

Omnilingual MT: Machine Translation for 1,600 Languages

Omnilingual MT Team, Belen Alastruey[†], Niyati Bafna[†], Andrea Caciolai[†], Kevin Heffernan[†], Artyom Kozhevnikov[†], Christophe Ropers[†], Eduardo Sánchez[†], Charles-Eric Saint-James[†], Ioannis Tsiamas[†], Xiang “Tony” Cao[§], Chierh Cheng[§], Joe Chuang[§], Paul-Ambroise Duquenne[§], Mark Duppenhaler[§], Nate Ekberg[§], Cynthia Gao[§], Pere Lluís Huguet Cabot[§], João Maria Janeiro[§], Jean Maillard[§], Gabriel Mejia Gonzalez[§], Holger Schwenk[§], Edan Toledo[§], Arina Turkatenko[§], Albert Ventayol-Boada[§], Ravi Arikapudi[‡], Rashel Moritz[‡], Alexandre Mourachko[‡], Surya Parimi[‡], Mary Williamson[‡], Shireen Yates[‡], David Dale[⊥], Marta R. Costa-jussa[⊥]

FAIR at Meta

costajussa@meta.com, daviddale@meta.com

[†]Core Contributors [§]Other Contributors [‡]Project Management [⊥]Technical Leadership

Abstract

We introduce **Omnilingual MT**, a family of multilingual translation (MT) systems covering more than 1,600 languages—the broadest benchmarked coverage of any MT system to date. We make contributions in several key research questions in the field: strategies for the use of monolingual data, curating seed training data, tokenization and suitable architectures, the competitiveness of small specialized MT models as opposed to general purpose large models, and evaluation protocols. Our data strategy combines large-scale multilingual corpora with backtranslation for 2,000+ languages and MeDLEy, a grammatically-diverse seed dataset for 109 low-resource languages. We explore two complementary architectures: OMT-LLAMA, a decoder-only model (1B/3B/8B) built on LLAMA3 with multilingual continual pretraining, and OMT-NLLB, a 3B encoder-decoder model with a novel three-stage training strategy exploiting non-parallel data via auto-encoding. For evaluation, we extend the BOU-QuET benchmark to 275 languages, develop Met-BOUQuET with human annotations of machine quality in 161 directions, and BLASER 3, a reference-free quality estimation metric. Our specialized 1B–8B models match or exceed 70B general-purpose LLMs, demonstrating that specialization is a more reliable path to multilingual translation than scale alone. Omnilingual MT achieves “passable” translation quality as judged by human annotators for more than 400 languages, doubling it over the previous state of the field at ~200 languages, and translates in zero-shot mode into ~1,200 languages versus 300–400 for the baselines. With Omnilingual MT, we showcase a path encompassing the full pipeline of data and system development towards designing models that are multilingual from the onset.

1 Introduction

The success of No Language Left Behind (NLLB; NLLB Team et al., 2024) demonstrated that high-

quality machine translation (MT) could extend to 200 languages, setting a new standard for linguistic inclusion. Despite its name, the project left the vast majority of the world’s ~7,000 languages without any direct support or mechanism of including new languages into an existing model. Early explorations toward 1,000-language MT (Siddhant et al., 2022) demonstrated its feasibility but did not yield sustained or extensible systems. Meanwhile, large language models (LLMs) have improved MT for high-resource languages (Alves et al., 2024; Dang et al., 2024; Gemma Team et al., 2025), but their multilingual gains remain dominated by well-resourced languages, and they are inefficient—often requiring tens of billions of parameters to match much smaller specialized models. Pretraining LLMs with Internet-scale data enabled including more languages than ever before in the training mix of these models. However, the evaluation of whether this data coverage actually results in improved translation from and into underserved languages has not been addressed systematically in the recent literature.

Since NLLB’s FLORES-200 benchmark (introduced in Goyal et al. (2022) and later extended), there has been no attempt to radically increase the language coverage of MT evaluation. But even FLORES-200 exposed a structural asymmetry: modern models can often *interpret* languages beyond their declared coverage through cross-lingual transfer, but they struggle to *generate* them reliably.

We argue that foundational models should aspire to provide at least a base level of support for *every* language, rather than concentrating capacity on a privileged subset. In this work, we introduce **Omnilingual MT**, a family of multilingual translation systems that extend support to more than 1,600 languages, the broadest benchmarked coverage to date. Our contributions span data, modeling, and evaluation (see Table 1). In order to produce them, we explored several research directions:

Table 1: **Summary of Omnilingual MT:** Each row lists the artifact, its type, and its primary claim.

Name	Type	Description	Claim
OMT-LLaMA	Model	LLaMA3-based decoder-only MT models (1B, 3B, 8B)	LLM for MT, capable of understanding and generating over 1,000 languages
OMT-NLLB	Model	Encoder-decoder MT built on OMNISONAR	A specialized MT system supporting 1,600-to-250 languages
BLASER 3	Model	Reference-free MT quality estimation	Highest multilingual coverage (1,600+ source languages)
OmniTOX	Model	Toxicity detection	Largest language coverage (1,600 languages), ROC AUC 0.86
MeDLEy	Dataset	MT training data	109 extremely low-resource languages, domain and grammatical diversity
BOUQuET	Dataset	Extension of an MT evaluation benchmark	275 languages, 56 language families, 8 domains
Met-BOUQuET	Dataset	Human judgments of MT quality	161 language directions; largest quality estimation dataset by language coverage

Monolingual data: How can monolingual data be leveraged to boost MT performance for low-resourced languages? We validate its utility with traditional methods (backtranslation, bitext mining and language modeling continual pretraining) as well as a more innovative mid-training autoencoding objective.

Seed data construction: How can we optimize diversity in a multiway parallel seed dataset created from scratch, given tight size constraints? We construct MeDLEy, a manually curated, grammatically-diverse, multicentric, and multidomain seed dataset for 109 low-resource languages.

Architectures: Are encoder-decoder Transformers still a viable option for NMT? We show that this is the case with the OMT-NLLB model that employs a cross-lingually aligned sentence encoder (Omnilingual SONAR Team et al., 2026) in an encoder-decoder architecture with a novel three-stage training strategy exploiting non-parallel data and supporting 1,600 and 250 languages on the source and target sides, respectively.

Model size: Can small translation-specific LLMs compete with larger general-purpose language models? We develop OMT-LLAMA, a series LLaMA3-based decoder-only models (1B/3B/8B) with a multilingual continual pretraining (CPT) recipe, post-training via supervised fine-tuning (SFT) and reinforcement learning (RL), covering 1,600+ languages. All three models outperform their much larger non-specialized counterpart, LLaMA3 70B.

Adaptability: How can we boost LLM-based MT quality for barely seen languages? We show that both fine-tuning and retrieval augmented gen-

eration work and their effects add up but cannot compensate for a weaker base model.

MT metrics: How can we scale the language coverage of scarce training data for quality estimation metrics? With BLASER 3, we combine human annotations with multiple protocols, synthetic data, and a strong base encoder to develop an omnilingual quality estimation model.

With these strategies combined, we were able to extend the number of languages that a multilingual model can “understand well enough for practical use” from about 200 to over 400, and the number of languages with better-than-random generation ability, from 300-400 to about 1200. Despite conventional wisdom, small specialised models (<8B) match or exceed larger general-purpose models. We hope that our work opens the way to *omnilinguality* in machine translation: models and resources aiming to provide better support to every language.

2 Related Work

Scaling multilingual MT. NLLB (NLLB Team et al., 2024) demonstrated high-quality MT for 200 languages using encoder-decoder Transformers with large-scale data curation and established FLoRes+ as the de facto evaluation standard. Google’s massively multilingual translation work (Siddhant et al., 2022) explored models covering up to 1,000 languages but did not produce sustained, extensible systems. Recent specialized systems such as MADLAD (Kudugunta et al., 2023) and TranslateGemma (Google Translate Research Team et al., 2026) have improved coverage beyond the NLLB-200 language set. More recently, Omnilingual SONAR Team et al. (2026) built a sentence

encoder with unlimited language support evaluated with translation from 1,560 languages, but stuck to the NLLB-200 language set on the target side.

LLMs for MT. Large decoder-only models exhibit surprising zero-shot and few-shot translation abilities (Zhu et al., 2024) and adaptability for the MT task (Xu et al., 2024), and specialized variants such as TowerLLM (Alves et al., 2024) validate LLM-based MT recipes. However, LLMs remain inefficient for MT: even 70B-parameter models lag behind smaller specialized systems on low-resource pairs. Recent work on MT-specific RL (Feng et al., 2025) and reasoning-based translation (Wang et al., 2025) has shown promise for high-resource pairs.

Data for long-tail languages. Creating training data for thousands of languages requires diverse strategies. Backtranslation (Sennrich et al., 2016; NLLB Team et al., 2024) and bitext mining (Schwenk et al., 2021) have been extended to massively multilingual settings. Manually curated seed datasets—NLLB Seed (NLLB Team et al., 2024), SMOL (Caswell et al., 2025), and Gatitos (Jones et al., 2023)—provide high-quality examples that bootstrap synthetic data pipelines.

Evaluation at scale. MT evaluation beyond 200 languages remains limited by the availability of the data and reliability of the metrics. The most widely used multilingual MT benchmark, FLORES-200 (Goyal et al., 2022; NLLB Team et al., 2024), has been extended to FLORES+ (Burchell et al., 2024; Dale et al., 2025) covering, as of today, 228 language varieties.¹ Few benchmarks go beyond that, with a notable exception of ChiKhaPo (Chang and Bafna, 2026) focused on word translation evaluation with thousands of languages.

Metrics are also a bottleneck. Surface-level metrics such as ChrF++, BLEU (Papineni et al., 2002; Popović, 2015) fail to capture meaning adequacy; neural metrics such as xCOMET (Guerreiro et al., 2024) and MetricX-24 (Juraska et al., 2024) are not guaranteed to work with target languages not included in their fine-tuning data. BLASER 2 (Dale and Costa-jussà, 2024) demonstrates cross-lingual

¹By “language variety” we denote the combination of language (e.g. Mandarin Chinese), writing system (e.g. Cyrillic script), and, in some cases, dialect (e.g. the Aranese dialect of the Occitan language included in FLORES+). Throughout this paper, when reporting a “number of languages”, we count the distinct language varieties.

Type	Source	# Texts	# Languages
Mono.	CC-2000-WEB	20M	~2000
	CC-2000-PDF	5M	~1700
Parallel	Bible	600M	~1600
	Panlex	2B	~1000
	CC-NLLB-200	500M	~200
	Other sources	90M	~500
Synthetic	Backtranslation	135M	~2000
	Mined paragraphs	100K	~60

Table 2: CPT data sources.

generalization, but is limited to the 200 languages of the underlying sentence encoder.

3 Data

3.1 Continual Pretraining Data

Our data strategy prioritizes both breadth (number of languages) and depth (domain diversity). Table 2 summarizes the CPT data sources.

Monolingual data. We collect multilingual monolingual corpora based on web documents from Common Crawl. After quality filtering and language identification with GlotLID (Kargaran et al., 2023), inspired by FineWeb2 (Penedo et al., 2025), we create CC-2000-WEB and CC-2000-PDF covering 2,000+ languages.

Parallel data. Bible texts (~1,600 languages)² provide the broadest coverage. Panlex dictionaries (Kamholz et al. (2014), ~1,000 languages) provide word-level alignment. CC-NLLB-200 reproduces the NLLB-200 composition of primary data sources. We further include Tatoeba (Tiedemann, 2012), SMOL (Caswell et al., 2025), Gatitos (Jones et al., 2023), and data contributed via the Language Technology Partnership Program³, enabling partners to contribute multilingual speech recordings and written text with the purpose of expanding the support of underserved languages in AI models, as well as a collection of over 50 miscellaneous publicly available datasets.

3.2 Backtranslation

For languages with monolingual text but little or no parallel data, we generate synthetic bi-text via backtranslation: (1) segment monolingual

²The majority of Bible texts were downloaded using the eBible tool: <https://github.com/BibleNLP/ebible>.

³<https://about.fb.com/news/2025/02/announcing-language-technology-partner-program/>

	En-YY				XX-En			
	H	M	L	VL	H	M	L	VL
No BT	44.7	24.0	17.0	20.0	57.0	42.3	35.7	38.2
+BT (LQ)	46.1	26.1	20.3	23.6	58.0	43.8	38.1	40.9
+BT (MQ)	46.3	26.6	20.8	24.3	58.2	44.0	38.3	41.3
+BT (HQ)	46.8	26.1	20.7	24.1	58.5	43.6	37.7	40.9
languages	28	89	71	24	28	89	71	24

Table 3: Effect of backtranslation on ChrF++ (FLo-Res+). MQ = medium quality filtering. H/M/L/VL = resource levels, defined in Section 3.5.

text into sentences using a fine-tuned SAT-12L-SM model (Frohmman et al., 2024); (2) filter via GlotLID with language-specific thresholds calibrated using Bible texts; (3) translate using NLLB-200 or an earlier version of OMT-LLAMA depending on the language support; and (4) estimate quality using OMNISONAR sentence embedding similarity and model-free signals: length ratio, unique character ratio to exclude repetitive hallucinations, edit distance to exclude excessive source copying.

We define quality filtering presets based on embedding similarity. A sentence pair is classified as “medium-quality” (MQ) if its similarity exceeds a reference threshold calibrated to the mean similarity of Bible texts for that target language. High-quality (HQ) and low-quality (LQ) presets adjust this threshold by adding or subtracting one standard deviation, respectively. We find that medium-quality filtering provides the best balance across resource levels (Table 3). The final dataset contains $\sim 230\text{M}$ sentences across $\sim 2,000$ languages.

3.3 Bitext mining

We adopt global mining using the OMNISONAR multilingual embedding space (Omnilingual SONAR Team et al., 2026) to discover parallel paragraphs in monolingual corpora. Source data includes CC-2000-Web (non-English text) and a filtered version of FineWeb-Edu (Lozhkov et al., 2024) (English text). We encode all data into OMNISONAR embeddings and apply product quantization using FAISS (Johnson et al., 2019) to enable efficient billion-scale similarity search. Quantization is applied to batches of up to 50 million embeddings per language. We then perform KNN search with $k = 3$ nearest neighbours using the stopes library (Andrews et al., 2022), retrieving English candidates for each non-English sentence and vice versa. Candidate pairs are scored using margin-based cosine similarity

in the OMNISONAR space, and pairs above language-specific thresholds are retained as mined bitexts.

3.4 MeDLEy: Seed Data for Low-Resource Languages

Existing multilingual datasets exhibit systematic gaps in domain diversity and coverage of marginalized languages. We create MeDLEy (Multicentric, Domain-diverse, Linguistically-diverse, Easy-to-translate seed dataset), a manually curated parallel corpus for 109 extremely low-resource languages. MeDLEy is the first seed data that focuses on covering grammatical phenomena in an arbitrary language. This is achieved by first curating 61 cross-linguistic grammatical *functions* (e.g., case, aspect, evidentiality) with high coverage over grammatical phenomena across human languages. Next, we asked linguists to manually create a set of 605 paragraphs ($\sim 2,200$ sentences) that employ these functions in 5 source languages across 5 domains. Finally, these texts are multiway-parallelized across 8 pivot languages and professionally translated into 109 low-resourced languages, with the goal of realizing language-specific forms of the underlying phenomena in each language. MeDLEy covers 92 languages not in prior datasets such as SMOL (Caswell et al., 2025) or NLLB Seed. In controlled experiments, MeDLEy matches or outperforms existing seed datasets and shows particular gains for into-English translation; it also proved essential for the post-training recipe (§4.2).

3.5 Evaluation Data

BOUQuET. We use BOUQuET originally introduced by Omnilingual MT Team et al. (2025) as a manually created benchmark covering 55 languages across 8 domains with diverse registers. In this work, we extend it to 275 languages through additional professional translation, making it the largest multidomain MT evaluation benchmark. The extended benchmark and the associated leaderboard are publicly available.⁴ BOUQuET is organized into paragraphs, but in this paper, we report the results of sentence-level translation.

Bible benchmark. We use the Gospel of John as an evaluation set (split into development and test parts), covering 1,560 languages. While Bible evaluation has limitations (narrow domain, formal reg-

⁴<https://huggingface.co/spaces/facebook/bouquet>

ister), it provides the broadest available reference-based evaluation coverage.

FLORES+. We report MT results on the FLORES+ dataset (version 2.1) with 213 language varieties in its devtest split.

Resource levels. When reporting MT evaluation results, we sometimes group them by language resource levels: “high” **H** (over 50M parallel sentences in the training data) “mid” **M** (over 1M), “low” **L** (over 40K), and “very low” **VL** (less than 40K parallel sentences). The coverage of language resource levels by the benchmarks is reported in Table 4.

Directions. With each benchmark dataset, we evaluate translation from all language varieties into English (XX-En) and from English into each language variety (En-YY).

4 Modeling

4.1 Vocabulary Extension

A critical prerequisite for omnilingual translation is adequate vocabulary coverage. The original LLAMA3 tokenizer produces demonstrates unacceptably high fertility for underserved scripts. We address this by: (1) adjusting pre-tokenization to better handle rare writing systems; and (2) extending from 128K to 256K tokens via continued byte-pair encoding (BPE) merging on balanced multilingual data. This reduces the average number of tokens per sentence on the FLoRes+ dataset almost twice, from 80.7 to 44.8. New token embeddings are initialized by averaging their sub-token decompositions (Gee et al., 2022).

Table 5 shows controlled ablation results on LLAMA3 1B Instruct (30K steps of CPT, identical data). The extended tokenizer achieves 26% relative ChrF++ improvement for En-YY (17.8→22.5) and 7% for XX-En.

4.2 OMT-LLaMA: Decoder-Only Models

Continual pretraining (CPT). OMT-LLAMA extends LLAMA3 3.1 8B (Grattafiori et al., 2024) with the 256K tokenizer via CPT on language modeling and translation tasks for 50K steps on 256 GPUs, preceded by a 10K-step embedding warmup on 80 GPUs with the decoder frozen. We use packed token-based batching at 8192 tokens per data-parallel rank (~2.1M tokens per optimizer step, no gradient accumulation), totalling ~105B

tokens. We apply the same recipe at 1B and 3B scales.⁵

Post-training (SFT + RL). We apply supervised fine-tuning (SFT) on a mixture of ~600K diverse instruction-tuning examples (covering 10 languages) and translation data from SMOL and MeDLEy. SFT recovers instruction-following capabilities and biases the model towards high-quality translations, as seen in Table 6.

We further apply reinforcement learning using DAPO⁶ (Yu et al., 2025) with a combined ChrF++/MetricX reward signal and Kullback–Leibler (KL) regularization. Table 6 shows that SFT and RL yield cumulative improvements, with RL providing additional gains on hard subsets without degrading any direction.

Retrieval-augmented generation (RAG). At inference time, we optionally retrieve parallel examples from our training data using a combination of full-text search and OMNISONAR embedding similarity, inserting them as few-shot examples. RAG provides complementary gains to fine-tuning, particularly for XX-En translation (§6.5).

4.3 OMT-NLLB: Encoder-Decoder Model

We train OMT-NLLB, a compact 3B-parameter encoder-decoder model that translates from 1,600 source languages into 250 target languages. Our key innovation is a three-stage training strategy that exploits non-parallel data through a cross-lingually aligned encoder.

Stage 1: Decoder enhancement with auto-encoding. We initialize from OMNISONAR, which maps 1,600+ languages into a shared sentence embedding space. Keeping the encoder frozen, we train the decoder on parallel MT data and monolingual data via auto-encoding—the decoder reconstructs monolingual input from the OMNISONAR sentence representations. This substantially increases training data for low-resource languages and boosts the generation fluency.

Stage 2: Removing the sentence-level bottleneck. The original OMNISONAR architecture uses pooled sentence-level representations for

⁵While the current paper experiments only with the LLAMA3 family, all the vocabulary extension and training components of the OMT-LLAMA recipe are applicable to any family of subword-level autoregressive language models, regardless of their architecture. We leave that for future work.

⁶Decoupled Clip and Dynamic sAmpling Policy Optimization

Benchmark	Domains	Sentences	Languages	by resource level			
				H	M	L	VL
BOUQuET	8	854	275	23	72	79	101
FLoRes+	3	1012	213	24	82	73	34
Bible	1	479	1,560	25	71	172	1,293

Table 4: Summary of the MT evaluation benchmarks. All three benchmarks are fully parallel with English. The number of sentences is indicated for the test split.

	En-YY					XX-En				
	H	M	L	VL	All	H	M	L	VL	All
128K	39.0	17.0	12.2	12.7	17.8	54.1	35.3	31.0	31.5	35.9
256K	41.2	23.2	16.7	15.5	22.5	54.8	39.1	33.6	33.1	38.7

Table 5: Tokenizer ablation (ChrF++ on FLoRes+). Extending from 128K to 256K tokens yields 26% relative improvement for En-YY.

	BOUQuET		FLoRes+		Bible	
	XX-En	En-YY	XX-En	En-YY	XX-En	En-YY
CPT	.725	.621	.735	.523	.683	.609
+SFT	.732	.611	.735	.539	.685	.656
+RL	.741	.616	.747	.541	.689	.661

Table 6: Post-training results (xCOMET). SFT and RL yield cumulative improvements across benchmarks.

cross-attention, limiting fine-grained information transfer. We remove this bottleneck and introduce token-level encoder-decoder attention. A decoder warm-up phase (encoder frozen) allows adaptation to full-sequence attention without training instabilities.

Stage 3: End-to-end fine-tuning. Both encoder and decoder are jointly updated on parallel data, enabling full optimization of token-level interactions.

Table 7 shows progressive improvements at each stage over the OMNISONAR baseline.

5 Proposed Evaluation

XSTS+R+P human evaluation protocol. We extend the established XSTS (cross-lingual semantic text similarity) protocol (Licht et al., 2022) with two additions designed to address systematic biases in previous MT evaluation protocols. The register (R) component adds scoring criteria that account for pragmatic discrepancies such as formality mismatches between source and translation. The paragraph-level (P) component adds consistency checks that penalize translations with sentence-level semantic equivalence but paragraph-level inconsistencies, such as inconsistent grammatical gender assignment to referents, shifting formality

	FLoRes+			Bible		
	ChrF	xCOM	MtrX	ChrF	xCOM	MtrX
OMNISONAR	47.5	.65	5.66	40.6	.52	8.79
+Stage 1	49.7	.65	5.48	45.9	.53	8.59
(w/o autoencoding)	49.6	.65	5.51	44.8	.52	8.61
+Stage 2	50.1	.65	5.36	46.4	.51	8.58
OMT-NLLB	50.9	.66	5.32	46.3	.50	8.61

Table 7: OMT-NLLB 3-stage training results. Each stage improves over OMNISONAR. ChrF stands for ChrF++, xCOM for xCOMET, MtrX for MetricX.

registers within a paragraph, or contradictory temporal framing across sentences. XSTS+R+P uses a 5-point scale: 1 (no useful information preserved), 2 (some information preserved but translation is largely unusable), 3 (core meaning conveyed, moderate errors), 4 (good translation with minor issues), and 5 (perfect translation). XSTS+R+P achieves Krippendorff’s α of 0.80, indicating substantially higher inter-annotator agreement than XSTS (0.57) or a simplified version of MQM (Multidimensional Quality Metrics by Lommel et al. (2013)) penalty (0.44).

Met-BOUQuET. To train and evaluate automatic metrics, we collect the largest human annotation dataset for MT quality: Met-BOUQuET, comprising annotations across 161 language directions and 119 unique languages using the XSTS+R+P protocol. Round 1 covers 104 directions with diverse system outputs; Round 2 compares OMT-LLAMA against the strongest baseline per direction (chosen on the BOUQuET dev split with automatic metrics) across 57 challenging, mostly low-resourced directions.

BLASER 3: Reference-Free Quality Estimation

Existing reference-free QE metrics (xCOMET, MetricX) use encoders with limited cross-lingual coverage. BLASER 3 is a small dense model (inspired by Dale and Costa-jussà (2024)) built on top of frozen OMNISONAR embeddings (1,600+ languages) and is trained multi-task on 6 evaluation protocols (such as DA and MQM) using human

annotations from the dev split of Met-BOUQuET, annotated data from [NLLB Team et al. \(2024\)](#) and WMT evaluation campaigns, and synthetic data scored with reference-based MetricX.

On Met-BOUQuET, BLASER 3 achieves average per-direction item-level Spearman correlation of 0.55 with human judgments, compared to 0.47 for MetricX and 0.43 for xCOMET, making it our primary reference-free metric.

OmniTOX OmniTOX extends MuTox ([Costa-jussà et al., 2024](#)) using OMNISONAR embeddings, providing binary toxicity classification for 1,600 languages—an eight-fold increase in language coverage over its predecessor, which supported approximately 200 languages. The model applies a classification head on top of OMNISONAR sentence embeddings, trained on toxicity annotations spanning multiple languages and domains. OmniTOX achieves a mean per-language ROC AUC of 0.86 across 30 diverse evaluation languages, outperforming MuTox by +0.06 mean per-language ROC AUC.

6 Results

We evaluate on BOUQuET (275 languages), FLoRes+ (213 languages), and Bible (1,560 languages), using ChrF++ for reference-based and BLASER 3 for reference-free evaluation, complemented by xCOMET ([Guerreiro et al., 2024](#)) and MetricX-24 ([Juraska et al., 2024](#)). Baselines include NLLB-200 ([NLLB Team et al., 2024](#)), LLAMA3 ([Grattafiori et al., 2024](#)), Gemma 3 ([Gemma Team et al., 2025](#)), GPT-OSS 20B/120B ([OpenAI, 2025](#)), MADLAD ([Kudugunta et al., 2023](#)), Aya-101 ([Üstün et al., 2024](#)), Tower+ ([Alves et al., 2024](#)), TranslateGemma ([Google Translate Research Team et al., 2026](#)), and Hunyuan-MT ([Zheng et al., 2025](#)).⁷

6.1 Standard Benchmarks

Table 8 reports ChrF++ on BOUQuET by resource level. For XX-En, both OMT-LLAMA and OMT-NLLB outperform all baselines at every level (except the high-resource languages where a larger specialized Tower+ model leads). For En-YY, OMT-NLLB excels on high/mid-resource languages while OMT-LLAMA is stronger on low/very-low-resource languages.

⁷For LLMs that do not provide their own translation prompt template, we used a simple prompt template defined in Appendix E.

	XX-En				En-YY			
	H	M	L	VL	H	M	L	VL
OMT-LLAMA 8B	65.1	53.5	38.7	27.6	60.7	45.8	30.8	18.8
OMT-NLLB	64.0	56.1	37.4	28.4	61.1	47.6	27.3	17.5
OMT-LLAMA 3B	64.4	54.6	38.8	28.0	59.4	45.0	28.7	16.8
GPT-OSS 120B	64.6	49.9	33.6	25.4	63.6	43.7	26.0	16.0
Gemma 3	65.3	51.3	34.8	27.4	62.6	40.9	24.0	14.8
NLLB-200	61.5	52.4	32.1	25.5	62.9	46.8	24.6	17.5
LLAMA 3 70B	65.0	47.6	33.3	26.2	60.2	37.2	23.7	16.2
OMT-LLAMA 1B	63.0	51.5	34.7	24.1	56.5	42.2	25.1	14.9
Tower+ 72B	66.5	43.3	31.7	25.2	58.9	26.9	15.2	11.3
LLAMA 3 8B	61.0	39.8	27.5	21.6	52.7	26.0	13.1	9.5

Table 8: BOUQuET results (ChrF++) by resource level. OMT models lead across nearly all categories. Selected baselines; full results in Appendix C.

Evaluation with FLoRes+ (Table 13 in the appendix) provides qualitatively similar results.

6.2 Long-Tail Understanding and Generation

We evaluate on the Bible benchmark (1,560 languages) to assess performance on the long tail of languages.

Understanding (XX-En). Figure 1 (as well as Table 14 in the appendix) shows that OMT-LLAMA and OMT-NLLB consistently outperform all baselines across 1,560 languages. OMT-LLAMA 8B strictly outperforms all baselines in 1,045 languages. OMT-NLLB outperforms all baselines in XX-En across all 1,560 evaluated languages despite its smaller size.

To understand how many languages our models understand “sufficiently well” for practical use, we define two quality thresholds: average XSTS+R+P score of 3.5 (“good”) and 2.5 (“passable”). We use the BOUQuET test set to map these values to the more neutral MetricX metric,⁸ obtaining thresholds of around 4.0 and 6.7, respectively. With them, we estimate that OMT-LLAMA achieves “passable” quality for 440 languages and OMT-NLLB for 416—nearly double the 221 for NLLB-200. Note that due to the potential fragility of double extrapolation, these numbers should be treated as orders of magnitude, not as exact values.

Generation (En-YY). Figure 2 reveals a striking generation bottleneck: baseline models degrade to near-random quality at 300–400 languages, while OMT-LLAMA maintains non-trivial quality for ~1,200. Determining the “passable” generation

⁸Appendix D explains why we find the extrapolations between BOUQuET and Bible datasets and between MetricX and XSTS+R+P scores reliable enough.




Figure 1: Long-tail **understanding** (XX-En, ChrF++ on Bible). OMT models (blue/orange) outperform baselines across 1,560 languages. Languages sorted by average performance; curves smoothed by EWM with a delay of 10.




Figure 2: Long-tail **generation** (En-YY, ChrF++ on Bible). Baselines degrade to random at 300–400 languages; OMT-LLAMA (orange) sustains quality for ~1,200 languages.

quality levels across such diverse output languages would require more human evaluation and is deferred to future work.

6.3 Model Size vs. Performance

Figure 3 shows the Pareto frontier of model size vs. quality on BOUQuET. The 1B OMT-LLAMA matches baselines $3\times$ – $8\times$ larger, and the 8B OMT-LLAMA outperforms the 70B LLaMA3—demonstrating that *specialization* outperforms scale for multilingual MT.

6.4 Added toxicity

Using OmniTOX, we systematically benchmarked added toxicity—the increase in toxicity between source text and translation—across 207 languages




Figure 3: Model size vs. translation quality (ChrF++ on BOUQuET, average of En-YY and XX-En). OMT models define the Pareto frontier—specialized 1B–8B models match or exceed 70B general-purpose LLMs.

for evaluated systems. We compute debiased toxicity logit differences between source and translation, with per-language baselines established on non-toxic BOUQuET sentences and flagging thresholds calibrated at a 5% false positive rate. None of the evaluated systems exhibit meaningful added toxicity, with flagging rates below 1.5% across all model configurations and language directions.

6.5 Extensibility

An *omnilingual* model should provide broad coverage, but especially allow for targeted improvement for specific languages where needed. We evaluate fine-tuning (FT) and RAG with 8B LLAMA3 and OMT-LLAMA on 5 low-resource (L) and 20 very low-resource (VL) languages using available parallel data (11K–300K examples per language, paired with English), as well as fine-tuning OMT-NLLB with the same data.

As Table 9 shows, fine-tuning OMT-NLLB improves the results in all translation direction groups (except the 5 higher-resourced languages into English, where the scores are already very good), and makes the fine-tuned model the winner in 3 of them. Fine-tuning OMT-LLAMA improves En-YY by +4.8 ChrF++ but slightly degrades XX-En.⁹ RAG improves XX-En by +6.3 ChrF++. Combining both yields the best results: +5.1 En-YY and +6.8 XX-En. Crucially, even with both techniques, the

⁹Following a similar observation in the SMOL paper (Caswell et al., 2025), we conjecture that this degradation is caused by the repetitive English target texts in the training dataset. Our experiments indicate that the problem should be mitigated by non-English-centric training data composition.

System	20 VL. languages		5 L. languages	
	En-YY	XX-En	En-YY	XX-En
OMT-LLaMA + FT + RAG	29.3	36.7	46.7	53.2
OMT-LLaMA + FT	28.9	29.0	47.5	42.7
OMT-NLLB + FT	28.2	37.9	52.4	56.7
OMT-LLaMA + RAG	26.8	36.2	44.8	54.0
OMT-NLLB	24.4	35.8	51.4	56.8
OMT-LLaMA	24.1	29.9	43.0	48.4
LLaMA-base + FT + RAG	20.1	27.9	33.2	40.9
LLaMA-base + FT	19.4	26.4	31.2	38.8
LLaMA-base + RAG	17.5	22.9	25.0	31.1
LLaMA-base	9.3	20.7	16.1	28.7

Table 9: Results of extending the OMT-LLaMA, LLaMA-base, and OMT-NLLB models (BOUQuET, dev split, ChrF++) with a sample of low-resourced and very low-resourced languages.

Direction group	Win%	OMT	Base	N
high → high	0.75	4.03	3.45	16
high → low	0.83	2.99	1.61	18
low → high	0.80	3.26	2.69	15
low → low	0.62	2.96	2.66	8
Total	0.77	3.38	2.67	57

Table 10: Human evaluation (Met-BOUQuET Round 2). Win rate = proportion of directions where OMT-LLaMA’s mean XSTS+R+P score exceeds the baseline’s. Scores are on a 1–5 scale; N denotes the number of translation directions. In this table, “high” denotes H+M languages, and “low”, L+VL.

base LLaMA3 does not outperform unadapted OMT-LLaMA, confirming that OMT-LLaMA is the superior starting point for language-specific adaptation.

6.6 Human Evaluation

We conduct human evaluation as the Round 2 of Met-BOUQuET, comparing OMT-LLaMA against the strongest baseline per direction across 57 directions and 80 languages (Table 10). OMT-LLaMA achieves a 77% win rate overall, with the largest gains for high→low directions (83%), where baselines fail to produce meaningful translations. The improvement from 2.67 to 3.38 mean XSTS+R+P score represents a jump from “minimally useful” to “useful.”

7 Conclusion

Omnilingual MT demonstrates that scaling multilingual translation to 1,600+ languages requires coordinated advances across data, modeling, and evaluation. Our data strategy—combining massive public corpora with backtranslation for 2,000+ languages and the manually curated MeDLEy dataset—addresses the generation bottle-

neck for long-tail languages. Our two architectural approaches—OMT-LLaMA (decoder-only, 1B/3B/8B) and OMT-NLLB (encoder-decoder, 3B with novel three-stage training)—deliver complementary strengths: OMT-LLaMA covers more output languages while OMT-NLLB excels at cross-lingual understanding.

Large-scale evaluation reveals that Omnilingual MT doubles the number of “passable quality” language varieties from ~200 to over 440, and maintains non-trivial generation quality for ~1,200 languages where baselines support 3–4 times less. This highlights that genuinely inclusive MT is a challenge that requires targeted data creation (such as MeDLEy and backtranslation) and innovative use of monolingual data (bottlenecked autoencoding) rather than simply scaling model parameters.

Our specialized 1B–8B models match or exceed the comparable 70B general-purpose LLM in massively multilingual MT, offering a clear Pareto advantage. Therefore, model specialization through continual pretraining and post-training provides a more efficient path to multilingual translation quality than “simply” increasing model and data size.

The human evaluation confirms a 77% win rate of OMT against strongest baselines across 57 diverse translation directions. Our evaluation suite, BOUQuET (275 languages) and Met-BOUQuET (161 directions), provides the community with tools to assess MT at unprecedented scale and develop even better translation systems and metrics.

Finally, we show that post-training techniques targeting multilingual extension of models do not compensate the lack of large improvements obtained by embracing massively multilingual training data and vocabulary in the earlier stages of training. This should motivate model developers interested in boosting multilingual tasks performance to develop models that are highly multilingual by design and are well equipped to extend their support to any additional language, if necessary.

Looking forward, we believe that foundational models should aim to provide at least a base level of support for *every* language—not only those with abundant data and rich markets. Omnilingual MT positions large-scale linguistic inclusion as an ongoing challenge demanding continued investment in data creation, architecture design, and evaluation methodology. The recipes presented here can in principle be reproduced on top of diverse base models, and we hope they inspire further research toward genuinely universal language technology.

Limitations and Ethical Considerations

Omnilingual MT has several limitations. OMT-NLLB generates only 250 target languages vs. $\sim 1,200$ for OMT-LLAMA. Both models are based on the Llama 3 family, now two years old, which makes it stale by the LLM standards (our early experiments indicate that the OMT recipe applies to more modern LLM families, but we leave the full exploration for future work). The training data overrepresents formal registers, and religious domains, and, probably, contains some mislabeled texts and obsolete orthographies. We did not exhaustively optimize the prompts of the baseline models.¹⁰ The extensibility results cover only 25 languages.

Our evaluation is also limited. The long-tail evaluation relies on Bible text (domain-limited, formal register), and human evaluation covers only 57 directions. While we expand the number of languages with “passable” understanding from ~ 200 to ~ 440 , this still covers only a fraction of the world’s languages, and these numbers themselves are subject to debate. And determining the target languages for which our models meet the “passable” bar of generation quality would require either more human evaluation or extensive research into automated multilingual metrics of fluency evaluation.

Because of the limited domain and style coverage of the test datasets and the fragility of automatic metrics, we advise against using the OMT models “out of the box” in high-stakes applications, without first evaluating their quality with domain- and language-specific test sets and, if possible, manual evaluation by native speakers. We believe that targeted fine-tuning of the OMT models will be effective for adapting them to specific domains and languages, but the results are still to be evaluated on a case-by-case basis.

At the moment of writing this text, some of the artifacts described in the paper are not yet publicly available, and some might never be released. To enable at least partial reproducibility, Appendix A presents a mapping between the private resources that we have used and the best publicly available analogs.

Due to the limited paper size, some parts of

¹⁰Our preliminary experiments showed that zero-shot translation with LLM has two main failure modes: translation into a wrong target language and providing verbose explanations alongside the translations. Our chosen prompt template, reported in Appendix E, mitigates both issues for the models evaluated in the paper.

the work were presented here only at a high level, which limits their reproducibility. A longer and more detailed technical report (Omnilingual MT Team et al., 2026) can be found at <https://arxiv.org/abs/2603.16309>.

The authors of this paper used AI assistants for writing (preparing a draft of the paper based on the longer technical report) and for creating a part of the research codebase. All generated text and code has been reviewed by the authors.

The human annotators were paid a fair rate.

References

- Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024. *Tower: An open multilingual large language model for translation-related tasks*. Preprint, arXiv:2402.17733.
- Pierre Andrews, Guillaume Wenzek, Kevin Heffernan, Onur Çelebi, Anna Sun, Ammar Kamran, Yingzhe Guo, Alexandre Mourachko, Holger Schwenk, and Angela Fan. 2022. *stopes - modular machine translation pipelines*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 258–265, Abu Dhabi, UAE. Association for Computational Linguistics.
- Laurie Burchell, Jean Maillard, Antonios Anastasopoulos, Christian Federmann, Philipp Koehn, and Skyler Wang. 2024. *Findings of the WMT 2024 shared task of the open language data initiative*. In *Proceedings of the Ninth Conference on Machine Translation*, pages 110–117, Miami, Florida, USA. Association for Computational Linguistics.
- Isaac Caswell, Elizabeth Nielsen, Jiaming Luo, Colin Cherry, Geza Kovacs, Hadar Shemtov, Partha Talukdar, Dinesh Tewari, Baba Mamadi Diane, Koulako Moussa Doumbouya, Djibrila Diane, Solo Farabado Cissé, Edoardo Ferrante, Alessandro Guasoni, Mamadou K. Keita, Sudhamoy DebBarma, Ali Kuzhuget, David Anugraha, Muhammad Ravi Shulthan Habibi, and 3 others. 2025. *SMOL: Professionally translated parallel data for 115 under-represented languages*. Preprint, arXiv:2502.12301.
- Emily Chang and Niyati Bafna. 2026. *ChiKhaPo: A large-scale multilingual benchmark for evaluating lexical comprehension and generation in large language models*. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33720–33754, San Diego, California, United States. Association for Computational Linguistics.
- Marta R. Costa-jussà, Mariano Coria Meglioli, Pierre Andrews, David Dale, Prangthip Hansanti, Elahe

- Kalbassi, Alex Mourachko, Christophe Ropers, and Carleigh Wood. 2024. [MuTox: Universal Multilingual audio-based TOXicity dataset and zero-shot detector](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5725–5734, Bangkok, Thailand. Association for Computational Linguistics.
- David Dale, Laurie Burchell, Jean Maillard, Idris Abdulmumin, Antonios Anastasopoulos, Isaac Caswell, and Philipp Koehn. 2025. [Findings of the WMT 2025 shared task of the open language data initiative](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 495–502, Suzhou, China. Association for Computational Linguistics.
- David Dale and Marta R. Costa-jussà. 2024. [BLASER 2.0: a metric for evaluation and quality estimation of massively multilingual speech and text translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16075–16085, Miami, Florida, USA. Association for Computational Linguistics.
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya expand: Combining research breakthroughs for a new multilingual frontier](#). *Preprint*, arXiv:2412.04261.
- Paul-Ambroise Duquenne, Holger Schwenk, and Benoit Sagot. 2023. [SONAR: sentence-level multimodal and language-agnostic representations](#).
- Zhaopeng Feng, Shaosheng Cao, Jiahao Ren, Jiayuan Su, Ruizhe Chen, Yan Zhang, Zhe Xu, Yao Hu, Jian Wu, and Zuozhu Liu. 2025. [Mt-r1-zero: Advancing llm-based machine translation via r1-zero-like reinforcement learning](#). *Preprint*, arXiv:2504.10160.
- Markus Frohmann, Igor Sterner, Ivan Vulić, Benjamin Minixhofer, and Markus Schedl. 2024. [Segment any text: A universal approach for robust, efficient and adaptable sentence segmentation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11908–11941, Miami, Florida, USA. Association for Computational Linguistics.
- Leonidas Gee, Andrea Zugarini, Leonardo Rigutini, and Paolo Torroni. 2022. [Fast vocabulary transfer for language model compression](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 409–416, Abu Dhabi, UAE. Association for Computational Linguistics.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 Technical Report](#). *Preprint*, arXiv:2503.19786.
- Google Translate Research Team, Mara Finkelstein, Isaac Caswell, Tobias Domhan, Jan-Thorsten Peter, Juraj Juraska, Parker Riley, Daniel Deutsch, Cole Dilanni, Colin Cherry, Eleftheria Briakou, Elizabeth Nielsen, Jiaming Luo, Sweta Agrawal, Wenda Xu, Erin Kats, Stephane Jaskiewicz, Markus Freitag, and David Vilar. 2026. [TranslateGemma Technical Report](#).
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- 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.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Alexander Jones, Isaac Caswell, Orhan Firat, and Ishank Saxena. 2023. [GATITOS: Using a new multilingual lexicon for low-resource machine translation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 371–405, Singapore. Association for Computational Linguistics.
- Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. 2024. [MetricX-24: The Google submission to the WMT 2024 metrics shared task](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.
- David Kamholz, Jonathan Pool, and Susan Colowick. 2014. [PanLex: Building a resource for panlingual lexical translation](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 3145–3150, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Amir Hossein Kargaran, Ayyoob Imani, François Yvon, and Hinrich Schuetze. 2023. [GlotLID: Language](#)

- identification for low-resource languages. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6155–6218, Singapore. Association for Computational Linguistics.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. Madlad-400: A multilingual and document-level large audited dataset. *Advances in Neural Information Processing Systems*, 36:67284–67296.
- Hynek Kydlíček, Guilherme Penedo, and Leandro von Werra. 2025. Finepdfs. <https://huggingface.co/datasets/HuggingFaceFW/finepdfs>.
- Daniel Licht, Cynthia Gao, Janice Lam, Francisco Guzman, Mona Diab, and Philipp Koehn. 2022. Consistent human evaluation of machine translation across language pairs. In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 309–321, Orlando, USA. Association for Machine Translation in the Americas.
- Arle 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.
- Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. 2024. Fineweb-edu: the finest collection of educational content.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, and 20 others. 2024. Scaling neural machine translation to 200 languages. *Nature*, 630(8018):841–846.
- Dayyán O’Brien, Bhavitvya Malik, Ona de Gibert, Pinzhen Chen, Barry Haddow, and Jörg Tiedemann. 2025. DocHPLT: A massively multilingual document-level translation dataset. *arXiv preprint*.
- Omnilingual MT Team, Belen Alastruey, Niyati Bafna, Andrea Caciolai, Kevin Heffernan, Artyom Kozhevnikov, Christophe Ropers, Eduardo Sánchez, Charles-Eric Saint-James, Ioannis Tsiamas, Xiang "Tony" Cao, Chierh Cheng, Joe Chuang, Paul-Ambroise Duquenne, Mark Duppenhaler, Nate Ekberg, Cynthia Gao, Pere Lluís Huguet Cabot, João Maria Janeiro, and 13 others. 2026. Omnilingual mt: Machine translation for 1,600 languages. *Preprint*, arXiv:2603.16309.
- Omnilingual MT Team, Pierre Andrews, Mikel Artetxe, Mariano Coria Meglioli, Marta R. Costa-jussà, Joe Chuang, David Dale, Mark Duppenhaler, Nathaniel Paul Ekberg, Cynthia Gao, Daniel Edward Licht, Jean Maillard, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Eduardo Sánchez, Ioannis Tsiamas, Arina Turkatenco, Albert Ventayol-Boada, and Shireen Yates. 2025. BOUQuET : dataset, benchmark and open initiative for universal quality evaluation in translation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27515–27535, Suzhou, China. Association for Computational Linguistics.
- Omnilingual SONAR Team, João Maria Janeiro, Pere-Lluís Huguet Cabot, Ioannis Tsiamas, Yen Meng, Vivek Iyer, Guillem Ramírez, Loic Barrault, Belen Alastruey, Yu-An Chung, Marta R. Costa-Jussa, David Dale, Kevin Heffernan, Jaehyeong Jo, Artyom Kozhevnikov, Alexandre Mourachko, Christophe Ropers, Holger Schwenk, and Paul-Ambroise Duquenne. 2026. Omnilingual SONAR: Cross-lingual and cross-modal sentence embeddings bridging massively multilingual text and speech. *Preprint*, arXiv:2603.16606.
- OpenAI. 2025. gpt-oss-120b & gpt-oss-20b model card. *Preprint*, arXiv:2508.10925.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Guilherme Penedo, Hynek Kydlíček, Amir Hossein Kargaran, and Leandro von Werra. 2026. Finetranslations. <https://huggingface.co/datasets/HuggingFaceFW/finetranslations>.
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargaran, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. 2025. Fineweb2: One pipeline to scale them all – adapting pre-training data processing to every language. *Preprint*, arXiv:2506.20920.
- 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.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. 2021. CCMatrix: Mining billions of high-quality parallel sentences on the web. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6490–6500, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96,

Berlin, Germany. Association for Computational Linguistics.

Mexico City, Mexico. Association for Computational Linguistics.

Aditya Siddhant, Ankur Bapna, Orhan Firat, Yuan Cao, Mia Xu Chen, Isaac Caswell, and Xavier Garcia. 2022. [Towards the next 1000 languages in multilingual machine translation: Exploring the synergy between supervised and self-supervised learning](#). *arXiv preprint arXiv:2201.03110*.

Gemma Team, Sherif El Abd, Vaibhav Aggarwal, Robin Algayres, Alek Andreev, Olivier Bachem, Ian Ballantyne, Cormac Brick, Victor Cărbune, Michelle Casbon, and 1 others. 2026. [Gemma 4 technical report](#). *arXiv 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).

Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D'souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction fine-tuned open-access multilingual language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15894–15939, Bangkok, Thailand. Association for Computational Linguistics.

Jiaan Wang, Fandong Meng, and Jie Zhou. 2025. [Deeptrans: Deep reasoning translation via reinforcement learning](#). *Preprint*, arXiv:2504.10187.

Haoran Xu, Young Jin Kim, Amr Mohamed Nabil Aly Aly Sharaf, and Hany Awadalla. 2024. A paradigm shift in machine translation: Boosting translation performance of large language models. In *International Conference on Learning Representations*, volume 2024, pages 2747–2767.

Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *Preprint*, arXiv:2503.14476.

Mao Zheng, Zheng Li, Bingxin Qu, Mingyang Song, Yang Du, Mingrui Sun, and Di Wang. 2025. [Hunyuan-mt technical report](#). *Preprint*, arXiv:2509.05209.

Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2024. [Multilingual machine translation with large language models: Empirical results and analysis](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2765–2781,

A Resources availability and substitutes

Unfortunately, we cannot promise neither complete reproducibility nor releasing all the resources described in the paper. To enable at least partial reproducibility, Table 11 presents a mapping between the private resources that we have used and the best publicly available analogues.

OMT resource	Analogue resource	Main differences
BOUQuET extension and Met-BOUQuET	The data is publicly available at hf.co/datasets/facebook/bouquet , with the accompanying evaluation code in github.com/facebookresearch/bouquet .	Not applicable.
OMT-Llama	TranslateGemma (Google Translate Research Team et al., 2026), Gemma4 (Team et al., 2026)	No direct analogues, but TranslateGemma is also MT-specialized, and the Gemma4 family of models demonstrates impressive coverage of languages.
OMT-NLLB	NLLB-200 (NLLB Team et al., 2024)	The main difference is language coverage: 200-to-200 vs 1600-to-250. OMT-NLLB is also better at longer texts than NLLB-200.
Bible benchmark	github.com/BibleNLP/ebible data	The ebible dataset would require additional filtering, and it lacks a portion of translations that we collected from other sources.
MeDLEy	SMOL (Caswell et al., 2025)	Different language coverage; different construction (focus on lexical vs grammatical diversity).
CC-2000-WEB	FineWeb 2 (Penedo et al., 2025)	Different filtering of similar underlying web data.
CC-2000-PDF	FinePDFs (Kydliček et al., 2025)	Different filtering of similar underlying web data.
CC-NLLB-200	NLLB-200 primary datasets collection (see the documentation in the “nllb” branch of the “fairseq” public repository); massively multilingual datasets from OPUS.	The versions and availability of some of the datasets might have evolved over time, so their exact re-compilation might be impossible.
OMT back-translated data	Finetranslations (Penedo et al., 2026)	A different language coverage (OMT backtranslated into and out of English, Finetranslations, only into English).
Mined paragraphs	DocHPLT (O’Brien et al., 2025); the “NLLB” dataset by AllenAI.	OMT mined data covers more languages than DocHPLT and is paragraph-level, unlike the NLLB mined dataset.
OmniSONAR	SONAR (Duquenne et al., 2023)	OmniSONAR covers more input languages than SONAR and operates in a different representation space.

Table 11: OMT resources and their closest public analogues.

B Languages at a glance

The next Table B lists the languages covered by the evaluation datasets that we contribute with this work (275 in BOUQuET, and 120 in Met-BOUQuET), our seed training dataset MeDLEy and our previous and ongoing evaluation community effort of FLoRes+.

ISO 639-3	ISO 15924	Language family	BOUQuET	MeDLEy	met-BOUQuET	FLoRes+
aar	Latn	Afro-Asiatic	✓	×	✓(R1)	×
abl	Latn	Austronesian	✓	×	✓(R2)	×
abz	Arab	Timor-Alor-Pantar	×	×	✓(R2)	×
ace	Arab	Austronesian	×	×	×	✓
ace	Latn	Austronesian	×	×	×	✓
acm	Arab	Afro-Asiatic	×	×	×	✓
acq	Arab	Afro-Asiatic	×	×	×	✓
aeb	Arab	Afro-Asiatic	×	×	×	✓
afr	Latn	Indo-European	✓	×	✓(R2)	✓
agr	Latn	Chicham	✓	×	✓(R1)	×
ahk	Thai	Sino-Tibetan	×	✓	×	×

aiq	Arab	Indo-European	✓	×	×	×
akb	Latn	Austronesian	×	✓	✓ (R2)	×
als	Latn	Indo-European	✓	×	×	✓
amh	Ethi	Afro-Asiatic	✓	×	×	✓
ami	Latn	Austronesian	✓	×	✓ (R1)	×
ane	Latn	Austronesian	✓	×	×	×
any	Latn	Atlantic-Congo	×	✓	✓ (R2)	×
apc	Arab	Afro-Asiatic	✓	×	✓ (R2)	✓
arb	Arab	Afro-Asiatic	×	✓	×	✓
arb	Latn	Afro-Asiatic	×	×	×	✓
arg	Latn	Indo-European	×	×	×	✓
arh	Latn	Chibchan	✓	×	×	×
arn	Latn	Araucanian	✓	×	×	×
ars	Arab	Afro-Asiatic	×	×	×	✓
ary	Arab	Afro-Asiatic	×	×	×	✓
arz	Arab	Afro-Asiatic	✓	×	✓ (R1)	✓
arz	Latn	Afro-Asiatic	✓	×	✓ (R1)	×
asm	Beng	Indo-European	✓	×	✓ (R2)	✓
ast	Latn	Indo-European	×	×	×	✓
ati	Latn	Atlantic-Congo	×	✓	×	×
awa	Deva	Indo-European	×	×	✓ (R2)	✓
ayr	Latn	Aymaran	✓	×	×	✓
ayz	Latn	Maybratic	✓	×	×	×
azb	Arab	Turkic	✓	✓	×	✓
azj	Latn	Turkic	✓	×	×	✓
azm	Latn	Otomanguean	✓	×	×	×
azz	Latn	Uto-Aztecan	✓	✓	×	×
bak	Cyrl	Turkic	✓	×	×	✓
bam	Latn	Mande	✓	✓	✓ (R2)	✓
ban	Latn	Austronesian	×	×	×	✓
bas	Latn	Atlantic-Congo	✓	×	×	×
bba	Latn	Atlantic-Congo	✓	✓	×	×
bbc	Latn	Austronesian	×	✓	×	×
bcc	Arab	Indo-European	×	✓	✓ (R2)	×
bel	Cyrl	Indo-European	✓	×	✓ (R2)	✓
bem	Latn	Atlantic-Congo	×	×	✓ (R2)	✓
ben	Beng	Indo-European	✓	×	✓ (R1, R2)	✓
ben	Latn	Indo-European	✓	×	×	×
bfa	Latn	Nilotic	×	✓	×	×
bft	Arab	Sino-Tibetan	✓	×	×	×
bhb	Deva	Indo-European	✓	×	✓ (R2)	×
bho	Deva	Indo-European	✓	×	×	✓
bib	Latn	Mande	×	✓	×	×
bjn	Arab	Austronesian	×	×	×	✓
bjn	Latn	Austronesian	×	×	×	✓
blt	Latn	Tai-Kadai	×	✓	×	×
bod	Tibt	Sino-Tibetan	✓	×	×	✓
bom	Latn	Atlantic-Congo	×	✓	×	×
bos	Latn	Indo-European	✓	×	×	✓
bre	Latn	Indo-European	✓	×	✓ (R2)	×
brh	Arab	Dravidian	✓	×	×	×
brx	Deva	Sino-Tibetan	✓	×	×	✓
bsh	Arab	Indo-European	✓	×	×	×
bsk	Arab	Burushaski	✓	×	×	×
bsq	Latn	Kru	×	✓	✓ (R2)	×
bug	Latn	Austronesian	×	×	×	✓
bul	Cyrl	Indo-European	✓	×	×	✓
cak	Latn	Mayan	✓	×	×	×
cat	Latn	Indo-European	✓	×	✓ (R2)	✓
ceb	Latn	Austronesian	✓	×	✓ (R2)	✓
ces	Latn	Indo-European	✓	×	✓ (R1)	✓
che	Cyrl	Nakh-Daghestanian	✓	×	×	×
chr	Cher	Iroquoian	✓	×	×	×
chv	Cyrl	Turkic	✓	×	×	✓
cja	Arab	Austronesian	✓	×	×	×
cjk	Latn	Atlantic-Congo	✓	✓	×	✓
ckb	Arab	Indo-European	✓	×	✓ (R2)	✓
ckl	Latn	Afro-Asiatic	✓	×	×	×
cmn	Hans	Sino-Tibetan	✓	×	✓ (R1)	✓
cmn	Hant	Sino-Tibetan	✓	×	✓ (R1)	✓

crh	Latn	Turkic	×	×	×	✓
crk	Cans	Algic	✓	×	✓(R1)	×
crk	Latn	Algic	✓	×	×	×
cux	Latn	Otomanguean	✓	×	×	×
cym	Latn	Indo-European	✓	×	×	✓
dan	Latn	Indo-European	✓	×	×	✓
daq	Deva	Dravidian	✓	×	×	×
dar	Cyrl	Nakh-Daghestanian	×	×	×	✓
deu	Latn	Indo-European	✓	×	✓(R1)	✓
dga	Latn	Atlantic-Congo	×	✓	✓(R2)	×
dgo	Deva	Indo-European	✓	×	×	✓
dik	Latn	Nilotic	✓	✓	✓(R2)	✓
diq	Latn	Indo-European	✓	×	×	×
div	Thaa	Indo-European	✓	×	×	×
djc	Latn	Dajuic	✓	×	×	×
dje	Latn	Songhay	✓	×	✓(R1)	×
dnj	Latn	Mande	×	✓	×	×
dtm	Latn	Dogon	✓	×	×	×
ds	Latn	Dogon	✓	×	×	×
dua	Latn	Atlantic-Congo	✓	×	×	×
dyu	Latn	Mande	×	✓	×	✓
dzo	Tibt	Sino-Tibetan	✓	✓	×	✓
ekk	Latn	Uralic	✓	×	✓(R2)	✓
ell	Grek	Indo-European	✓	×	✓(R1)	✓
enb	Latn	Nilotic	✓	×	×	×
eng	Latn	Indo-European	✓	✓	✓(R1, R2)	✓
enl	Latn	Lengua-Mascoy	✓	×	×	×
epo	Latn	Esperanto	×	×	×	✓
eto	Latn	Atlantic-Congo	✓	×	×	×
eus	Latn	Basque	✓	×	✓(R2)	✓
ewe	Latn	Atlantic-Congo	×	✓	×	✓
ewo	Latn	Atlantic-Congo	✓	×	×	×
fao	Latn	Indo-European	✓	×	×	✓
fia	Copt	Nubian	✓	×	×	×
fij	Latn	Austronesian	×	✓	×	✓
fil	Latn	Austronesian	×	×	×	✓
fin	Latn	Uralic	✓	×	×	✓
fon	Latn	Atlantic-Congo	×	✓	×	✓
fra	Latn	Indo-European	✓	✓	✓(R1, R2)	✓
fry	Latn	Indo-European	✓	×	✓(R2)	×
fuc	Latn	Atlantic-Congo	✓	×	×	×
fur	Latn	Indo-European	×	×	×	✓
fuv	Latn	Atlantic-Congo	✓	✓	✓(R2)	✓
fvr	Latn	Furan	✓	×	×	×
gax	Latn	Afro-Asiatic	✓	×	×	×
gaz	Latn	Afro-Asiatic	✓	×	✓(R1)	✓
gil	Latn	Austronesian	✓	×	✓(R1)	×
gkp	Latn	Mande	✓	×	×	×
gla	Latn	Indo-European	✓	×	×	✓
gle	Latn	Indo-European	✓	×	×	✓
glg	Latn	Indo-European	✓	×	×	✓
gom	Deva	Indo-European	✓	×	×	✓
gor	Latn	Austronesian	×	✓	×	×
grt	Latn	Sino-Tibetan	×	✓	×	×
guc	Latn	Arawakan	✓	×	✓(R1)	×
gug	Latn	Tupian	✓	×	✓(R2)	✓
guj	Gujr	Indo-European	✓	×	×	✓
guz	Latn	Atlantic-Congo	✓	×	×	×
gxx	Latn	Kru	✓	×	×	×
hat	Latn	Indo-European	✓	×	✓(R2)	✓
hau	Latn	Afro-Asiatic	✓	×	✓(R2)	✓
heb	Hebr	Afro-Asiatic	✓	×	✓(R2)	✓
heh	Latn	Atlantic-Congo	✓	✓	×	×
hig	Latn	Afro-Asiatic	×	✓	×	×
hin	Deva	Indo-European	✓	✓	✓(R1, R2)	✓
hin	Latn	Indo-European	✓	×	✓(R1)	×
hne	Deva	Indo-European	✓	×	×	✓
hrv	Latn	Indo-European	✓	×	✓(R1)	✓
hun	Latn	Uralic	✓	×	✓(R1)	✓
hve	Latn	Huavean	✓	×	×	×

hye	Armn	Indo-European	✓	×	✓ (R2)	✓
ibo	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
ijc	Latn	Ijoid	✓	×	×	×
ilo	Latn	Austronesian	✓	×	×	✓
ind	Latn	Austronesian	✓	✓	✓ (R1, R2)	✓
irk	Latn	Afro-Asiatic	✓	✓	×	×
isl	Latn	Indo-European	✓	×	×	✓
ita	Latn	Indo-European	✓	×	✓ (R1, R2)	✓
jav	Latn	Austronesian	✓	×	✓ (R1)	✓
jmc	Latn	Atlantic-Congo	✓	✓	×	×
jnj	Latn	Ta-Ne-Omoti	✓	×	×	×
jpn	Jpan	Japanic	✓	×	✓ (R1)	✓
kaa	Cyrl	Turkic	✓	×	✓ (R1)	✓
kab	Latn	Afro-Asiatic	×	×	×	✓
kac	Latn	Sino-Tibetan	✓	✓	×	✓
kai	Latn	Afro-Asiatic	✓	×	×	×
kal	Latn	Eskimo-Aleut	✓	×	✓ (R1)	×
kam	Latn	Atlantic-Congo	✓	✓	×	✓
kan	Knda	Dravidian	✓	×	×	✓
kas	Arab	Indo-European	×	×	×	✓
kas	Deva	Indo-European	×	×	×	✓
kat	Geor	Kartvelian	✓	×	✓ (R2)	✓
kaz	Cyrl	Turkic	✓	×	×	✓
kbp	Latn	Atlantic-Congo	×	✓	×	✓
kde	Latn	Atlantic-Congo	×	✓	×	×
kdj	Latn	Nilotic	×	✓	✓ (R2)	×
kea	Latn	Indo-European	✓	×	×	✓
kek	Latn	Mayan	✓	✓	×	×
khk	Cyrl	Mongolic-Khit	✓	×	×	✓
khm	Khmr	Austroasiatic	✓	×	✓ (R1)	✓
khq	Latn	Songhay	✓	×	×	×
khw	Latn	Indo-European	✓	×	×	×
kik	Latn	Atlantic-Congo	×	×	×	✓
kin	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
kir	Cyrl	Turkic	✓	×	×	✓
klx	Arab	Indo-European	✓	×	×	×
kmb	Latn	Atlantic-Congo	✓	✓	×	✓
kmr	Latn	Indo-European	✓	×	×	✓
knc	Arab	Saharan	✓	✓	✓ (R2)	✓
knc	Latn	Saharan	×	×	×	✓
knw	Latn	Kxa	✓	×	×	×
kor	Kore	Koreanic	✓	×	✓ (R1)	✓
krt	Latn	Saharan	✓	×	×	×
kru	Deva	Dravidian	✓	×	✓ (R1)	×
ksf	Latn	Atlantic-Congo	✓	×	×	×
ktu	Latn	Atlantic-Congo	✓	✓	×	✓
kuj	Latn	Atlantic-Congo	✓	×	×	×
kus	Latn	Atlantic-Congo	×	✓	×	×
kwy	Latn	Atlantic-Congo	✓	×	×	×
kxp	Arab	Indo-European	✓	×	×	×
lao	Lao	Tai-Kadai	✓	×	×	✓
led	Latn	Central Sudanic	✓	×	×	×
lgg	Latn	Central Sudanic	✓	✓	×	×
lia	Latn	Atlantic-Congo	×	✓	×	×
lij	Latn	Indo-European	✓	×	✓ (R1)	✓
lim	Latn	Indo-European	✓	×	×	✓
lin	Latn	Atlantic-Congo	✓	×	✓ (R1, R2)	✓
lir	Latn	Pidgin	✓	×	✓ (R2)	×
lit	Latn	Indo-European	✓	×	×	✓
lld	Latn	Indo-European	×	×	×	✓
lmo	Latn	Indo-European	×	×	×	✓
loa	Latn	North Halmahera	✓	×	×	×
loh	Latn	Surmic	✓	×	×	×
lon	Latn	Atlantic-Congo	×	✓	×	×
ltg	Latn	Indo-European	×	×	×	✓
ltz	Latn	Indo-European	×	×	×	✓
lua	Latn	Atlantic-Congo	×	✓	✓ (R2)	✓
lug	Latn	Atlantic-Congo	✓	✓	×	✓
luo	Latn	Nilotic	✓	✓	✓ (R2)	✓
lus	Latn	Sino-Tibetan	×	×	×	✓

lvs	Latn	Indo-European	✓	×	×	✓
mad	Latn	Austronesian	×	✓	✓ (R2)	×
maf	Latn	Afro-Asiatic	✓	×	×	×
mag	Deva	Indo-European	×	×	×	✓
mah	Latn	Austronesian	×	✓	×	×
mai	Deva	Indo-European	✓	×	×	✓
mak	Latn	Austronesian	×	✓	×	×
mal	Mlym	Dravidian	✓	×	✓ (R2)	✓
mam	Latn	Mayan	✓	✓	×	×
mar	Deva	Indo-European	✓	×	✓ (R2)	✓
mas	Latn	Nilotic	✓	×	×	×
men	Latn	Mande	×	✓	×	×
mey	Latn	Afro-Asiatic	✓	×	×	×
mfe	Latn	Indo-European	×	×	×	✓
mhr	Cyrl	Uralic	×	×	×	✓
mie	Latn	Otomanguean	✓	×	×	×
mim	Latn	Otomanguean	×	✓	✓ (R2)	×
min	Arab	Austronesian	✓	×	✓ (R2)	✓
min	Latn	Austronesian	×	×	×	✓
mio	Latn	Otomanguean	×	✓	×	×
miq	Latn	Misumalpan	✓	✓	×	×
mkd	Cyrl	Indo-European	✓	×	×	✓
mlt	Latn	Afro-Asiatic	✓	×	×	✓
mni	Beng	Sino-Tibetan	×	×	×	✓
mni	Mtei	Sino-Tibetan	×	✓	✓ (R2)	✓
mos	Latn	Atlantic-Congo	✓	✓	✓ (R2)	✓
mri	Latn	Austronesian	✓	×	×	✓
mrw	Latn	Austronesian	×	✓	×	×
mtq	Latn	Austroasiatic	✓	×	×	×
mya	Mymr	Sino-Tibetan	✓	×	✓ (R1)	✓
myv	Cyrl	Uralic	×	×	×	✓
myx	Latn	Atlantic-Congo	×	✓	×	×
mzl	Latn	Mixe-Zoque	✓	×	×	×
mzm	Latn	Atlantic-Congo	×	✓	×	×
naq	Latn	Khoe-Kwadi	✓	×	×	×
nga	Latn	Atlantic-Congo	×	✓	×	×
ngl	Latn	Atlantic-Congo	×	✓	×	×
ngu	Latn	Uto-Aztecan	×	✓	×	×
nhe	Latn	Uto-Aztecan	✓	×	×	×
nia	Latn	Austronesian	×	✓	×	×
nij	Latn	Austronesian	×	✓	×	×
nim	Latn	Atlantic-Congo	×	✓	×	×
nld	Latn	Indo-European	✓	×	✓ (R1, R2)	✓
nlv	Latn	Uto-Aztecan	✓	×	✓ (R2)	×
nno	Latn	Indo-European	✓	×	×	✓
nob	Latn	Indo-European	×	×	×	✓
npi	Deva	Indo-European	✓	×	✓ (R2)	✓
nqo	Nkoo	N'Ko	×	×	×	✓
nso	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
nuj	Latn	Atlantic-Congo	×	✓	×	×
nus	Latn	Nilotic	✓	✓	✓ (R2)	✓
nya	Latn	Atlantic-Congo	✓	×	×	✓
nyy	Latn	Atlantic-Congo	×	✓	×	×
oci	Latn	Indo-European	×	×	×	✓
ory	Orya	Indo-European	✓	×	×	✓
osi	Latn	Austronesian	×	×	✓ (R2)	×
pag	Latn	Austronesian	×	×	×	✓
pam	Latn	Austronesian	×	✓	×	×
pan	Guru	Indo-European	✓	×	×	✓
pap	Latn	Indo-European	×	×	×	✓
pbs	Latn	Otomanguean	✓	×	×	×
pbt	Arab	Indo-European	✓	×	✓ (R2)	✓
pcm	Latn	Indo-European	✓	✓	✓ (R2)	×
pes	Arab	Indo-European	✓	×	✓ (R1, R2)	✓
plt	Latn	Austronesian	✓	×	×	✓
pol	Latn	Indo-European	✓	×	✓ (R1, R2)	✓
por	Latn	Indo-European	✓	×	✓ (R1)	✓
prs	Arab	Indo-European	×	×	×	✓
quc	Latn	Mayan	✓	✓	×	×
quh	Latn	Quechuan	✓	×	×	×

quy	Latn	Quechuan	×	×	×	✓
quz	Latn	Quechuan	✓	×	×	×
rhg	Rohg	Indo-European	×	✓	×	×
rim	Latn	Atlantic-Congo	×	✓	×	×
rmy	Latn	Indo-European	×	✓	×	×
rob	Latn	Austronesian	✓	×	×	×
roh	Latn	Indo-European	✓	×	×	×
ron	Latn	Indo-European	✓	×	✓ (R1)	✓
run	Latn	Atlantic-Congo	×	×	×	✓
rus	Cyrl	Indo-European	✓	✓	✓ (R1, R2)	✓
sag	Latn	Atlantic-Congo	×	×	×	✓
san	Deva	Indo-European	×	×	×	✓
sat	Olck	Austroasiatic	✓	✓	×	✓
sba	Latn	Central Sudanic	✓	✓	✓ (R1)	×
scn	Latn	Indo-European	✓	×	×	✓
sgc	Latn	Nilotic	✓	×	×	×
shk	Latn	Nilotic	×	✓	×	×
shn	Mymr	Tai-Kadai	✓	✓	×	✓
sif	Latn	Siamou	✓	×	×	×
sin	Sinh	Indo-European	✓	×	×	✓
skr	Arab	Indo-European	✓	×	×	×
slk	Latn	Indo-European	✓	×	×	✓
slv	Latn	Indo-European	✓	×	×	✓
sme	Latn	Uralic	✓	×	×	×
smo	Latn	Austronesian	×	×	×	✓
sna	Latn	Atlantic-congo	✓	×	×	✓
snd	Arab	Indo-European	✓	×	×	✓
snd	Deva	Indo-European	×	×	×	✓
som	Latn	Afro-Asiatic	✓	×	×	✓
sot	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
spa	Latn	Indo-European	✓	✓	✓ (R1, R2)	✓
sro	Latn	Indo-European	✓	×	×	×
srd	Latn	Indo-European	×	×	×	✓
srp	Cyrl	Indo-European	✓	×	×	✓
ssw	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
sun	Latn	Austronesian	✓	×	×	✓
swe	Latn	Indo-European	✓	×	✓ (R1)	✓
swh	Latn	Atlantic-Congo	✓	✓	✓ (R1, R2)	✓
syl	Beng	Indo-European	×	✓	×	×
szl	Latn	Indo-European	✓	×	✓ (R2)	✓
taj	Deva	Sino-Tibetan	×	✓	×	×
tam	Latn	Dravidian	✓	×	×	×
tam	Taml	Dravidian	✓	×	✓ (R2)	✓
taq	Latn	Afro-Asiatic	✓	✓	×	✓
taq	Tfng	Afro-Asiatic	✓	✓	×	✓
tat	Cyrl	Turkic	✓	×	✓ (R2)	✓
tda	Latn	Songhay	✓	×	✓ (R2)	×
tel	Latn	Dravidian	✓	×	×	✓
tel	Telu	Dravidian	✓	×	×	×
tem	Latn	Atlantic-Congo	×	✓	×	×
teo	Latn	Nilotic	×	✓	×	×
tgk	Cyrl	Indo-European	✓	×	×	✓
tgl	Latn	Austronesian	✓	×	✓ (R1, R2)	×
tha	Thai	Tai-Kadai	✓	×	✓ (R1)	✓
tir	Ethi	Afro-Asiatic	✓	✓	✓ (R1)	✓
toc	Latn	Totonacan	✓	×	×	×
tpi	Latn	Indo-European	✓	×	×	✓
tpl	Latn	Otomanguean	✓	×	×	×
tsg	Latn	Austronesian	✓	×	✓ (R2)	×
tsn	Latn	Atlantic-Congo	✓	✓	✓ (R2)	✓
tso	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
tsz	Latn	Tarascan	✓	✓	×	✓
tui	Latn	Atlantic-Congo	✓	×	×	×
tuk	Latn	Turkic	×	×	×	✓
tum	Latn	Atlantic-Congo	×	✓	×	✓
tur	Latn	Turkic	✓	×	✓ (R1)	✓
twi	Latn	Atlantic-Congo	✓	×	×	✓
tyv	Cyrl	Turkic	×	×	×	✓
tzh	Latn	Mayan	✓	✓	×	×
tzm	Tfng	Afro-Asiatic	✓	✓	✓ (R2)	×

tzo	Latn	Mayan	×	✓	✓ (R2)	×
uig	Arab	Turkic	✓	×	×	✓
ukr	Cyrl	Indo-European	✓	×	✓ (R1)	✓
umb	Latn	Atlantic-Congo	✓	✓	×	✓
urd	Arab	Indo-European	✓	×	✓ (R1)	✓
urd	Latn	Indo-European	✓	×	×	×
uzn	Latn	Turkic	✓	×	×	✓
uzs	Arab	Turkic	×	×	×	✓
vec	Latn	Indo-European	×	×	×	✓
ven	Latn	Atlantic-congo	✓	×	×	×
vie	Latn	Austroasiatic	✓	×	✓ (R1)	✓
wlv	Latn	Mataguanan	✓	×	×	×
vmw	Latn	Atlantic-Congo	✓	✓	×	✓
war	Latn	Austronesian	✓	×	✓ (R2)	✓
wol	Latn	Atlantic-Congo	✓	✓	✓ (R2)	✓
wsg	Deva	Dravidian	×	✓	×	×
wuu	Hans	Sino-Tibetan	✓	×	×	✓
xho	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓
xon	Latn	Atlantic-congo	×	✓	×	×
xsr	Deva	Sino-Tibetan	×	✓	×	×
xuu	Latn	Khoe-Kwadi	✓	×	×	×
yao	Latn	Atlantic-Congo	×	✓	×	×
ybb	Latn	Atlantic-Congo	×	✓	×	×
ydd	Hebr	Indo-European	✓	×	×	✓
ydg	Arab	Indo-European	✓	×	×	×
yor	Latn	Atlantic-Congo	✓	✓	✓ (R1, R2)	×
yua	Latn	Mayan	✓	✓	×	×
yue	Hant	Sino-Tibetan	✓	×	×	×
zai	Latn	Otomanguean	✓	×	×	×
zgh	Tfng	Afro-Asiatic	×	×	×	✓
zsm	Latn	Austronesian	✓	×	✓ (R1)	✓
zne	Latn	Atlantic-Congo	✓	✓	×	×
zul	Latn	Atlantic-Congo	✓	×	✓ (R2)	✓

C Additional evaluation results

Tables 12, 13, and 14 report the ChrF++ scores on the BOUQuET, FLoRes+ and Bible datasets respectively, translated between each language and English, aggregated by the non-English language resource level.

Resource level	XX-En						En-YY						All total
	high	mid	low	v. low	zero	total	high	mid	low	v. low	zero	total	
OMT-LLaMA 8B	65.1	53.5	38.7	27.6	21.6	40.6	60.7	45.8	30.8	18.8	12.6	32.8	36.7
OMT-NLLB	64.0	56.1	37.4	28.4	23.2	41.3	61.1	47.6	27.3	17.5	11.5	31.9	36.6
OMT-LLaMA 3B	64.4	54.6	38.8	28.0	22.3	41.0	59.4	45.0	28.7	16.8	11.3	31.3	36.2
NLLB-200	61.5	52.4	32.1	25.5	22.7	37.9	62.9	46.8	24.6	17.5	13.2	31.3	34.6
GPT-OSS 120B	64.6	49.9	33.6	25.4	24.3	38.3	63.6	43.7	26.0	16.0	13.4	30.8	34.5
Gemma 3	65.3	51.3	34.8	27.4	24.6	39.4	62.6	40.9	24.0	14.8	11.6	28.9	34.1
TranslateGemma	63.6	50.9	35.0	27.7	24.9	39.3	59.5	42.0	22.1	15.5	12.4	28.5	33.9
OMT-LLaMA 1B	63.0	51.5	34.7	24.1	19.9	37.9	56.5	42.2	25.1	14.9	12.3	29.1	33.5
LLaMA 3 70B	65.0	47.6	33.3	26.2	24.0	37.6	60.2	37.2	23.7	16.2	14.3	28.2	32.9
MADLAD	64.7	47.8	31.0	24.3	20.5	36.0	62.2	35.5	21.5	15.3	11.5	26.6	31.3
GPT-OSS 20B	62.5	47.0	30.7	23.2	22.4	35.7	59.2	36.7	20.2	13.1	12.1	26.1	30.9
Aya-101	59.6	48.3	30.9	25.5	23.0	36.3	52.8	33.7	16.6	11.4	10.9	23.1	29.7
Tower+	66.5	43.3	31.7	25.2	24.4	36.0	58.9	26.9	15.2	11.3	10.3	21.3	28.7
Tiny Aya Global	62.7	38.1	24.6	18.8	20.6	30.5	58.5	29.5	13.3	9.0	10.3	21.0	25.8
LLaMA 3 8B	61.0	39.8	27.5	21.6	20.6	32.1	52.7	26.0	13.1	9.5	8.2	19.1	25.6
Hunyuan-MT	56.6	33.9	24.4	20.3	21.0	29.0	47.8	21.5	13.9	12.4	10.8	18.5	23.8

Table 12: Translation performance on BOUQuET (ChrF++) by the non-English language resource level. Here, we separate the very low- and zero-resourced languages by the threshold of 1K examples in the training data.

Resource level	XX-En					En-YY					All total
	high	mid	low	v. low	total	high	mid	low	v. low	total	
OMT-NLLB	65.0	58.3	54.6	52.3	57.3	57.6	46.0	41.0	34.6	44.6	50.9
OMT-LLaMA 8B	64.6	56.4	50.6	48.3	54.6	57.4	46.5	39.8	33.8	44.2	49.4
NLLB-200	63.5	55.3	50.9	49.3	54.2	57.2	44.3	39.4	32.5	43.0	48.6
Gemma 3	64.2	52.3	46.4	48.1	51.4	58.6	40.1	31.5	26.4	38.1	44.8
TranslateGemma	62.5	51.2	45.5	47.5	50.4	55.4	39.5	30.5	27.5	37.2	43.8
GPT-OSS 120B	62.5	49.5	42.9	44.7	48.5	58.2	40.9	31.7	29.2	38.8	43.6
LLaMA 3 70B	64.1	49.8	44.8	45.7	49.5	56.7	37.9	31.1	26.3	36.8	43.2
MADLAD	63.1	45.4	43.8	44.3	47.1	57.2	34.1	31.4	29.3	35.7	41.4
GPT-OSS 20B	61.2	47.1	39.9	40.7	45.8	54.8	33.9	25.0	24.1	32.6	39.2
Aya-101	58.2	48.5	42.1	43.5	47.1	48.5	32.3	22.9	20.3	29.9	38.5
Tower+	65.0	45.8	43.1	45.8	47.4	54.5	27.3	22.6	22.9	28.8	38.1
LLaMA 3 8B	61.0	44.1	38.9	39.7	44.1	51.4	27.5	21.7	20.6	27.9	36.0
Tiny Aya Global	61.3	39.9	34.6	35.6	40.4	53.6	27.6	19.5	19.7	27.4	33.9
Hunyuan-MT	57.3	38.1	35.5	37.7	39.7	45.6	22.2	20.5	23.4	24.9	32.3
Aya Expanse	59.0	34.8	34.5	38.2	38.3	48.0	22.3	20.1	21.8	24.9	31.6
EuroLLM	62.5	33.7	34.2	38.2	38.2	55.7	16.6	18.8	22.3	23.2	30.7

Table 13: Translation performance on FLoRes+ (ChrF++) by the non-English language resource level.

Resource level	XX-En					En-YY					All total
	high	mid	low	v. low	total	high	mid	low	v. low	total	
OMT-LLaMA 8B	63.3	61.1	52.3	35.1	38.8	53.1	47.5	40.0	28.3	31.0	34.9
OMT-NLLB	62.8	59.9	58.5	43.4	46.3	54.4	44.5	22.3	10.5	14.3	30.3
NLLB-200	54.8	51.1	33.3	20.4	24.0	50.9	49.6	24.4	11.6	15.6	19.8
LLaMA 3 70B	63.8	47.3	30.6	21.6	24.7	48.7	30.0	18.8	10.7	13.3	19.0
Gemma 3	61.6	48.0	30.7	21.5	24.5	48.4	31.8	18.0	9.0	11.9	18.2
TranslateGemma	58.4	45.3	29.6	20.8	23.7	45.1	31.6	16.7	8.0	10.8	17.3
GPT-OSS 120B	54.1	40.5	27.1	20.0	22.5	44.5	31.1	17.4	9.0	11.7	17.1
Aya-101	53.5	41.6	27.9	19.4	22.1	40.8	24.3	12.0	7.6	9.5	15.8
Aya Expanse	54.5	28.9	21.7	18.3	19.9	39.1	17.4	11.8	9.3	10.5	15.2
GPT-OSS 20B	51.9	37.5	23.9	17.2	19.6	41.2	25.9	12.7	8.4	10.4	15.0
Hunyuan-MT	49.8	29.8	21.3	17.5	19.1	36.4	16.5	11.7	9.2	10.4	14.7
Tiny Aya Global	51.3	29.4	20.1	16.5	18.2	42.7	20.9	9.1	6.8	8.4	13.3
LLaMA 3 8B	55.6	35.3	23.2	17.1	19.4	42.4	20.2	10.0	5.1	7.1	13.3
EuroLLM	62.2	27.1	19.1	13.0	15.3	47.2	10.7	7.7	5.9	7.2	11.2

Table 14: Translation performance on Bible (ChrF++) by the non-English language resource level.

D Extrapolations across evaluation metrics and datasets

Beyond the relative performance of models in understanding hundreds of languages, we seek to determine how many languages the OMT models understand "well enough" in absolute terms.

Extrapolation between metrics. We define a passing quality threshold as an average XSTS+R+P score above 2.5. Based on the definition of XSTS+R+P scores (see Section 5), this roughly means that the system is capable of conveying the core meaning of a sentence in the majority of cases. We rely on MetricX (reference-based) to estimate whether a translation meets this criterion. The primary motivation for using MetricX is that it is a well-established external metric that effectively leverages both target and reference translations (and is therefore relatively reliable even for very difficult source languages, as well as it understands the target language). Furthermore, unlike BLASER 3, it is not based on OMNISONAR, thereby avoiding potential bias toward the OMT-NLLB model, which shares the same encoder. We apply a monotonic regression over Met-BOUQuET (restricted to the XX-En directions for direct compatibility with the Bible evaluation setup) to map between MetricX and XSTS+R+P scores on the item level (see Figure 4). On the item level, cross-validated RMSE of this regression is 1.008 XSTS+R+P points. However, most of this noise gets canceled out by aggregation: the cross-validated RMSE between the predicted and the actual scores averaged on the language and system level is only 0.36 points.




Figure 4: A curve to predict XSTS+R+P scores from reference-based MetricX (on Met-BOUQuET, XX-En subset; to visualize XSTS+R+P distribution, a random jitter has been added to the Y axis).

We can use this formula to estimate the proportion of languages in each bouquet level that would pass the “good” and “passable” into-English translation quality bars. They are reported in Table 15 for the OMT-LLAMA-8B and OMT-NLLB models.

		OMT-LLAMA			OMT-NLLB		
high	27	4.24	1.00	1.00	4.23	1.00	1.00
mid	75	3.52	0.61	0.92	3.79	0.80	0.96
low	83	2.50	0.25	0.45	2.44	0.28	0.37
very low	38	1.90	0.05	0.21	1.98	0.13	0.21
zero	51	1.74	0.04	0.10	1.75	0.04	0.18

Table 15: Projected proportions of “good” (above 3.5) and “passable” (above 2.5) quality levels for into-English translation on BOUQuET test set, estimated from the MetricX scores. In this table, we make a distinction between very-low-resourced languages (less than 40K parallel data examples) and zero-resourced ones (less than 1K).

Extrapolation between datasets. Even though BOUQuET contains a representative sample of languages with diverse language families, writing systems, geographies, and resource levels, it covers “only” 275 language varieties, an order of magnitude fewer than the Bible dataset. Therefore, the question of correlation of evaluation results between these two benchmarks is relevant. Indeed, our analysis suggests that there is a high degree of equivalence between the Bible- and BOUQuET-based evaluation results. We computed averages (per each source language and MT system combination) of the scores of 19 diverse translation systems for the 159 language varieties shared between the Bible and BOUQuET benchmarks (into-English translation, to keep the scores comparable), and for most languages, they seem to be relatively close to each other (with well-performing languages even demonstrating an advantage on BOUQuET), as Figure 5 illustrates. Statistical analysis quantifies this equivalence:

- For MetricX, Pearson correlation is 0.92 (OLS fit is $BOUQuET = -0.580 + 1.0153 \times Bible$)
- For ChrF++, Pearson correlation is 0.87 (OLS fit is $BOUQuET = 3.3140 + 0.9798 \times Bible$)

For both fitted lines, the slope is statistically barely distinguishable from 1, but the intercept is statistically different from 0, indicating that the BOUQuET sentences are on average slightly easier to translate than the Bible ones. Therefore, we can treat the automatic XX-En scores on the Bible dataset as sufficiently accurate and conservative estimates of their BOUQuET counterparts.




Figure 5: Correlation of MetricX values, averaged over 19 MT systems and shared 159 XX-En directions, between BOUQuET and Bible datasets.

Bible-based count of supported languages. We use the Bible benchmark (comprising 1,560 languages) to estimate the average XSTS+R+P score using the MetricX proxy for each source language translated into English. Figure 6 presents the results of this extrapolation for OMT-LLAMA, OMT-NLLB, and NLLB-200 as a baseline. All three models cover approximately the same number of Bible languages (around 130) at the "good" quality threshold of at least 3.5 extrapolated XSTS+R+P points on average. However, the number of Bible languages for which the "passable" quality threshold of 2.5 points is exceeded is substantially larger for the OMT models: 440 languages for the OMT-LLAMA 8B model and 416 languages for OMT-NLLB —nearly double the 221 varieties for which NLLB-200 surpasses this threshold. We therefore conclude that, while we remain far from completely "solving" machine translation for the long tail of languages, the OMT models double the number of reasonably well-understood source languages compared to previous massively multilingual models.




Figure 6: Predicted XSTS+R+P machine translation quality score for the Bible languages into English using MetricX extrapolation on the test set (languages sorted individually for each of the evaluated models) .

E Prompt template

The following box defines the prompt template that we used as the default for evaluating third-party LLMs for the translation task. The prompt uses the English names of the language, followed, if necessary, by the name of the writing system (such as "Mandarin Chinese (Traditional script)").

Translate the following text from {source language} into {target language}.
Please write only its translation to {target language}, without any additional comments.
Make sure that your response is a translation to {target language} and not the original text.
{source language}: {source text}
{target language}: