 35 and 36 report per-class precision and recall at the type and subtype levels respectively. The 49 evaluated pairs are listed in Table 37 (one pair excluded as the variant was not a genuine terminological link).

Type	Support	Precision	Recall	Accuracy
GRAPH	14	67%	100%	—
MORPH	1	100%	100%	—
RED	4	100%	100%	—
EXP	2	100%	100%	—
LEX	19	95%	100%	—
COMBO	9	100%	11%	—
Total	49	94%	85%	84%

Table 35: Type-level evaluation of the LLM-based variation classifier on 49 manually annotated (preferred term, variant) pairs. Per-class precision and recall are reported; accuracy is a global metric (total correct / total) reported at the Total row only.

Results. The classifier achieves 84% overall type recall and 77% subtype accuracy across the 49 evaluable pairs (one pair was excluded as the variant was not a genuine terminological link). Performance is near-perfect for single-type variations: GRAPH, LEX, RED, EXP, and MORPH all reach 100% recall at the type level, with subtype accuracy above 85% for the two most frequent classes (GRAPH: 86%, LEX: 89%).

The only systematic failure concerns compound variations (COMBO): the classifier recovers only 1 out of 9 gold COMBO instances (recall 11%), con-

sistently predicting a single type instead of a combination. All 8 missed COMBO cases follow the same pattern, a multi-word term is abbreviated to an acronym that is also an ellipsis of the original (e.g., *topic detection / TDT, model / LM*), which the classifier labels as pure GRAPH (acronym) rather than the compound GRAPH+EXP that gold annotation assigns. This over-simplification depresses GRAPH precision to 67%, since seven COMBO pairs are absorbed into the GRAPH class.

The remaining subtype errors are minor: two GRAPH1 predictions should be GRAPH5 (partial abbreviation, e.g., *ASR system*), and two LEX1 predictions should be LEX2 (near-synonymy with head-word change, e.g., *similarity function, parallel text*).

Overall, the classifier is reliable for the large majority of variation types encountered in practice. The COMBO class is both the rarest in the sample and the hardest to predict, but its impact on downstream analyses is limited since compound variations are a small fraction of the full dataset.

Type	Support	Precision	Recall	Accuracy
GRAPH	14	57%	86%	—
MORPH	1	100%	100%	—
RED	4	100%	100%	—
EXP	2	100%	100%	—
LEX	19	85%	89%	—
COMBO	9	100%	11%	—
Total	49	90%	81%	76%

Table 36: Subtype-level evaluation of the LLM-based variation classifier on 49 manually annotated pairs, grouped by gold type. Precision is computed over all instances predicted as a given type; recall over gold instances of that type. Accuracy is global.

C Manual Evaluation of Alignment Accuracy

The alignment step (Section 3) links each source term occurrence to its target span using Bertalign and SimAlign, with an LLM fallback for low-confidence spans. To assess the reliability of these alignments, we conducted a manual evaluation on a sample of 100 term alignments drawn from the full set of pipeline outputs across all documents. To ensure balanced coverage, the sample was stratified across the nine system outputs, that is, the four MT systems under

both the baseline and base+terms conditions plus the human reference, each contributing eleven alignments (twelve for the reference); alignments with an empty target span were excluded. For each sampled alignment, one of the authors, a native French speaker fluent in English with expertise in computational linguistics, judged whether the aligned target span is the correct realization of the source term in the given sentence pair, and, when incorrect, recorded the error type. Table 38 reports accuracy per system and Table 39 lists the incorrect alignments.

System	Condition	#	Acc.
Llama	baseline	11	91%
Llama	base+terms	11	91%
Qwen	baseline	11	100%
Qwen	base+terms	11	100%
Euro9	baseline	11	100%
Euro9	base+terms	11	100%
Euro22	baseline	11	100%
Euro22	base+terms	11	91%
Ref	—	12	92%
Total		100	96%

Table 38: Per-system accuracy of the manual alignment evaluation on 100 alignments stratified across the nine system outputs. Each row reports the number of judged alignments and the proportion judged correct.

Results. The aligner is correct on 96% of the sampled alignments (96/100). Accuracy is high and comparable across conditions, at 98% for baseline (43/44), 95% for base+terms (42/44), and 92% for the human reference (11/12); the small per-system differences are not meaningful given that each system contributes only 11–12 judgements. Accuracy is likewise stable across variation types, with no error on the ten sampled dynamic variants.

The four errors fall into two categories. Three are lexical-choice errors, where the span captures an adjacent target word rather than the term itself: *corpus* aligned to *Récolte* (from *Corpus Harvesting / Récolte de corpus*), *indexing* to *discours* (from *speech indexing / indexation des discours*), and *future* echoed as *Future* where no distinct French realization appears in the sentence. The remaining case is a span-boundary error, where *automatic speech recog-*

nition was aligned to the truncated *reconnaissance de la parole* instead of *reconnaissance automatique de la parole*. All four affect single occurrences and none reflects a systematic failure mode, consistent with the low overall error rate.

D Prompts

This appendix collects the full prompts used by the LLM components of our pipeline: alignment disambiguation (Figure 2), baseline translation (Figure 3), glossary-guided translation (Figure 6), and variation classification (Figure 9).

E Reproducibility Statement

The variation classifier and the alignment-disambiguation step in our pipeline use OpenAI's

gpt-4.1-mini (Section 3; prompts in Appendix D); this usage is intrinsic to the contribution and is fully described in the methodology. All MT inference was run on a single NVIDIA H100 96GB GPU. Total GPU time across 4 systems and 2 conditions (baseline and base+terms) for the 42 documents from ParaNLP and IWSLT2023 is approximately 10 minutes using vLLM (Kwon et al., 2023) with pre-generated prompts. The labeling pipeline issued approximately 5,000 requests to gpt-4.1-mini (snapshot gpt-4.1-mini-2026-04-01), for a total cost of approximately 15 USD. Reliance on a closed commercial LLM introduces a versioning risk: OpenAI silently retires model snapshots, so re-running the pipeline more than a few months from publication may result in slightly different labels.

```

"""You are a professional {source_language}-{target_language} translator, teaching the students the course
on technical translation. You are checking a student's translation of a sentence that contains a
technical term. You are given an {source_language} term (it can be a word or an expression), a source
{source_language} sentence containing this term (it may be cased differently or contain additional
punctuation), and a student's {target_language} translation. You need to find how the student has
translated the term in question in {target_language}, and return only that term.

Important: do not change the translated term anyhow, copy it straight from the sentence! For example, keep
the casing and the grammar form of the translated term as is.

When completing the task, follow the examples below:

{source_language} sentence: This paradigm based on ant colony algorithms for the exploration of the graph
removes the need to dynamically expand the graph: the memory footprint becomes independant of the
language model size.
{source_language} term: language model
{target_language} translation: Ce paradigme basé sur les algorithmes de colonie de fourmis pour l'
exploration du graphe supprime le besoin d'étendre dynamiquement le graphe : l'empreinte mémoire
devient indépendante de la taille du modèle de langage.
Translated term: modèle de langage
.
.
.
{source_language} sentence: We used our set of manually-written chronologies as a training corpus to
perform machine learning experiments.
{source_language} term: training corpus
{target_language} translation: Nous avons utilisé notre ensemble de chronologies rédigées manuellement
comme corpus d'apprentissage pour effectuer des expériences d'apprentissage automatique.
Translated term: corpus d'apprentissage

{source_language} sentence: {src_segment}
{source_language} term: {src_term}
{target_language} translation: {tgt_segment}
Translated term: """

```

Figure 2: Alignment prompt: maps a source term to its translation in the target sentence.

```

system prompt:
"""
You are a good translator! Translate the following text from English into French. Do not include any
extraneous note, commentary, explanations, or annotations. You must reply only with the translated
text in French.
"""
user prompt:
"""
English: $src\nFrench:
"""

```

Figure 3: Baseline translation prompt for Llama3.1-8B-Instruct.

```

system prompt:
'''
Translate the following English source text to French:
'''
user prompt:
'''
English: $src\nFrench:
'''

```

Figure 4: Baseline translation prompt for EuroLLM.

```

'''
You are a good translator! Translate the following text from English into French.\nEnglish: ${src}\nFrench:
'''

```

Figure 5: Baseline translation prompt for Qwen3-8B.

```

<|begin_of_text|><|start_header_id|>system<|end_header_id|>

Cutting Knowledge Date: December 2023
Today Date: 26 Jul 2024

You are a good translator! Consider the provided terms in English and their French translations:
Terms: present = présent - borrowing = emprunt - corpus = corpus - annotated corpus = corpus annot  
Translate the following text from English into French. Do not include any extraneous note, commentary,
explanations, or annotations. You must reply only with the translated text in French.<|eot_id|><|
start_header_id|>user<|end_header_id|>

English: Hi, this is Elena and I'm going to be presenting our work, Detecting Unassimilated Borrowings in
Spanish: An Annotated Corpus and Approaches to Modeling.
French:<|eot_id|><|start_header_id|>assistant<|end_header_id|>

```

Figure 6: Term-injection translation prompt for Llama3.1-8B-Instruct.

```

<|im_start|>system
Translate the following English source text to French, considering the provided terms in English and their
French translations:
Terms: present = pr  sent - borrowing = emprunt - corpus = corpus - annotated corpus = corpus annot  <|im_end
|>
<|im_start|>user
English: Hi, this is Elena and I'm going to be presenting our work, Detecting Unassimilated Borrowings in
Spanish: An Annotated Corpus and Approaches to Modeling.
French:<|im_end|>
<|im_start|>assistant

```

Figure 7: Term-injection translation prompt for EuroLLM.

```

<|im_start|>user
You are a good translator! Consider the provided terms in English and their French translations:
Terms: present = présent - borrowing = emprunt - corpus = corpus - annotated corpus = corpus annoté
Translate the following text from English into French.
English: Hi, this is Elena and I'm going to be presenting our work, Detecting Unassimilated Borrowings in
        Spanish: An Annotated Corpus and Approaches to Modeling.
French: <|im_end|>
<|im_start|>assistant
<think>

</think>

```

Figure 8: Term-injection translation prompt for Qwen3-8B.

Source	We dealt with time-stamped events governed by a reported speech verb or otherwise appearing in reported speech .
Reference	Nous avons traité des événements datés régis par un verbe de parole rapportée ou apparaissant dans un discours rapporté .
Llama	Nous avons traité des événements datés régis par un verbe de discours rapporté ou apparaissant autrement dans le discours rapporté .
Euro22	Nous avons traité des événements marqués par un verbe de discours rapporté ou apparaissant autrement dans un discours rapporté .
Euro9	Nous avons traité des événements marqués par un horodatage régis par un verbe de discours rapporté ou apparaissant autrement dans un discours rapporté .
Qwen	Nous avons traité des événements datés de temps, régis par un verbe de discours direct ou autrement apparaissant dans un discours rapporté .

Table 15: A sentence from the ParaNLP corpus where the reference alternates between *parole rapportée* and *discours rapporté* for the same source term *reported speech*, both standard French renderings of the same concept. The four MT systems produce a single repeated form (Qwen substitutes *discours direct*, a different concept).

Source	Reference
... via Reinforcement Learning (RL).	... via l'apprentissage par renforcement (RL). <i>GRAPH</i> → <i>GRAPH</i> ✓
... a data-driven method such as RL .	... des méthodes d'apprentissage statistique telles que l'apprentissage par renforcement. <i>GRAPH</i> → <i>NO-VAR</i> ✗
An RL agent compares different ...	Un agent apprenant par renforcement compare différentes ... <i>GRAPH</i> → <i>COMBO</i> ✗

Table 16: Cross-term variation for *reinforcement learning* → *apprentissage par renforcement*: *RL* is preserved or expanded, illustrating type-dependent variability.

Table 18: Average proportion of terms covered by no evaluation method per document (*occ* = occurrences; *dist* = distinct terms; never evaluated $\equiv$ not in glossary $\wedge$ hapax ($k=1$) $\wedge$ not ITC-eligible)

Category	Avg occ/doc		Avg dist/doc	
	Avg <i>N</i>	%	Avg <i>N</i>	%
All terms	273.9	—	68.7	—
Eval. by ≥ 1 method	265.5	96.8%	60.3	87.9%
Never evaluated	8.4	3.2%	8.4	12.1%
Hapax ($k=1$, excl. by TLC)	34.2	12.5%	34.2	49.8%
Not in glossary (excl. by EM)	29.5	10.8%	6.7	9.8%

Averaged over 31 documents. Never evaluated $\equiv$ not in glossary $\wedge$ hapax $\wedge$ not ITC-eligible.

Category	Avg occ/doc		Avg dist/doc	
	Avg <i>N</i>	%	Avg <i>N</i>	%
All terms	120.3	—	35.2	—
Eval. by ≥ 1 method	117.3	97.5%	32.2	91.9%
Never evaluated	3.0	2.5%	3.0	8.1%
Hapax ($k=1$, excl. by TLC)	16.5	13.7%	16.5	46.9%
Not in glossary (excl. by EM)	7.4	6.2%	2.2	6.2%

Averaged over 10 documents. Never evaluated $\equiv$ not in glossary $\wedge$ hapax $\wedge$ not ITC-eligible.

Table 19: Average proportion of terms covered by no evaluation method per document (*occ* = occurrences; *dist* = distinct terms; never evaluated $\equiv$ not in glossary $\wedge$ hapax ($k=1$) $\wedge$ not ITC-eligible)

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	197.7	72.2%	85.31	83.15	83.77	81.54	81.80	88.4	73.5%	72.63	72.47	73.07	74.87	70.85
RED	24.6	9.0%	87.09	85.13	82.00	86.15	88.11	9.9	8.2%	67.92	65.42	65.42	68.75	66.67
MORPH	1.7	0.6%	0.00	0.00	0.00	0.00	0.00	0.3	0.2%	0.00	0.00	0.00	0.00	0.00
EXP	3.2	1.2%	6.45	6.45	23.12	2.15	23.12	0.7	0.6%	0.00	0.00	0.00	0.00	0.00
GRAPH	20.6	7.5%	45.76	43.11	42.51	45.08	42.00	11.4	9.5%	73.78	58.15	59.26	65.87	56.35
LEX	22.0	8.0%	40.35	47.79	49.63	38.59	45.24	7.4	6.2%	34.66	45.98	55.24	45.77	44.25
COMBO	4.0	1.5%	0.00	0.00	0.00	0.00	0.00	2.2	1.8%	0.00	0.00	0.00	0.00	0.00

Table 21: EM accuracy (%) by source variation category under the baseline condition. *N*: avg. occurrences per document; %: share of total occurrences. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	64.9%	88.96	90.82	90.67	88.20	90.11	13.2	70.6%	95.41	85.58	88.26	88.35	87.71
RED	3.8	11.0%	89.95	89.62	84.55	88.54	91.99	2.4	12.8%	94.24	84.86	82.94	82.38	87.44
MORPH	0.0	0.1%	1.61	1.61	1.61	1.61	1.61	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.5	1.3%	26.87	22.65	25.84	18.62	24.48	0.1	0.5%	8.00	6.00	10.00	10.00	6.00
GRAPH	3.5	10.0%	55.89	65.18	67.74	65.19	67.68	1.3	7.0%	88.86	65.36	73.71	80.91	58.38
LEX	3.9	11.3%	87.66	78.66	85.06	81.45	84.19	1.3	7.0%	52.95	52.50	49.92	53.41	49.47
COMBO	0.5	1.3%	19.97	20.89	21.89	21.12	24.19	0.4	2.1%	19.17	12.50	12.50	14.17	14.17

Table 22: TCR pseudo-ref (%) by source variation category under the baseline condition. *N*: avg. distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	64.9%	87.98	90.79	89.69	87.50	88.71	13.2	70.6%	94.81	86.60	86.63	85.78	84.85
RED	3.8	11.0%	84.97	86.31	82.64	84.23	89.23	2.4	12.8%	89.43	72.43	72.29	70.84	80.71
MORPH	0.0	0.1%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.5	1.3%	24.20	18.00	20.83	13.86	24.13	0.1	0.5%	5.00	0.00	10.00	10.00	5.00
GRAPH	3.5	10.0%	54.49	60.30	63.54	64.16	67.23	1.3	7.0%	88.24	55.36	61.23	73.32	53.82
LEX	3.9	11.3%	84.00	74.87	81.65	78.19	77.61	1.3	7.0%	51.73	47.07	46.61	49.64	44.53
COMBO	0.5	1.3%	13.30	14.29	17.19	15.44	19.88	0.4	2.1%	17.78	6.67	6.67	8.89	8.89

Table 23: TCR temporal (%) by source variation category under the baseline condition. *N*: avg. distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	80.7%	86.57	84.57	85.17	82.95	83.33	13.2	80.0%	74.01	73.92	74.05	76.70	72.28
RED	1.4	5.1%	87.82	86.36	83.19	88.58	89.21	1.2	7.3%	72.08	69.58	69.58	75.42	70.83
MORPH	0.0	0.1%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.2	0.7%	1.61	1.61	7.53	0.54	7.53	0.1	0.6%	0.00	0.00	0.00	0.00	0.00
GRAPH	2.3	8.1%	44.78	41.81	41.06	44.45	40.71	1.3	7.9%	76.61	60.36	61.46	68.66	58.49
LEX	1.4	5.0%	37.90	46.82	47.36	38.14	42.47	0.7	4.2%	13.89	25.35	25.03	25.00	23.76
COMBO	0.1	0.2%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00

Table 24: TCR_{EM_{macro}} (%) by source variation category under the baseline condition. *N*: avg. in-glossary distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
GRAPH	9.1	34.2%	57.95	69.92	72.06	67.73	77.83	1.0	33.3%	90.00	85.71	85.71	50.00	85.71
RED	8.4	31.5%	87.36	84.85	81.03	84.65	83.95	1.1	36.7%	81.82	95.65	100.00	100.00	100.00
EXP	1.3	4.8%	40.00	32.43	29.73	23.68	28.21	0.2	6.7%	100.00	100.00	50.00	0.00	0.00
LEX	7.9	29.5%	56.97	70.35	72.12	70.76	69.71	0.7	23.3%	85.71	60.00	60.00	50.00	60.00

Table 25: CTV variation preservation (%) by source variation category under the baseline condition. *N*: avg. eligible occurrences per document (variants whose head term is translated as NO-VAR); %: share of eligible total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	N	%	Ref _S	Euro22	Euro9	Qwen	Llama	N	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	197.7	72.2%	85.31	93.23	90.38	93.96	93.63	88.4	73.5%	72.63	87.12	80.69	90.22	89.31
RED	24.6	9.0%	87.09	88.75	86.14	89.39	89.56	9.9	8.2%	67.92	70.00	65.42	72.50	72.50
MORPH	1.7	0.6%	0.00	0.00	0.00	0.00	0.00	0.3	0.2%	0.00	0.00	0.00	0.00	0.00
EXP	3.2	1.2%	6.45	12.90	12.90	24.73	22.58	0.7	0.6%	0.00	0.00	0.00	0.00	0.00
GRAPH	20.6	7.5%	45.76	44.40	40.98	43.41	48.61	11.4	9.5%	73.78	63.41	58.57	78.41	76.27
LEX	22.0	8.0%	40.35	47.31	48.12	37.01	44.01	7.4	6.2%	34.66	46.54	48.08	47.71	42.59
COMBO	4.0	1.5%	0.00	0.00	0.00	0.00	0.00	2.2	1.8%	0.00	0.00	0.00	0.00	0.00

Table 26: EM accuracy (%) by source variation category under the base+terms condition. N : avg. occurrences per document; %: share of total occurrences. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	N	%	Ref _S	Euro22	Euro9	Qwen	Llama	N	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	64.9%	88.96	95.28	93.22	94.98	94.01	13.2	70.6%	95.41	90.70	91.49	94.98	93.47
RED	3.8	11.0%	89.95	91.14	87.83	91.61	92.15	2.4	12.8%	94.24	82.44	85.44	80.89	86.06
MORPH	0.0	0.1%	1.61	1.61	1.61	1.61	1.61	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.5	1.3%	26.87	22.61	25.69	23.46	24.01	0.1	0.5%	8.00	10.00	10.00	6.00	6.00
GRAPH	3.5	10.0%	55.89	58.55	60.92	59.85	61.30	1.3	7.0%	88.86	74.46	73.60	81.46	79.89
LEX	3.9	11.3%	87.66	85.32	88.17	79.56	80.29	1.3	7.0%	52.95	54.32	52.65	49.32	48.11
COMBO	0.5	1.3%	19.97	21.43	23.35	22.81	23.66	0.4	2.1%	19.17	12.50	12.50	12.50	13.33

Table 27: TCR pseudo-ref (%) by source variation category under the base+terms condition. N : avg. distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	N	%	Ref _S	Euro22	Euro9	Qwen	Llama	N	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	64.9%	87.98	94.87	92.76	95.10	93.72	13.2	70.6%	94.81	89.86	91.53	94.15	92.52
RED	3.8	11.0%	84.97	88.03	87.68	88.56	89.14	2.4	12.8%	89.43	75.71	73.95	71.24	82.05
MORPH	0.0	0.1%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.5	1.3%	24.20	20.45	23.71	21.89	18.36	0.1	0.5%	5.00	10.00	10.00	0.00	5.00
GRAPH	3.5	10.0%	54.49	54.51	54.78	56.08	56.27	1.3	7.0%	88.24	63.51	61.22	71.10	69.85
LEX	3.9	11.3%	84.00	81.54	86.06	74.26	77.26	1.3	7.0%	51.73	49.42	48.51	46.29	40.64
COMBO	0.5	1.3%	13.30	17.19	17.86	17.36	18.26	0.4	2.1%	17.78	7.78	8.89	6.67	8.89

Table 28: TCR temporal (%) by source variation category under the base+terms condition. N : avg. distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	N	%	Ref _S	Euro22	Euro9	Qwen	Llama	N	%	Ref	Euro22	Euro9	Qwen	Llama
NO-VAR	22.4	80.7%	86.57	94.15	91.32	94.59	94.38	13.2	80.0%	74.01	88.09	81.32	91.39	89.68
RED	1.4	5.1%	87.82	89.80	86.90	90.41	90.24	1.2	7.3%	72.08	74.17	69.58	76.67	76.67
MORPH	0.0	0.1%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00
EXP	0.2	0.7%	1.61	1.61	1.61	5.91	6.99	0.1	0.6%	0.00	0.00	0.00	0.00	0.00
GRAPH	2.3	8.1%	44.78	43.58	39.93	43.16	48.63	1.3	7.9%	76.61	66.21	61.35	71.21	67.64
LEX	1.4	5.0%	37.90	46.66	45.74	33.94	39.35	0.7	4.2%	13.89	25.91	27.73	16.94	24.60
COMBO	0.1	0.2%	0.00	0.00	0.00	0.00	0.00	0.0	0.0%	0.00	0.00	0.00	0.00	0.00

Table 29: TCR_{EM_{macro}} (%) by source variation category under the base+terms condition. N : avg. in-glossary distinct terms per document with $k \geq 2$ occurrences; %: share of total. Best system per corpus in bold.

Var.	ParaNLP							IWSLT2023						
	<i>N</i>	%	Ref _S	Euro22	Euro9	Qwen	Llama	<i>N</i>	%	Ref	Euro22	Euro9	Qwen	Llama
GRAPH	9.1	34.2%	57.95	60.23	65.47	63.28	66.22	1.0	33.3%	90.00	85.71	85.71	90.00	90.00
RED	8.4	31.5%	87.36	82.98	80.29	83.57	81.42	1.1	36.7%	81.82	83.33	95.65	85.71	86.67
EXP	1.3	4.8%	40.00	37.25	32.00	62.50	36.54	0.2	6.7%	100.00	50.00	0.00	100.00	100.00
LEX	7.9	29.5%	56.97	64.23	66.93	63.24	62.55	0.7	23.3%	85.71	60.00	60.00	60.00	60.00

Table 30: CTV variation preservation (%) by source variation category under the base+terms condition. *N*: avg. eligible occurrences per document (variants whose head term is translated as NO-VAR); %: share of eligible total. Best system per corpus in bold.

Type	Subtype	Description and Examples
<i>Single-Type Variations</i>		
GRAPH	GRAPH 1	Acronyms/Initialisms (e.g., <i>World Health Organization</i> / <i>WHO</i>)
	GRAPH 2	Symbols/Formulas (e.g., <i>carbon dioxide</i> / <i>CO₂</i>)
	GRAPH 3	Spelling changes (e.g., <i>micro-organism</i> / <i>micro organism</i>)
	GRAPH 4	Multiple spelling variants (e.g., <i>micro-organism</i> / <i>microorganism</i>)
	GRAPH 5	Partial abbreviation (e.g., <i>chemotherapy</i> / <i>chemo</i>)
MORPH	MORPH 1	Word order (e.g., <i>automatic translation system</i> / <i>translation system automatic</i>)
	MORPH 2	Article presence/absence (e.g., <i>collection cost</i> / <i>cost of the collection</i>)
	MORPH 3	Inflection (e.g., <i>adaptation</i> / <i>adaptations</i>)
	MORPH 4	Derivation/Affixation (e.g., <i>white</i> / <i>whiten</i>)
	MORPH 5	Syntactic structure (e.g., <i>carbon emission</i> / <i>emission of carbon</i>)
RED	RED 1	Base reduction (e.g., <i>automatic translation system</i> / <i>translation system</i>)
	RED 2	Extension reduction (e.g., <i>automatic translation system</i> / <i>automatic system</i>)
	RED 3	Other reductions (e.g., <i>medium-sized enterprise</i> / <i>enterprise</i>)
EXP	EXP 1	Semantic addition (e.g., <i>translation system</i> / <i>automatic translation system</i>)
	EXP 2	Explicit form (e.g., <i>stats</i> / <i>statistics</i>)
	EXP 3	Lexical insertion (e.g., <i>dark urine</i> / <i>dark colored urine</i>)
	EXP 4	Abbreviation expansion (e.g., <i>ex.</i> / <i>example</i>)
LEX	LEX 1	Simple unit substitution (e.g., <i>residues</i> / <i>waste</i>)
	LEX 2	Base change (e.g., <i>translation system</i> / <i>translation model</i>)
	LEX 3	Extension change (e.g., <i>binary classification model</i> / <i>logistic regression model</i>)
	LEX 4	Base and extension change (e.g., <i>coronavirus disease 2019</i> / <i>Wuhan pneumonia</i>)
<i>Multiple-Type Variations</i>		
CM	GRAPH +LEX	Graphical and lexical (e.g., <i>alpha particle</i> / <i>α ray</i>)
	LEX +RED	Lexical and reduction (e.g., <i>translation system</i> / <i>model</i>)
	GRAPH +RED	Graphical and reduction (e.g., <i>micro-organism pathogen</i> / <i>microorganism</i>)
	GRAPH +MORPH	Graphical and morphosyntactic (e.g., <i>micro-organisms</i> / <i>microorganism</i>)
	MORPH +LEX	Morphosyntactic and lexical (e.g., <i>system of translation</i> / <i>translation model</i>)
	EXP +LEX	Expansion and lexical (e.g., <i>model</i> / <i>automatic translation system</i>)
	Other	Other combinations of two or more types (format: CODE1+CODE2+...)

Table 34: Terminological variation typology with subtypes and examples. GRAPH: Graphical, MORPH: Morphosyntactic, RED: Reduction, EXP: Expansion, LEX: Lexical, CM: Combined Multiple.

#	Head term	Variant	Pred. type	Pred. subtype	Gold type	Gold subtype
1	word error rate	WER	GRAPH	GRAPH 1	GRAPH	GRAPH 1
2	hidden Markov model	HMM	GRAPH	GRAPH 1	GRAPH	GRAPH 1
3	automatic speech recognition system	ASR system	GRAPH	GRAPH 1	GRAPH	GRAPH 5
4	data set	dataset	GRAPH	GRAPH 3	GRAPH	GRAPH 3
5	tokenization	tokenisation	GRAPH	GRAPH 3	GRAPH	GRAPH 3
6	named entity recognition	NER	GRAPH	GRAPH 1	GRAPH	GRAPH 1
7	topic detection	TDT	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
8	conditional random field	CRF	GRAPH	GRAPH 1	GRAPH	GRAPH 1
9	Vector space model	VSM	GRAPH	GRAPH 1	GRAPH	GRAPH 1
10	optimization	optimisation	GRAPH	GRAPH 3	GRAPH	GRAPH 3
11	phrase	PP	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
12	information retrieval model	IR model	GRAPH	GRAPH 1	GRAPH	GRAPH 5
13	cluster	IC	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
14	reinforcement learning	RL	GRAPH	GRAPH 1	GRAPH	GRAPH 1
15	context-free grammar	CFG	GRAPH	GRAPH 1	GRAPH	GRAPH 1
16	Bayesian network	DBN	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
17	n-gram	n gram	GRAPH	GRAPH 3	GRAPH	GRAPH 3
18	speech recognition	TTS	GRAPH	GRAPH 1	NOT LINK	
19	finite state automata	DFA	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
20	syntactic category	POS	GRAPH	GRAPH 1	COMBO	GRAPH 1 + LEX 4
21	model	LM	GRAPH	GRAPH 1	COMBO	GRAPH 1 + EXP 1
22	self-organizing map	SOM	GRAPH	GRAPH 1	GRAPH	GRAPH 1
23	input representation	input sequence	LEX	LEX 3	LEX	LEX 3
24	structural annotation	segmentation	LEX	LEX 1	LEX	LEX 1
25	validation data	development data	LEX	LEX 3	LEX	LEX 3
26	elision	deletion	LEX	LEX 1	LEX	LEX 1
27	inference	entailment	LEX	LEX 1	LEX	LEX 1
28	phrase	VP	LEX	LEX 1	COMBO	GRAPH 1 + EXP 1
29	similarity measure	similarity function	LEX	LEX 1	LEX	LEX 2
30	parallel corpus	parallel text	LEX	LEX 1	LEX	LEX 2
31	syntactic category	word class	LEX	LEX 1	LEX	LEX 1
32	parser	analyzer	LEX	LEX 1	LEX	LEX 1
33	lexeme	lexical item	LEX	LEX 1	LEX	LEX 1
34	interjection	exclamation	LEX	LEX 1	LEX	LEX 1
35	syntactic transfer	reordering	LEX	LEX 1	LEX	LEX 1
36	alignment	matching	LEX	LEX 1	LEX	LEX 1
37	true negative rate	specificity	LEX	LEX 1	LEX	LEX 1
38	Markov model	Markov chain	LEX	LEX 2	LEX	LEX 2
39	syntactic category	part-of-speech	LEX	LEX 1	LEX	LEX 1
40	training data	training dataset	LEX	LEX 2	LEX	LEX 2
41	foundation model	base model	LEX	LEX 3	LEX	LEX 3
42	lexeme	lexical unit	LEX	LEX 1	LEX	LEX 1
43	multi-word term	multi-word	RED	RED 1	RED	RED 1
44	grammatical subject	subject	RED	RED 2	RED	RED 2
45	programming language C++	C++	RED	RED 1	RED	RED 1
46	Apache OpenNLP	OpenNLP	RED	RED 1	RED	RED 1
47	speech recognition	automatic speech recognition	EXP	EXP 1	EXP	EXP 1
48	future	future tense	EXP	EXP 3	EXP	EXP 3
49	syntactic parsing	syntactical analysis	COMBO	LEX 3+MORPH 4	COMBO	LEX 3+MORPH 4
50	similarity measure	measure of similarity	MORPH	MORPH 5	MORPH	MORPH 5

Table 37: The 50 manually evaluated (preferred term, variant) pairs with predicted and gold variation labels.

Incorrect alignment in context (source term and aligned target span in bold)	
Llama (baseline), lexical choice: <i>corpus</i> → <i>Récolte</i> (should align <i>corpus</i>)	
EN	Corpus Harvesting
FR	Récolte de corpus
Llama (base+terms), lexical choice: <i>future</i> → <i>Future</i> (source form echoed; no distinct French realization)	
EN	Conclusion and Future Work
FR	Conclusion et Travail à Faire
Euro22 (base+terms), lexical choice: <i>indexing</i> → <i>discours</i> (should align <i>indexation</i>)	
EN	speech indexing
FR	indexation des discours
Ref, span boundary: <i>automatic speech recognition</i> → <i>reconnaissance de la parole</i> (drops <i>automatique</i>)	
EN	The work presented here aims at expanding this paradigm to automatic speech recognition .
FR	Les travaux présentés ici visent à étendre ce paradigme à la reconnaissance automatique de la parole .

Table 39: The four incorrect alignments among the 100 manually evaluated pairs, shown in source (EN) and target (FR) context. The source term is bold in the English segment and the aligned target span is bold in the French segment; for the span-boundary error, the non-bold word (*automatique*) is the token wrongly excluded from the span.

""You are an expert in scientific terminology and linguistic variation, teaching a course on terminology. A student has identified a potential variant of a technical term and needs your help to classify the relationship between the head term and its variant according to the established typology.

You receive:

- A head term (the preferred form of a technical term)
- A variant (an alternative form that may or may not be related to the head term)
- Context examples showing how these terms appear in real scientific texts (when available)

Your task: Determine the type of variation between the head term and the variant. Classify according to the variation typology. Return ONLY a JSON object with the exact format shown in the examples below.

Important rules:

- 1) First verify that the variant is actually related to the head term
- 2) Count the TYPES of changes, not the number of words: one type -> \VG/\VMS/\VR/\VE/\VL, multiple types -> CM
- 3) For CM (Multiple Changes), combine codes with + (e.g., "\VG+\VL" or "\VMS+\VR+\VL")
- 4) Copy the exact JSON format from the examples - no markdown, no extra text
- 5) Keep justifications brief (1-2 sentences maximum)

Reference of variation types:

\VG (Graphical): Written form changes without semantic change (acronyms, spelling, symbols)

- > \VG1: Acronym (European Parliament/EP), \VG2: Symbol (CO2), \VG3: Spelling (trade mark/trademark), \VG4: Multiple spelling, \VG5: Partial abbreviation

\VMS (Morphosyntactic): Structure or inflection changes (word order, articles, inflection, derivation)

- > \VMS1: Order, \VMS2: Article, \VMS3: Inflection, \VMS4: Derivation, \VMS5: Structure

\VR (Reduction): Element deletion

- > \VR1: Base reduction, \VR2: Extension reduction, \VR3: Other reduction

\VE (Expansion): Element addition

- > \VE1: Semantic addition, \VE2: Explicit form, \VE3: Lexical insertion, \VE4: Abbreviation development

\VL (Lexical): Lexical substitution (synonyms, near-synonyms, translations validated as variant in context)

- > \VL1: Simple unit, \VL2: Base change, \VL3: Attribute/extension change, \VL4: Base+extension change

CM (Multiple Changes): Combination of two or more types above

When completing the task, follow the examples below:

Head term: European Parliament
Variant: EP
{"categorie": "\VG", "sous_type": "\VG1", "justification": "Acronym formed from the initials of the head term."}

.

.

.

Head term: coronavirus disease 2019
Variant: COVID-19
{"categorie": "CM", "sous_type": "\VG+\VR", "justification": "Acronym formation (\VG1) combined with structural reduction."}}

Head term: airborne dust particle
Variant: airborne dust
{"categorie": "\VR", "sous_type": "\VR1", "justification": "Base reduction: deletion of 'particle' from the base noun phrase."}}

{context_section}

Head term: {head_term}
Variant: {variant}""

Figure 9: Variation classification prompt.