

Quantifying the Impact of Translation Errors on Multilingual LLM Evaluation

Klaudia-Doris Thellmann¹, Bernhard Stadler^{1,2}, Michael Färber¹, Jens Lehmann^{3,1,2}
firstname.lastname@tu-dresden.de

¹TUD Dresden University of Technology and ScaDS.AI,

²InfAI e.V., ³Amazon (work done outside of Amazon)

Correspondence: klaudia-doris.thellmann@tu-dresden.de

Abstract

Machine-translated benchmarks are widely used to assess the multilingual capabilities of large language models (LLMs), yet translation errors in these benchmarks remain underexplored, raising concerns about the reliability and comparability of multilingual evaluation. We address two practical gaps: (i) how well automatic MQM-style error spans from LLM judges and a span-aware QE baseline (xCOMET-XXL) match expert human span annotations on benchmark translations, and (ii) how strongly translation errors (as opposed to source-side issues in the English original) explain accuracy drops on translated benchmarks. We find that span agreement is non-trivial on naturally occurring benchmark translations, and that target-side translation errors are consistently associated with measurable, percentage-point drops in translated accuracy even after controlling for English correctness and source-side anomalies.

1 Introduction

In multilingual evaluation, machine-translated datasets are widely used as reference data, yet translation quality is often overlooked, undermining reliability and comparability (Choenni et al., 2024; Artetxe et al., 2020; Plaza et al., 2024). Human protocols such as MQM (Lommel et al., 2013, 2024) and Error Span Annotation (ESA) (Kocmi et al., 2024) provide increasingly diagnostic assessments (Freitag et al., 2021).

More recently, researchers have treated LLMs themselves as translation judges (“LLM-as-a-judge”; Kocmi and Federmann, 2023), using zero- or few-shot prompting to tag MQM-style error spans, while recent fine-grained quality-estimation work has likewise moved toward span-aware error localization, most prominently in xCOMET-XXL (Guerreiro et al., 2024); however, its behavior on naturally occurring benchmark translations and

the implications for benchmark reliability remain underexplored.

This trend is exemplified by GEMBA and GEMBA-ESA (Freitag et al., 2024), and by GPT-based evaluators such as AutoMQM (Huang et al., 2024) for inline span detection, or MQM-APE (Lu et al., 2025), which uses automatic post-editing to refine translations.

Prior work investigating the effects of translation artifacts on model performance relies either on manual inspection of small samples (Artetxe et al., 2020; Plaza et al., 2024), which provides qualitative insights but does not scale, or on heuristics (Park et al., 2024; Choenni et al., 2024) such as sentence length ratios or learned quality estimation scores (e.g., COMET-QE (Rei et al., 2020)), both of which lack precision in identifying the type and location of translation errors. In addition, most of these studies are limited to single benchmarks or languages (e.g. Spanish MMLU (Plaza et al., 2024) or XNLI (Artetxe et al., 2020)).

Table 1 illustrates why source-side anomalies matter: some are corrected in translation, while others are preserved or exacerbated, so source- and target-side errors can independently affect both benchmark outcomes and the construct being tested.

Two practical questions remain under-addressed: (i) how well do LLM-produced MQM-style spans match expert human spans on real benchmark translations, and (ii) how strongly do translation errors (vs. source-side issues) explain accuracy drops on translated benchmarks?

We study these questions on the EU20 suite (Thellmann et al., 2024) and annotate translation errors with four LLM annotators using an MQM-inspired prompt (§3). To strengthen the comparison, we also include xCOMET-XXL as a baseline, complement thresholded span matching with threshold-free character-overlap analyses, and test in a counterfactual analysis whether removing translation-error effects changes downstream

Source	EN key phrase	Translation key phrase	Why it matters for evaluation
<i>Source-side anomalies (S) – flagged on the English source</i>			
ARC (DA), corrected	rate of speed	samme hastighed	Source oddness corrected in target; not a target-side translation error.
ARC (DE), preserved	terrestrial plants	terrestrische Pflanzen	Source anomaly (plants instead of planets) preserved in target; may be mistaken for a target-side error.
<i>Target-side translation errors (T) – English solvable, translation not</i>			
MMLU (FR)	Wrong	Faux	Moral judgment → factual judgment. EN ✓ / FR ×
ARC (DA)	shorter distance	kortere tid	Stopping distance → stopping time. EN ✓ / DA ×
GSM8K (SL)	feet / cubic feet	metra / kubična metra	Unit shift invalidates the calculation. EN ✓ / SL ×
HellaSwag (DA)	liner	foderet	Cleaning action → feeding action. EN ✓ / DA ×

Table 1: Illustrative examples of source-side anomalies (S) and target-side translation errors (T). S cases are flagged on the English source and either corrected or preserved in the translation; they do *not* count as target-side translation errors. T cases change the intended meaning of the target item; in our sample, the English original was solved correctly (EN ✓) while the translation was not (target ×).

system rankings.

We make three contributions:

1. Human reference and LLM span agreement on EU20 (§3). We release a professional span-level MQM reference for an EU20 subset (225 items, nine target languages) and evaluate four LLM annotators and xCOMET-XXL against it. GPT-5.2 shows the highest agreement with human annotations under strict span matching (mean span-level F1 **0.55** under position-overlap matching) and complementary character-overlap metrics (Char-F1 / Char-F1w).

2. Meta-evaluation with Span-ACES_{Ref} (§4). We evaluate span localization on Span-ACES_{Ref} (1,407 items), a cleaned projection of SPAN-ACES (Moghe et al., 2025), and validate the transformation on 178 records where manual review and GPT-5.2 agree on 165/178 (0.93), enabling controlled classic vs. tolerant span metrics.

3. Performance impact with source controls (§5). Using our annotation-and-evaluation setup, we estimate the impact of target-side translation errors (T) and source-side issues (S) on translated accuracy via logistic regression (fixed effects; English controls; bootstrap CIs): across annotators, T is associated with drops of about **6–8 pp** (full model) and **6–11 pp** among English-solvable items. A counterfactual T=0 ranking analysis further shows that translation errors mainly shift absolute scores rather than relative system rankings.

The annotated datasets, Span-ACES_{Ref} resources, and our codebases and prompts are available under [hf.co/btqe](https://huggingface.co/btqe) and github.com/btqe.

2 Related Work

Translation artifacts and their effects. Several studies have shown that translation artifacts can undermine the reliability of model evaluation: Choenni et al. (2024) found that MT-generated test sets may overestimate model capabilities, especially in low-resource languages; Artetxe et al. (2020) demonstrated that subtle “translationese” can bias cross-lingual benchmarks like XNLI; Plaza et al. (2024) reported that mistranslations in Spanish MMLU data cause 6–13% accuracy loss for GPT-4, with up to 60% of failures directly linked to translation errors; and Park et al. (2024) observed similar effects for VQA models. Agrawal et al. (Agrawal et al., 2024) further show that translation errors can substantially affect cross-lingual learning outcomes, especially in lower-resource settings, reinforcing that benchmark translation quality is not merely a surface issue but can alter downstream conclusions. While these findings underscore the need for rigorous quality control, prior work remains limited in scale and granularity. Quality estimation and automated translation evaluation have also moved toward span-aware error localization, for example with xCOMET-XXL (Guerreiro et al., 2024) and recent WMT shared tasks on fine-grained evaluation (Blain et al., 2023; Zerva et al., 2024; Lavie et al., 2025).

Our work differs in focusing on naturally occurring benchmark translations and on linking localized translation errors to downstream benchmark accuracy.

Multilingual benchmarks. Recent multilingual benchmarks range from carefully curated, manually translated datasets (e.g., SuperGLEBer (Pfister and Hotho, 2024), ScandEval (Nielsen, 2023), Iber-

oBench (Baucells et al., 2025), FrenchBench (introduced as part of CroissantLLM; Faysse et al., 2025), BenCzechMark (Fajcik et al., 2025)) to large-scale resources generated via machine translation. While they offer high quality, manually constructed benchmarks are costly and difficult to scale, prompting the use of machine translation for broader coverage, e.g. Global MMLU (Singh et al., 2025), XNLI (Conneau et al., 2018), OKAPI (Lai et al., 2023), and LAMBADA (Paperno et al., 2016). However, many such resources lack transparent quality control. Our work advances this line of research by combining automated span-level error annotation, benchmarked against expert human annotations and a strong baseline, with statistical analysis to assess how translation errors affect downstream LLM performance on multilingual benchmarks.

3 Human and LLM MQM Annotation

In this section, we present our MQM-based annotation setup for capturing translation errors and assess the span-level annotation accuracy of LLM-based annotators. We compare LLM-generated MQM annotations against a professionally produced human reference on an EU20 subset and additionally evaluate LLM annotators on Span-ACES_{Ref}.

3.1 Methodology

Human Annotation. Using EU20 (Thellmann et al., 2024) (DeepL translations¹ of MMLU (Hendrycks et al., 2021), ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), GSM8K (Cobbe et al., 2021), and TruthfulQA (Lin et al., 2022) into 20 official EU languages), we create a span-level human reference for translation error analysis by manually annotating nine target languages spanning Germanic (DE, DA), Romance (FR, IT, RO), Uralic (ET, HU), Slavic (SL), and Baltic (LT), covering a range of high- to low-resource settings.

For each language, we randomly selected 25 segments (225 total) using fixed per-dataset quotas to balance task types; a minimum-length filter excluded trivially short instances. Dataset statistics are reported in Appendix A.1 (Table 5).

Annotation was carried out by professional translators/linguists with one annotator per language using an MQM-inspired, span-based protocol implemented in a custom Argilla interface (Appendix A.1). For segments with errors, annotators

marked erroneous spans and assigned an MQM label and severity (Major/Minor), distinguishing *Accuracy* errors (faithfulness/meaning transfer; Table 6) from *Fluency/Style* errors (grammaticality and naturalness; Table 7). They highlighted target-side error spans and, where applicable, the corresponding source spans (Figure 2), and then submitted a minimally post-edited corrected translation and optional clarification comment (Figure 3).

After collection, we performed systematic consistency checks (e.g., presence of error labels/severity, and source–target span alignment where applicable) and investigated irregular cases such as source–target span mismatches by consulting annotator comments and the original segment context. Where necessary, we corrected format inconsistencies and obvious annotation-entry issues to obtain a clean, machine-readable reference. In addition, we conducted a targeted manual cross-check by comparing human-annotated spans with LLM-proposed annotations as a complementary sanity check.

LLM-based Annotation. To generate automatic, span-level MQM annotations, we use four instruction-tuned LLMs: GPT-5.2, GPT-4o-mini, LLaMA-4, and Mistral-Large. Model references are listed in Appendix A.1 (Table 8). For comparison, we additionally include xCOMET-XXL (src+mt, reference-free/QE) as a strong non-generative baseline with span-level error predictions. We adopt a variant of the GEMBA-ESA prompting approach² and extend it with curated multilingual few-shot examples, covering a broad range of error types and both structured content (e.g., multiple-choice questions) and general-purpose text. Our structured JSON output also aligns with the broader move toward explicit span-tagging formats for translation error detection in reasoning-sensitive settings, such as tagged span annotation (Yeom et al., 2025). To ensure comparability with the reference, the prompt restricts labels to the same MQM inventory (Accuracy vs. Fluency/Style, see Table 6 and 7 from Appendix A.1). Each model is prompted to produce a JSON-structured response containing a list of annotated error spans with MQM label and severity.

In addition, we annotate source-side anomalies on the same EU20 subset using GPT-5.2 to account for cases where the source has irregularities that may affect downstream TQE. These annotations follow an MQM-inspired, span-based protocol for flu-

¹developers.deepl.com/docs

²github.com/MicrosoftTranslator/GEMBA

ency/coherence anomalies and distinguish surface issues (e.g., *typo*, *grammar*, *punctuation*, *awkward*) from semantic oddness (e.g., *contradiction*, *broken_coreference*, *implausible_logic*), with severity (major/minor) indicating potential impact on interpretability or answerability.

Metrics for Span-Level Agreement. Because span annotation is a unitizing task—annotators define both the presence and the boundaries of spans, rather than labeling pre-segmented units—we complement pairwise span-matching scores with reliability analyses that serve as robustness checks, including Krippendorff’s unitized alpha (α_u) for free-span annotation settings (Krippendorff, 1995; Krippendorff et al., 2016). This unitizing setting is also reflected in NLP annotation tooling such as DKPro Agreement (Meyer et al., 2014).

We report span-level Recall and F1 computed with greedy 1:1 matching with fixed thresholds and micro-aggregation over items (TP/FP/FN summed per language). For the position-based overlap coefficient (OC), a candidate pair is matchable if $OC \geq 0.8$, where OC is overlap length divided by minimum span length. For the string-based metric (SIM), we use raw character 3-gram Dice similarity (no normalization) with threshold $SIM \geq 0.6$. SIM is computed on spans deduplicated by exact target-span text (on both sides), whereas OC operates on the full span lists and requires valid target offsets. Spans with missing offsets cannot be matched under OC and therefore contribute to FP/FN in the OC-based Span-F1 computation.

Beyond thresholded OC/SIM matching, we also report threshold-free character-overlap metrics. Our severity-aware character-overlap metric Char-F1w follows the overlap logic of the WMT25 span-level evaluation setting, giving full credit to severity-matched overlaps and partial credit to severity mismatches (Lavie et al., 2025). We additionally report a binary Char-F1 variant that ignores severity labels. These character-level views complement the thresholded span-matching scores used in the main text and help test whether annotator ranking is robust to metric choice.

The SIM-based results are reported in Appendix A.2 (Table 9). Additional unitizing-reliability analyses are reported in Appendix A.4 (Table 13); threshold-free character-overlap results are reported in Appendix A.5 (Table 14), with complementary threshold-sweep ranges in Table 15.

Source-overlap. To quantify the tendency to annotate source-side oddness as target errors, we compute the fraction of an annotator’s target spans (with valid offsets) that overlap target regions linked to GPT-5.2 source-anomaly annotations (OC threshold 0.8, each annotator span counted at most once). For each source anomaly, GPT-5.2 provides an English anchor span and, when possible, a target-side anchor in the translation. We ground these anchors conservatively: provided offsets are validated by exact substring checks, and if offsets are missing or don’t match, we attempt exact substring search to recover offsets only when the match is unambiguous. Otherwise, the record remains unlinked and is excluded from offset-based computations. We then compute source-overlap with OC threshold 0.8, counting each annotator span at most once via its best overlap.

3.2 Results and Discussion

Annotator	OC Recall	OC F_1	F_1 range (lang.)	Char- F_1 [95% CI]
GPT-5.2	.44	.55	.34–.78	.48 [.437, .529]
xCOMET-XXL	.32	.28	—	.26 [.223, .301]
Mistral-Large	.17	.23	.19–.32	.28 [.235, .317]
GPT-4o-mini	.09	.15	.08–.25	.21 [.163, .261]
LLaMA-4	.08	.13	.03–.23	.14 [.104, .183]

Table 2: Span-level agreement with the professional human MQM reference on the EU20 subset. OC columns use greedy one-to-one matching on target character offsets ($OC \geq 0.8$), averaged over 9 languages (25 items each). F_1 range gives the per-language min–max Span- F_1 . Char- F_1 is a threshold-free character-level span-overlap metric on the binary target-side error mask, reported as global micro with 95% bootstrap CIs (per-language ranges are unavailable for xCOMET-XXL).

Span-level Agreement. Table 2 shows that GPT-5.2 yields the highest agreement with the human reference under OC-based span matching (mean OC Span-F1 = .55), outperforming the other automatic annotators. Among automatic baselines, xCOMET-XXL is stronger (mean OC Span-F1 = .28) than GPT-4o-mini and LLaMA-4 and broadly comparable to or slightly above Mistral-Large under our OC-based matching, but remains below GPT-5.2 on naturally occurring benchmark translations. Additional OC-based pairwise comparisons using GPT-5.2 and xCOMET-XXL as reference annotators are reported in Appendix A.2 (Table 10). Appendix A.3 provides worked examples for OC and SIM, while Appendix A.7 collects addi-

tional qualitative source-side and target-side example cases. Appendix A.5 (Table 16) further shows that xCOMET-XXL is more fine-grained than the other automatic annotators, producing more and shorter spans on average.

Besides GPT-5.2, Mistral-Large generally achieves higher agreement scores than GPT-4o-mini and LLaMA-4, with LLaMA-4 typically lowest. This qualitative ranking is robust across alternative evaluation views: threshold-free character-overlap metrics and pooled bootstrap confidence intervals support the same overall conclusion, with GPT-5.2 remaining strongest and the remaining LLM annotators ordered similarly (Appendix A.5, Table 14; Appendix A.6, Table 17). Appendix A.2 further reports SIM-based agreement (Table 9) and per-language Span-F1 for OC and SIM (Table 11), showing that relative annotator rankings are stable even though absolute agreement varies across language batches. Appendix A.6 (Table 17) further reports pooled bootstrap confidence intervals over items, confirming the same qualitative ranking for the four LLM annotators.

The low agreement between the automatic annotators and the human reference indicates that span boundaries are often chosen differently and that annotators may also disagree on what should count as a translation error in the target text.

Annotator	Mean overlap rate	Range
GPT-5.2	.06	.00–.13
Human	.08	.04–.13
xCOMET-XXL	.08	.04–.13
GPT-4o-mini	.10	.00–.25
LLaMA-4	.13	.06–.29
Mistral-Large	.14	.09–.20

Table 3: Source-overlap: fraction of an annotator’s target spans with valid offsets that overlap target regions linked to GPT-5.2 source-anomaly annotations. Values are means over 9 languages; ranges denote min–max across languages. Higher values indicate a stronger tendency to mark source-linked target regions as target-side errors ($OC \geq 0.8$; each annotator span counted at most once via its best overlap).

Source-overlap indicator. Table 3 provides a diagnostic view of source-linked confounds. GPT-5.2 has the lowest source-overlap rate, consistent with a stricter target-only annotation tendency. xCOMET-XXL is close to the human rate, while Mistral-Large and LLaMA-4 show higher overlap, suggesting that they more often mark source-linked regions

as target-side translation errors.

Human annotations are non-zero, reflecting that some source issues are preserved in the translation. Per-language source-overlap rates are reported in Appendix A.2 (Table 12) and help contextualize language-dependent differences in span agreement. More broadly, source-overlap helps explain part of the agreement differences across annotators: systems that frequently mark regions tied to source-side irregularities can diverge from annotators that focus more strictly on target-only errors. Additional source-side and target-side example cases are collected in Appendix A.7.

Qualitative inspection. To contextualize the agreement scores, we manually spot-checked a small set of items across DE/DA/IT/RO. Although limited, the inspection revealed recurring divergence patterns consistent with the quantitative results. GPT-4o-mini, LLaMA-4, and Mistral-Large sometimes label answer-option issues as *Mistranslation* even when the translation is plausible, especially in multiple-choice settings. Such cases are less frequent in the human reference and GPT-5.2. Both human and LLM annotators also occasionally treat source-side ill-formedness as target-side translation errors, consistent with non-zero source-overlap rates (Appendix A.2, Table 12), while occasional data irregularities (e.g., source–target mismatches) can reduce span agreement independently of annotation quality. We also observed model-specific tendencies: GPT-5.2 sometimes reverts acceptable human post-edits and misses subtle meaning differences, whereas GPT-4o-mini can produce broader spans. These observations suggest that disagreement is driven not only by span boundary variation but also by differing judgments about what should count as a target-side translation error, especially in the presence of source anomalies. These findings motivate treating source anomalies as a confounder and reporting agreement both with and without source-linked regions in future work.

4 Span-ACES_{Ref}

To complement the EU20 study with an external gold-span benchmark, we evaluate the same LLM annotators on Span-ACES_{Ref}, a cleaned projection of Span-ACES.

4.1 Methodology

Span-ACES_{Ref} construction. A key limitation of SPAN-ACES is that it labels only the introduced

Aggregation	GPT-4o-mini				LLaMA-4				Mistral-Large			
	F1	R	F1 _t	R _t	F1	R	F1 _t	R _t	F1	R	F1 _t	R _t
<i>Accuracy (20 EU20 target languages)</i>												
mean (N)	.10/.21	.17/.23	.33/.48	.57/.55	.13/.10	.17/.12	.34/.31	.43/.36	.24/.34	.36/.42	.47/.60	.70/.71
mean (cap)	.17/.28	.25/.31	.31/.46	.49/.51	.21/.20	.26/.22	.34/.35	.43/.40	.28/.39	.39/.45	.41/.57	.59/.65
<i>Fluency / Style (German-only, four anaphoric phenomena)</i>												
mean (N)	.05/.32	.09/.37	.41/.57	.68/.67	.15/.16	.19/.17	.42/.41	.59/.47	.45/.46	.62/.54	.55/.56	.82/.67
mean (cap)	.04/.32	.08/.36	.47/.61	.74/.71	.12/.14	.16/.15	.40/.41	.58/.49	.47/.48	.62/.56	.58/.59	.83/.70

Table 4: Span-ACES_{Ref} span localization across subsets. Classic matching (Span-F1 $F1$, Span-Recall R) requires exact span equality; tolerant matching ($F1_t$, R_t) allows gold-centered boundary slack up to $k=3$ tokens per side. Each cell reports *Baseline/Ours* (initial GEMBA-ESA vs. our updated prompt).

error span per item, leaving any additional translation errors unlabeled. Consequently, span-based annotators that correctly flag such extra errors are counted as false positives, biasing precision and Span-F1 downward. Each item consists of a human reference, a good translation, and an incorrect translation containing exactly one targeted phenomenon. We construct Span-ACES_{Ref} by extracting the single contentful token-level difference between the good and incorrect translations, mapping that edit to the corresponding location in the human reference, and discarding items with multiple diffs or ambiguous matches. This reduces unlabeled noise and yields more reliable span-level precision and Span-F1.

Mapping to MQM categories. Using the ACES authors’ released mapping, we collapse phenomena into the two coarse MQM categories *Accuracy* and *Fluency/Style*, yielding 1,155 *Accuracy* instances across the 20 EU20 target languages and 252 *Fluency/Style* instances for German only in Span-ACES_{Ref}. Dataset composition follows the final Span-ACES_{Ref} split described above, and the MQM mapping is reported in Appendix B.1 (Table 20).

Dataset validation (human vs. GPT-5.2). To validate the transformation pipeline and identify projection errors, we manually reviewed 178 Span-ACES_{Ref} records from a stratified sample using a fixed checklist and compared the resulting pass/fail judgments with GPT-5.2. The two assessments agreed on 165/178 items (0.93), suggesting that most cases are straightforward under these criteria and that GPT-5.2 can provide a useful auxiliary signal for scalable dataset validation. The remaining 13 disagreements were mostly borderline cases, such as type drift after projection or subtle side effects.

Span metrics (classic vs. tolerant). We report classic span metrics (Span-F1 $F1$ and Span-Recall R) and a tolerant variant ($F1_t$, R_t) that allows small boundary mismatches without overlap thresholds. Classic matching requires exact equality between predicted and gold spans. Tolerant matching is gold-centered: a prediction is counted as correct if it contains the gold span and may extend by up to k tokens on either side. Metrics are first averaged within each phenomenon and then aggregated per MQM category using either sample-count weighting (**mean (N)**) or capped weighting (**mean (cap)**, $C=25$) to limit dominance by large phenomena.

Baseline vs. our prompt. We evaluate two prompting setups on Span-ACES_{Ref}: (i) the initial GEMBA-ESA baseline prompt, and (ii) our updated MQM-style prompting (Section 3.1, paragraph “LLM-based Annotation”). Unless stated otherwise, tables report each metric as *Baseline/Ours*.

4.2 Results and Discussion

Accuracy and Fluency/Style. Table 4 reports span localization results for *Accuracy* and *Fluency/Style*. On *Accuracy*, Mistral-Large performs best overall, and tolerant matching substantially improves scores for all models, indicating frequent boundary near-misses. Under mean (N), the updated prompt raises classic F1 for GPT-4o-mini (**.10→.21**) and Mistral-Large (**.24→.34**), but slightly lowers it for LLaMA-4 (**.13→.10**), a pattern that also holds under tolerant matching and therefore suggests prompt–model sensitivity. On *Fluency/Style*, scores are higher overall, but because this subset is German-only and limited in scope, we view it mainly as a controlled diagnostic. Here too, prompt refinements help GPT-4o-mini most strongly (classic F1 **.05→.32**, mean (N)), while gains for Mistral-Large (**.45→.46**) and LLaMA-

4 (.15→.16) are marginal, indicating that prompt improvements do not transfer uniformly across architectures. Across models and categories, tolerant scores (F_{1_t} , R_t) are consistently higher than classic scores, indicating frequent boundary near-misses rather than complete detection failures. This matches a common pattern in LLM annotations: different systems often point to the same error region but choose slightly different span boundaries (e.g., broader or shifted spans).

5 Performance Impact Analysis

To assess how translation quality relates to downstream LLM performance, we link span-level error annotations to whether models answer translated instances correctly and analyze these associations with logistic regression.

5.1 Methodology

Data and annotations. For our performance analysis, we use MMLU, ARC, and GSM8K from the EU20 benchmark suite and annotate target-side translation errors with an MQM-inspired prompting scheme (Section 3) using four LLM-based annotators: GPT-5.2, GPT-4o-mini, LLaMA-4, and Mistral-Large. For robustness, we additionally include xCOMET-XXL as non-generative baseline in the same regression pipeline. To control for issues already present in the original English items (e.g., inconsistencies or ambiguities that may affect model accuracy independent of translation), we additionally annotate the English source for a paired subset of 651 translated instances (item–language pairs) using GPT-5.2. Together with the corresponding target-side annotations, this paired subset enables a direct comparison of source-side issues (S) versus translation-induced errors (T). For the performance-impact regressions, we exclude HellaSwag and TruthfulQA: in HellaSwag, many “odd” endings are intentionally constructed distractors (task design rather than source noise), which would blur the interpretation of S , and for TruthfulQA we could not reliably derive binary correctness labels from the released evaluation logs. The resulting paired subset contains 651 item–language pairs and 4,993 translated observations across the evaluated models. Detailed dataset statistics, including translated and English correctness counts and annotation sparsity, are reported in Appendix C.1 (Table 22).

Model evaluation. We evaluate eight instruction-tuned multilingual LLMs (Appendix C.2, Table 23)

using a modified version of EleutherAI’s LM Evaluation Harness³ and record a binary outcome (correct/incorrect) for each item in English and in each translated variant.

Analysis dataset. We merge per-instance correctness with the corresponding span-level error annotations to construct an analysis dataset that links (i) the presence of source-side issues and/or translation errors to (ii) each evaluated LLM’s correctness on the English original and its translated variants. We analyze this dataset with logistic regression, focusing on the paired-annotation subset when comparing translation errors against source-side issues.

Regression setup. We use logistic regression to estimate how the presence of translation errors (T) and source-side issues (S) predicts correctness on translated items, while controlling for whether the model solves the English original and for systematic differences across languages, datasets, and evaluation models. A key property of our data is that for each underlying item and evaluation model we observe both (i) performance on the English original and (ii) performance on its EU20 translations, enabling a clean separation between English-level ability and translation-associated failures.

Instances and variables. Our unit of analysis is a translated instance: we record whether evaluation model m answers an item in target language ℓ correctly ($y \in \{0, 1\}$; English excluded). For the same item and model, we record whether the English original is answered correctly, denoted $y^{EN} \in \{0, 1\}$. From the span-level annotations, we derive two binary indicators: $T = 1$ if the translation contains at least one annotated target-side error (annotator-specific), and $S = 1$ if at least one issue was annotated in the English source. Since S refers to the English original, it is shared across all translations of the same item.

Specifications. We fit logistic regressions with fixed effects for target language, dataset, and evaluation model, denoted by C .

(A) Full model with English control. This specification asks: *How strongly are translation errors associated with lower accuracy after controlling for whether the model solves the English original?*

$$\Pr(y = 1) = \sigma(\beta_T T + \beta_S S + \beta_{EN} y^{EN} + C). \quad (1)$$

(B) Only items solved in English. Estimated on the subset with $y^{EN} = 1$, this specification

³github.com/EleutherAI/lm-evaluation-harness

asks: *Among items the model answers correctly in English, how are translation errors (T) and source issues (S) associated with accuracy in translation?*

$$\Pr(y = 1) = \sigma(\beta_T T + \beta_S S + C). \quad (2)$$

Ablations (omit S). To assess sensitivity to omitted-variable bias, we refit (A) and (B) without S . If translation errors correlate with source-side issues that also depress accuracy, omitting S can make the estimated translation-error association more negative.

Effect reporting and uncertainty. For each target annotator and specification, we report average marginal effects (AMEs) of T and S in probability points (more interpretable than log-odds). We compute confidence intervals using a block bootstrap over underlying items: we resample items with replacement, include all associated rows (across target languages and evaluation models) for each resampled item, refit the model, and take percentile intervals from the bootstrap distribution. We also run a counterfactual ranking analysis by predicting each evaluation model’s translated accuracy under $T=0$ and comparing the resulting ranking to the observed ranking (Appendix C.4).

Counterfactual ranking. To assess whether translation errors affect not only absolute scores but also relative system ordering, we additionally predict each evaluation model’s translated accuracy under the counterfactual setting $T=0$ and compare the resulting ranking to the observed ranking.

5.2 Results and Discussion

Figure 1 summarizes the estimated effect of target-side translation errors on translated accuracy under the two main specifications. xCOMET-XXL yields the same qualitative pattern as baseline annotator: its estimated translation-error effect is also negative, borderline in the full model (Spec. A: $\text{AME}(T) = -4.60$ pp, 95% CI [-9.07, 0.00]) but clearly below zero in the English-correct subset (Spec. B: -6.20 pp, 95% CI [-10.88, -1.57]). This indicates that the main result does not depend on a single LLM-as-a-judge annotation approach. Across all four annotators, target-side translation errors (T) are consistently associated with lower correctness in translation, even when controlling for source-side issues (S), English correctness, and fixed effects. In the full model with an English control (Spec. A), $\text{AME}(T)$ ranges from **-5.59 to -7.51** pp and all 95%

bootstrap CIs lie entirely below zero. Conditioning on items solved in English (Spec. B; $y^{EN}=1$) yields larger drops, with $\text{AME}(T)$ between **-6.12 and -10.71** pp, again with CIs fully below zero. Together, A and B provide complementary evidence that translation errors are associated with accuracy losses and that these losses persist even when the underlying item is solvable in English.

Source-side issues (S) are directionally negative but weaker and less stable. In Spec. A, $\text{AME}(S)$ is around **-3.5 to -3.8** pp, but the corresponding CIs generally overlap zero, consistent with y^{EN} absorbing much of the item-difficulty signal. In Spec. B, $\text{AME}(S)$ is somewhat larger and becomes significant for some annotators, suggesting that source anomalies can additionally depress translated performance once we restrict to English-solvable items. Omitting S changes $\text{AME}(T)$ only marginally (typically about **0.2–0.5** pp), so the translation-error effect is not simply explained away by source-side issues. Logit coefficients and odds ratios mirror these patterns (Appendix Tables 24 and 26), with $\text{OR}(T)$ consistently below 1 (roughly .65–.72 in A and .54–.69 in B). Full AME estimates for all specifications, including xCOMET-XXL, are reported in Appendix C.3 (Tables 27 and 25).

Practical impact. Figure 1 reports effects in probability points (pp), which can be read directly as accuracy drops. Across annotators, a translation with at least one annotated error ($T=1$) is associated with a **6–8 pp** lower probability of a correct answer in the full model (Spec. A) and a **6–11 pp** lower probability when restricted to items solved in English (Spec. B; $y^{EN}=1$). Thus, translation errors are associated with a noticeable loss in measured accuracy even when the underlying item is solvable in English.

To translate these per-item drops into overall impact, we use the approximation $\text{overall loss} \approx \Pr(T=1) \times |\text{AME}(T)|$. With GPT-5.2 annotations, $T=1$ occurs in **63%** of translated observations in Spec. A; combined with $\text{AME}(T) = -6.76$ pp, this implies an overall loss of about **4.3 pp**. In the English-correct subset of Spec. B, the same calculation gives larger impacts (e.g., GPT-5.2: **62% × 8.98 pp ≈ 5.6 pp**). These figures are intended as a readable approximation of the observed association, not a separate causal estimand.

For intuition, every additional **10%** of items with a translation error corresponds to about $|\text{AME}(T)|/10$ points of accuracy loss: roughly

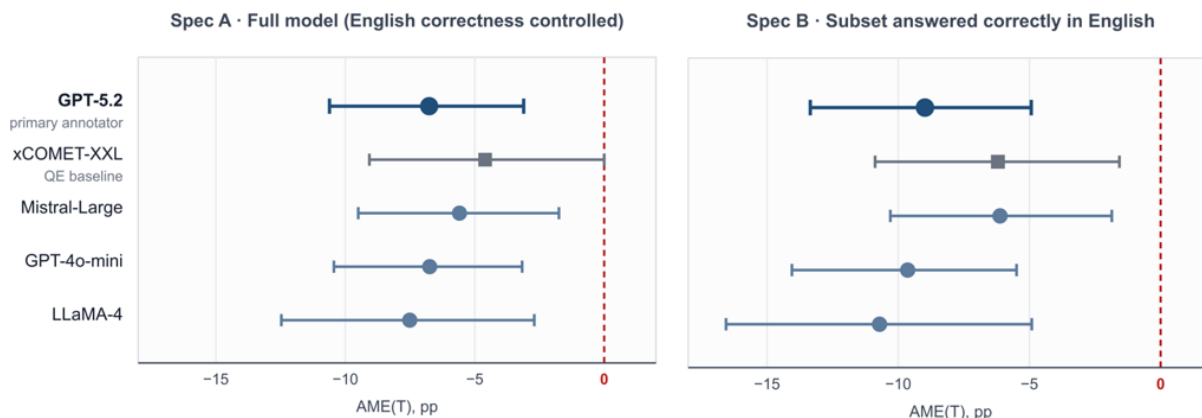


Figure 1: Average marginal effects (AMEs) of target-side translation errors T on translated correctness. Panel A shows the full model controlling for English correctness; Panel B is restricted to items answered correctly in English ($y^{\text{EN}}=1$). Points show $\text{AME}(T)$ in probability points (pp), bars 95% block-bootstrap confidence intervals over underlying items, and the dashed vertical line marks zero. Across annotators, translation errors are associated with lower translated accuracy; xCOMET-XXL shows the same negative direction and is clearly below zero in Spec. B.

0.6–0.8 pp in Spec. A and **0.6–1.1 pp** in Spec. B. This approximation ignores overlap between T and source-side issues S beyond what is already captured by the regression controls.

Counterfactual ranking. To assess whether translation errors affect not only absolute scores but also relative system ordering, we compare the observed ranking of evaluation models to a counterfactual ranking obtained by predicting translated accuracy under $T=0$. Across the three annotators included in this auxiliary analysis (Mistral-Large, LLaMA-4, and GPT-5.2), the observed and counterfactual rankings are nearly identical: Spearman’s ρ is essentially 1.0 in all cases, with equally tight Kendall τ estimates (Appendix C.4, Table 28). This suggests that translation errors behave like an approximately uniform penalty across systems: they shift absolute scores downward, but usually do not induce rank swaps. At the same time, the corresponding counterfactual score uplifts remain non-trivial—roughly 1–4.6 percentage points depending on annotator and evaluation model (Appendix C.4, Table 29)—so ranking stability does not remove the broader problem of biased absolute measurements.

6 Conclusion

We study how translation errors in machine-translated benchmarks affect multilingual LLM evaluation by combining span-level MQM-style error annotation with downstream performance modeling. On an EU20 subset with a professional human reference (225 items, nine target languages),

GPT-5.2 shows the highest agreement with humans, while xCOMET-XXL provides the strongest non-generative baseline. This ranking is supported not only by thresholded span matching, but also by threshold-free character-overlap and unitizing-reliability analyses, indicating that the main agreement result is not a metric artifact.

We further evaluate span localization on Span-ACES_{Ref} (1,407 items), where tolerant span metrics are substantially higher than classic ones, suggesting that many residual disagreements are boundary near-misses rather than complete detection failures. Finally, in regressions with English controls and explicit source-side issue indicators, target-side translation errors are consistently associated with lower translated accuracy (about 6–8 pp in the full model and 6–11 pp when restricting to items solved in English), while source-side issues have smaller and less stable effects. A counterfactual ranking analysis further shows that translation errors mainly shift absolute scores rather than relative system rankings: leaderboard orderings are usually preserved, but absolute performance estimates remain biased downward. Our results suggest that translated leaderboards can remain internally stable while still understating absolute multilingual performance. Future work should extend this analysis to finer-grained error types and multi-span phenomena, and develop more boundary-robust span evaluation and validation protocols for translation-aware benchmark release.

Limitations

First, high-quality span-annotated references remain scarce, particularly for lower-resource languages and for fine-grained MQM error inventories. Our EU20 human reference covers nine target languages and 225 items, with one professional annotator per language, so human-human inter-annotator agreement is not available. To partially address this, we report a shared-subset reliability analysis on the 225-segment EU20 subset using the available automatic annotators plus the human reference; however, we do not present this as a replacement for human-human IAA. Span-ACES_{Ref} provides an additional external benchmark, but it covers controlled phenomena rather than naturally occurring error distributions.

Second, span-level MQM annotation has no widely accepted cross-lingual gold standard, and expert annotation is costly; while we validate parts of our setup (human reference checks and a Span-ACES_{Ref} validation subset), residual subjectivity in span boundaries and in what counts as a target-side error is unavoidable.

Third, our performance-impact analysis is correlational: although we control for English correctness, fixed effects, and source-side issues, unobserved confounders and model/specification choices can still affect effect sizes, and some signals (especially source-side issues) are sparse and therefore estimated with greater uncertainty.

Fourth, automatic annotation quality depends on the chosen models and prompting setup. We mitigate this by comparing multiple LLM annotators, including another span-aware baseline (xCOMET-XXL), and by reporting robustness analyses across alternative agreement views; however, conclusions about the strongest annotator may still not transfer to other model families, prompting strategies, or annotation interfaces.

Finally, our results reflect a specific selection of benchmarks (EU20 tasks), MT system, source and target languages, and evaluated LLMs; translation artifacts and their downstream impact may differ in other languages, domains, genres, or translation pipelines.

Acknowledgments

This work was funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) through the OpenGPT-X project (no. 68GX21007D). The authors further acknowl-

edge support from the Federal Ministry of Research, Technology and Space of Germany (BMFTR), the Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus through ScaDS.AI (Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig), and the BMFTR under grant no. 01IS24077A. We also thank the Gauss Centre for Supercomputing e.V. for computing time on JUWELS at Jülich Supercomputing Centre (JSC), and ZIH at TU Dresden for infrastructure used in automatic evaluation computations. GNU parallel (Tange, 2024) supported parts of the data processing.

References

- Ashish Agrawal, Barah Fazili, and Preethi Jyothi. 2024. [Translation errors significantly impact low-resource languages in cross-lingual learning](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 319–329, St. Julian’s, Malta. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2020. [Translation artifacts in cross-lingual transfer learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7674–7684, Online. Association for Computational Linguistics.
- Irene Baucells, Javier Aula-Blasco, Iria de Dios-Flores, Silvia Paniagua Suárez, Naiara Perez, Anna Salles, Susana Sotelo Docio, Júlia Falcão, Jose Javier Saiz, Robiert Sepulveda Torres, Jeremy Barnes, Pablo Gamallo, Aitor Gonzalez-Agirre, German Rigau, and Marta Villegas. 2025. [IberoBench: A benchmark for LLM evaluation in Iberian languages](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10491–10519, Abu Dhabi, UAE. Association for Computational Linguistics.
- Frederic Blain, Chrysoula Zerva, Ricardo Rei, Nuno M. Guerreiro, Diptesh Kanojia, José G. C. de Souza, Beatriz Silva, Tânia Vaz, Yan Jingxuan, Fatemeh Azadi, Constantin Orasan, and André Martins. 2023. [Findings of the WMT 2023 shared task on quality estimation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 629–653, Singapore. Association for Computational Linguistics.
- Rochelle Choenni, Sara Rajae, Christof Monz, and Ekaterina Shutova. 2024. [On the evaluation practices in multilingual NLP: Can machine translation offer an alternative to human translations?](#) *Preprint*, arXiv:2406.14267.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#). *Preprint*, arXiv:1803.05457.

- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training Verifiers to Solve Math Word Problems](#). *Preprint*, arXiv:2110.14168.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Martin Fajcik, Martin Docekal, Jan Dolezal, Karel Ondrej, Karel Beneš, Jan Kapsa, Pavel Smrz, Alexander Polok, Michal Hradis, Zuzana Neverilova, Ales Horak, Radoslav Sabol, Michal Stefanik, Adam Jirkovsky, David Adamczyk, Petr Hyner, Jan Hula, and Hynek Kydlicek. 2025. [BenCzechMark: A Czech-centric multitask and multimetric benchmark for large language models with duel scoring mechanism](#). *Preprint*, arXiv:2412.17933.
- Manuel Faysse, Patrick Fernandes, Nuno M. Guerreiro, António Loison, Duarte M. Alves, Caio Corro, Nicolas Boizard, João Alves, Ricardo Rei, Pedro H. Martins, Antoni Bigata Casademunt, François Yvon, André F. T. Martins, Gautier Viaud, Céline Hudelot, and Pierre Colombo. 2025. [CroissantLLM: A truly bilingual French-English language model \(introducing FrenchBench\)](#). *Preprint*, arXiv:2402.00786.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, Errors, and Context: A Large-Scale Study of Human Evaluation for Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Markus Freitag, Nitika Mathur, Daniel Deutsch, Chi-Kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Frederic Blain, Tom Kocmi, Jiayi Wang, David Ifeoluwa Adelani, Marianna Buchicchio, Chrysoula Zerva, and Alon Lavie. 2024. [Are LLMs Breaking MT Metrics? Results of the WMT24 Metrics Shared Task](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 47–81, Miami, Florida, USA. Association for Computational Linguistics.
- Nuno M. Guerreiro, Ricardo Rei, Daan Van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. [xCOMET: Transparent Machine Translation Evaluation through Fine-grained Error Detection](#). *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Xu Huang, Zhirui Zhang, Xiang Geng, Yichao Du, Jijun Chen, and Shujian Huang. 2024. [Lost in the source language: How large language models evaluate the quality of machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3546–3562, Bangkok, Thailand. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Tom Kocmi, Vilém Zouhar, Eleftherios Avramidis, Roman Grundkiewicz, Marzena Karpinska, Maja Popović, Mrinmaya Sachan, and Mariya Shmatova. 2024. [Error Span Annotation: A Balanced Approach for Human Evaluation of Machine Translation](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1440–1453, Miami, Florida, USA. Association for Computational Linguistics.
- Klaus Krippendorff. 1995. [On the reliability of unitizing continuous data](#). *Sociological Methodology*, 25:47–76.
- Klaus Krippendorff, Yann Mathet, Stéphane Bouvry, and Antoine Widlöcher. 2016. [On the reliability of unitizing textual continua: Further developments](#). *Quality & Quantity*, 50(6):2347–2364.
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. [Okapi: Instruction-tuned Large Language Models in Multiple Languages with Reinforcement Learning from Human Feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327, Singapore. Association for Computational Linguistics.
- Alon Lavie, Greg Hanneman, Sweta Agrawal, Diptesh Kanojia, Chi-Kiu Lo, Vilém Zouhar, Frederic Blain, Chrysoula Zerva, Eleftherios Avramidis, Sourabh Deoghare, Archchana Sindhujan, Jiayi Wang, David Ifeoluwa Adelani, Brian Thompson, Tom Kocmi, Markus Freitag, and Daniel Deutsch. 2025. [Findings of the WMT25 shared task on automated translation evaluation systems: Linguistic diversity is challenging and references still help](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 436–483, Suzhou, China. Association for Computational Linguistics.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring How Models Mimic Human Falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Arle Lommel, Serge Gladkoff, Alan Melby, Sue Ellen Wright, Ingemar Strandvik, Katerina Gasova, Angelika Vaasa, Andy Benzo, Romina Marazzato Sparano,

- Monica Foresi, Johani Innis, Lifeng Han, and Goran Nenadic. 2024. [The multi-range theory of translation quality measurement: MQM scoring models and statistical quality control](#). In *Proceedings of the 16th Conference of the Association for Machine Translation in the Americas (Volume 2: Presentations)*, pages 75–94, Chicago, USA. Association for Machine Translation in the Americas.
- Arlé Richard Lommel, Aljoscha Burchardt, and Hans Uszkoreit. 2013. [Multidimensional quality metrics: A flexible system for assessing translation quality](#). In *Proceedings of Translating and the Computer 35*, London, UK. Aslib.
- Qingyu Lu, Liang Ding, Kanjian Zhang, Jinxia Zhang, and Dacheng Tao. 2025. [MQM-APE: Toward high-quality error annotation predictors with automatic post-editing in LLM translation evaluators](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5570–5587, Abu Dhabi, UAE. Association for Computational Linguistics.
- Christian M. Meyer, Margot Mieskes, Christian Stab, and Iryna Gurevych. 2014. [DKPro agreement: An open-source Java library for measuring inter-rater agreement](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 105–109, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Nikita Moghe, Arnisa Fazla, Chantal Amrhein, Tom Kocmi, Mark Steedman, Alexandra Birch, Rico Senrich, and Liane Guillou. 2025. [Machine Translation Meta Evaluation through Translation Accuracy Challenge Sets](#). *Computational Linguistics*, pages 1–65.
- Dan Nielsen. 2023. [ScandEval: A benchmark for Scandinavian natural language processing](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 185–201, Tórshavn, Faroe Islands. University of Tartu Library.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. [The LAMBADA dataset: Word prediction requiring a broad discourse context](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany. Association for Computational Linguistics.
- ChaeHun Park, Koanho Lee, Hyesu Lim, Jaeseok Kim, Junmo Park, Yu-Jung Heo, Du-Seong Chang, and Jaegul Choo. 2024. [Translation deserves better: Analyzing translation artifacts in cross-lingual visual question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5193–5221, Bangkok, Thailand. Association for Computational Linguistics.
- Jan Pfister and Andreas Hotho. 2024. [SuperGLEBer: German language understanding evaluation benchmark](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7904–7923, Mexico City, Mexico. Association for Computational Linguistics.
- Irene Plaza, Nina Melero, Cristina Pozo, Javier Conde, Pedro Reviriego, Marina Mayor-Rocher, and María Grandury. 2024. [Spanish and LLM benchmarks: Is MMLU lost in translation?](#) *Preprint*, arXiv:2406.17789.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David I. Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Wei-Yin Ko, Sebastian Ruder, Madeline Smith, Antoine Bosselut, Alice Oh, Andre F. T. Martins, Leshem Choshen, and 5 others. 2025. [Global MMLU: Understanding and Addressing Cultural and Linguistic Biases in Multilingual Evaluation](#). *Preprint*, arXiv:2412.03304v2.
- Ole Tange. 2024. [GNU Parallel 20240722 \('assange'\)](#).
- Klaudia Thellmann, Bernhard Stadler, Michael Fromm, Jasper Schulze Buschhoff, Alex Jude, Fabio Barth, Johannes Leveling, Nicolas Flores-Herr, Joachim Köhler, René Jäkel, and Mehdi Ali. 2024. [Towards Multilingual LLM Evaluation for European Languages](#). *Preprint*, arXiv:2410.08928.
- Taemin Yeom, Yonghyun Ryu, Yoonjung Choi, and Jinyeong Bak. 2025. [Tagged span annotation for detecting translation errors in reasoning LLMs](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 878–886, Suzhou, China. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a Machine Really Finish Your Sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Chrysoula Zerva, Frederic Blain, José G. C. De Souza, Diptesh Kanojia, Sourabh Deoghare, Nuno M. Guerreiro, Giuseppe Attanasio, Ricardo Rei, Constantin Orasan, Matteo Negri, Marco Turchi, Rajen Chatterjee, Pushpak Bhattacharyya, Markus Freitag, and André Martins. 2024. [Findings of the quality estimation shared task at WMT 2024: Are LLMs closing the gap in QE?](#) In *Proceedings of the Ninth Conference on Machine Translation*, pages 26–76, Miami, Florida, USA. Association for Computational Linguistics.

A MQM Annotation

A.1 Annotation Guide and Interface

This subsection summarizes the manual protocol used to create our human reference and documents the annotation interface. Our guidelines are based on the Multidimensional Quality Metrics (MQM) framework,⁴ using a reduced, task-focused typology. Concretely, we group error labels into two high-level dimensions—*Accuracy* (semantic faithfulness) and *Fluency/Style* (well-formedness and naturalness)—and annotate errors as spans. For each error, annotators highlight the erroneous span in the target translation and, where applicable, the corresponding source span to support span-level alignment analyses. We treat *Addition* and *Omission* as asymmetric cases: *Addition* is annotated only on the target side (no source span), whereas *Omission* is annotated only on the source side (no target span). Table 5 summarizes the sampled EU20 subset, and Tables 6 and 7 list the reduced MQM inventory used in the manual annotation.

Dataset	k/lang	Ex.	Tok.	Tok. min–max	Sent.
ARC	4	36	56.6	31–98	5.0
GSM8K	2	18	129.5	93–191	8.3
HellaSwag	9	81	166.2	112–253	12.3
MMLU	6	54	153.5	62–402	8.7
TruthfulQA	4	36	99.3	66–201	10.5

Table 5: Aggregated statistics for the manually annotated EU20 subset across nine languages. **k/lang** is the sampling quota per language, yielding **Ex.** total instances per dataset across all languages. **Tok.** and **Sent.** report the mean number of tokens and sentences per instance, respectively. **Tok. min–max** gives the observed token-length range within each dataset.

Label	Description
Addition	Adds content not supported by source
Omission	Drops source content
Mistranslation	Meaning is changed or incorrect
Under-translation	Meaning is too vague; nuance is lost
Over-translation	Adds unwarranted specificity or detail
Reordering	Changes attachment or meaning
Untranslated	Leaves a source fragment untranslated
Wrong language	Uses tokens from unintended language
Do-not-translate	Content that should remain unchanged

Table 6: Reduced MQM label set for *Accuracy* errors used in manual annotation.

Annotators assign one of two severity levels to each annotated span. *Major* errors change mean-

⁴<https://themqm.org/the-mqm-full-typology/>

Label	Description
Grammar	Grammatical error (agreement, tense, case)
Spelling	Spelling or character error
Punctuation	Punctuation error affecting readability
Inconsistent	Inconsistent terminology or naming
Awkward	Unnatural or non-idiomatic phrasing
Unintelligible	Output is not reliably interpretable

Table 7: Reduced MQM label set for *Fluency/Style* errors used in manual annotation.

ing, can mislead, or substantially harm understanding/faithfulness, whereas *Minor* errors largely preserve meaning and primarily affect fluency/style or constitute a limited local defect. If no label fits or annotators are uncertain, they mark the span as *Other/Unknown* (used sparingly).

System	Type	Reference
GPT-5.2	LLM	OpenAI model page
GPT-4o-mini	LLM	OpenAI model page
LLaMA-4 Scout 17B	LLM	Hugging Face model card
Mistral-Large-Instruct-2411	LLM	Hugging Face model card
xCOMET-XXL	QE / span baseline	Guerreiro et al. (2024b)

Table 8: Automatic annotators used in the appendix analyses. The four LLMs are evaluated throughout the appendix; xCOMET-XXL is additionally included whenever the auxiliary analysis covers the non-generative span baseline.

Figure 2 and Figure 3 illustrate the Argilla-based interface and the annotation workflow. Concretely, annotators proceeded as follows:

1. Decide whether the translation is error-free (*Yes/No/Unsure*).
2. If errors are present, mark erroneous target spans and assign an MQM label and severity.
3. Mark corresponding source spans where applicable; *Addition* is annotated only on the target side, whereas *Omission* is annotated only on the source side.
4. Provide a minimally post-edited corrected translation.
5. Answer a control question indicating whether all important errors were captured.
6. Optionally add comments for clarification.

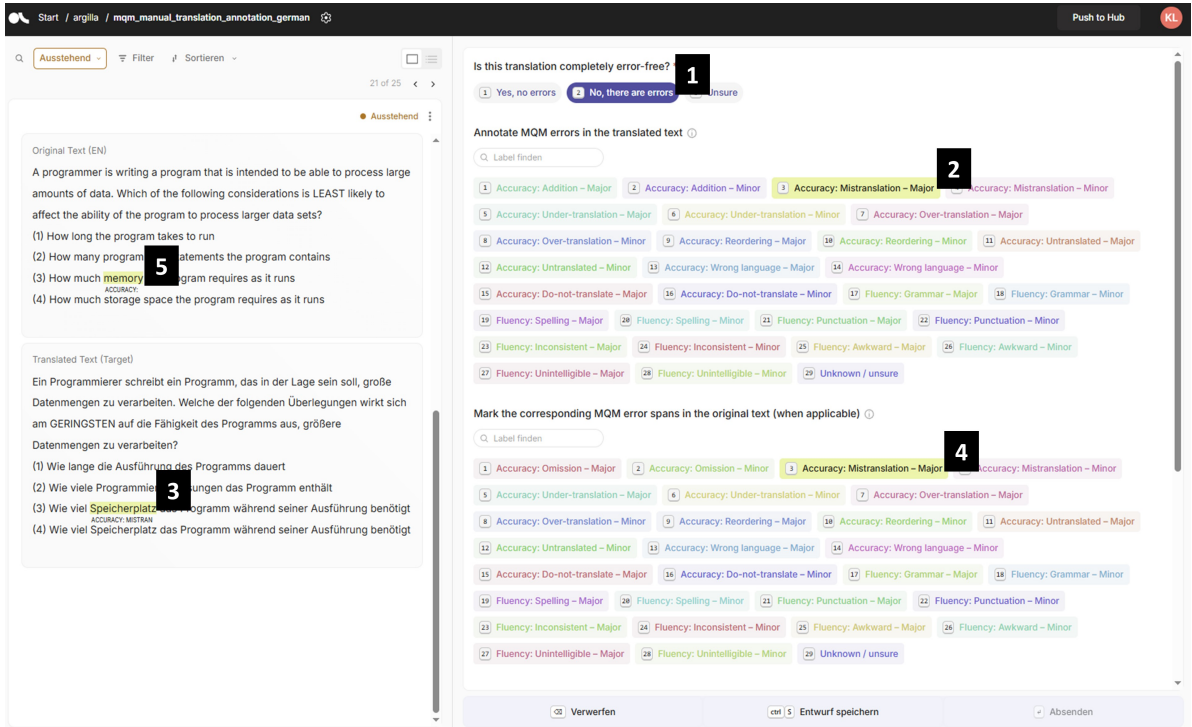


Figure 2: Argilla interface for span-based MQM annotation. The numbered markers indicate the main interaction elements: (1) error-free decision, (2) target-side MQM label selection, (3) target-span highlighting, (4) source-span highlighting, and (5) the source/context panel.

A.2 Metrics and Matching

We compare span sets using (i) a position-based overlap coefficient (OC) computed from target character offsets, and (ii) a string-based similarity metric (SIM) defined as raw character 3-gram Dice similarity, without text normalization. OC and SIM capture complementary notions of span agreement: OC is offset-based and rewards positional overlap, whereas SIM is text-based and remains applicable when offsets are missing or boundaries drift. Offsets are validated against the raw target string. If missing or inconsistent, we attempt to recover them via exact substring search and leave offsets missing in ambiguous cases. Not all auxiliary analyses in this appendix cover the same system set: where relevant, we report xCOMET-XXL alongside the LLM annotators, but some robustness analyses are restricted to the four LLM-based comparisons.

For character-offset spans a and b , we define the overlap coefficient as

$$\text{OC}(a, b) = \frac{|a \cap b|}{\min(|a|, |b|)},$$

i.e., the character-overlap length divided by the shorter span length. For OC-based matching, a gold span and a predicted span are matchable if $\text{OC} \geq \tau_{\text{OC}}$; in the main text we use $\tau_{\text{OC}} = 0.8$.

For span texts s_1 and s_2 , let $\mathcal{G}_3(s)$ denote the multiset of raw character trigrams of s . We define

$$\text{SIM}(s_1, s_2) = \frac{2 \sum_g \min(c_g(s_1), c_g(s_2))}{\sum_g c_g(s_1) + \sum_g c_g(s_2)},$$

i.e., Sørensen–Dice similarity over character trigrams, where $c_g(s)$ is the count of trigram g in s . For SIM-based matching, a gold span and a predicted span are matchable if $\text{SIM} \geq \tau_{\text{SIM}}$; in the main text we use $\tau_{\text{SIM}} = 0.6$.

OC matching requires valid target offsets. Spans without offsets cannot be matched and contribute to FP/FN. For SIM-based matching, we apply the metric to spans deduplicated by exact target-span text on both sides.

To compute precision/recall/F1, we perform greedy one-to-one matching: all matchable span pairs are sorted by score (descending) and selected if neither span has been matched before. For each comparison, we aggregate TP/FP/FN across samples to obtain micro-precision, micro-recall, and micro-F1. For sensitivity analysis, we additionally sweep thresholds over $\text{OC} \in \{0.7, 0.8, 0.9\}$ and $\text{SIM} \in \{0.4, 0.5, 0.6\}$.

To complement the human-reference comparison in the main text, Table 10 reports additional

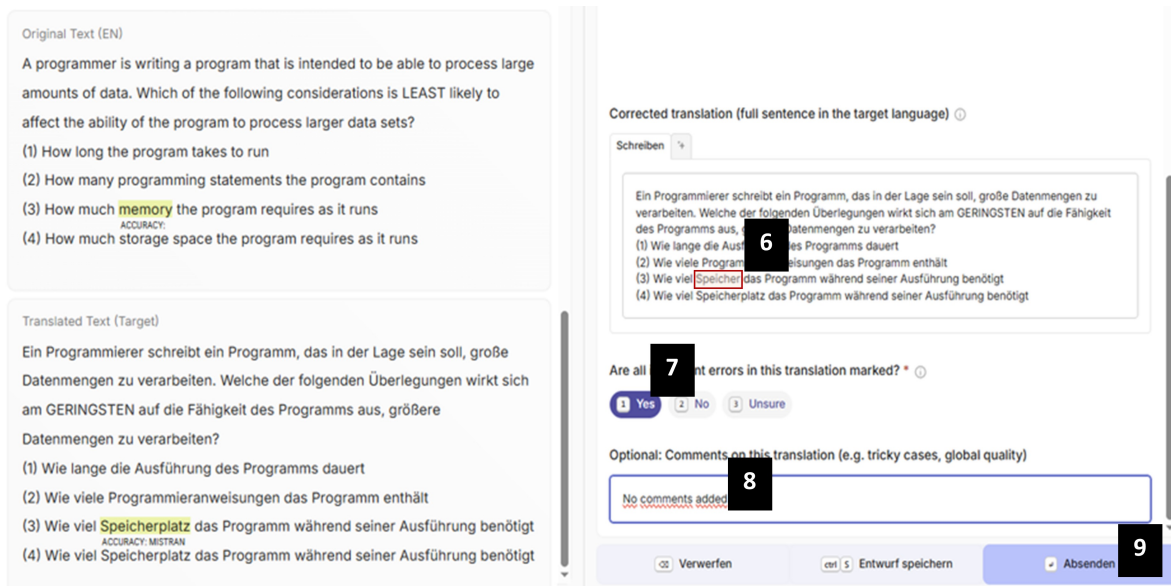


Figure 3: Argilla interface for post-editing and metadata. The numbered markers indicate (1) the minimally corrected translation field, (2) the completeness control question, (3) optional annotator comments, and (4) final submission.

Comparison	SIM Recall	SIM F1
Human vs GPT-5.2	.38	.45 (.30–.78)
Human vs Mistral-Large	.14	.17 (.10–.26)
Human vs GPT-4o-mini	.11	.15 (.09–.25)
Human vs LLaMA-4	.06	.10 (.02–.19)
GPT-5.2 vs Mistral-Large	.16	.17 (.09–.24)
GPT-5.2 vs GPT-4o-mini	.15	.21 (.11–.30)
GPT-5.2 vs LLaMA-4	.09	.12 (.03–.19)

Table 9: Span-level agreement based on string similarity (SIM) averaged over nine languages (25 items each). We report mean Span-Recall and Span-F1 across languages. Parentheses denote the min–max Span-F1 over languages. SIM uses raw character 3-gram Dice with threshold 0.6 and greedy one-to-one matching on spans deduplicated by exact target-span text.

OC-based pairwise span agreement using GPT-5.2 and xCOMET-XXL as reference annotators. Because Span-F1 is symmetric but recall is directional, these auxiliary views help contextualize agreement patterns beyond the human-reference setting.

The per-language values in Table 11 are descriptive and should be interpreted cautiously, given the small per-language sample size (25 items per language).

A.3 Worked Examples for OC and SIM

We illustrate the two span-matching criteria from §3 on concrete examples drawn from our annotations.

Example 1: OC with partial overlap. Consider a human span H and a candidate match M on the same target text:

(A) Human as reference		
Comparison	OC Recall	OC Span-F1
Human vs GPT-5.2	.44	.55
Human vs xCOMET-XXL	.32	.28
Human vs Mistral-Large	.17	.23
Human vs GPT-4o-mini	.09	.15
Human vs LLaMA-4	.08	.13
(B) GPT-5.2 as reference		
Comparison	OC Recall	OC Span-F1
GPT-5.2 vs xCOMET-XXL	.42	.26
GPT-5.2 vs Mistral-Large	.24	.25
GPT-5.2 vs GPT-4o-mini	.15	.21
GPT-5.2 vs LLaMA-4	.12	.17
(C) xCOMET-XXL as reference		
Comparison	OC Recall	OC Span-F1
xCOMET-XXL vs GPT-5.2	.19	.26
xCOMET-XXL vs Mistral-Large	.16	.22
xCOMET-XXL vs GPT-4o-mini	.08	.13
xCOMET-XXL vs LLaMA-4	.07	.11

Table 10: Extended OC-based pairwise span agreement on the EU20 human-reference subset. Span-F1 is symmetric across comparison direction, while recall is directional and depends on the chosen reference annotator.

H: heißen Zehentrenner nehmen und ihn auf Ihr Gesicht auftragen
M: einen heißen Zehentrenner nehmen und ihn

Lang	Metric	GPT-5.2 vs			Human vs			GPT-5.2
		G4o	L4	Mis	G4o	L4	Mis	
DA	OC	.23	.17	.25	.12	.13	.19	.48
DE	OC	.16	.14	.19	.16	.13	.20	.61
ET	OC	.20	.20	.24	.14	.14	.20	.57
FR	OC	.17	.12	.24	.10	.09	.19	.34
HU	OC	.28	.23	.28	.15	.12	.22	.44
IT	OC	.21	.18	.36	.16	.14	.29	.53
LT	OC	.23	.05	.18	.16	.04	.25	.49
RO	OC	.09	.18	.25	.08	.14	.21	.68
SL	OC	.30	.26	.31	.25	.23	.32	.78
DA	SIM	.23	.08	.10	.09	.08	.15	.30
DE	SIM	.14	.12	.19	.16	.10	.14	.39
ET	SIM	.20	.13	.13	.19	.09	.15	.44
FR	SIM	.18	.13	.19	.09	.06	.10	.34
HU	SIM	.16	.14	.20	.10	.10	.18	.49
IT	SIM	.20	.12	.24	.18	.10	.25	.51
LT	SIM	.23	.03	.11	.16	.02	.26	.39
RO	SIM	.11	.17	.12	.09	.09	.12	.42
SL	SIM	.30	.19	.24	.25	.19	.22	.78

Table 11: Span- F_1 by language for OC and SIM for the four LLM annotators and the human-reference comparison. These values are descriptive rather than inferential, given the small per-language sample size. G4o = GPT-4o-mini, L4 = LLaMA-4, Mis = Mistral-Large.

Lang	G4o	L4	Mis	GPT-5.2	Human
DA	.00	.07	.10	.00	.04
DE	.17	.09	.20	.06	.13
ET	.09	.09	.18	.04	.05
FR	.10	.17	.09	.07	.07
HU	.25	.15	.14	.09	.04
IT	.15	.22	.18	.13	.13
LT	.00	.06	.15	.04	.09
RO	.04	.07	.12	.03	.07
SL	.14	.29	.09	.12	.11

Table 12: Source-overlap rate by language for the four LLM annotators and the human reference (Span-Recall-style micro aggregation): fraction of target error spans (with valid offsets) that overlap target regions linked to GPT-5.2 source-anomaly annotations (OC threshold 0.8).

With target character offsets $H = [811, 871)$ and $M = [805, 845)$, the span lengths and overlap evaluate to

$$|H| = 60, \quad |M| = 40, \quad |H \cap M| = 34,$$

yielding

$$\text{OC}(H, M) = \frac{|H \cap M|}{\min(|H|, |M|)} = \frac{34}{40} = 0.85.$$

Since $0.85 \geq 0.8$, the spans are matchable under the OC criterion used in the main text. A readable view of the same overlap is:

Human: ... einen [H: heißen
Zehentrenner nehmen und ihn auf
Ihr Gesicht auftragen] ...
Match: ... [M: einen heißen
Zehentrenner nehmen und ihn] auf
Ihr Gesicht auftragen ...

This example shows that OC can count a genuine partial overlap as a match even when the span boundaries differ substantially.

Example 2: SIM with character-trigram Dice.

Let A and B be two candidate spans over the same sentence,

A: sich stärker verfestigen
B: die Mitglieder der Gruppe A sich stärker verfestigen

and let $\mathcal{G}_3(s)$ denote the multiset of raw character trigrams of a string s . Then $|\mathcal{G}_3(A)| = 22$ and $|\mathcal{G}_3(B)| = 50$. Because A occurs as a contiguous substring of B , every trigram in A also occurs in B , so the trigram-overlap count equals

$$\sum_g \min(c_g(A), c_g(B)) = 22,$$

where $c_g(\cdot)$ is the count of trigram g . The Dice similarity is therefore

$$\text{SIM}(A, B) = \frac{2 \cdot 22}{22 + 50} = \frac{44}{72} \approx 0.611.$$

Since $0.611 \geq 0.6$, the two spans are matchable under the SIM criterion. Example shared trigrams include `sic`, `ich`, `stä`, `ver`, `tig`, and `gen`. This example shows that SIM can recover boundary-expanded textual matches even without relying on offsets.

A.4 Systems-only Multi-Rater Reliability

To make span-level annotation variability explicit, we additionally report systems-only multi-rater reliability on the shared subset using Krippendorff’s unitized alpha (α_u), which is designed for free-span annotation settings. Because the human reference contains one annotator per language, human–human unitizing agreement is unavailable for this subset. Table 13 summarizes the resulting reliability estimates.

Metric	Value
Krippendorff α (nominal, any-error)	0.037
Krippendorff α_u (unitizing, target-span localization)	0.128
α_u range across languages (9×25 segments)	0.075–0.186

Table 13: Systems-only multi-rater reliability on the shared subset (225 segments, 9 languages, 5 raters). α_u is the global systems-only unitizing reliability for target-text error spans, measuring agreement on span locations and boundaries. The reported range is the min–max of the nine per-language α_u values and is not a confidence interval.

A.5 Threshold-free Character-overlap Metrics

To complement thresholded OC/SIM span matching, we also report threshold-free character-level overlap metrics. Char-F1 is a binary character-overlap measure that ignores severity labels. Char-F1w is severity-aware and follows a WMT25-style overlap logic, giving full credit to severity-matched overlaps and partial credit (0.5) to severity-mismatched overlaps, with critical grouped with major. Table 14 reports the corresponding character-level results, and Table 15 reports complementary threshold-sweep robustness ranges for OC- and SIM-based matching.

xCOMET-XXL tends to produce more numerous and shorter spans, which helps explain its more recall-oriented behavior under greedy one-to-one matching. Table 16 summarizes these annotator statistics.

A.6 Bootstrap Confidence Intervals

To quantify uncertainty despite the relatively small per-language sample size (25 items per language), we compute pooled stratified bootstrap confidence intervals over items for the four LLM annotators. This auxiliary analysis excludes xCOMET-XXL,

System	Any-error F1	Char-F1w [95% CI]	Char-F1 [95% CI]
GPT-5.2	.95	.42 [.383, .469]	.48 [.437, .529]
xCOMET-XXL	.88	.20 [.169, .230]	.26 [.223, .301]
Mistral-Large	.88	.21 [.173, .242]	.28 [.235, .317]
GPT-4o-mini	.77	.17 [.129, .217]	.21 [.163, .261]
LLaMA-4	.63	.11 [.078, .142]	.14 [.104, .183]

Table 14: Threshold-free character-overlap results on the shared subset. Any-error F1 is segment-level error/no-error agreement. Char-F1 is threshold-free character-level span-overlap F1 on the binary target-side error mask (global micro, 95% bootstrap CI via segment resampling). Char-F1w is a severity-aware WMT25-style character-overlap variant with full credit for severity-matched overlaps and partial credit for severity mismatches. Point estimates rounded to two decimals; CI bounds to three.

System	OC F1 range	SIM F1 range
GPT-5.2	.585–.592	.426–.506
xCOMET-XXL	.294–.304	.154–.229
Mistral-Large	.267–.271	.157–.208
GPT-4o-mini	.169–.171	.122–.148
LLaMA-4	.138–.144	.074–.114

Table 15: Threshold-sweep robustness on the shared subset. OC/SIM ranges denote min–max micro-F1 across threshold sweeps (OC: 0.7/0.8/0.9; SIM: 0.4/0.5/0.6). Ranges reported at three decimals to preserve the sweep variation.

Annotator	Spans / sample	Median span len.	Coverage (union)
xCOMET-XXL	6.95	7.56	0.16
Human	6.05	9.94	0.11
GPT-5.2	3.21	19.22	0.10
Mistral-Large	2.74	21.17	0.11
GPT-4o-mini	1.40	20.33	0.04
LLaMA-4	1.29	23.17	0.04

Table 16: Annotator statistics (mean over languages): average number of spans per sample, median span length in characters, and union coverage of target characters. xCOMET-XXL is markedly more fine-grained than the other automatic annotators, which helps explain its relatively higher recall but lower precision under greedy one-to-one matching.

since the bootstrap runs were prepared only for the human-vs.-LLM comparisons. In each replicate, we resample 25 items with replacement per language (9 languages; 225 pooled items total), micro-aggregate TP/FP/FN across resampled items, and recompute span-level precision, recall, and F1. We use 2,500 bootstrap iterations and report percentile-based 95% confidence intervals.

Table 17 reports the resulting pooled confidence intervals.

Comparison Human vs.	OC micro-F_1 [95% CI]	SIM micro-F_1 [95% CI]
GPT-5.2	.526 [.484, .568]	.437 [.397, .482]
Mistral-Large	.228 [.198, .259]	.176 [.149, .204]
GPT-4o-mini	.141 [.116, .167]	.144 [.119, .171]
LLaMA-4	.120 [.097, .145]	.086 [.066, .107]

Table 17: Pooled bootstrap confidence intervals for span-level agreement against the human reference, computed for the four LLM annotators. CIs are based on 2,500 stratified bootstrap replicates over items (25 items per language, sampled with replacement; 9 languages, 225 pooled items per replicate). Scores are micro- F_1 values obtained by aggregating TP/FP/FN across resampled items.

A.7 Additional Qualitative Example Cases

To complement the quantitative agreement tables, we include additional target-side (Table 18) example cases drawn from the manually inspected EU20 subset. The source-side cases (Section 1; Table 1) are diagnostic only and are not counted as target-side translation errors. They motivate our separate treatment of source-linked confounds, while the target-side cases illustrate how translation errors can change the construct or answerability of benchmark items.

Source	EN snippet	Translation snippet	Distortion type and effect
MMLU / FR prof_medicine/150	“A 14-year-old boy ... He is embarrassed because he has grown breasts ... this patient’s condition”	“... <i>il s’est fait pousser des seins</i> ... <i>cette patiente</i> ”	Gender flip / clinical meaning shift: the male patient becomes female in the translation, changing the diagnostic interpretation from male gynecomastia to female breast development.
MMLU / SL moral_scenarios/349	Option (3): “Not wrong, Wrong” Option (4): “Not wrong, Not wrong”	Option (3): “ <i>Ni napačno, ni napačno</i> ” Option (4): “ <i>Ni napačno, ni napačno</i> ”	Structural corruption: two distinct answer options collapse into the same translation, making the item partially unsolvable.
MMLU / FR hs_world_hist/84	“600 C.E.–1450 C.E.”	“ <i>600 à 1450 avant notre ère</i> ”	Time-period inversion: C.E. becomes B.C.E., shifting the historical scope by roughly two millennia.
MMLU / RO prof_law/115	“The estates of the two decedents sued the company”	“ <i>Avea celor doi decedați a dat în judecată compania</i> ”	Legal-semantic shift: legal estates (the juridical entities representing the decedents) becomes personal wealth/property, removing the entity that can sue.
MMLU / SL hs_world_hist/80	“in conformity without resolve”	“ <i>V skladu z odločitvijo</i> ” (with decision)	Polarity reversal: the rhetorical meaning flips from “without resolve” to “with decision”, changing the interpretation of the historical statement.
MMLU / SL, ET formal_logic/69, /84	“(Cx . Ox)”	“(Cx - Ox)”	Symbol corruption: the logical conjunction symbol is replaced by a hyphen, making the formal expression syntactically invalid.
MMLU / DA prof_law/1334	“200-acre tract”	first “ <i>200 acres</i> ”, later “ <i>200 hektar</i> ”	Measurement distortion: the translation inconsistently renders acres and then inflates the area to 200 hectares (~2.5× larger), altering the legal facts of the case.
GSM8K / DA main/test/150	“Tim lives 2 miles away from the school”	“ <i>Tim bor 3 km væk fra skolen</i> ”	Number + unit change: both the distance unit and the quantity are changed, so the translated math problem no longer matches the source.
GSM8K / RO main/test/659	“A pound of almonds costs \$10 while a pound of walnuts costs \$15”	“ <i>Un kilogram de migdale costă 10 dolari ... un kilogram de nuci costă 15 dolari</i> ”	Unit distortion: pound becomes kilogram throughout, changing the price basis and invalidating the original calculation chain.
GSM8K / ET main/test/130	“Nissa hires 60 seasonal workers to play elves in her department store’s Santa village”	“ <i>mängivad tontu</i> ”	Concept mistranslation: Estonian <i>tontu</i> means ghosts/goblins rather than Santa’s elves (<i>päkapikud</i>), distorting the cultural concept being tested.

Table 18: Additional target-side translation cases from our example collection, beyond the four main-text examples. The selected cases illustrate diverse distortion types, including structural corruption, time-period inversion, legal-semantic shifts, polarity reversals, symbol corruption, and unit or concept changes.

B Span-ACES_{Ref}

This appendix expands the construction details for Span-ACES_{Ref} and summarizes its dataset composition and phenomenon-to-MQM mapping. A key limitation of the original SPAN-ACES setup is that each item targets exactly one introduced phenomenon, so the released span corresponds only to that targeted error. If an incorrect translation contains additional issues beyond the targeted one, those extra issues remain unlabeled and can therefore count as false positives during span evaluation. Span-ACES_{Ref} reduces this source of noise by projecting the targeted good→incorrect edit onto the human reference and discarding cases where that projection is not unique or not reliable.

B.1 Construction, Mapping, and Validation Details

Operationally, we construct Span-ACES_{Ref} as follows:

1. Start from a SPAN-ACES item consisting of a human reference, a *good* translation, and an *incorrect* translation.
2. Compute a case-sensitive token diff between the good and incorrect translations.
3. Keep only items with a single *contentful diff*, i.e., one non-empty contiguous token-level edit that changes lexical content rather than only punctuation or formatting.
4. Project this edit into the human reference by locating the corresponding reference-side token sequence and use that projected span as the gold span.
5. Discard items with multiple diffs, no reference-side match, or ambiguous repeated matches in the reference.

This procedure is intentionally conservative: *projection* here does not attempt full semantic alignment, but a simple grounded transfer of the targeted edit into the human reference whenever the match is unique enough to support reliable span evaluation. In this sense, the original SPAN-ACES data can be viewed as a targeted one-error-per-item resource, whereas Span-ACES_{Ref} converts that targeted edit into a cleaner reference-side span benchmark.

Table 19 summarizes the final dataset composition. For consistency with our EU20 annotation

Split	Items	Langs.	MQM types	Phenomena
Accuracy	1,155	20	4	14
Fluency/Style	252	1	1	4
Total	1,407	20	5	18

Table 19: Final composition of Span-ACES_{Ref}. The *Fluency/Style* portion is German-only and consists of four anaphoric phenomena; the *Accuracy* portion spans the 20 EU20 target languages.

setup, we map SPAN-ACES phenomena to MQM types using the ACES authors’ released mapping⁵ and collapse the resulting MQM types into two coarse categories: *Accuracy* and *Fluency/Style*. In total, Span-ACES_{Ref} contains 1,155 *Accuracy* instances spanning the 20 EU20 target languages and 252 *Fluency/Style* instances (German only), for 1,407 items overall. Table 20 lists the phenomenon-to-MQM mapping used in Span-ACES_{Ref}.

MQM type	ACES phenomenon
<i>Accuracy</i>	
mistranslation	pleonastic_it:substitution coreference-based-on-commonsense hallucination-date-time overly-literal-vs-ref-word overly-literal-vs-explanation overly-literal-vs-synonym real-world-knowledge-hypernym-vs-distractor real-world-knowledge-entailment real-world-knowledge-synonym-vs-antonym ordering-mismatch
untranslated	untranslated-vs-ref-word untranslated-vs-synonym
no-translate	do-not-translate
addition	addition
<i>Fluency/Style</i>	
grammar	anaphoric_intra_non-subject_it:substitution anaphoric_intra_subject_it:substitution anaphoric_intra_they:substitution anaphoric_group_it-they:substitution

Table 20: Phenomenon-to-MQM mapping used in Span-ACES_{Ref}. Accuracy covers 4 MQM types and 14 distinct phenomena; Fluency/Style covers 1 MQM type (*grammar*) and 4 anaphoric phenomena.

To sanity-check the transformation pipeline, we manually reviewed a stratified subset of 178 Span-ACES_{Ref} records using a fixed checklist: (i) whether the projected edit matches the original good→incorrect change, (ii) whether the intended error type is

⁵github.com/EdinburghNLP/ACES

Stage		Illustrative toy example
Good translation		Der Ausschuss genehmigte den Vorschlag.
Incorrect translation		Der Ausschuss lehnte den Vorschlag ab .
Case-sensitive token diff		Single contentful edit: <code>genehmigte</code> → <code>lehnte...ab</code>
Human reference	refer-	Der Ausschuss genehmigte den Vorschlag.
Projected span	gold	<code>genehmigte</code>
Decision		Keep the item: one contentful diff, unique reference-side match.

Table 21: Toy illustration of the good→incorrect diff and projection step used to derive a gold span in Span-ACES_{Ref}. The example is schematic and explains the mechanics only. If the human reference did not contain a unique reference-side counterpart of the targeted edit, the item would be discarded.

preserved in the reference context, and (iii) whether the projection introduces side effects elsewhere in the sentence. We applied the same checklist with GPT-5.2 and compared final `pass/fail` verdicts. The two assessments agree on 165/178 items (0.93), indicating that most projected cases are straightforward; the remaining disagreements are mostly borderline cases such as type drift after projection or subtle contextual side effects.

For evaluation, we report both classic and tolerant span metrics. Classic matching requires exact equality between a predicted span and the gold span. Tolerant matching is gold-centered and counts a prediction as correct if it contains the gold span and optionally up to k tokens of boundary slack on either side. We aggregate scores first within each phenomenon and then per MQM category using two weighting schemes: **mean (N)**, which weights phenomena by sample count, and **mean (cap)**, which caps each phenomenon at $C=25$ to reduce domination by a few large phenomena.

C Performance Analysis

This appendix collects supporting material for Section 5. Section C.1 summarizes the paired regression dataset, Section C.2 lists the downstream evaluation models, Section C.3 reports the detailed pooled fixed-effects regressions, and Section C.4 provides an exploratory counterfactual ranking diagnostic. Unless noted otherwise, the appendix tables below correspond directly to the confirmatory analyses discussed in the main text.

C.1 Regression Dataset

Table 22 summarizes the paired regression dataset used in Section 5. The paired subset contains 651 unique (item, target-language) pairs and 4,993 translated observations across evaluation models after excluding HellaSwag and TruthfulQA. Besides translated and English correctness counts, we also report how often source-side and target-side annotation lists are empty or non-empty; these annotation counts are over (pair, annotator) records.

Statistic	MMLU	GSM8K	ARC	Total
N_{pair}	237	215	199	651
N_{obs}	1896	1505	1592	4993
$y=1$	1127	847	966	2940
$y=0$	769	658	626	2053
$y^{EN}=1$	1373	1017	1255	3645
$y^{EN}=0$	523	488	337	1348
SET= \emptyset	814	673	722	2209
SET $\neq\emptyset$	134	187	74	395
TET= \emptyset	472	347	478	1297
TET $\neq\emptyset$	476	513	318	1307

Table 22: Dataset statistics for the regression dataset used in Section 5 (HellaSwag and TruthfulQA excluded). N_{pair} counts unique (item, target-language) pairs and N_{obs} translated observations across evaluation models. y and y^{EN} denote translated versus English correctness (counts over N_{obs}). SET/TET indicate empty vs. non-empty source-side and target-side annotation lists, respectively; these counts are over (pair, annotator) records.

C.2 Evaluation Models

Table 23 lists the eight multilingual instruction-tuned LLMs used as downstream evaluation models in Section 5. We report short names in the main text and the corresponding Hugging Face model identifiers here for reproducibility. The additional xCOMET-XXL system used in the regression pipeline is a span annotator baseline rather than a

downstream evaluation model and is therefore not included in this table.

Short name	HF model identifier	#Params
Aya	CohereLabs/aya-expanse-32b	32.3B
Command-A	CohereLabs/c4ai-command-a-03-2025	111B
Gemma	google/gemma-3-27b-it	27.4B
Mistral	mistralai/Mistral-Small-3.1-24B-Instruct-2503	24B
Pharia	Aleph-Alpha/Pharia-1-LLM-7B-control-aligned-hf	7.0B
Phi	microsoft/Phi-4-multimodal-instruct	5.6B
Qwen	Qwen/Qwen2.5-32B	32.8B
Salamandra	BSC-LT/salamandra-7b-instruct	7.8B

Table 23: Downstream evaluation models used in the performance analysis. The table lists the short names used in the paper, the corresponding Hugging Face model identifiers, and approximate parameter counts.

C.3 Detailed Regression Results

Tables 24–25 report the pooled fixed-effects regressions underlying Section 5. Tables 24 and 26 provide the same four LLM annotator models in log-odds and odds-ratio form, respectively. The corresponding average marginal effects (AMEs) are the most directly interpretable view and are reported separately per specification: Table 27 covers the full model (Spec. A, with English correctness control), and Table 25 covers the subset restricted to items solved in English (Spec. B). Both AME tables include xCOMET-XXL as an additional span annotator baseline.

C.4 Counterfactual Ranking Analysis

This subsection reports an exploratory transparency analysis for the question of whether translation errors affect only absolute translated scores or also downstream leaderboard orderings. Using the fitted regression models, we predict each evaluation model’s translated accuracy under the counterfactual setting $T=0$ (no target-side translation errors) and compare the resulting ranking to the observed ranking. Table 28 reports bootstrapped rank correlations for three representative annotators (Mistral-Large, LLaMA-4, and GPT-5.2), and Table 29 reports the per-model counterfactual uplift Δ , defined as the predicted increase in translated accuracy (in percentage points) when moving from the observed setting to $T=0$.

Annotator	Spec.	coef(T)	coef(S)
GPT-4o-mini	A	-.40 [-.62, -.19]	-.21 [-.45, .04]
	$A_{\neg S}$	-.41 [-.64, -.20]	–
	B	-.58 [-.84, -.34]	-.27 [-.55, .01]
	$B_{\neg S}$	-.60 [-.86, -.36]	–
GPT-5.2	A	-.41 [-.64, -.19]	-.21 [-.45, .03]
	$A_{\neg S}$	-.42 [-.66, -.20]	–
	B	-.56 [-.86, -.31]	-.28 [-.57, -.01]
	$B_{\neg S}$	-.57 [-.86, -.31]	–
LLaMA-4	A	-.44 [-.72, -.16]	-.21 [-.45, .04]
	$A_{\neg S}$	-.45 [-.73, -.20]	–
	B	-.61 [-.95, -.29]	-.25 [-.54, .03]
	$B_{\neg S}$	-.64 [-.97, -.34]	–
Mistral-Large	A	-.33 [-.58, -.10]	-.22 [-.47, .03]
	$A_{\neg S}$	-.35 [-.58, -.12]	–
	B	-.38 [-.66, -.12]	-.29 [-.58, -.01]
	$B_{\neg S}$	-.40 [-.69, -.13]	–

Table 24: Logit coefficients (log-odds) for translation errors T and source-side issues S , with 95% block-bootstrap confidence intervals. Spec. A includes English correctness y^{EN} ; Spec. B restricts to $y^{EN}=1$. $A_{\neg S}$ and $B_{\neg S}$ omit S . All models include fixed effects for target language, dataset, and evaluation model.

Annotator	Spec.	AME(T)	AME(S)
GPT-4o-mini	B	-9.64 [-14.05, -5.49]	-4.34 [-8.75, 0.19]
	$B_{\neg S}$	-10.01 [-14.52, -5.86]	–
GPT-5.2	B	-8.98 [-13.35, -4.93]	-4.64 [-9.19, -0.10]
	$B_{\neg S}$	-9.26 [-13.46, -5.03]	–
LLaMA-4	B	-10.71 [-16.56, -4.91]	-4.15 [-8.66, 0.42]
	$B_{\neg S}$	-11.23 [-17.10, -5.57]	–
Mistral-Large	B	-6.12 [-10.30, -1.86]	-4.74 [-9.29, -0.20]
	$B_{\neg S}$	-6.48 [-10.74, -2.18]	–
xCOMET-XXL	B	-6.20 [-10.88, -1.57]	-4.91 [-9.45, -0.51]
	$B_{\neg S}$	-6.48 [-11.06, -1.77]	–

Table 25: Average marginal effects (AMEs; probability points) of target-side translation errors T and source-side issues S on correctness in translation under **Spec. B** (subset with $y^{EN}=1$, i.e., items the model solves correctly in English). $B_{\neg S}$ omits S . Point estimates are shown with 95% block-bootstrap CIs below. All models include fixed effects for target language, dataset, and evaluation model.

Annotator	Spec.	OR(T)	OR(S)
GPT-4o-mini	A	.67 [.54, .83]	.81 [.64, 1.04]
	$A_{\neg S}$.66 [.53, .82]	–
	B	.56 [.43, .71]	.77 [.58, 1.01]
	$B_{\neg S}$.55 [.42, .70]	–
GPT-5.2	A	.67 [.53, .83]	.81 [.64, 1.03]
	$A_{\neg S}$.66 [.52, .82]	–
	B	.57 [.43, .74]	.75 [.56, .99]
	$B_{\neg S}$.56 [.43, .73]	–
LLaMA-4	A	.65 [.49, .85]	.81 [.63, 1.04]
	$A_{\neg S}$.63 [.48, .82]	–
	B	.54 [.39, .75]	.78 [.58, 1.03]
	$B_{\neg S}$.53 [.38, .71]	–
Mistral-Large	A	.72 [.56, .90]	.80 [.63, 1.03]
	$A_{\neg S}$.71 [.56, .88]	–
	B	.69 [.52, .89]	.75 [.56, .99]
	$B_{\neg S}$.67 [.50, .88]	–

Table 26: Odds ratios (OR) for translation errors T and source-side issues S , with 95% block-bootstrap confidence intervals. Spec. A includes English correctness y^{EN} ; Spec. B restricts to $y^{EN}=1$. $A_{\neg S}$ and $B_{\neg S}$ omit S . All models include fixed effects for target language, dataset, and evaluation model.

Annotator	Spec.	AME(T)	AME(S)
GPT-4o-mini	A	-6.74 [-10.44, -3.17]	-3.50 [-7.46, 0.61]
	$A_{\neg S}$	-6.99 [-10.62, -3.22]	–
GPT-5.2	A	-6.76 [-10.61, -3.11]	-3.54 [-7.47, 0.57]
	$A_{\neg S}$	-7.01 [-10.80, -3.26]	–
LLaMA-4	A	-7.51 [-12.47, -2.70]	-3.47 [-7.56, 0.62]
	$A_{\neg S}$	-7.83 [-12.50, -3.33]	–
Mistral-Large	A	-5.59 [-9.50, -1.69]	-3.69 [-7.66, 0.44]
	$A_{\neg S}$	-5.82 [-9.47, -2.04]	–
xCOMET-XXL	A	-4.60 [-9.07, 0.00]	-3.82 [-7.69, 0.19]
	$A_{\neg S}$	-4.82 [-9.10, -0.44]	–

Table 27: Average marginal effects (AMEs; probability points) of target-side translation errors T and source-side issues S on correctness in translation under **Spec. A** (full model, controlling for English correctness y^{EN}). $A_{\neg S}$ omits S . Point estimates are shown with 95% block-bootstrap CIs below. All models include fixed effects for target language, dataset, and evaluation model.

Metric	Mistral-Large	LLaMA-4	GPT-5.2
Spearman ρ	0.9992	0.9995	0.9995
95% CI	[0.9762, 1.0]	[1.0, 1.0]	[1.0, 1.0]
Kendall τ	0.9976	0.9984	0.9986
95% CI	[0.9286, 1.0]	[1.0, 1.0]	[1.0, 1.0]
Boot	2500 / 0	2500 / 0	2500 / 0

Table 28: Rank correlation between observed evaluation-model rankings and counterfactual rankings predicted under $T=0$. Correlations are bootstrapped over underlying items ($B = 2500$). “Boot” reports successful / failed bootstrap runs. Near-perfect correlations indicate that the relative ordering of evaluation models is highly stable under the $T=0$ intervention in this paired subset.

(a) Annotator = Mistral-Large		
Eval model	Δ mean	95% CI
Phi-4-multimodal-instr.-6B	4.57	[1.44, 7.71]
aya-expanse-32b-Instr.	4.39	[1.39, 7.37]
Salamandra-7b-instr.	4.29	[1.32, 7.33]
Qwen-2.5-32B-Instr.	4.13	[1.30, 6.96]
Mistral-Small-3.1-24B-Instr.-2503	4.10	[1.29, 6.86]
c4ai-command-a-03-111B-Instr.	3.85	[1.21, 6.44]
Pharia-1-LLM-7B-contr.-aligned	3.84	[1.18, 6.59]
Gemma-3-27B-Instr.	3.59	[1.12, 6.01]
(b) Annotator = LLaMA-4		
Eval model	Δ mean	95% CI
Qwen-2.5-32B-Instr.	1.36	[-0.21, 2.91]
Phi-4-multimodal-instr.-6B	1.27	[-0.20, 2.73]
aya-expanse-32b-Instr.	1.27	[-0.19, 2.73]
Mistral-Small-3.1-24B-Instr.-2503	1.20	[-0.18, 2.60]
Salamandra-7b-instr.	1.18	[-0.19, 2.50]
c4ai-command-a-03-111B-Instr.	1.14	[-0.17, 2.46]
Gemma-3-27B-Instr.	1.09	[-0.16, 2.36]
Pharia-1-LLM-7B-contr.-aligned	1.06	[-0.17, 2.24]
(c) Annotator = GPT-5.2		
Eval model	Δ mean	95% CI
Phi-4-multimodal-instr.-6B	2.45	[-0.46, 5.44]
aya-expanse-32b-Instr.	2.36	[-0.45, 5.22]
Salamandra-7b-instr.	2.25	[-0.42, 5.03]
Qwen-2.5-32B-Instr.	2.25	[-0.42, 4.96]
Mistral-Small-3.1-24B-Instr.-2503	2.22	[-0.42, 4.89]
c4ai-command-a-03-111B-Instr.	2.09	[-0.39, 4.55]
Pharia-1-LLM-7B-contr.-aligned	2.01	[-0.38, 4.48]
Gemma-3-27B-Instr.	1.95	[-0.37, 4.32]

Table 29: Per-model counterfactual uplift Δ (percentage points), defined as the predicted increase in translated accuracy when moving from the observed setting to the counterfactual setting $T=0$. Confidence intervals are bootstrapped over underlying items ($B = 2500$). Although the ranking changes are negligible, the corresponding score uplifts remain non-trivial, indicating that translation errors still bias absolute multilingual performance estimates downward.

Taken together, the near-perfect rank correlations and non-zero score uplifts suggest that, in this paired subset, translation errors behave mostly like an approximately shared downward shift across systems: relative leaderboard orderings are largely preserved, but absolute translated scores remain biased downward. This is why we treat ranking stability as a useful diagnostic rather than as evidence that translation quality is unimportant.