

SalamandraTA at WMT 2026 Terminology Shared Task: Hard Examples Are Better Teachers

Xixian Liao

Barcelona Supercomputer Center
xixian.liao@bsc.es

Maite Melero

Barcelona Supercomputer Center
maite.melero@bsc.es

Abstract

Terminology-aware translation asks for more than a correct translation: the output must use the exact terms a glossary prescribes. The standard recipe, fine-tuning on glossary-annotated translation pairs, hides an inefficiency: for most examples the glossary prescribes exactly what the model would have produced anyway, so they teach nothing about following a glossary. We therefore keep only the examples where the model’s own translation contradicts the glossary. In a controlled study at fixed data volume, this selection alone raises term accuracy from 78.7% to 89.9%. The filtered data, built by a two-way synthetic pipeline on open models, is part of the instruction-tuning mixture of our public release SALAMANDRATA-7B-INSTRUCT v3.0, which, used exactly as released and wrapped in a document-level inference pipeline, forms the BSC submission to the WMT26 Terminology Shared Task Track 1. At the official WMT26 evaluation, our system achieves 94.2% term success at 74.6 chrF++, with only two of the twenty-two submissions outperforming it on both metrics. On last year’s benchmark, it also surpasses our GRPO-based system, despite being trained solely with ordinary supervised fine-tuning.

1 Introduction

A translation can preserve the meaning perfectly and still fail its audience. A Spanish-speaking developer who opens a *pull request* will call it exactly that, yet a generic machine translation (MT) system may render the phrase as *solicitud de extracción*. Both are understandable, but only the English term matches what developers read every day in their interfaces and documentation. Terminology-aware translation addresses this gap: beyond producing a correct translation, it must produce the prescribed one, whether the prescription comes from a client’s glossary or from the usage of a community.

A common way to teach this behaviour is to fine-tune the model on glossary-annotated translation



Figure 1: Data selection: train only on what SALAMANDRATA-7B-INSTRUCT gets wrong.

pairs. Much of this supervision, however, is redundant: for a large share of training examples the glossary prescribes exactly what the model would have produced anyway, so the example can be answered correctly without consulting the glossary at all. The useful signal is concentrated in the cases where the prescribed term differs from the model’s default choice. This observation suggests a simple data selection rule: keep for training only those where at least one term is hard for the model and comes out wrong (Figure 1). Our controlled experiment confirms that the rule matters. Holding the training budget and every other ingredient fixed, we vary only the proportion of these hard examples in the terminology data, and term accuracy climbs from 78.7% with none of them to 89.9% with all of them (Figure 2).

This paper describes the Barcelona Supercomputing Center (BSC) submission to the WMT26 Terminology Translation Shared Task Track 1,

document-level translation with an explicit dictionary. Our system is our latest public release, SALAMANDRATA-7B-INSTRUCT v3.0,¹ used exactly as released: its instruction-tuning mixture already includes terminology data built with this recipe, generated by a two-way synthetic pipeline on Gemma-4-31B (Team, 2026) and filtered against our previous release. At inference time, documents are translated chunk by chunk with per-chunk glossaries and lightly post-edited under a guard that protects prescribed terms. At the official WMT26 evaluation, our system achieves 94.2% term success at 74.6 chrF++, with only two of the twenty-two submissions outperforming it on both metrics. On last year’s benchmark, it also surpasses our GRPO-based system, despite being trained solely with ordinary supervised fine-tuning.

In summary, we make three contributions. First, we show that a lot of glossary-annotated training data is redundant, and that a simple remedy, keeping only the examples where the glossary contradicts the model’s default translation, yields an eleven-point term-accuracy gain at fixed data volume. Second, we describe a two-way synthetic pipeline built on open models, which turns glossaries into documents and documents into glossaries, and whose output trains the terminology component of the public release SALAMANDRATA-7B-INSTRUCT v3.0. Third, we present the BSC Track 1 submission: the released model, unmodified, wrapped in a document-level inference pipeline with per-chunk glossaries and a post-editing guard that protects prescribed terms.

2 Related work

Making models follow glossaries. The terminology constraint was earlier imposed at decoding time, through lexically constrained beam search (Hokamp and Liu, 2017), at a cost in speed and fluency. Dinu et al. (2019) moved it into training, annotating source terms with their prescribed translations so the model learns to copy them in context. The WMT terminology shared tasks (Alam et al., 2021; Semenov et al., 2023, 2025) record the field’s shift to LLMs. Our own previous submission moved the constraint into the reward, optimising models with GRPO under a joint adherence and quality objective (Garcia Gilabert et al., 2025). This year the recipe is simpler: ordinary

supervised fine-tuning, targeting longer document inputs. What is new is not how we train but what we train on: only the examples whose terminology the model gets wrong.

Synthetic terminology data. LLM-generated parallel text now rivals web-crawled corpora in quality (Finkelstein et al., 2024), and strong WMT25 terminology systems already relied on curated synthetic data and LLM post-editing (Jaswal, 2025; Semenov et al., 2025). Our pipeline differs in shape and in openness: it runs in both directions, expanding glossaries into documents and mining documents for glossaries, and is built entirely on openly licensed models and data.

Choosing what to train on. That some examples teach more than others is a recurring observation. In MT, selection began as domain adaptation (Axelrod et al., 2011), later made dynamic during training (van der Wees et al., 2017); curriculum learning orders examples by difficulty instead of filtering them (Bengio et al., 2009; Kocmi and Bojar, 2017; Platanios et al., 2019). The instruction-tuning era sharpened the point that a small, well-chosen set can beat a large undifferentiated one (Zhou et al., 2023; Chen et al., 2024; Liu et al., 2024), and, closest to us, Kocmi et al. (2025) filter MT data by difficulty before preference tuning. Concentrating supervision where the current model falls short is also the animating idea of active learning (Settles, 2009), long used in MT to decide what to annotate (Haffari et al., 2009; Zhao et al., 2020). In our case, we select from an already-labelled pool, so no annotator is queried; and where prior selection relies on a continuous score, however obtained, our notion of hardness is self-referential and binary: an example is hard exactly when the model’s own unconstrained output contradicts the prescribed term. The filter itself has no scores to threshold, and §4 measures its individual contribution.

3 Data: two synthetic pipelines

The terminology component of SALAMANDRATA v3’s instruction-tuning mixture is built by two synthetic pipelines that run in opposite directions. One starts from terminology and generates text: bilingual medical glossaries (MeSpEn; Villegas et al., 2018) seed Gemma-4-31B² (Team, 2026), which writes parallel texts around them. The other starts from text and extracts terminology: Gemma-4-31B

¹<https://huggingface.co/BSC-LT/salamandraTA-7b-instruct>

²<https://huggingface.co/google/gemma-4-31b>

extracted aligned term pairs from the EMEA parallel corpus³ (Tiedemann, 2012), whose sentences we then assemble into documents. Both pipelines share one selection step, described in §3.3, that keeps only what our own model gets wrong. Together they yield 33,615 instances covering 94 directed pairs across 29 languages (12,479 from MeSpEn, 21,136 from EMEA); Appendix A documents cleaning rules, generation prompts, and filtering thresholds (Figure 3 gives an overview).

3.1 MeSpEn: terminology → text

After cleaning the raw glossaries, we keep only the hard terms, those our model mistranslates in isolation (§3.3), and use them as seeds. Gemma-4-31B writes a parallel text around small groups of seeds, grouped by semantic similarity so that co-occurring terms are plausible together, at lengths ranging from a single sentence to multi-paragraph documents. In document-length samples the main term must appear at least twice and be translated identically each time, mirroring the consistency requirement of the shared task. A sample survives only if every annotated term actually appears on both sides, and term and sentence pairs exceed LaBSE (Feng et al., 2022) similarity thresholds.

3.2 EMEA: text → terminology → documents

Gemma-4-31B extracts aligned term pairs from each EMEA sentence pair, and sentences on which our model already produces every gold term are discarded (§3.3). After quality filtering, deduplication, and LaBSE filtering, consecutive sentences are concatenated into a mixture of granularities: single sentences, paragraphs, and multi-paragraph documents with literal newline breaks, each document’s glossary being the union of its sentences’ term pairs. Finally, glossary entries are reduced to lemma form while the text is left untouched, so the model must learn to inflect prescribed terms in context rather than copy surface strings.

3.3 The hard-example filter

Both pipelines rely on the same selection step, and it is the step this paper is about. Every candidate, a glossary term in MeSpEn or a sentence in EMEA, is translated by SALAMANDRATA-7B-INSTRUCT v2.0 *without* any terminology in the prompt, and the output is checked against the gold terms with a lenient matcher that forgives casing, punctuation,

Harder terminology examples → higher term accuracy




Figure 2: Term accuracy against the share of hard examples in the terminology data.

and spacing differences. If every required term already appears, the model needs no help on this example and it is discarded; only candidates with at least one missed term survive. We call the discarded examples *easy* and the surviving ones *hard*. The filter is computed once, against our previous public release.

4 Controlled study: does hardness matter?

The filter of §3.3 discards nearly half of the corpus (47.5%; Appendix B). This section asks whether that is actually a good idea. We hold everything about training fixed and vary a single quantity, the share of hard examples in the terminology data, from 0% to 100% in five steps plus a random baseline, and trace term accuracy along the way (Figure 2).

4.1 Setup

Training runs. In Figure 2, each point on the curve is a complete instruction-tuning run. Starting from the SALAMANDRATA-7B base model (the continually pre-trained translation checkpoint), we fine-tune on the fixed-background mixture of the public SALAMANDRATA-7B-INSTRUCT v2 release recipe (see data details in Appendix C). All runs use identical hyperparameters: one epoch, learning rate 10^{-5} with cosine decay and 3% warmup, context length 8,192, and an effective batch size of 32.

The six mixtures. Into this background we insert a terminology component, the only ingredient that changes between runs. It spans the 19 directions the EMEA branch produces, with a fixed per-direction budget of up to 4,000 examples, identical across mixtures. Six compositions are tested: 0%,

³<https://opus.nlpl.eu/datasets/EMEA>

25%, 50%, 75%, and 100% hard examples, plus a *random* baseline sampled from the unfiltered pool, which inherits the corpus’s natural per-direction proportion of hard examples. Easy examples and hard examples pass the same cleaning, quality, and semantic-similarity gates, so the two pools differ only in difficulty.

Evaluation. We evaluate on the WMT2025 Terminology Track 1 test set (Semenov et al., 2025): en→de, en→es, and en→ru.⁴ None of the three occurs as a training direction in the terminology component, so any gain on them already reflects transfer. We report the official term accuracy, micro-averaged over all 1,589 term occurrences, together with chrF (Popović, 2015) as a check that translation quality does not silently degrade; confidence intervals and *p*-values come from a clustered paired bootstrap over source segments (95%, 10k resamples).

Relation to the released model. These runs do not reproduce the v3.0 recipe. The background mixture follows the earlier v2.0 release recipe and does not include the additional tasks introduced in v3.0 (structured-text translation, register control, paraphrasing, and further post-editing and gender data). The terminology component also differs from the released one: v3.0 is trained on MeSpEn and EMEA data together, whereas each run here uses EMEA data only, at a fixed volume of 76,000 instances, so that the hard-example ratio is the only moving part across runs. The curve should accordingly be read as evidence for the selection principle, not as a decomposition of v3.0’s gains.

4.2 Results

Selection alone is worth eleven points. Figure 2 shows the result, with per-direction numbers in Appendix D: averaged over the three directions, term accuracy climbs from 78.7% with no hard examples to 89.9% with only hard examples, a gain of 11.2 points (95% CI [+9.5, +12.8], $p < 0.0001$) at identical data volume. The easy examples are not noise; they passed exactly the same cleaning, quality, and similarity gates as the hard ones.

Accuracy generally follows the dose of difficulty across every mixture we tested. Hard examples are potent: a quarter dose already delivers 71% of the full effect. The small dip at 50%, traced to the

en→ru direction, lies within noise on the average (25% vs. 50%: +0.3 points, 95% CI [−0.6, +1.3], $p = 0.55$, paired bootstrap). A forced 50/50 split trails the fully filtered set by 3.5 points. An unfiltered random sample of the same size comes closer, but for a revealing reason: EMEA is naturally rich in hard sentences, so a random draw is itself a moderately hard mixture, and it lands where its hard share predicts. Even against this lucky corpus, filtering adds a significant +1.3 points ($p = 0.003$); on a corpus with fewer hard cases, random sampling would inherit correspondingly less, while the filter turns the dose of difficulty into a guarantee.

The strongest evidence that difficulty is what does the teaching comes from the direction the training data covers worst. For en→ru, with no Russian anywhere in the terminology data, the easy-to-hard gain is the largest of the three (+15.7 points). The en→es curve climbs strictly with the hard share (+10.8), and en→de, the direction closest to the training languages, saturates earliest, with 75% hard and the random mixture edging past pure-hard (93.0 vs. 92.3).

Hardness costs nothing in quality. Across the five mixtures that contain hard examples, average chrF stays within half a point (67.0–67.4), and the all-easy mixture is the weakest on chrF as well (66.5). The gains are not bought with translation quality.

5 Document-level submission pipeline

The WMT26 Track 1 test sets are full documents, each paired with a corpus-level glossary that mixes entries relevant to the document with distractors from the whole collection. The 8,192-token context window of our system must hold the document, its glossary, and the translation as it is generated, so full documents invite truncation and repetition; and oversized glossaries dilute attention over entries that never occur in the text. The submission therefore translates each document in three stages.

1. Layout-preserving chunking. A document is split only at line breaks into chunks of roughly 500 source words, translated chunk by chunk with the terminology prompt of the model card (beam search, beam size 5), and re-joined into documents. A chunk spans several lines, and the model reproduces its line structure directly, having been trained on exactly these granularities (§3). An automatic check verifies every chunk, and the few that call for it are re-translated line by line, so the submitted

⁴<https://github.com/wmt-conference/wmt25-terminology>

layout matches the source by construction.

2. Glossary filtering. Each chunk receives only the glossary entries whose source term actually occurs in it. A term counts as present if it appears verbatim or if its lemmas match. This shrinks the corpus-level glossary to the handful of entries the model must act on.

3. QE-guided post-editing. Chunk translation leaves occasional local defects: dropped clauses, agreement errors, spelling noise. We therefore score every segment with CometKiwi (Rei et al., 2022) and let the same released model post-edit the low-scoring ones, using the post-editing prompt from its instruction tuning, so that the submission comes entirely from one model.

The risk is that fluency and terminology disagree about repetition. The test documents often name one concept several ways while the glossary maps every name to the same target, so faithful translations repeat themselves: *a steel wheel (or steelie)* becomes *koto stalowe (lub koto stalowe)*. The post-editor does what any writer would and smooths the repetition away. Sometimes this is harmless, because the document keeps enough copies of the term; sometimes it deletes an occurrence that term accuracy still needed, or replaces the prescribed wording altogether. CometKiwi rewards all three alike. An edit is therefore applied only if it raises CometKiwi and does not lower document-level term accuracy. Most edits that survive fix grammar, spelling, or punctuation, or remove a genuinely surplus repetition.

6 Results

6.1 Official WMT26 results

Table 1 places SALAMANDRATA-7B-INSTRUCT v3.0, used exactly as released and wrapped in the pipeline of §5, among the WMT26 Track 1 submissions (Charkiewicz et al., 2026). The system reaches 94.2% lemmatised term success at 74.6 document chrF++, only 2.2 points of term success below the human reference. Of the twenty-two submissions, only COZY and Agentic-OPUS improve on both axes at once; four further systems reach higher term success at lower chrF++.

6.2 Comparison with our GRPO submission

Our WMT25 system optimised the same model family with GRPO under a joint adherence and

quality reward (Garcia Gilabert et al., 2025). It operates only at the sentence level, so the two can only be compared on last year’s benchmark (Table 2). The released v3.0 model improves on both axes, from 67.3 to 69.4 average chrF and from 91.3% to 94.0% average term accuracy, with no reinforcement learning, no task-specific adaptation, and no access to the test-time domain. The comparison is not fully controlled, however. The two systems differ not only in their terminology data, but also in the broader instruction mixture. We therefore view this as a system-level comparison rather than a controlled ablation. Nevertheless, the result does show that a general-purpose translation model, given our terminology data as part of ordinary instruction tuning, can outperform a dedicated RL pipeline on this task.

7 Discussion

Many-to-one glossaries make quality and adherence disagree. The test glossaries are not injective: *boot lid* and *trunk lid* share the single prescribed target *klapa bagażnika*. When a source document deliberately enumerates synonyms, a faithful translation must repeat that target: “both the boot lid and the trunk lid” can only become *zarówno klapy bagażnika, jak i klapy bagażnika*, text any human editor would immediately smooth out. Quality estimation sides with the editor: CometKiwi scored the smoothed variant, ... *jak i pokrywa bagażnika*, 0.13 higher, and across our post-editing stage it consistently assigned positive deltas to the very edits that damage adherence. We suspect this tension generalises: any system that optimises a quality metric on top of a terminology task needs an explicit adherence constraint, because the two objectives are locally adversarial on some many-to-one entries. What the translations of these sentences should ideally look like, and how they should be scored in the terminology task, remain open questions.

What the filter buys. The controlled study suggests the value of hard examples is not the term pairs they contain but the behaviour they force: reading the glossary when the model’s default would have been something else. This also suggests the filter should transfer to other adherence-flavoured tasks (style guides, register constraints, do-not-translate lists) where most naturally occurring supervision is similarly redundant. Whether this behavioural account holds up mechanistically,

System	Term success			chrF++		
	Avg	en-pl	es-eu	Avg	en-pl	es-eu
<i>gold</i>	<i>96.4</i>	<i>93.6</i>	<i>99.1</i>	<i>100.0</i>	<i>100.0</i>	<i>100.0</i>
Agentic-OPUS	95.7	95.4	96.0	75.3	77.1	73.5
COZY _{flash}	95.4	94.4	96.3	74.4	76.5	72.2
HW-TSC	95.3	95.8	94.8	73.8	75.6	72.1
CoZY	95.0	93.6	96.4	75.6	75.9	75.3
Agentic-Sonnet	94.7	95.2	94.3	73.1	75.6	70.7
TaT	94.5	97.2	91.7	73.7	73.6	73.8
SalamandraTA v3.0 (ours)	94.2	93.2	95.2	74.6	75.9	73.3
CUNI-UFAL	93.4	91.6	95.2	74.4	74.2	74.5

Table 1: WMT26 Track 1, proper mode: lemmatised exclusive term success (%) and document chrF++, per direction and averaged. The eight submissions with the highest average term success, out of twenty-two; best among these in bold. Full table in the findings paper (Charkiewicz et al., 2026).

System (WMT25 Track 1)	chrF	Term
o3-term-guide	71.0	99.1
duterm	70.1	98.2
SalamandraTA v3.0 (ours)	69.4	94.0
salamandrata (ours, GRPO)	67.3	91.3

Table 2: WMT25 Track 1, averaged over en→es/de/ru: our released model against our previous GRPO submission and the two top-ranked systems (Semenov et al., 2025).

and what in the model implements it, we leave to future work.

8 Conclusion

We described the BSC submission to WMT26 Terminology Translation Track 1: the public release SALAMANDRATA-7B-INSTRUCT v3.0, whose terminology training data is built by a two-way synthetic pipeline on open models and filtered by a single rule, keep only what the model gets wrong, together with a document-level inference pipeline with per-chunk glossaries and QE-guided post-editing. A controlled study shows the filter is responsible for an eleven-point term-accuracy gain at fixed data volume, and the released system reaches 94.2% term success at 74.6 chrF++ in the WMT26 official evaluation. Hard examples, it turns out, are better teachers.

Limitations

The WMT25 test set provides per-sentence terminology, so its errors are all genuine adherence failures. The WMT26 setting differs: a corpus-level glossary must first be narrowed to each chunk by surface and lemma matching (§5), and a term whose inflected form escapes the match never

reaches the model at all. Our submission’s errors will therefore mix adherence failures with retrieval failures. The controlled study is deliberately narrow: one data source, three test directions, one data volume, and a single reference model defining hardness. The filter is also static, computed once against v2.0 before training; a dynamic variant that re-estimates difficulty during training might select better but was out of scope. Our terminology data is biased toward the medical domain of its seed resources, while the test domains this year differ. Finally, the base model was continually pre-trained on sentence-level parallel data only. Document-level translation behaviour comes primarily from instruction tuning, and very long documents must be chunked at inference, which can lose cross-chunk context such as antecedents for consistent term choice. We are currently extending continual pre-training to document-level parallel corpora, which we expect will improve document-level context modeling.

Acknowledgments

This work/research has been promoted and financed by the Government of Catalonia through the Aina project.

This work is supported by MLLM4TRA (PID2024-158157OB-C32) funded by MCIN/AEI/10.13039/501100011033/FEDER, UE.

This work was supported by the TaMTAS project (PCI2025-167117-2), funded by MICIU/AEI/10.13039/501100011033 and the European Union under the CHIST-ERA Call 2025, *Science in Your Own Language*.

References

- Md Mahfuz Ibn Alam, Ivana Kvapilíková, Antonios Anastasopoulos, Laurent Besacier, Georgiana Dinu, Marcello Federico, Matthias Gallé, Kweonwoo Jung, Philipp Koehn, and Vassilina Nikoulina. 2021. [Findings of the WMT shared task on machine translation using terminologies](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 652–663, Online. Association for Computational Linguistics.
- Amittai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. [Domain adaptation via pseudo in-domain data selection](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 355–362, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 41–48, New York, NY, USA. Association for Computing Machinery.
- Adrian Charkiewicz, Pinzhen Chen, Thierry Etchegoyhen, Harritxu Gete Ugarte, Kamil Guttman, Xu Huang, David Ponce, Artur Nowakowski, Frédéric Odeh, Arturo Oncevay, Kirill Semenov, Dawei Zhu, and Vilém Zouhar. 2026. Findings of the WMT26 terminology translation task. In *To appear in Proceedings of the Eleventh Conference on Machine Translation*. Association for Computational Linguistics.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. 2024. [Alpagasus: Training a better alpaca with fewer data](#). Preprint, arXiv:2307.08701.
- Georgiana Dinu, Prashant Mathur, Marcello Federico, and Yaser Al-Onaizan. 2019. [Training neural machine translation to apply terminology constraints](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3063–3068, Florence, Italy. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ari-vazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Mara Finkelstein, David Vilar, and Markus Freitag. 2024. [Introducing the NewsPaLM MBR and QE dataset: LLM-generated high-quality parallel data outperforms traditional web-crawled data](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1355–1372, Miami, Florida, USA. Association for Computational Linguistics.
- Javier Garcia Gilabert, Carlos Escolano, Xixian Liao, and Maite Melero. 2025. [Terminology-constrained translation from monolingual data using GRPO](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1335–1343, Suzhou, China. Association for Computational Linguistics.
- Gholamreza Haffari, Maxim Roy, and Anoop Sarkar. 2009. [Active learning for statistical phrase-based machine translation](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 415–423, Boulder, Colorado. Association for Computational Linguistics.
- Chris Hokamp and Qun Liu. 2017. [Lexically constrained decoding for sequence generation using grid beam search](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- Akshat Jaswal. 2025. [It takes two: A dual stage approach for terminology-aware translation](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1344–1350, Suzhou, China. Association for Computational Linguistics.
- Tom Kocmi, Arkady Arkhangorodsky, Alexandre Berard, Phil Blunsom, Samuel Cahyawijaya, Théo Dehaze, Marzieh Fadaee, Nicholas Frosst, Matthias Galle, Aidan Gomez, Nithya Govindarajan, Wei-Yin Ko, Julia Kreutzer, Kelly Marchisio, Ahmet Üstün, Sebastian Vincent, and Ivan Zhang. 2025. [Command-a-translate: Raising the bar of machine translation with difficulty filtering](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 789–799, Suzhou, China. Association for Computational Linguistics.
- Tom Kocmi and Ondřej Bojar. 2017. [Curriculum learning and minibatch bucketing in neural machine translation](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 379–386, Varna, Bulgaria. INCOMA Ltd.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024. [What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning](#). Preprint, arXiv:2312.15685.
- Audrey Mash, Ella Paulina Bohman, and Maite Melero. 2026. ParaCLEAN: Improving Translation Quality through Systematic Parallel Data Cleaning. *Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026)*, pages 6630–6640.
- Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabas Poczos, and Tom Mitchell. 2019. [Competence-based curriculum learning for neural machine translation](#). In *Proceedings of the 2019*

- Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1162–1172, Minneapolis, Minnesota. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Kirill Semenov, Xu Huang, Vilém Zouhar, Nathaniel Berger, Dawei Zhu, Arturo Oncevay, and Pinzhen Chen. 2025. [Findings of the WMT25 terminology translation task: Terminology is useful especially for good MTs](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 554–576, Suzhou, China. Association for Computational Linguistics.
- Kirill Semenov, Vilém Zouhar, Tom Kocmi, Dongdong Zhang, Wangchunshu Zhou, and Yuchen Eleanor Jiang. 2023. [Findings of the WMT 2023 shared task on machine translation with terminologies](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 663–671, Singapore. Association for Computational Linguistics.
- Burr Settles. 2009. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- Gemma Team. 2026. [Gemma 4 technical report](#). Preprint, arXiv:2607.02770.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Marlies van der Wees, Arianna Bisazza, and Christof Monz. 2017. [Dynamic data selection for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1400–1410, Copenhagen, Denmark. Association for Computational Linguistics.
- Marta Villegas, Ander Intxaurreondo, Aitor Gonzalez-Agirre, Montserrat Marimon, and Martin Krallinger. 2018. The mespen resource for english-spanish medical machine translation and terminologies: Census of parallel corpora, glossaries and term translations. *LREC MultilingualBIO: multilingual biomedical text processing*.
- Yuekai Zhao, Haoran Zhang, Shuchang Zhou, and Zhihua Zhang. 2020. [Active learning approaches to enhancing neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1796–1806, Online. Association for Computational Linguistics.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [Lima: Less is more for alignment](#). Preprint, arXiv:2305.11206.

A Synthetic-data pipeline details

Figure 3 gives an overview of the two pipelines. MeSpEn starts from terminology and generates text around it, while EMEA pipeline starts from the EMEA corpus after cleaning with ParaCLEAN (Mash et al., 2026) and extracts terminology out of it. Both contain the same two quality gates: a *hard-case selection* step, where SALAMANDRATA itself decides what is worth training on, and a *semantic filter* based on LaBSE similarity.

A.1 MeSpEn: generating text from terminology

1. Glossary cleaning. The raw MeSpEn glossary files contain formatting noise: numbered variant lists, domain markers, phonetic transcriptions, part-of-speech tags, and occasional swapped columns. A single cleaning pass fixes these, splits multi-variant entries into pipe-separated alternatives, and deduplicates rows.

2. Hard-term selection. SALAMANDRATA-7B-INSTRUCT v2.0 translates every glossary term on its own, and the output is compared to the gold entry with a lenient matcher, so that casing, punctuation, or spacing differences do not count as errors. Terms the model already translates correctly are discarded; only genuine mismatches are used as seeds.

3. Semantic grouping. Each synthetic text is generated around a small group of glossary terms, so the terms in a group must plausibly belong to the same text: sampling them at random produces absurd combinations (e.g., an obstetrics term next to a dental one). We therefore cluster the glossary into medical sub-domains with LaBSE, and draw the terms of each group mostly (80% of the time) from a single cluster.

4. Text generation. Gemma-4-31B receives a term group and writes the source text and its translation in one pass, at lengths ranging from a single sentence (25%) over short passages and paragraphs (58%) to multi-paragraph documents (17%); short texts carry 1–3 terms, documents more. Two tricks keep the output realistic. First, the model must invent a one-sentence scenario before writing (e.g., “a discharge summary after knee surgery”), which anchors the text in a concrete setting instead of a vague list of facts. Second, in document-length samples the main term has to appear at least twice and be translated the same way each time. The

model finally reports which surface form each term actually took (say, the plural *células* rather than the glossary form *célula*); the next step depends on these.

5. Filtering. Two kinds of errors slip through generation. Sometimes a reported term does not actually appear in the text it claims to be in; simple rule checks catch these, and also drop degenerate samples (extreme length ratios, source copied as target, duplicates). Sometimes both source and target read fluently, yet they quietly say different things; LaBSE similarity catches these, by requiring each term pair and the full text pair to score at least 0.80. Surviving records are formatted as prompt–answer pairs, with document-form records oversampled ($\approx 10\times$) to strengthen multi-paragraph layout.

A.2 EMEA: extracting terminology from text

1. Term extraction. Gemma-4-31B reads each aligned sentence pair and extracts the domain-specific term pairs it contains. Numeric “terms” (dosages such as *720 IU/kg*) are removed.

2. Normalization. EMEA text carries detokenization artifacts that would otherwise end up inside glossary terms: “*Parkinson’ s*”, “*13, 6 mg*”, “*mg/ ml*”. A conservative pass repairs these (“*Parkinson’s*”, “*13,6 mg*”, “*mg/ml*”).

3. Hard-sentence selection. SALAMANDRATA translates every source sentence *without* any glossary in the prompt, and we check which gold terms appear in its output. If all of them do, the model needs no help on this sentence, so only sentences with at least one missed term are kept.

4. Filtering and audit. The same two error types as in MeSpEn are checked. Sometimes Gemma reports a term that is not literally in the sentence (e.g., one it inferred rather than read), so we look up every extracted term in its own sentence and drop the record when it is missing, along with truncated or length-distorted samples. Sometimes the two sides quietly diverge in meaning; LaBSE similarity catches these, by requiring the sentence pair to score at least 0.88.

5. Document construction. EMEA sentences arrive in corpus order, so consecutive sentences from the same document section can simply be concatenated back into longer units. Each training instance is sampled from a length profile, that is, a single sentence, a short space-joined paragraph, or




Figure 3: The two synthetic-data pipelines.

a document of several hundred to $\approx 1,000$ source words with newline layout, so that the model sees every granularity. A document’s glossary is the union of its sentences’ term pairs.

6. Glossary lemmatization. Finally, every glossary entry is reduced to its lemma form with Stanza (Qi et al., 2020) and deduplicated. The glossary thus advertises *apple* $\rightarrow$ *manzana* even where the text contains *apples/manzanas*, and the model must learn to inflect prescribed terms in context rather than copy surface strings. We verified that the longest formatted instances stay within the 8,192-token training context.

A.3 Data profile

Figure 4 summarizes the two corpora. MeSpEn contributes short, term-dense texts in 79 language pairs (median 45 source words, usually 1–2 glossary entries), while EMEA contributes long documents in 19 pairs (median 645 source words, median 83 glossary entries).

B Hard/easy pool sizes per direction

Table 3 gives, for every direction, the size of the hard pool (v2.0 misses at least one term) and the easy pool (v2.0 already produces every term), counted after the quality and LaBSE audit gates and before any subsampling. Every ratio-curve run of §4 draws 4,000 records per direction from these pools; the *random* baseline draws them from the

Direction	Hard	Easy	%Hard
cs-de	7,291	7,929	47.9
cs-en	15,187	30,291	33.4
de-en	14,835	30,886	32.4
en-et	17,348	12,457	58.2
en-pl	20,404	18,558	52.4
es-bg	32,219	32,839	49.5
es-cs	36,281	28,473	56.0
es-da	34,639	30,472	53.2
es-el	23,646	23,587	50.1
es-fi	22,734	13,322	63.1
es-hu	31,301	17,805	63.7
es-lt	24,345	19,242	55.9
es-mt	33,659	27,723	54.8
es-nl	36,629	36,781	49.9
es-pl	26,656	22,981	53.7
es-ro	42,265	29,028	59.3
es-sk	38,744	28,557	57.6
es-sl	33,429	28,357	54.1
es-sv	30,664	32,429	48.6
Total	522,276	471,717	52.5

Table 3: Per-direction pool sizes after quality filtering. The hard share ranges from 32.4% (de-en) to 63.7% (es-hu).

union of the two, which is why its hard share tracks the corpus’s natural one.

C Background instruction-tuning mixture

All training runs share the same background mixture: the instruction-tuning recipe of the pub-




Figure 4: Data profile of the two synthetic datasets.

lic SALAMANDRATA-7B-INSTRUCT v2 release,⁵ which covers general, multi-reference, paragraph-level, and document-level translation, post-editing, grammar correction, gender-fair translation, and named-entity recognition. Table 5 lists every dataset with its sampling cap and instance count; per-dataset language coverage, licences, and references are given in the release model card. High-resource corpora enter the mixture through a fixed per-direction sampling cap (the “Sample/dir.” column); smaller datasets are used in full. The background totals 682,431 instances and is identical in every run. Only the terminology-aware component added on top of it changes (§A).

D Ablation: full numbers

Mixture	en→de	en→es	en→ru	Avg
0% hard	85.1 (68.1)	80.1 (74.0)	70.5 (57.4)	78.7 (66.5)
25%	90.8 (68.2)	87.5 (74.9)	81.3 (58.9)	86.7 (67.4)
50%	91.0 (68.4)	88.1 (74.7)	79.6 (58.0)	86.3 (67.0)
random	93.0 (68.5)	88.3 (74.4)	84.1 (58.5)	88.5 (67.1)
75%	93.0 (68.6)	89.9 (74.5)	84.3 (59.0)	89.2 (67.4)
100%	92.3 (67.9)	90.9 (74.6)	86.2 (59.0)	89.9 (67.2)

Table 4: Term accuracy (chrF in parentheses) per direction for every point of Figure 2. Accuracy averages are micro-averaged over all 1,589 term occurrences; chrF averages are means over the three directions.

⁵<https://huggingface.co/BSC-LT/salamandraTA-7b-instruct/tree/v2.0>

Task	Dataset	Coverage	Sample/dir.	Instances
General MT	Tatoeba	hundreds of directions	500	121,361
	NTEU	ca/es/en $\leftrightarrow$ EU languages	500	36,000
	EuPress	ca/es/en $\leftrightarrow$ EU languages	500	36,000
	WMT++	en $\rightarrow$ 32	500	16,000
	FLEURS	en/ca/es $\rightarrow$ XX	100	12,900
Multi-reference Paragraph-level	Tatoeba	hundreds of directions	500	100,973
	News-Commentary	56 directions	1,000	56,000
	ACADData (train)	96 directions	500	27,852
	NewsPalm	en $\leftrightarrow$ de	5,000	10,000
Document-level	ACADData (bench)	12 directions	full	5,944
	Europarl	210 directions	500	105,000
	Project Gutenberg (filtered)	39 pairs	full	8,959
Post-editing	LangMark	en $\rightarrow$ 7	5,000	35,000
	Q21	en $\leftrightarrow$ de/cs/lv	5,000	30,000
	ApeQuest	en $\rightarrow$ fr/nl/pt	5,000	15,000
Grammar	Synthetic noise (NTEU, EuPress)	25 languages	500	12,500
Gender	EuroGEST	en $\rightarrow$ 21	full	29,076
	GeNTE	en-it	full	3,000
	GLITTER	en-de	full	1,202
NER	AnCora	ca	full	10,629
	SLI NERC	gl	full	6,483
	EIEC	eu	full	2,552
Total background				682,431

Table 5: The fixed background mixture shared by all instruction-tuning runs: the recipe of the SALAMANDRATA-7B-INSTRUCT v2 release. “Sample/dir.” is the per-translation-direction cap applied to high-resource corpora; “full” means the entire dataset is used.