

DialUp! Modeling the Language Continuum by Adapting Models to Dialects and Dialects to Models

Niyati Bafna¹, Emily Chang², Nathaniel R. Robinson¹,
David R. Mortensen³, Kenton Murray¹, David Yarowsky¹, and Hale Sirin¹

¹Johns Hopkins University, Center for Language and Speech Processing;

²University of Virginia; ³Language Technologies Institute, Carnegie Mellon University
{nbafna1, nrobin38}@jhu.edu, ec5ug@virginia.edu

Abstract

Most of the world’s languages and dialects are low-resource, and lack support in mainstream machine translation (MT) models. However, many of them have a closely-related high-resource language (HRL) neighbor, and differ in linguistically regular ways from it. This underscores the importance of model robustness to dialectal variation and cross-lingual generalization to the HRL dialect continuum. We present DialUp, consisting of a training-time technique for adapting a pretrained **model** to dialectal **data** (M→D), and an inference-time intervention adapting dialectal **data** to the **model** expertise (D→M). M→D induces model robustness to potentially unseen and unknown dialects by exposure to synthetic data exemplifying linguistic mechanisms of dialectal variation, whereas D→M treats dialectal divergence for known target dialects. These methods show considerable performance gains for several dialects from four language families, and modest gains for two other language families. We also conduct feature and error analyses, which show that language varieties with low baseline MT performance are more likely to benefit from these approaches.¹

1 Introduction

Recent years have seen advancement in MT quality and language coverage, with models such as M2M100 (Fan et al., 2021) and NLLB (Team et al., 2022) covering 100 and 200 languages respectively, as well as multi-purpose generative language models such as the GPT, BLOOM and Llama (Muenighoff et al., 2023; Dubey et al., 2024) series expanding to dozens of languages. These models have displayed cross-lingual generalization capabilities for some unseen or low-resource languages, but they tend to suffer significant performance degradation for most (Jiao et al., 2023; Robinson




Figure 1: Two paradigms for robustness to dialects on a continuum of distances from an HRL. DialUp involves M→D: training the **model** on artificial dialectal variation, and D→M: bringing dialectal **data** closer to model expectations (HRL-like input) at inference.

et al., 2023; Ziems et al., 2023; Cahyawijaya et al., 2024; Joshi et al., 2024).

The roughly 7000 languages in the world can be grouped into a few hundred families. Closely-related languages (CRLs) and dialects within a family largely exhibit continuous and structured differences along phonological, morphological, and lexical dimensions, rather than being discrete monoliths. (Hovy and Purschke, 2018; Bergman and Diab, 2022). Not only is it unfeasible to collect training data for every lect, but language is fluid and dynamic across and between dialect boundaries, calling for principled general approaches to dialectal variation.

Many language families have at least one HRL member that is supported by state-of-the-art models. We seek to expand models proficient in HRLs to their lower-resource CRLs, by inducing robustness to unseen varieties across language continua. We do this with DialUp: adapting models from an HRL to its language relatives, and adapting those relatives’ data inversely towards the HRL model, as in Figure 1.

We first aim to adapt the model to the language continuum, or **model-to-data (M→D) adaptation**

¹<https://github.com/niyatibafna/dialup/>

Turkic ["boils"]	kaynar (tur)	qaynar (azj)	gaýnar (Tuk)	qaynar (crh)
Romance ["tree"]	albero (ita)	árboře (glg)	arbře (oci)	árvulu (scn)
Creole ["today"]	jodi a (hat)	jórdi (lou)	ozordi (crs)	zordi (mfe)

Table 1: Cognates differ in predictable ways, via **sound change** patterns, vowel changes, and new **suffixal paradigms**.

(yellow arrows in Figure 1). The goal is robustness to unseen CRLs (both known CRLs absent in pre-training data and yet undocumented CRL varieties), using only HRL bitext. We do this by simulating dialectal variation patterns over HRL fine-tuning text, teaching the model to generalise to realistic variants such as actual CRL inputs.

Next, we aim to adapt data to the HRL model, or **data-to-model (D→M) adaptation** (red arrows in Figure 1), in the case that a HRL-CRL bilingual lexicon is available or can be induced. In this circumstance, we pull CRL data towards the model distribution at inference time by interchanging words for their HRL counterparts.

Note that cross-lingual generalization can be viewed as a train-test data mismatch problem: we want our models to work on dialectal data which differs in various ways from the training data of the model (HRL data). Our two approaches above then represent two broad paradigms for this general problem in machine learning (see Figure 1). Model-to-data adaptation has roots in existing cross-lingual transfer approaches, which approximate CRLs using an HRL, relying on linguistic similarities for transfer; however, this does not train the model to handle dialectal divergences. M→D innovates on this by synthetically approximating such divergences. By approximating language continua, we may avoid pressing questions raised by discrete language paradigms, e.g. which transfer language to use and how much training data to collect (Dalmia et al., 2018). D→M adaptation has roots in prior approaches to approximate a model’s expected domain in data (Nie et al., 2023).

These two approaches have different advantages. M→D does not require any CRL resources or data and does not require partitioning the continuum into discrete dialects. D→M is train-free and can be directly applied even to closed-source models. These approaches can also be used in tandem, by applying D→M at inference time to a model adapted via M→D. We expect these methods to be most beneficial for unseen CRLs and CRLs that have high linguistic overlap with the HRL, i.e. languages that depend on existing HRL representations in the

pre-trained model. In sum, we contribute:

- DialUp, a principled and inexpensive method to induce robustness over language family continua via M→D and D→M adaptation.
- Evaluations of DialUp’s benefits for X→eng MT with two models for 49 CRLs across six language (sub)families.
- Consistent gains via M→D for low-baseline varieties across 4/6 language families (up to mean +11.4 BLEU for Romance languages)
- Gains from D→M adaptation over baselines for certain families and CRLs (yielding up to mean +12 BLEU for Indic dialects), showing for the first time that adapting dialectal *function words* is more beneficial than adapting content words
- Evidence that M→D and D→M combine advantageously, and provide a recipe for increasing the flexibility of existing MT models to general dialectal variation.

2 DialUp

2.1 Model-to-data (M→D)

In this paradigm, given an HRL-proficient model and bitext in that HRL, we adapt the model to unseen CRLs. We generate varied artificial CRLs of the HRL by simulating mechanisms of dialectal variation over the HRL bitext, and fine-tune on the resulting synthetic data to induce model robustness to dialectal variation.

Artificial language generation We employ the approach of Bafna et al. (2024b) to simulate language variation along **phonological**, **morphological**, and **lexical** (function word and (non-cognate) content word) dimensions. Each constituent process operates on the relevant units of the inputs: phonemes, suffixes, and function and content lexemes, and “noises” a given unit with probability θ^n , $n \in \{p, m, f, c\}$, replacing all instances of that unit with a plausible alternative in the constructed language (i.e. phones in the original unit randomly replaced with alternate phones of high phonetic proximity, or suffixes swapped with similar-sounding alternatives).² θ^n serves as a proxy for linguistic distance in each of these dimensions: increase or “dialing up” of θ^n results in more

²We provide further details on each noiser in Appendix B.

Using HRL text (top) to generate artificial dialect text (bottom) via noising	
Indic	मित्रों , तीन आइडिया ऐसे हैं जिन पर मैं अधिकतम चर्चा और भागीदारी का आग्रह करना चाहता हूँ। मित्रों , तीम आइडिया ऐषआ हों शिन फर नैं अधिखतम चर्चा और भागीडईरी गा आग्रह करना चाहडईई।
Austronesian	Beberapa proyek akan melalui langkah 2, 3 dan 4 beberapa kali. Veverofu broyek aqen meyalui langkih 2, 3 dun 4 veverofu kali
Romance	Paul accetta con riluttanza, sapendo ciò che sta per accadere. Paul acceddu gen riluttanza, sapendo ciò che zta ber accadere.
Turkic	Şükür sadakası olarak da bilinir. Şükür jadakası oyarüg ta bilinir.
French-Creole	Kijan nou ka montre gratitud nou pou sa Sovè a fè pou nou? Kijan ngou ka mondre gratitud ngou pou sa Sovè ò fè pou ngou?

Table 2: Linguistically-motivated noisers mimic **function word**, **suffixal**, and regular **sound** changes.

divergent artificial varieties. Hence, artificial languages exhibit various kinds of noise, at respective θ^n -dials. (See Table 2 for examples of generated artificial language text.) Each artificial language is defined as the map of changes made from the HRL. We present two alternative ways to distribute these artificial languages: **M→D-shell** generates them on a single hypersphere at a fixed $\theta^{p,m,f,c}$ -radius from the HRL³, while **M→D-cloud** generates them on multiple hyperspheres around the HRL, populating the hypothesized dialect continuum.⁴ (Figure 1 depicts M→D-cloud; M→D-shell would show all yellow arrows extending to the same blue band.)

2.2 Data-to-model (D→M)

In this paradigm, given CRL-HRL bilingual lexicons, we swap divergent parts of CRL input with known HRL equivalents to pull data towards the model’s proficiency distribution. Languages from the same family are generally syntactically similar or monotonic (Posner and Sala, 2024); hence this switching should largely maintain comprehensible grammatical structure.

We explore three D→M settings: **func**, **cont**, and **all**; in which we swap out only function words,⁵ only content words,⁶ and both, respectively. Function and content word classes differ both in their role in language, with the former crucial for grammaticality and coherence, and the latter providing semantics, as well as in the extent to which they

³We use $\theta_{0.05,0.3,0.5,0.001}^{p,m,f,c}$, as an average CRL-HRL distance given θ^n posteriors over several CRL-HRL pairs (Bafna et al., 2024b). See § E.1 for more details.

⁴We use $K = 10$ hyperspheres at uniform intervals in each dimension from the HRL and $\theta^{n_i} = \theta_{max}^n \cdot \frac{i}{K}$. Each hypersphere is therefore used for noising $1/K^{th}$ of the total data. We set θ_{max} as $\theta_{0.07,0.5,0.8,0.001}^{p,m,f,c}$, as reasonable maximum CRL-HRL distances in each dimension. See § E.1 for more details and tuning experiments for both variants.

⁵determiners, adpositions, auxiliaries, conjunctions, pronouns, and determiners

⁶all other word classes

are affected by language change: function words diverge quickly and often opaquely across dialects as compared to less frequent content words (Ellis, 2008). It is therefore relevant to isolate the effects of D→M on these word classes.

Note that isolating these settings requires function word identification in CRL input. We achieve this by collecting HRL function words from the Universal Dependencies corpus (Nivre et al., 2016) and annotation projection using the HRL-CRL lexicons. Any word that is not identified as a function word is considered a content word.

2.3 M↔D and baselines

We also combine M→D (–cloud fine-tuning) with D→M (inference-time adaptation) into M↔D.

We compare all proposed approaches to three baselines: (1) evaluating the model on CRL→eng MT without any adaptation (**off-the-shelf**), (2) evaluating on CRL→eng after fine-tuning on ordinary HRL→eng bitext without any simulated linguistic variation (**fthr1**), and (3) fine-tuning and evaluation after augmenting the HRL→eng bitext with completely random (i.e. not linguistically motivated or plausible) variations (**randaug**) For the final baseline, we randomly swap out characters in the HRL text with uniformly sampled target same-script characters, and words with different words from the source language vocabulary at probability (θ^{rc} , θ^{rw}). This is similar to prior work (Blinkov and Bisk, 2018; Heigold et al., 2018) that introduces random character and word perturbations in the HRL as a method of data augmentation. We implement shell and cloud variants of this baseline analogously to M→D, maintain analogous parametrization for a fair comparison⁷, and only report the better of the two.

⁷See § E.1 for details.

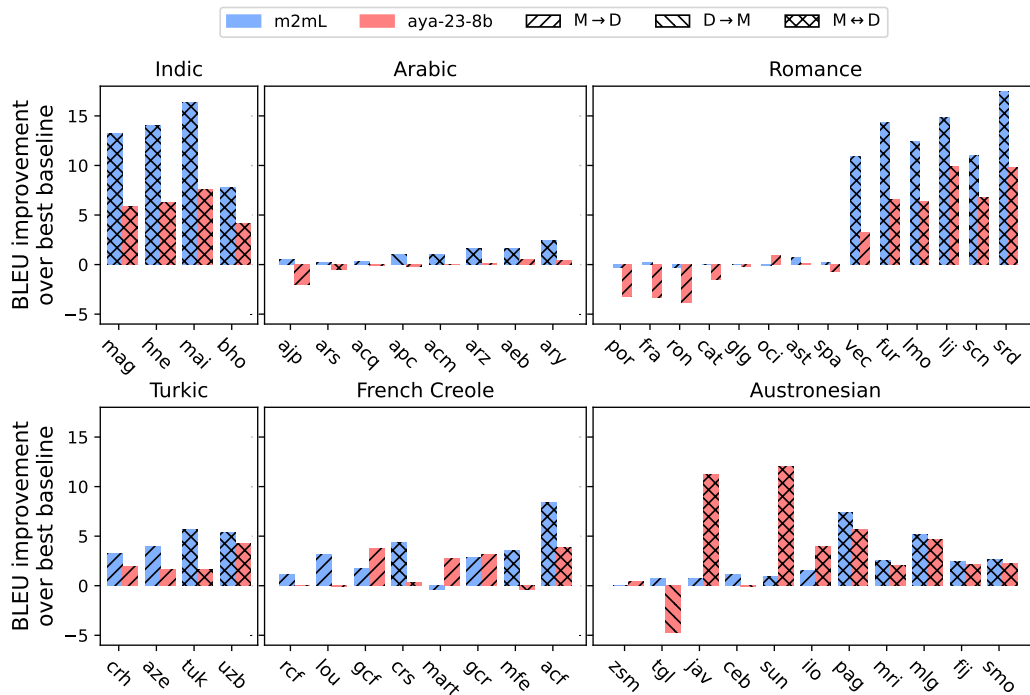


Figure 2: BLEU point improvement of the best DialUp method (M→D, D→M, or M↔D) over the best baseline (off-the-shelf, fthrl, or randaug). Languages are ordered by their M2M off-the-shelf performance.

3 Experimental Setup

We work with X→eng MT for 49 languages in six (sub)families.

Languages and Datasets We include the following language groups and designate one HRL in each: Austronesian (with HRL Indonesian), Indic (Hindi), Turkic (Turkish), Romance (Italian), Arabic (Standard Arabic), and French-related Creole languages (Haitian). (See Table 6 in the Appendix for a full list.) We use Wiki text bitext (Schwenk et al., 2021) in the HRLs as training data; we do not use any CRL bitext. All compared methods use the same amount of total HRL bitext (100K sentences). We included CRLs per language family according to availability of evaluation data in the FloRes-200 (Team et al., 2022) dataset, maintaining a single script within each family. We use Kreyòl-MT evaluation sets for Creole CRLs (absent in FloRes) (Robinson et al., 2024). Our set of CRLs includes languages on a spectrum of relatedness to each HRL, as well as variety of resource levels. Some of the CRLs we include, such as French, are high-resource themselves, while most are low-resource. The Turkic, Austronesian, and Arabic languages we included vary widely in HRL proximity; some have a high degree of mutual intelligibility with the HRL (Azerbaijani, Malay, Saudi

Arabic), and some are quite distant (Uzbek, Tagalog, Moroccan Arabic) (Hammarström et al., 2024; Nouri and Yangarber, 2014).

Bilingual lexicons For each CRL-HRL pair, we use a combination of PanLex (Kamholz et al., 2014), and Swadesh lexicons.⁸ We also added Art Dieli’s Sicilian-Italian dictionary (Dieli, 2011) and Indic dialect lexicons from Bafna et al. (2024a). However, most of these have little coverage of function words. We therefore additionally used (function word-inclusive) lexicons obtained from performing statistical alignment (Dyer et al., 2013) on FloRes dev data.⁹ The Creole CRLs are not included in any of the above lexicon datasets; we used statistically aligned lexicons from JHU Bible translations (McCarthy et al., 2020), available for acf, crs, and mfe. We therefore only report D→M and M↔D results for these three Creole CRLs.

Models We employ two MT models in our experiments: one multilingual supervised seq2seq model, M2M-100-1.2B (Fan et al., 2020); and one

⁸https://en.wiktionary.org/wiki/Category:Swadesh_lists_by_language

⁹See Appendix D for our survey of available lexicons, details and statistics of used lexicons. We also report the type coverage of CRL text for each lexicon, and note that this is generally low for PanLex and Swadesh.

multilingual instruction-tuned LLM, Aya-23-8B (Aryabumi et al., 2024). We selected these models because they are highly multilingual (supporting all our selected HRLs) but lack support for enough of our CRLs to warrant legitimate evaluation of adaptation to unseen languages. M2M only supports 15 of our 49 selected CRLs, while Aya-23 supports only 4. We selected one seq2seq model and one LLM in order to evaluate our methods in these two settings. We curated translation instructions from our bitexts to fine-tune Aya-23. (See § C.1 for details.) We use LoRA (Hu et al., 2021) for a single epoch for all fine-tuning processes.

4 Results

Results across all languages are in Figure 2. Our approaches give gains across the board, with varied trends by language family. Mean gains of low-performing languages (< 25 BLEU for off-the-shelf) per method and language family are in Table 3. See Appendix K for detailed results and COMET (Rei et al., 2020) scores. We found that these show trends consistent with BLEU.

We observe that for both models, languages with poorer baseline scores, which tend to be the low-performing varieties we aim to assist in this work, benefit more from DialUp, while languages with better baselines show small or negative gains. This trend is visible in Figure 2, where languages are ordered by their off-the-shelf BLEU score, and especially pronounced for Romance languages. We also see that M2M typically benefits more. This may be because Aya-23 was likely exposed to some CRLs in pre-training, despite not being trained explicitly on them, given the heterogeneity and poor documentation of LLM pre-training sets. Two exceptions to this trend are Javanese (jav) and Sundanese (sun) (gaining +13.3 BLEU with Aya-23, and little with M2M), which Aya-23 translates poorly off-the-shelf, and which M2M explicitly supports in pre-training. Proximity to the HRL also appears to play a role: CRLs distant from their HRL, such as Samoan and Tagalog, as well as those extremely close to the HRL, such as Malay (zsm), Azerbaijani (aze), and Saudi Arabic (ars), appear to benefit little from DialUp.

M→D gives mean gains over best baselines for all language families, demonstrating the efficacy of linguistically-motivated (as opposed to random) synthetic data augmentation. Table 4 shows examples of improvements on cognate words, includ-

ing new inflections and function words. We often see that while off-the-shelf models transliterated them or left them as are instead of translating, M→D-tuned models decode them correctly.¹⁰ D→M introduces gains for many low-performing varieties—up to +11.4 BLEU for Aya-23 on Sundanese and +13.7 BLEU for M2M on Chhattisgarhi (hne)—where replacing dialectal words with HRL equivalents is helpful. However, it is also damaging for several, especially high-performing CRLs like French (fra), Spanish (spa), and Tagalog (tgl). Conceivably, if the base model is already proficient in a CRL, D→M code-switching introduces counterproductive unnaturalness. Since D→M is an inference-time method, practitioners may activate or deactivate it according to language needs. See Table 5 for examples of D→M-func input and output, demonstrating the brittleness of baseline models on dialectal function words, as well as the impact of treating them.

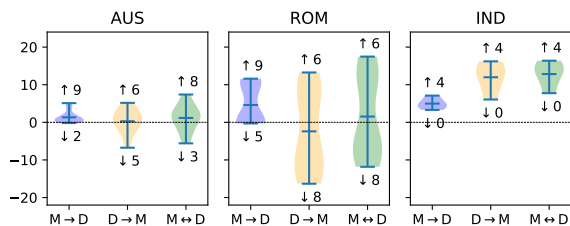


Figure 3: BLEU score improvements over the best baseline with M2M for three language families. ↑ and ↓: # CRLs with positive/negative gains. M→D gives more consistent positive gains.

So which paradigm performs better? That depends on the model and the language. See Figure 3 for a distribution of score improvements for two families, and number of wins over the best baseline. Most families behave similarly to Austronesian and Romance (see Appendix H for similar plots for all languages/models). In general M→D gives lower-variance gains, with wins over the baseline for most CRLs. D→M often shows similar or higher mean gains but with higher variance, and consistently fewer wins. Indic languages show a different distribution, with a unanimous preference for D→M. Notably, M↔D generally performs best across languages. (See Table 3.) Holistically, results suggest that **M→D provides gains with low-risk consistency, and that D→M provides significant benefits for a select set of CRLs.**

We note that D→M-func consistently outper-

¹⁰See full-length examples in Appendix G.

		M2M							Aya-23						
		AUS (9)	ARA (7)	ROM (6)	TUR (4)	IND (4)	CRE (8)	Mean	AUS (9)	ARA (1)	ROM (5)	TUR (4)	IND (4)	CRE (8)	Mean
Baselines	off-the-shelf	10.8	20.2	9.6	5.0	12.6	5.9	10.7	<u>7.6</u>	<u>23.5</u>	17.1	7.9	<u>20.4</u>	7.3	14.0
	fthrln	<u>+0.3</u>	<u>+0.3</u>	<u>+0.5</u>	-0.1	<u>+0.4</u>	<u>+0.9</u>	<u>+0.4</u>	-0.6	-1.8	+1.4	<u>+1.2</u>	-1.8	<u>+2.3</u>	<u>+0.1</u>
	ftrandaug	+0.2	+0.2	<u>+0.6</u>	<u>+0.2</u>	<u>+0.4</u>	<u>+0.9</u>	<u>+0.4</u>	-0.4	-2.3	<u>+1.9</u>	+1.1	-1.8	+2.1	<u>+0.1</u>
M→D	-shell	<u>+1.9</u>	<u>+1.0</u>	<u>+1.5</u>	<u>+2.5</u>	+5.0	<u>+2.7</u>	<u>+4.1</u>	+2.1	<u>+0.4</u>	<u>+7.9</u>	+2.9	<u>+2.8</u>	+3.0	<u>+3.2</u>
	-cloud	+1.3	+0.9	+9.1	+1.6	<u>+5.4</u>	<u>+2.7</u>	+3.5	<u>+2.2</u>	+0.1	+7.7	<u>+3.2</u>	+2.7	<u>+3.5</u>	<u>+3.2</u>
D→M	-cont	-0.4	-2.2	+1.6	+1.4	+0.3	+2.5	+0.5	+0.8	-6.2	-1.0	-0.6	-1.1	-3.8	-1.9
	-func	-0.1	<u>+0.8</u>	+8.0	+1.0	<u>+12.0</u>	+3.7	+4.3	+2.4	<u>-1.8</u>	<u>+4.7</u>	<u>+0.3</u>	<u>+5.6</u>	<u>-3.3</u>	<u>+1.3</u>
	-all	<u>+0.0</u>	-1.4	+9.8	<u>+3.0</u>	+11.4	<u>+5.6</u>	<u>+4.7</u>	<u>+3.7</u>	-6.9	+3.1	+0.0	+4.4	-3.5	+0.2
M↔D	-cloud-cont	<u>+1.4</u>	-1.5	+9.0	+2.6	+5.1	+3.7	+3.4	+3.3	-4.3	+5.3	+1.8	+1.6	+2.6	+1.7
	-cloud-func	+1.1	<u>+1.3</u>	<u>+14.1</u>	+2.3	<u>+13.0</u>	+4.6	<u>+6.1</u>	+3.8	<u>-0.8</u>	<u>+9.8</u>	<u>+3.1</u>	<u>+6.0</u>	<u>+2.8</u>	<u>+4.1</u>
	-cloud-all	+0.9	-1.2	+13.1	<u>+3.9</u>	+12.2	<u>+6.1</u>	+5.8	<u>+5.1</u>	-4.5	+7.4	+2.1	+4.6	+2.5	+2.9

Table 3: BLEU score performance gains relative to off-the-shelf, for low-performing CRLs (off-the-shelf score < 25), averaged by language family. # of such CRLs per family provided in parentheses. The overall best score is **bolded** and the best score in each paradigm is underlined. Language families are abbreviated with first three letters (e.g. Indic (IND)). M→D and D→M both outperform best baselines; M↔D is generally the best.

lang	Input CRL phrase	HRL equivalent	Baseline (CRL→eng)	M→D (CRL→eng)
bho	अध्ययन कइला पर	(hin) अध्ययन करने पर	the study on Kila revealed	studies have shown
lij	da reduce a mortalitæ do Ebola	(ita) da ridurre la mortalità dell'Ebola	da reduce a mortalitæ do Ebola	to reduce Ebola mortality
crh	açıkqlama	(tur) açıklama	[empty]	announcement

Table 4: M→D improves M2M on input with cognate processes such as **function words**, **new suffixes**, and **phonological/orthographic** change. The baseline often transliterates the input or leaves it as is.

lang	Input CRL Phrase	D→M modification	Baseline (CRL→eng)	D→M (CRL→eng)
mag	जौरे एगो मुकदमो	साथ एक मुकदमो	Jure Ego Prosecution	a litigation with
scn	ntà cademia svidisi	nella cademia svidisi	ntà cademia svidisi	at the Swedish Academy
jav	kaya sawetara ahli liyane	seperti beberapa ahli liyane	rich a handful of other experts	like some other experts

Table 5: D→M improves M2M output by replacing **function words** in the input (underlined in output) with HRL equivalents.

forms D→M-cont, and is often better or close to D→M-all. This highlights the importance of treating dialectal function words in enabling model comprehension, and the utility of collecting CRL-HRL function word maps for low-resource languages. This insight is particularly convenient given that function words form a small closed class and are much easier to collect comprehensive lexicons for than open-class content words; they are also more likely to be accurately aligned with statistical alignment than relatively rarer content words.

In fact, exchanging content words is largely unhelpful across the board, including from large high-quality lexicons like Art Dieli for Sicilian (scn). We attribute this to the higher degree of context dependence in handling content words; synonymy, word sense variation, naturalness, and lexicon noise likely contribute to this problem.

See Appendix E for our additional M→D experiments in using multiple HRLs per family, training with more data or for epochs, different source datasets, and more aggressive noising. These mildly improve or maintain current trends. We also show that switching CRL words directly into En-

glish instead of the HRL for D→M degrades performance, presumably due to unnaturalness of CRL-eng (as compared to CRL-HRL)code-mixing.

5 Discussion

When does M→D help? To test our hypotheses that baseline performance and proximity to the HRL interact meaningfully with score improvement, and to understand which circumstances render DialUp more or less effective, we analyze how language features correlate with gains. We computed Spearman’s ρ coefficients between BLEU improvement from M→D over the off-the-shelf baseline and features indicating both baseline CRL support (off-the-shelf BLEU score, whether the model explicitly supports the CRL, and number of CRL Wikipedia pages¹¹) and HRL-CRL relatedness (character F-score (Popović, 2015) between HRL and CRL FloRes dev sets and average token-count ratio in CRL dev lines to HRL dev lines when using the model’s HRL tokenizer¹²). For

¹¹per Robinson et al. (2023)

¹²In the case of the Creole CRLs, which lack FloRes dev sets, we used JHU Bible translations, available for acf, crs, and mfe.

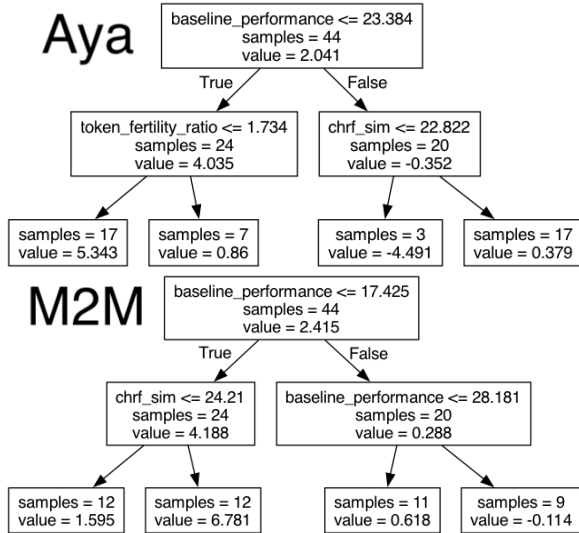


Figure 4: Decision trees indicate that the languages benefiting most from adaptation are low-baseline languages with less than 1.75 times HRL token fertility for Aya, and low-baseline languages with more than 24.2 chrF proximity to the HRL.

both Aya and M2M the only such feature that correlated significantly ($p < 0.01$) was `off-the-shelf BLEU`, with $\rho = -0.52$ and -0.45 , respectively—indicating a **moderate-to-strong negative correlation between baseline performance and improvement from M→D**, as hypothesized in § 4.

Further analysis also suggests that baseline score is more predictive of M→D success than other features; this is confirmed by random forests fitted over the same features for both models (see decision trees in Figure 4), as well as feature weights from a linear regressor. (See Table 35.)

Measurements of HRL-CRL relatedness had the second-highest forest feature importance for both models. Though our measurements of HRL-CRL relatedness do not correlate with BLEU score improvement, we suspect there may be a non-linear relationship between dev-set character-F score and method effectiveness. When plotted in Figure 6, the points suggest the contour of a downward-facing parabola, indicating that M→D is most effective for CRLs that are not too close but not too far from the HRL, as we hypothesized. This has an intuitive interpretation: CRLs that are too close may either already be well-performing, or may not display enough dialectal divergence to benefit from DialUp, whereas CRLs that are too far may have a higher amount of non-cognate divergences from the HRL that DialUp does not help with. However, note that language family is possibly confounded with




Figure 5: Gains in BLEU points for different values of θ^p (1-dimensional noiser) for Indic languages, with dotted lines showing the performance of M→D-c, using the the 3-dimensional noiser $\theta^{p,m,f}$ with default parameters. Tuning only θ^p for Indic is competitive with M→D-c.

character-F score, with only Austronesian CRLs covering a wide range on our observed contour; it would be difficult to disentangle the effect of language family without many more languages.

Setting $\theta^{p,m,f}$ -dials Although M→D uses phonological, morphological, and function word noising to model cognate processes in CRLs, the latter two are specialized versions of the phonological noiser, applying it at a fixed high dial to HRL suffixes/function words (Appendix B). We show that this 3-dimensional noiser can be approximated solely by the phonological noiser at a suitable θ^p -dial. Setting all other $\theta^n = 0$, we sweep over θ^p for Indic (Figure 5). We observe that with an optimal θ^p , the phonological noiser achieves competitive gains as M→D-c, which uses all noisers. Our default M→D θ -dials are eminently reasonable, yielding the best possible performance achievable by the above procedure; however, these defaults may be more suited to particular language families than others. Given that tuning $\theta^{p,m,f}$ in combination is intractable, we therefore recommend simply tuning θ^p for a given language family as a good proxy, along with trying our defaults. The above results indicate that **model gains from M→D largely arise from learning to decode sound changes in dialects**, regardless of their placement in words.

M→D error analysis Arabic varieties gain less than other groups from M→D and M↔D. This may indicate that our M→D method does not approximate the Arabic linguistic variation well. Arabic varieties do exhibit phonological differences (which M→D simulates). However, these are often not expressed in the orthography, making lexical, syntactic, and register differences more prominent differentiators in text. We experimented simulating lexical choice variation in CRLs using HRL

WordNets as part of our artificial language generation method for 5 language families; however, this did not give us significant gains (Appendix E). DialUp in its current form may not be well suited for Arabic. Additionally, we note that FloRes sets are extremely close to standard Arabic, with a mean dev "dialectness" score (Keleg et al., 2023) of only 28%.¹³ This may also limit the scope of improvements from DialUp methods on this test set.

French-related Creole languages, which also had low baseline scores and small M→D gains, may have also suffered from genre mismatch (given their heterogeneous test sets (Robinson et al., 2024)). We initially hypothesized that their poor performance may be due to the HRL’s (Haitian) poor baseline performance with Aya-23. We attempted curriculum training with fine-tuning first on Haitian bitext and then with M→D; however this did not materially help. It is still possible that given Haitian’s low-resource nature (and its extremely low-resource CRLs), training token frequencies are too low to risk obfuscating them with DialUp noise. We hope future work will find more concrete conclusions on this trend.

D→M error analysis We investigate why D→M-func degrades the performance of Romance and Austronesian HRLs, while benefiting several CRLs in these families and Indic. We perform a manual evaluation of 100 words from our function word lexicons for 4 languages from these families (Table 9)¹⁴ to characterize the extent of noise in the D→M-func pipeline: a CRL word may be misidentified as a function word (mean function word identification accuracy: 83.6%), or it may be mistranslated (mean function word translation accuracy: 56.5%, general translation accuracy: 59.7%). The introduced noise is naturally particularly damaging when the baseline translation is good (as in for many high-resource CRLs). In fact, we find that D→M swaps hurt for high-resource CRLs even with ideal-case accurate swaps from clean lexicons. On the other hand, for a number of CRLs such as Maithili (mai), we observe that the benefits of making key function words swaps outweighs the negative impact of this noise; we find that even noisy automatically collected lexicons contribute to D→M

¹³Comparatively, MADAR-26 dev sets (Bouamor et al., 2018) for the same varieties achieve 59%. We initially attempted to use this dataset, but the genre mismatch with train data rendered both baseline and challenger scores too low.

¹⁴See Appendix I for details.




Figure 6: CRL-HRL proximity, measured as character F-score between dev sets, plotted against BLEU increase from M→D, suggests a peak from 20-30 chrF. Arabic varieties are outliers because the FloRes dialectal Arabic sets are unusually close to standard Arabic.

improvements in some languages (e.g. excluding relatively noisy Bafna et al. (2024a) lexicons results in a drop of 9 BLEU for Maithili with M2M). D→M benefits also depend on the model and HRL: in the case of Haitian, D→M performs well with M2M (yielding +7 BLEU for acf), but degrades baseline performance for Aya-23. This can be explained by the fact that Aya-23 has <10 BLEU performance on Haitian itself, making it unhelpful to switch other CRL words into Haitian; on the other hand, M2M baseline is strong for Haitian but weak for related CRLs, rendering D→M a useful strategy. **We therefore recommend D→M for low-baseline CRLs, with a model that performs well on the language family HRL.**

Note that D→M shows small gains for Arabic and Turkic CRLs. This is likely caused by a scarcity of word exchanges, which D→M relies on for increasing CRL input comprehensibility. In fact, D→M-func results in only 16.3% and 15% of words being swapped for Arabic and Turkic, respectively, compared to 43.1% for Austronesian, 38.5% for Romance, and 33.4% for Indic. (See Appendix F). This is potentially a result of Arabic and Turkic languages’ comparatively complex morphology, in addition to above noted issues with Arabic test sets. Grammatical morphemes in these languages are frequently affixes, clitics, or attached to affixes and seldom occur on their own; these will be missed by our whitespace-reliant word-switching technique. Future work can investigate integrating morphological analysis technologies when available into our methods.

6 Prior Work

DialUp has theoretical bases in a large body of prior work of noise pattern induction for robustness. Similar to M→D previous researchers have developed methods of injecting noise into training data, including orthographic variants or typos (Belinkov

and Bisk, 2018; Heigold et al., 2018), grammatical errors (Anastasopoulos et al., 2019), learned noise types (Brahma et al., 2023), and bilingual lexicon-induced code-mixing (Jones et al., 2023; Xia et al., 2019). These approaches require monolingual data or bilingual lexicons for from-scratch training. $M \rightarrow D$ innovates by inducing general dialectal robustness in a pre-trained model to unseen dialects without knowledge or resources in target varieties. $D \rightarrow M$ also draws from prior work, namely input processing techniques for LLM in-context learning via CRL-English lexicons, morphological analyzers, and grammar references (Zhang et al., 2024a; Ghazvininejad et al., 2023; Zhang et al., 2024b; Tanzer et al., 2024; Dimakis et al., 2024). $D \rightarrow M$ is likewise an inference-time intervention, but it is applicable to both traditional MT models and LLMs. Notably, our exploration of $D \rightarrow M$ is the first to validate the insight that model comprehensibility of often-ignored function words is beneficial, and to show considerable gains using small, automatically collected lexicons.

7 Conclusion

We present DialUp, a general recipe for expanding pretrained model coverage from its high-resource training languages to their dialect continua. This involves a finetuning-based technique involving artificial dialect generation using only HRL data, as well as an inference-time technique for rendering dialectal input more comprehensible to the model. DialUp shows large gains on several low-resource languages, and introduces a promising paradigm for making NLP flexible to potentially unseen and undocumented dialectal variation.

Limitations

Language-family-dependent gains The benefits of $M \rightarrow D$ and $D \rightarrow M$ clearly vary by language family: Indic and Romance benefit much more than other families, and Arabic in particular benefits minimally. $M \rightarrow D$ relies on approximating realistic synthetic dialects of the HRL by simulating various mechanisms of language variation. While these are broadly similar across language families and can be grouped into phonological, morphological, and lexical processes, which our artificial language generation method seeks to model, each language family naturally varies in typology, intra-family linguistic diversity, and the kinds of relationship that its CRLs have with each other and with the HRL.

For example, certain variants may differ from the HRL in their extent of lexical influence from a third (colonial) language, as in the cases of some Arabic and Turkic languages. Certain typologies or scripts may also be more or less amenable to our methods; e.g. we observe that suffix stacking in Turkic languages is challenging for the morphological noiser, which performs suffix identification in a manner most suited for fusional or morphologically non-complex languages (Ismayilzada et al., 2024). Finally, some language families, such as Arabic, may exhibit syntactic variation, which is not modeled at all by our current method. Our current artificial language generation method does not cater to these family-specific characteristics, which may contribute to its poorer performance on these families. We leave it to future work to refine the current catch-all technique of artificial language generation for specific families.

Different-script CRLs $M \rightarrow D$ does not support artificial language generation across scripts, i.e. it cannot perturb an HRL in one script into an artificial language in a different script. This is relevant for language families with more than one script: for example, several Turkic low-resource languages use the Cyrillic script, and cannot be supported with our current methods. Future work can look into transliteration-based approaches to handle this issue.

Models We provide results on one representative model of the traditional NMT as well as LLM paradigm, choosing an appropriate model with respect to its training languages given our goal of handling unseen and unknown target dialects, and available evaluation data. Naturally, our method should be tested on a variety of models from each paradigm to gauge its usefulness as a general data augmentation technique for achieving dialectal robustness.

Availability of bilingual lexicons Our best-performing $D \rightarrow M$ -func approach depends on the availability of CRL-HRL function word mappings. These are much easier to collect than bilingual lexicons given that they are a small class of closed class words, and we show that using even noisy alignments from a small amount of bitext is effective for this purpose: however, this may still be unfeasible for particularly low-resource languages.

Long way to go While our methods show considerable improvements on cognates for CRLs (ex-

amples in Table 4), and improved scores using automated MT evaluation metrics such as BLEU and COMET, the improved translation may still not be usable or coherent. We provide examples of this in Table 7.

Ethics Statement

Our work seeks to expand the benefits of mainstream NLP (specifically, machine translation) in standard high-resource varieties of languages to their diverse, dialectal, colloquial, and/or low-resource neighbours. Our work introduces general techniques towards this goal; however, any particular language community may have specific and different needs, potentially not served by mainstream NLP, which our work does not take into account (Bird, 2020). While we attempt to use minimal resources in low-resource languages to achieve our goal in order to make our methods resource-light and widely applicable, using only HRL data in training risks reinforcing a culturally mainstream view of the world in our models. This may be particularly problematic for continua where dialect-speaking communities diverge in political or other views from communities speaking the standard variety.

References

- Rached Ben Abdelkader, Abdeljelil Ayed, and Aziz Naouar. 1977. [Peace corps english-tunisian arabic dictionary](#).
- Sotiris Anagnostidis and Jannis Bulian. 2024. [How susceptible are llms to influence in prompts?](#) *Preprint*, arXiv:2408.11865.
- Antonios Anastasopoulos, Alison Lui, Toan Q. Nguyen, and David Chiang. 2019. [Neural machine translation of text from non-native speakers](#). 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 3070–3080, Minneapolis, Minnesota. Association for Computational Linguistics.
- Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Kelly Marchisio, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick Frosst, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. [Aya 23: Open weight releases to further multilingual progress](#). *Preprint*, arXiv:2405.15032.
- Niyati Bafna, Cristina España-Bonet, Josef van Genabith, Benoît Sagot, and Rachel Bawden. 2024a. [When your cousin has the right connections: Unsupervised bilingual lexicon induction for related data-imbalanced languages](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 17544–17556, Torino, Italia. ELRA and ICCL.
- Niyati Bafna, Kenton Murray, and David Yarowsky. 2024b. [Evaluating large language models along dimensions of language variation: A systematic investigation of cross-lingual generalization](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18742–18762, Miami, Florida, USA. Association for Computational Linguistics.
- Özge Bakay, Özlem Ergelen, Elif Sarıms, Selin Yıldırım, Bilge Nas Arıcan, Atilla Kocabalcıoğlu, Merve Özçelik, Ezgi Samyay, Oğuzhan Kuyrukçu, Begüm Avar, and Olcay Taner Yıldız. 2021. [Turkish WordNet KeNet](#). In *Proceedings of the 11th Global Wordnet Conference*, pages 166–174, University of South Africa (UNISA). Global Wordnet Association.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and Natural Noise Both Break Neural Machine Translation](#). *Preprint*, arxiv:1711.02173.
- Martin Benjamin. 2019. [Teach you backwards: An in-depth study of google translate for 108 languages](#). Online.
- A. Bergman and Mona Diab. 2022. [Towards responsible natural language annotation for the varieties of Arabic](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 364–371, Dublin, Ireland. Association for Computational Linguistics.
- Pushpak Bhattacharyya. 2010. [IndoWordNet](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Steven Bird. 2020. [Decolonising speech and language technology](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Houda Bouamor, Nizar Habash, Mohammad Salameh, Wajdi Zaghواني, Owen Rambow, Dana Abdulrahim, Ossama Obeid, Salam Khalifa, Fadhil Eryani, Alexander Erdmann, and Kemal Oflazer. 2018. [The MADAR Arabic dialect corpus and lexicon](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Maharaj Brahma, Kaushal Maurya, and Maunendra Desarkar. 2023. [SelectNoise: Unsupervised noise injection to enable zero-shot machine translation for extremely low-resource languages](#). In *Findings of the*

- Association for Computational Linguistics: EMNLP 2023*, pages 1615–1629, Singapore. Association for Computational Linguistics.
- Samuel Cahyawijaya, Holy Lovenia, and Pascale Fung. 2024. [LLMs are few-shot in-context low-resource language learners](#). *arXiv preprint arXiv:2403.16512*.
- Siddharth Dalmia, Ramon Sanabria, Florian Metze, and Alan W. Black. 2018. [Sequence-based multi-lingual low resource speech recognition](#). *CoRR*, abs/1802.07420.
- Art Dieli. 2011. [The sicilian language](#). Online.
- Antonios Dimakis, Stella Markantonatou, and Antonios Anastasopoulos. 2024. [Dictionary-aided translation for handling multi-word expressions in low-resource languages](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2588–2595, Bangkok, Thailand. Association for Computational Linguistics.
- Sumanth Doddapaneni, Rahul Aralikkatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan, and Pratyush Kumar. 2023. [Towards leaving no Indic language behind: Building monolingual corpora, benchmark and models for Indic languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12402–12426, Toronto, Canada. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 644–648.
- Anas Assadek El Ouamari. 2024. [Moroccan darija arabic translator](#).
- Nick C Ellis. 2008. The dynamics of second language emergence: Cycles of language use, language change, and language acquisition. *The modern language journal*, 92(2):232–249.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Man-deep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. [Beyond english-centric multilingual machine translation](#). *Preprint*, arXiv:2010.11125.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- Jonathan Garrett, Greg Lastowka, Akgul Muhammetmuradova, Jahan Myradova, Kimberly Naahielua, Meena Pallipamu, Muhammed Rustamov, Meretgul Sharipova, and Aynabat Yaylymov. 1996. [Turkmen-english dictionary](#).
- Ronald Gatty. 2009. [Fijian-english dictionary with notes on fijian culture and natural history](#).
- Marjan Ghazvininejad, Hila Gonen, and Luke Zettlemoyer. 2023. [Dictionary-based phrase-level prompting of large language models for machine translation](#). *arXiv preprint arXiv:2302.07856*.
- Mike Green. 2020. [Egyptian arabic dictionary](#).
- Harald Hammarström, Robert Forkel, Martin Haspelmath, and Sebastian Bank. 2024. [Glottolog](#). Available online at <http://glottolog.org>, Accessed on 2024-12-14.
- Georg Heigold, Stalin Varanasi, Günter Neumann, and Josef van Genabith. 2018. [How Robust Are Character-Based Word Embeddings in Tagging and MT Against Word Scrambling or Random Noise?](#) In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 68–80, Boston, MA. Association for Machine Translation in the Americas.
- Dirk Hovy and Christoph Purschke. 2018. [Capturing regional variation with distributed place representations and geographic retrofitting](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4383–4394, Brussels, Belgium. Association for Computational Linguistics.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Mete Ismayilzada, Defne Circi, Jonne Sälevä, Hale Sirin, Abdullatif Köksal, Bhuwan Dhingra, Antoine Bosselut, Lonneke van der Plas, and Duygu Ataman. 2024. [Evaluating morphological compositional generalization in large language models](#). *Preprint*, arXiv:2410.12656.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is ChatGPT A Good Translator? Yes With GPT-4 As The Engine](#). *Preprint*, arxiv:2301.08745.
- 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.

- Aditya Joshi, Raj Dabre, Diptesh Kanojia, Zhuang Li, Haolan Zhan, Gholamreza Haffari, and Doris Dippold. 2024. Natural language processing for dialects of a language: A survey. *arXiv preprint arXiv:2401.05632*.
- 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).
- Amr Keleg, Sharon Goldwater, and Walid Magdy. 2023. **ALDi: Quantifying the Arabic level of dialectness of text**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10597–10611, Singapore. Association for Computational Linguistics.
- Dylan Kobayashi, Jason Leigh, John Maher, and James Metz. 2016. **Samoan / english dictionary**. Department of Indo-Pacific Languages and Literature, and Department of Information & Computer Sciences, University of Hawai'i at Manoa.
- Arya D. McCarthy, Rachel Wicks, Dylan Lewis, Aaron Mueller, Winston Wu, Oliver Adams, Garrett Nicolai, Matt Post, and David Yarowsky. 2020. **The Johns Hopkins University Bible corpus: 1600+ tongues for typological exploration**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2884–2892, Marseille, France. European Language Resources Association.
- John Moorfield. 2024. **Te aka māori dictionary**.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, et al. 2023. **Crosslingual generalization through multitask finetuning**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111.
- Ercong Nie, Sheng Liang, Helmut Schmid, and Hinrich Schütze. 2023. **Cross-lingual retrieval augmented prompt for low-resource languages**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8320–8340, Toronto, Canada. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. **Universal dependencies v1: A multilingual treebank collection**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.
- Nurri Hirfana Bte Mohamed Noor, Suerya Sapuan, and Francis Bond. 2011. **Creating the open Wordnet Bahasa**. In *Proceedings of the 25th Pacific Asia Conference on Language, Information and Computation*, pages 255–264, Singapore. Institute of Digital Enhancement of Cognitive Processing, Waseda University.
- Javad Nouri and Roman Yangarber. 2014. **Measuring language closeness by modeling regularity**. In *Proceedings of the EMNLP'2014 Workshop on Language Technology for Closely Related Languages and Language Variants*, pages 56–65, Doha, Qatar. 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.
- Rebecca Posner and Marius Sala. 2024. **Linguistic characteristics of the romance languages**.
- Yasser Rezagui, Lahsen Abouenour, Fettoum Krieche, Karim Bouzoubaa, and Paolo Rosso. 2016. **Arabic WordNet: New content and new applications**. In *Proceedings of the 8th Global WordNet Conference (GWC)*, pages 333–341, Bucharest, Romania. Global Wordnet Association.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. **Comet: A neural framework for mt evaluation**. *Preprint*, arXiv:2009.09025.
- Ayla Rigouts Terryn and Miryam de Lhoneux. 2024. **Exploratory study on the impact of English bias of generative large language models in Dutch and French**. In *Proceedings of the Fourth Workshop on Human Evaluation of NLP Systems (HumEval) @ LREC-COLING 2024*, pages 12–27, Torino, Italia. ELRA and ICCL.
- Nathaniel Robinson, Raj Dabre, Ammon Shurtz, Rasul Dent, Onenamiyi Onesi, Claire Monroc, Loïc Grobol, Hasan Muhammad, Ashi Garg, Naome Etori, Vijay Murari Tiyyala, Olanrewaju Samuel, Matthew Stutzman, Bismarck Odoom, Sanjeev Khudanpur, Stephen Richardson, and Kenton Murray. 2024. **Kreyòl-MT: Building MT for Latin American, Caribbean and colonial African creole languages**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3083–3110, Mexico City, Mexico. Association for Computational Linguistics.
- Nathaniel Robinson, Perez Ogayo, David R. Mortensen, and Graham Neubig. 2023. **ChatGPT MT: Competitive for high- (but not low-) resource languages**. In *Proceedings of the Eighth Conference on Machine Translation*, pages 392–418, Singapore. Association for Computational Linguistics.
- Adriana Roventini, Antonietta Alonge, Nicoletta Calzolari, Bernardo Magnini, and Francesca Bertagna. 2000. **ItalWordNet: a large semantic database for**

Italian. In *Proceedings of the Second International Conference on Language Resources and Evaluation (LREC'00)*, Athens, Greece. European Language Resources Association (ELRA).

Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [Wiki-Matrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.

Morris Swadesh. 2015. [Swadesh lists by language](#). Online.

Garrett Tanzer, Mirac Suzgun, Eline Visser, Dan Jurafsky, and Luke Melas-Kyriazi. 2024. [A Benchmark for Learning to Translate a New Language from One Grammar Book](#). *Preprint*, arxiv:2309.16575.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, 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, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.

Mengzhou Xia, Xiang Kong, Antonios Anastasopoulos, and Graham Neubig. 2019. Generalized data augmentation for low-resource translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5786–5796.

Chen Zhang, Xiao Liu, Jiuheng Lin, and Yansong Feng. 2024a. [Teaching Large Language Models an Unseen Language on the Fly](#). *Preprint*, arxiv:2402.19167.

Kexun Zhang, Yee Man Choi, Zhenqiao Song, Taiqi He, William Yang Wang, and Lei Li. 2024b. [Hire a Linguist!: Learning Endangered Languages with In-Context Linguistic Descriptions](#). *Preprint*, arxiv:2402.18025.

Caleb Ziems, William Held, Jingfeng Yang, Jwala Dhamala, Rahul Gupta, and Diyi Yang. 2023. [Multi-VALUE: A framework for cross-dialectal English NLP](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 744–768, Toronto, Canada. Association for Computational Linguistics.

A Languages

See Table 6.

B Details of artificial language generation

In § 2, we introduce the artificial language generation method for $M \rightarrow D$, which consists of phonological, morphological, function word, and content word “noising”, or corruption, taken from [Bafna et al. \(2024b\)](#).

- The **phonological** noiser changes phonemes in a given left and right phonological context to phonetically nearby targets (e.g. $t \rightarrow d$, which differ only in voicing). We use script-to-IPA mappings and vice versa to first convert our input into IPA. We define plausible target phoneme sets for each phoneme based on its phonetic characteristics. The noised phoneme is then converted back into the input script.
- The **morphological** noiser targets suffixes in the HRL. Suffixes are identified heuristically, collected by frequency over an HRL corpus. A given suffix is noised by applying a fixed high amount of phonological noise over it ($\theta^p=0.8$), and changed globally over all occurrences to its noised version.
- Similarly, the **function word** noiser targets function words in the HRL. Function word identification is done in the following way: we curate a set of POS tags for closed-class words, including determiners, pronouns, conjunctions, auxiliaries, and adpositions. We then use the tagged Universal Dependencies corpus ([Nivre et al., 2016](#)) to identify all words in a given language whose most frequent tag is one of these. Since function words are a relatively small set, this procedure should yield coverage even over a small corpus. A given function word is noised by applying a fixed high amount of phonological noise over it ($\theta^p=0.8$), and changed globally over all occurrences to its noised version.
- The non-cognate **content** word noiser generates non-words using an HRL character ngram model as replacements for HRL content words. Note that this is the only noiser that models non-cognate processes.

C Experimental details

C.1 Aya prompt

LLMs are known to be sensitive to the prompt used for the task ([Anagnostidis and Bulian, 2024](#)).

We evaluated Aya-23 off-the-shelf on the following prompts, varying in their specification of the source language:

1. Translate into English:
<source_text>
2. Translate from a dialect of <hrl> into English:
<source_text>
3. Translate from <flores_code_of_lrl> into English:
<source_text>

The first two prompts are chosen in accordance with our goal of translating potentially unknown dialects of the HRL. While prompts 2 and 3 generally give improvements on the baseline for known target dialects (+1.3, +2.5, +1.3, -0.5 mean BLEU improvements for Arabic, Romance, Turkic and Indic respectively) the best performing prompt differs by language. For the purpose of our study, using prompts 2 and 3 is non-ideal, since it introduces the confound of degradation stemming from a certain target being referred to as a dialect of another language, or model familiarity with the particular dialect name or code. This is relevant because our evaluation languages range over degrees of resourcedness and presence in Aya-23, as well as relatedness from the corresponding HRL. Further, note that also use our chosen prompt during instruction finetuning over noisy text for consistency, and it is unclear what language name or code we should use for the noised text for best transfer to all target dialects in this setting. In order to avoid the above confounds and for simplicity, we use the first prompt for all approaches and languages, both during training and evaluation.

C.2 M2M tokenizer

M2M requires specification of source and target language tokens; the source language token is appended to the end of the input text during tokenization, and the target language token is provided to the decoder as its first token to specify the output language. We use CRL language tokens when they

Language	Family	Subfamily	ISO code	FloRes/NLLB	In M2M?	In Aya-23?	Resource Level
Indonesian	Austronesian	Malayo-Polynesian	ind	ind_Latn	id	ind	HRL
Javanese	Austronesian	Malayo-Polynesian	jav	jav_Latn	-	-	<200K
Sundanese	Austronesian	Malayo-Polynesian	sun	sun_Latn	-	-	<100K
Samoan	Austronesian	Malayo-Polynesian, Polynesian	smo	smo_Latn	-	-	<10K
Maori	Austronesian	Eastern Polynesian	mri	mri_Latn	-	-	<50K
Cebuano	Austronesian	Malayo-Polynesian	ceb	ceb_Latn	-	-	<12M
Standard Malay	Austronesian	Malayo-Polynesian	zsm	zsm_Latn	-	-	<2M
Tagalog	Austronesian	Malayo-Polynesian, Philippine	tgl	tgl_Latn	-	-	<500K
Ilocano	Austronesian	Malayo-Polynesian, Philippine	ilo	ilo_Latn	-	-	<100K
Fijian	Austronesian	Eastern Malayo-Polynesian	fij	fij_Latn	-	-	<10K
Plateau Malagasy	Austronesian	Malayo-Polynesian, Western Indonesian	plt	plt_Latn	-	-	<500K
Pangasinan	Austronesian	Malayo-Polynesian, Philippine	pag	pag_Latn	-	-	<10K
Arabic (MSA)	Arabic	Modern Standard Arabic	arb	arb_Arab	ar	ara	HRL
Mesopotamian Arabic	Arabic	Mesopotamian Arabic	acm	acm_Arab	-	-	<1K
Ta'izzi-Adeni Arabic	Arabic	Southern Yemeni	acq	acq_Arab	-	-	<1K
Tunisian Arabic	Arabic	Maghrebi	aeb	aeb_Arab	-	-	<1K
South Levantine Arabic	Arabic	Levantine Arabic	ajp	ajp_Arab	-	-	<1K
North Levantine Arabic	Arabic	Levantine Arabic	apc	apc_Arab	-	-	<1K
Najdi Arabic	Arabic	Peninsular Arabic	ars	ars_Arab	-	-	<1K
Moroccan Arabic	Arabic	Maghrebi	ary	ary_Arab	-	-	<100K
Egyptian Arabic	Arabic	Egyptian Arabic	arz	arz_Arab	-	-	<3M
Italian	Romance	Italo-Western	ita	ita_Latn	it	ita	HRL
Spanish	Romance	Italo-Western	spa	spa_Latn	es	spa	HRL
French	Romance	Italo-Western	fra	fra_Latn	fr	fra	HRL
Portuguese	Romance	Italo-Western	por	por_Latn	pt	por	HRL
Romanian	Romance	Eastern Romance	ron	ron_Latn	ro	ron	<3M
Lombard	Romance	Italo-Western	lmo	lmo_Latn	-	-	<200K
Asturian	Romance	Italo-Western	ast	ast_Latn	ast	-	<500K
Galician	Romance	Italo-Western	glg	glg_Latn	gl	-	<500K
Venetian	Romance	Italo-Western	vec	vec_Latn	-	-	<200K
Catalan	Romance	Italo-Western	cat	cat_Latn	ca	-	<2M
Sicilian	Romance	Italo-Western	scn	scn_Latn	-	-	<100K
Sardinian	Romance	Southern Romance	srd	srd_Latn	-	-	<50K
Friulian	Romance	Italo-Western	fur	fur_Latn	-	-	<10K
Ligurian	Romance	Italo-Western	lij	lij_Latn	-	-	<50K
Occitan	Romance	Italo-Western	oci	oci_Latn	oc	-	<200K
Turkish	Turkic	Oghuz	tur	tur_Latn	tr	tur	HRL
Azerbaijani	Turkic	Oghuz	azj	azj_Latn	az	-	<1M
Crimean Tatar	Turkic	Kipchak	crh	crh_Latn	-	-	<50K
Turkmen	Turkic	Oghuz	tuk	tuk_Latn	-	-	<50K
Uzbek	Turkic	Karluk	uzb	uzn_Latn	uz	-	<1M
Hindi	Indic	Western	hin	hin_Deva	-	-	HRL
Awadhi	Indic	Eastern	awa	awa_Deva	-	-	CRL
Bhojpuri	Indic	Eastern	bho	bho_Deva	-	-	<100K
Chhattisgarhi	Indic	Eastern	hne	hne_Deva	-	-	<2M
Magahi	Indic	Eastern	mag	mag_Deva	-	-	<2M
Maithili	Indic	Eastern	mai	mag_Deva	-	-	<50K
Haitian	French-related Creole	Haitian Creole	hat	hat_Latn	ht	hat	HRL
Guadeloupean	French-related Creole	Antillean Creole	gcf	-	-	-	<1K
Martinican	French-related Creole	Antillean Creole	gcf	-	-	-	<1K
Saint Lucian Patois	French-related Creole	Antillean Creole	acf	-	-	-	<1K
French Guianese	French-related Creole	French Guianese Creole	gcr	-	-	-	<1K
Louisiana Creole	French-related Creole	Louisiana Creole	lou	-	-	-	<1K
Mauritian	French-related Creole	Bourbonnais Creoles	mfe	-	-	-	<1K
Réunion Creole	French-related Creole	Bourbonnais Creoles	lmo	-	-	-	<1K
Seychellois	French-related Creole	Bourbonnais Creoles	crs	-	-	-	<1K

Table 6: Languages studied organized by language family. We also report the languages supported by the models used, as well resource-level. We report the number of Wikipedia articles for a language as a proxy for its resource level.

are supported by M2M, and the HRL language token otherwise. Note that the tokenizer itself is shared across all languages, as well as encoder parameters. While decoder parameters have language-family-specific components depending on the target language, this is irrelevant for us as we always have English as our target language.

D Bilingual Lexicon Collection for D→M

Our D→M approach uses CRL-HRL bilingual lexicons for an inference-time intervention. Ideally, we would like the lexicons to cover fully inflected words and all parts-of-speech.

D.1 Survey of lexicons, and challenges in collection

See a list of the language-pair specific bilingual lexicons we considered in Table 29: as shown here, many lexicons list lemmas rather than inflected words, rendering them unusable for our purpose. This is naturally particularly a problem for morphologically rich languages.

APIs and web apps Many available sources for bilingual lexicons are APIs and web apps. These include but are not limited to Google Translate ¹⁵, Microsoft Bing ¹⁶, ModernMT API ¹⁷, Alibaba

¹⁵<https://translate.google.com/>

¹⁶<https://www.bing.com/translator>

¹⁷<https://www.modernmt.com/translate>

Translate¹⁸, tradukka¹⁹, freelang²⁰, From-to.io²¹, iTranslate²², and Glosbe²³. However:

- **Querying is often problematic.** APIs are often blocked by some paywall and mining the databases of web apps may be considered unethical or illegal. This is the case for The Living Arabic Project²⁴, a database of endangered Arabic dialects.
- **The quality of such resources can be poor.** A study from the Kamusi Project found that Google Translate performed poorly on some of the study’s languages, namely Plateau Malagasy, Ilocano, Samoan, and Maori (Benjamin, 2019). Glosbe is a crowdsourced database, meaning the translations may not be accurate.

Programmatic readability Some lexicons that exist as a PDF may have text that is difficult to extract. For example, the Peace Corps Tunisian Arabic Dictionary (Abdelkader et al., 1977) is grainy, making it difficult to perform OCR.

English-centric resources While we could find many lexicons that translated from a low/mid-resource language into English, we had trouble finding lexicons that translated from a low/mid-resource language into a related non-English resource language (see Table 29). This underscores the existing bias toward English-centric resources and technologies in natural language processing (Rigouts Terry and de Lhoneux, 2024).

D.2 Lexicon sizes

See Tables 30, 31, and 32 for lexicon sizes of the lexicons we aggregated and used for each CRL-HRL pair, with a breakdown by function and content word counts. In these tables, we also report the word type coverage, a measure of how extensively the lexicon documents a LRL. Coverage is calculated as the percentage of unique words in the FloRes dev set/JHU Bible test set that are documented in the associated LRL lexicon.

Note that we only use publicly available data resources.

¹⁸<https://translate.alibaba.com/>

¹⁹<https://tradukka.com/>

²⁰<https://www.freelang.net/>

²¹<https://from-to.io/>

²²<https://itranslate.com/>

²³<https://glosbe.com/>

²⁴<https://www.livingarabic.com/>

E Variant approaches and hyperparameter exploration

E.1 Choosing noising dials

The search space for $\theta^{p,m,f,c}$, i.e., the noising dials used to generate artificial dialectal data, is large. We conduct some small-scale experiments to understand the impact of this choice on the M→D approaches.

Default choices for $\theta^{p,m,f,c}$ We use the noiser framework as set up in Bafna et al. (2024b). Within the Bayesian generative framework of these noisers, it is possible to compute the posteriors (MLE estimates) of θ^n independently for each noiser, given a real CRL-HRL pair. Bafna et al. (2024b) compute and provide these posteriors over a number of CRL-HRL pairs for a few different language families, that include languages from four of the six language families that we work with. Our default choices for these parameters depend on the observed range of these θ -posteriors over 18 real language pairs. For M→D-shell, we take a simple average of these posteriors to give us a single $\theta^{p,m,f,c}$ -radius, which serves as a reasonable default distance between a random CRL and HRL. For M→D-cloud, given that we generate artificial languages on several hyperspheres up to a maximum $\theta^{p,m,f,c}$ -radius, we take maximums over the observed posteriors instead in each dimension, discarding some clear outliers. While this yields reasonable defaults for the phonological, morphological, and function word noisers (which all model cognate processes), we find that this is not ideal for the content word noiser, as discussed next.

Setting θ^c Our initial experiments indicate that setting θ^c , i.e. the content word noiser dial in accordance with the procedure for calculating other noiser dials is suboptimal, and much lower θ^c ’s are better. This makes sense: the content word noiser is the only noiser that models *non-cognate* processes, and introduces non-words with no connection to the source word into the noised text. This is naturally not helpful, since the introduced non-words have no systematic relationships with the target dialects, and there is no way for the model to generalize what it observes. Therefore, we set $\theta^c = 0.001$, i.e., very low, and largely rely on and discuss the other three noisers for this work (as in § 5).

Effect of hyperparameters on M→D-cloud We tried more aggressive noising for Indic, Turkic, and

Haitian, using $\theta_{0.2,0.8,0.8,0.001}^{p,m,f,c}$ as the max-radius for M→D-cloud; i.e. significantly increasing θ^p and θ^m (the θ^f default is already high at 0.8). This only gave minor variations on existing results for all language families and models. This is not very surprising: since M→D-cloud samples from many radii through the hypersphere up to the max-radius, it is relatively less sensitive to the choice of the max-radius $\theta^{p,m,f,c}$.

Effect of hyperparameters on M→D-shell On the other hand, M→D-shell only samples from the shell defined by the radius, and therefore might be more sensitive to this choice. We observe that it performs slightly better or worse than M→D-cloud (see Table 19 and Table 20) depending on the language and model. Given a specific target language family, it may make sense to conduct a hyperparameter search for θ for optimal M→D-shell performance. Our small-scale experiments with Indic with a few different θ 's indicate only minor improvements from this search over our default parameters; however, this may vary by language family and model. Given the intractability of a systematic hyperparameter search over 3-dimensional θ -radius (excluding the non-cognate lexical noiser, i.e. $\theta^{p,m,f}$), we instead show that it is possible to approximate all cognate processes solely with θ^p ; we discuss this in § 5.

Number of radii N for M→D-cloud For the approach M→D-cloud, we use N radii at uniform intervals from 0 to max-radius θ ; thus, N defines the density of the cloud given a max-radius. Of course, the total amount of data used remains constant regardless of N and θ . We compute results for Indic for the choice of number of radii $K \in \{5, 10, 20\}$ and observe very minor differences. We use $N = 10$ throughout the paper. Note that since θ is a probability, $N = 10$ is already enough to approximate an effectively continuous cloud up to the max-radius. While there is no cost to increasing N , we observe no gains from doing so beyond a certain point.

Parametrizing randaug-shell and -cloud We can consider randaug to consist of a character-level and a word-level noiser. The difference from our noisers lies only in the target generation: while our noisers choose linguistically plausible targets as described in § 2, the randaug noisers choose random equivalents in accordance with previous literature. The character-level noiser is analogous to the

M→D phonological noiser which affects characters, and the word-level noiser with the content word noiser which affects lexemes. For a fair comparison and in order to test the importance of linguistically-motivated target generation, we therefore set the dials for these noisers to be the same as the default dials for their analogous noisers, both for -shell and -cloud variants. We use $\theta_{rc,rw}^{0.05,0.001}$ for randaug-shell and $\theta_{rc,rw}^{0.07,0.001}$ as the max-radius for randaug-cloud. For the latter, we use $K = 10$ hyperspheres, consistent with M→D-cloud. Small tuning experiments over the Indic family yield minor variations in the baseline results obtained with these parameters. Note that the morphological and function word noisers do not have random baseline equivalents in the literature: these are inherently linguistically-motivated and inspired by patterns of dialectal variation.

E.2 Multiple HRLs

A language family may have more than one HRL, and the LRLs in the family may be closer to any of them. We experiment with using two HRLs as sources for noising, additionally using French for the Romance family, Uzbek for Turkic, and Tunisian for the Arabic family, along with the primary HRLs. We still maintain the total amount of finetuning data, but split it uniformly between our source HRLs. We choose Uzbek since it differs from Turkish in having a high amount of Russian-influenced and Persian-influenced vocabulary, similar to crh_Latn. Similarly, we choose Tunisian as a representative of the Maghrebi Arabic dialects, hoping to help with Moroccan (ary_Arab) which also belongs to this subfamily. See Table 21 for the results for Aya-23. We see that this helps for some Romance dialects that share high similarities with French and Spanish (such as oci_Latn and glg_Latn). Similarly, for Arabic, this gives a considerable performance boost for Moroccan Arabic. However, we don't observe any improvements for the Turkic family except Uzbek itself.

E.3 Different datasets

We experiment with different choices of base dataset for noising for Indic: namely, indiccorp (Doddapaneni et al., 2023) and the NLLB dataset (Team et al., 2022). We observed that the choice of dataset matters for the absolute performance of the finetuned model (fthr1): specifically, finetuning with indiccorp actually hurts the baseline performance of both models. However, the gains from

the noising method over this finetuned baseline are consistent across these three datasets, for the Indic family. We present results on indiccorp in Table 22); those on the NLLB dataset show similar trends (see Table 23).

Training hyperparameters We observe that increasing either the amount of data used or number of epochs worsens ordinary HRL fine-tuning results slightly for CRLs, while M→D approaches show small improvements for some languages. We fix these hyperparameters for all experiments (i.e. 100K sentences of bitext, a single epoch) for ease of comparison.

E.4 Using CRL-Eng swapping

We experiment with switching CRL words directly into English instead of the HRL: as expected, this degrades performance considerably since the model now receives unnatural CRL-English code-mixed input. See Table 39 and Table 38 for results.

E.5 Enhancing noisers to model lexical variation: Semantic noise

We would like to model different patterns of lexical usage among closely-related languages. For example, while the word “book” is “pustak” in Marathi and “kitab” in Hindi by common usage, “pustak” is also an entirely comprehensible Hindi word. We would like our noisers therefore to expose the model during finetuning to various lexical variants of a given concept as they might be realized in different dialects of a language family. We use WordNets in all our HRLs: IndoWordNet for Hindi (Bhattacharyya, 2010), TurkishWordNet (Bakay et al., 2021), Arabic WordNet (AWN v2) (Regragui et al., 2016), WordNet Bahasa (Noor et al., 2011), and Italian Wordnet (Roventini et al., 2000) to “noise” a given word with some probability θ^s to its (randomly chosen) synonym. Note that since these wordnets contain lemmas, whereas we typically encounter inflected forms in text, we first lemmatize the word to be noised, perform semantic noising, and then re-append its original inflections, to roughly maintain grammaticality. This process may introduce incidental noise at inflection boundaries, but is tolerable given our general goal of perturbing the data.

This noised and inflected synonym then undergoes the other noising processes as per usual, modeling the intuition that the diverging lexeme would still bear the effects of general phonological and

morphological change in a given dialect.

We set $\theta^s \in \{0.1, 0, 3\}$, keeping other noising parameters constant; however, we do not see that this helps much. See language family means for M2M in Table 24.

F Function word identification for D→M

HRL M→D requires function word identification in the HRL, for the function word lexical noiser. We use the same procedure for function word identification in the HRL as used by the lexical noiser in differential between functional and content words, consistent with Bafna et al. (2024b), summarized in Appendix B, to create a list of HRL function words.

CRLs The D→M-func approach only affects function words in the CRL input, and therefore requires function word identification for CRL text. Given that we have a set of HRL function words and HRL-CRL bilingual lexicons, we can identify function words in the CRL using projection of POS tags from the HRL.

Comparing against natural distribution The impact of our best-performing variant D→M-func depends on the identification of function words in the CRL input as described above. We would like to evaluate the coverage of this method; however, we lack annotated data for the CRLs for this evaluation. Instead, assuming that language families share general distributive properties of function words, we compare the percentage of identified function words in our CRL input against the natural percentage of function words in the HRL (as a representative of the language family) in the UD corpus. See Table 36. We observe that D→M affects an expected percentage of words for Romance, Turkic, and Indic, but there is over-identification for Austronesian and under-identification for Arabic, possibly due to noisier automatic alignments for these families.

See Table 37 for a summary of the number of (1) content, (2) functional, and (3) content and functional combined (represented as “all” in the table) words switched out in each D→M approach.

G More examples

See examples of M→D and D→M outputs in Table 7 and Table 8.

Input	ई छेत्र में एगो एंटीबॉडी कॉकटेल , जेकर नाम ZMapp बा , शुरू में अच्छा नतीजा देखइलस , लेकिन अध्ययन कइला पर पता चलल कि हई इबोला से हीवे वाला मौत रोके खातिर ओतना फायदेमंद नइखे जेतना एकरा बारे में बतावल गइल रहे .
Reference	One antibody cocktail, ZMapp, initially showed promise in the field, but formal studies indicated it had less benefit than sought in preventing death.
Baseline Output	The egocentric antibody cocktail, Jake named ZMapp Ba , initially had a good outcome, but the study on Kila revealed that the deaths caused by HIV/AIDS in order to prevent HIV/AIDS remained beneficial to the awareness of HIV/AIDS .
M→D Output	The antibody cocktail , ZMapp, was initially a good result , but studies have shown that it has been a very beneficial outcome in preventing deaths from Ebola .
Input	Ind ol test PALM, ZMapp a'l é stait drovad coma contròll, 'l saria a dir qe a i sienziats a i 'l han drovad coma basament e a i hann comparad a i olter tré trataments a quest qi.
Reference	In the PALM trial, ZMapp served as a control, meaning scientists used it as a baseline and compared the three other treatments to it.
Baseline Output	In the test PALM, ZMapp a'l é stait drovad coma contròll , 'l would be to say that to the sienziats to the l'han drovad coma basament and to the hann compared to the olter tré treatments to this qi .
M→D Output	In the PALM test, ZMapp's standing as a control would be to say that the subjects were found as a base and compared to the other treatments to this qi .
Input	Nen Lunes, si Sara Danius, say permanentin sekretaryo na Nobel Committee for Literature ed Swedish Academy et inyanunsio tod publiko diad sakey a programa ed Sveriges Radio diad Sweden, ya aga direktan nakontak na komite si Bob Dylan diad impanalo tod 2016 Nobel Prize in Literatura, kanian agla ra nansagpot ya sikatoy kontaken lamet.
Reference	On Monday, Sara Danius, permanent secretary of the Nobel Committee for Literature at the Swedish Academy, publicly announced during a radio program on Sveriges Radio in Sweden the committee, unable to reach Bob Dylan directly about winning the 2016 Nobel Prize in Literature, had abandoned its efforts to reach him.
Baseline Output	Nen Lunes, si Sara Danius, say permanentin secretary na Nobel Committee for Literature ed Swedish Academy et inyanunsio tod publiko diad sakey a programa ed Sveriges Radio diad Sweden, ya aga direktan nakontaktak na komite si Bob Dylan diad impanalo tod 2016 Nobel Prize in Literatura, kanian agla ra nansagpot ya sikatoy kontaken lamet.
M→D Output	On Monday, Sara Danius was the permanent secretary of the Nobel Committee for Literature and the Swedish Academy and the public in Sweden's Sveriges Radio program, and the director of Bob Dylan's committee for the 2016 Nobel Prize in Literature.

Table 7: Showing improvement of M→D over off-the-shelf

Original & D→M Input	Luneddi, Sara Danius, segretaria permanenti dū cumitatu Nobel pā littiratura ntā cademia svidisi, dissi pubbricamenti nna nu prugramma radiu di Radio Sveriges nna Svezzia ca lu svidisi, ha pubbricamenti in un prugramma radiu di Radio Sveriges in Svezzia che il cumitatu, nun putennu rintracciari direttamenti Bob Dylan e nfirmàrilu ca avia vinciutu lu cumitatu, non putennu rintracciari direttamente Bob Dylan e nfirmàrilu che aveva vinciutu il premiu Nobel pā littiratura 2016, smisi di pruvari a cuntattàrilu.
Reference	On Monday, Sara Danius, permanent secretary of the Nobel Committee for Literature at the Swedish Academy, publicly announced during a radio program on Sveriges Radio in Sweden the committee, unable to reach Bob Dylan directly about winning the 2016 Nobel Prize in Literature, had abandoned its efforts to reach him.
Original Output	On Monday, Sara Danius, permanent secretary due cumitatu Nobel pā littiratura ntā cademia svidisi , said publications nna nu prugramma radiu di Radio Sveriges nna Sweden ca lu cumitatu, nun putennu rintracciari direttamenti Bob Dylan and nfirmàrilu ca avia vinciutu lu premio Nobel pā littiratura 2016, ceased pruvari a cuntattàrilu .
D→M Output	On Monday, Sara Danius, permanent secretary of the Nobel laureate for literature at the Swedish Academy , has published in a radio prugram of Radio Sveriges in Sweden that the cumitatu, could not trace Bob Dylan and nfirmàrilu who had won the Nobel prize for literature 2016, stopped looking at cuntattàrilu
Original & D→M input	हुनका सभक अनुशासित रक्षा , गेंद संग खेलबाक कौशल आ बेस नीक टीम प्रयास हुनका सभ के आन से अलग टीम बनाबैत उन्हें के अनुशासित रक्षा , गेंद साथ खेलबाक कौशल और ने नीक टीम प्रयास उन्हें के को आन से अलग टीम बनाबैत छल आ ई स्पष्ट भ' गेल छल जे ई हराबए वला टीम अछि। था और यह स्पष्ट भ' गया था जो यह हराबए वला टीम है।
Reference	Their disciplined defence, ball handling skills and excellent team work made them stand out and it was clear that this was the team to beat.
Original Output	Jhfxjyfhjhjhkhhdtdtfnhjhjb hvydhbn hey nah gun hvybj hvybj hvybj him bun jvycstghnh tcj her n f n him g j t j him ubufybknknugyfybj hvybj much hbi b mhvghvhtdj gf
D→M Output	They had disciplined defense, backbone skills with the ball and made a nick team effort to make them a different team from K'u'lu and it was clear that this was a bad team .

Table 8: Showing improvement of D→M over off-the-shelf

H Wins and variance for approach paradigms

See Figure 7 and Figure 8 for a comparison of win rates over baseline and improvement distributions

for different paradigms for M2M and Aya-23 respectively, for all language families.

I Manual Evaluation of Function Word Lexicons

lang	FW acc. (%)	FW Id. (%)	Gen. acc. (%)
fra_ita	67.6	67.0	60.4
bho_hin	97.9	51.6	74.7
arz_arb	84.0	32.9	38.2
azj_tur	84.9	74.5	65.3
Mean	83.6	56.5	59.7

Table 9: Function word (FW) accuracy (%), function word identification accuracy (FW Id, %), general translation accuracy (%) in function word lexicons

The evaluation was conducted by three of our authors, fluent in Turkish (tur), Hindi (hin), Arabic (arb), Egyptian (ary), and French (fra), with a working understanding of the respective CRLs and Italian (ita). See Table 9.

J Training and evaluation details

J.1 Compute

Our noising approaches and baselines require finetuning: we had 2 models x 8 approaches x 6 HRLs = 96 finetuning runs. Each experiment required about 3 hours on a single A100, totaling 288 GPU hours.

We conducted evaluation for 54 languages (49 CRLs + 6 HRLs) x 15 total approaches = 810 evaluation runs. Each evaluation took 30 minutes on a single GTX 1080 machine, totaling 405 GPU hours.

These calculations do not include development and initial experiments, and ablation studies.

J.2 Evaluation

We tested our models on the devtest splits of the FloRes-200 dataset, using BLEU (Hugging-Face evaluate wrapper: mixed case, tokenization with sacrebleu 13a) as well as COMET (Unbabel/wmt22-comet-da).

K Results by language and approach

We summarize our results in the following.

- Table 10 and Table 11 detail the average BLEU and COMET scores respectively for each approach within the M→D, D→M, and

M↔D paradigms for each language family. Table 12 provides COMET score means for low-baseline CRLs (<25 BLEU) per language family.

- Table 13, Table 14, Table 15, and Table 16 present the BLEU and COMET scores of best-performing variant each of the baseline, M→D, D→M, and M↔D approaches, for M2M and Aya-23. Note that the best approach is chosen on a family-wide basis, in order to recommend a general strategy for a given language family, by comparing means over the language family for each variant, and may not reflect the best scores for an individual language in that paradigm (similarly for the best baseline). **This is different from how best variants are chosen for Figure 2, Figure 3, Figure 7, and Figure 8**, which consider best performing paradigms variants on a per-language basis.

The following tables provide a more detailed view of the performance of each paradigm.

- M→D: Table 17, Table 18, Table 19. and Table 20 detail the BLEU and COMET performance for each language in our language family families using each approach in the M→D paradigm. Table 7 provides a qualitative example of how translations improved following the M→D paradigm.
- D→M: Table 25, Table 26, Table 27, and Table 28 detail the BLEU and COMET performance for each language in our language family families using each approach in the D→M paradigm. Table 8 provides a qualitative example of how translations improved following the D→M paradigm.
- M↔D: Table 40, Table 41, Table 42, and Table 43 detail the BLEU and COMET performance for each language in our language family families using each approach in the M↔D paradigm.

L Use of AI assistants

We used GitHub copilot for coding assistance. No AI assistants were used for any writing purposes.

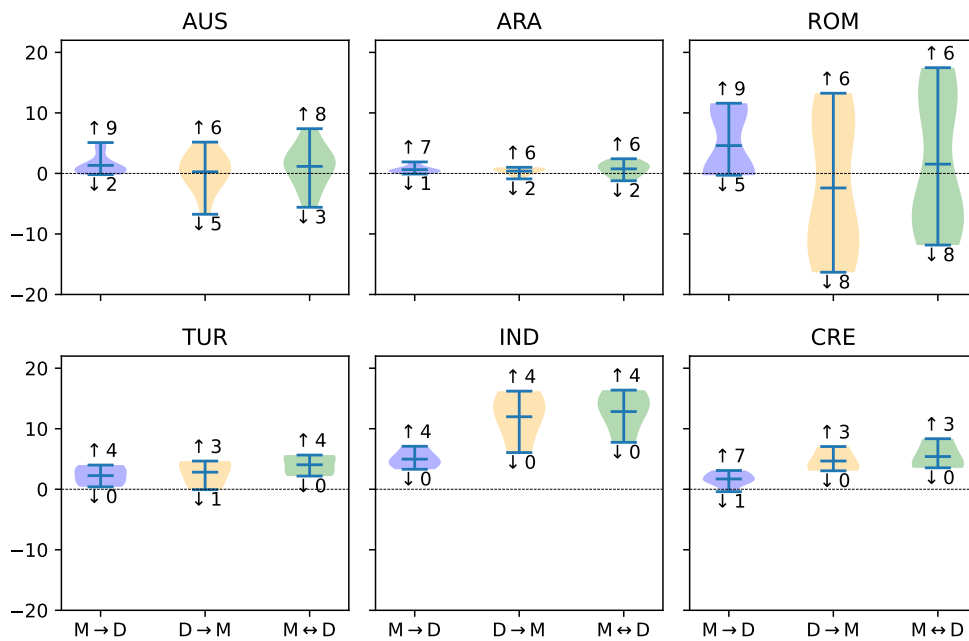


Figure 7: BLEU score improvements over the best baseline with M2M for all language families. ↑ and ↓: # CRLs with positive/negative gains. M→D gives more consistent positive gains.




Figure 8: BLEU score improvements over the best baseline with Aya-23 for all language families. ↑ and ↓: # CRLs with positive/negative gains. M→D gives more consistent positive gains.

		M2M							Aya-23						
		AUS (9)	ARA (7)	ROM (6)	TUR (4)	IND (4)	CRE (8)	Mean	AUS (9)	ARA (1)	ROM (5)	TUR (4)	IND (4)	CRE (9)	Mean
Baselines	off-the-shelf	14.9	21.2	26.6	5.0	12.6	5.9	14.3	<u>11.9</u>	31.5	29.5	7.9	<u>20.4</u>	7.3	<u>18.1</u>
	fthrln	<u>+0.3</u>	<u>+0.3</u>	<u>+0.4</u>	-0.1	<u>+0.4</u>	<u>+0.9</u>	<u>+0.4</u>	-1.0	-0.7	-0.1	<u>+1.2</u>	-1.8	<u>+2.3</u>	-0.0
	ftrandaug	+0.2	+0.2	+0.2	<u>+0.2</u>	<u>+0.4</u>	<u>+0.9</u>	+0.3	-0.9	-0.5	<u>+0.2</u>	+1.1	-1.8	+2.1	<u>+0.0</u>
M→D	-shell	<u>+1.7</u>	<u>+1.0</u>	<u>+5.1</u>	<u>+2.5</u>	+5.0	<u>+2.7</u>	<u>+3.0</u>	+1.1	-0.2	+2.4	+2.9	<u>+2.8</u>	+3.0	+2.0
	-cloud	+1.1	+0.7	+3.9	+1.6	<u>+5.4</u>	<u>+2.7</u>	+2.6	<u>+1.2</u>	-0.0	<u>+2.5</u>	<u>+3.2</u>	+2.7	<u>+3.5</u>	<u>+2.2</u>
D→M	-cont	<u>-1.5</u>	-2.0	-7.1	+1.4	+0.3	+2.5	-1.1	-0.1	-4.3	-7.6	-0.6	-1.1	-3.8	-2.9
	-func	<u>-1.5</u>	<u>+0.7</u>	<u>-2.8</u>	+1.0	<u>+12.0</u>	+3.7	<u>+2.2</u>	+0.7	<u>-1.1</u>	<u>-1.8</u>	<u>+0.3</u>	<u>+5.6</u>	<u>-3.3</u>	<u>+0.1</u>
	-all	-2.1	-1.2	-9.2	<u>+3.0</u>	+11.4	<u>+5.6</u>	+1.2	<u>+1.5</u>	-5.0	-8.4	+0.0	+4.4	-3.5	-1.8
M↔D	-cloud-cont	<u>+0.1</u>	-1.4	-2.7	+2.6	+5.1	+3.7	+1.2	+1.6	-3.6	-4.0	+1.8	+1.6	+2.6	+0.0
	-cloud-func	-0.3	<u>+1.1</u>	<u>+2.0</u>	+2.3	<u>+13.0</u>	+4.6	<u>+3.8</u>	+1.9	<u>-0.8</u>	<u>+0.5</u>	<u>+3.1</u>	<u>+6.0</u>	<u>+2.8</u>	<u>+2.2</u>
	-cloud-all	-1.3	-1.1	-4.3	<u>+3.9</u>	+12.2	<u>+6.1</u>	+2.6	<u>+2.8</u>	-4.0	-4.9	+2.1	+4.6	+2.5	+0.5

Table 10: Comparison of language family BLEU means by model. Performance gains/loss are relative to off-the-shelf. The overall score is **bolded** and the best score in each paradigm is underlined.

		M2M							Aya-23						
		AUS (9)	ARA (7)	ROM (6)	TUR (4)	IND (4)	CRE (8)	Mean	AUS (9)	ARB (1)	ROM (5)	TUR (4)	IND (4)	CRE (9)	Mean
Baselines	off-the-shelf	52.9	<u>72.6</u>	<u>67.9</u>	<u>47.6</u>	<u>64.2</u>	<u>47.6</u>	<u>58.8</u>	<u>55.2</u>	81.7	74.0	60.7	<u>76.5</u>	47.6	66.0
	fthrl	<u>+0.7</u>	-0.1	-0.6	-0.6	<u>+0.0</u>	-0.3	-0.1	-2.0	+0.2	+0.9	<u>+4.3</u>	-1.7	<u>+6.2</u>	+1.3
	randaug-shell	+0.1	-0.5	-0.2	-0.3	-0.0	-0.3	-0.2	-1.6	<u>+0.3</u>	<u>+1.5</u>	+4.3	-1.8	+5.7	<u>+1.4</u>
M→D	-shell	<u>+4.3</u>	<u>+1.3</u>	<u>+8.3</u>	<u>+8.0</u>	+8.0	<u>+3.6</u>	<u>+5.6</u>	+3.8	+1.0	+5.1	<u>+8.3</u>	+2.8	+7.5	+4.8
	-cloud	+2.9	+0.8	+7.0	+4.4	<u>+8.7</u>	+3.3	+4.5	<u>+4.1</u>	<u>+1.1</u>	<u>+5.2</u>	<u>+8.2</u>	<u>+3.0</u>	<u>+7.8</u>	<u>+4.9</u>
D→M	-cont	+0.3	-2.7	-8.7	+5.2	+0.8	-	+1.2	<u>+0.5</u>	-4.1	-8.6	-3.7	-1.3	-	+0.2
	-func	-0.6	<u>+0.5</u>	-1.2	+2.9	<u>+10.8</u>	-	+4.7	-1.7	<u>-0.7</u>	<u>-0.9</u>	<u>+0.1</u>	<u>+3.1</u>	-	<u>+3.6</u>
	-all	<u>+1.3</u>	-2.3	-10.1	<u>+7.6</u>	+10.0	-	+3.5	-0.3	-5.0	-10.1	-3.4	+1.7	-	+0.3
M↔D	-cloud-cont	<u>+4.7</u>	-1.9	-1.0	+8.3	+7.8	-	+5.8	<u>+5.7</u>	-1.7	-1.0	+5.9	+1.6	-	+5.8
	-cloud-func	+2.6	<u>+1.0</u>	<u>+5.9</u>	+6.5	<u>+12.8</u>	-	<u>+8.0</u>	+3.0	<u>+0.5</u>	<u>+4.0</u>	<u>+8.0</u>	<u>+4.1</u>	-	<u>+7.6</u>
	-cloud-all	+3.5	-1.8	-2.6	<u>+9.8</u>	+11.4	-	+6.3	+4.8	-2.3	-2.4	+5.6	+2.8	-	+5.4

Table 11: Comparison of language family COMET means by model. Performance gains/losses are relative to off-the-shelf. The overall score is **bolded** and best score in each paradigm is underlined.

		M2M							Aya-23						
		AUS (9)	ARA (7)	ROM (6)	TUR (4)	IND (4)	CRE (8)	Mean	AUS (9)	ARB (1)	ROM (5)	TUR (4)	IND (4)	CRE (9)	Mean
Baselines	off-the-shelf	46.9	<u>71.5</u>	<u>44.7</u>	<u>47.6</u>	64.2	<u>47.6</u>	<u>53.8</u>	<u>49.5</u>	74.0	58.4	60.7	<u>76.5</u>	47.6	61.1
	fthrl	<u>+0.7</u>	-0.1	-1.3	-0.6	<u>+0.0</u>	-0.3	-0.2	-2.0	+0.2	+2.8	<u>+4.3</u>	-1.7	<u>+6.2</u>	+1.6
	randaug-shell	+0.2	-0.5	-0.3	-0.3	-0.0	-0.3	-0.2	-1.6	<u>+0.6</u>	<u>+4.0</u>	+4.3	-1.8	+5.7	<u>+1.9</u>
M→D	-shell	<u>+5.1</u>	<u>+1.5</u>	<u>+19.4</u>	<u>+8.0</u>	+8.0	<u>+3.6</u>	<u>+7.6</u>	+5.3	+2.7	+12.7	<u>+8.3</u>	+2.8	+7.5	+6.6
	-cloud	+3.4	+1.0	+16.5	+4.4	<u>+8.7</u>	+3.3	+6.2	<u>+5.5</u>	<u>+2.9</u>	<u>+12.7</u>	<u>+8.2</u>	<u>+3.0</u>	<u>+7.8</u>	<u>+6.7</u>
D→M	-cont	+2.0	-3.1	+2.8	+5.2	+0.8	-	+2.8	+1.9	-6.5	-2.1	-3.7	-1.3	-	+0.4
	-func	+1.0	<u>+0.6</u>	+9.0	+2.9	<u>+10.8</u>	-	+6.1	-0.1	<u>-1.3</u>	<u>+4.5</u>	<u>+0.1</u>	<u>+3.1</u>	-	<u>+4.0</u>
	-all	<u>+4.4</u>	-2.6	<u>+11.5</u>	<u>+7.6</u>	+10.0	-	+7.4	<u>+2.4</u>	-7.7	+1.3	-3.4	+1.7	-	+1.6
M↔D	-cloud-cont	<u>+7.2</u>	-2.2	+15.3	+8.3	+7.8	-	+8.5	<u>+8.4</u>	-1.7	+9.2	+5.9	+1.6	-	+7.4
	-cloud-func	+4.4	<u>+1.2</u>	<u>+19.8</u>	+6.5	<u>+12.8</u>	-	<u>+10.2</u>	+5.4	<u>+1.6</u>	<u>+13.4</u>	<u>+8.0</u>	<u>+4.1</u>	-	<u>+9.2</u>
	-cloud-all	+6.6	-2.0	+17.3	<u>+9.8</u>	+11.4	-	+9.8	+7.8	-2.8	+10.0	+5.6	+2.8	-	+7.4

Table 12: Comparison of language family COMET means by model where baseline < 25. Performance gains/loss are relative to off-the-shelf. The overall best score for each language is **bolded** and the best score in each paradigm is underlined. Numbers in parenthesis represent the number of languages within the model and language family whose baseline off-the-shelf BLEU score was less than 25.

Language	Best Baseline	Best M→D	Best D→M	Best M↔D
Austronesian Language Family				
ind_Latn (Indonesian)	38.9	-0.5	-	-
jav_Latn (Javanese)	23.3	+0.7	-0.2	+0.7
sun_Latn (Sundanese)	23.5	+0.8	+0.4	+0.9
smo_Latn (Samoan)	2.4	-0.2	+2.0	+2.7
mri_Latn (Maori)	2.5	+0.4	+2.1	+2.6
ceb_Latn (Cebuano)	22.6	+1.1	-3.1	-2.3
zsm_Latn (Standard Malay)	39.0	+0.0	-1.8	-2.5
tgl_Latn (Tagalog)	28.3	+0.7	-6.8	-5.6
ilo_Latn (Ilocano)	14.3	+1.5	-0.4	+1.4
fij_Latn (Fijian)	2.5	+0.5	+1.7	+2.4
plt_Latn (Plateau Malagasy)	2.5	+5.1	+3.8	+5.2
pag_Latn (Pangasinan)	6.7	+4.0	+5.2	+7.4
Average	15.2	+1.3	+0.3	+1.2
Baseline < 25 Average	11.1	+1.5	+1.3	+2.3
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	29.9	-0.1	-	-
acm_Arab (Mesopotamian Arabic)	22.9	+0.4	+0.5	+1.0
acq_Arab (Ta'izzi-Adeni Arabic)	25.2	+0.3	+0.2	+0.2
aeb_Arab (Tunisian Arabic)	18.0	+0.8	+1.0	+1.7
ajp_Arab (South Levantine Arabic)	24.6	+0.5	-0.9	-1.2
apc_Arab (North Levantine Arabic)	21.7	-0.1	+1.0	+0.8
ars_Arab (Najdi Arabic)	28.7	+0.2	-0.1	-0.4
ary_Arab (Moroccan Arabic)	13.1	+1.9	+0.8	+2.4
arz_Arab (Egyptian Arabic)	18.3	+1.0	+0.6	+1.6
Average	21.6	+0.6	+0.4	+0.8
Baseline < 25 Average	20.5	+0.7	+0.5	+0.9
Romance Language Family				
ita_Latn (Italian)	31.7	+0.0	-	-
spa_Latn (Spanish)	27.1	+0.2	-7.5	-4.2
fra_Latn (French)	41.8	+0.2	-13.2	-8.7
por_Latn (Portuguese)	46.0	-0.3	-15.1	-8.8
ron_Latn (Romanian)	41.5	-0.3	-16.3	-11.8
glg_Latn (Galician)	38.2	+0.0	-11.8	-6.5
cat_Latn (Catalan)	42.9	+0.0	-13.0	-8.5
oci_Latn (Occitan)	44.7	-0.1	-5.2	-4.7
ast_Latn (Asturian)	34.7	+0.7	-8.0	-6.1
lmo_Latn (Lombard)	10.1	+11.6	+6.2	+12.4
vec_Latn (Venetian)	17.7	+10.0	+7.5	+10.9
scn_Latn (Sicilian)	9.4	+8.9	+7.9	+11.0
srd_Latn (Sardinian)	5.8	+11.5	+13.3	+17.5
fur_Latn (Friulian)	8.6	+10.9	+10.8	+14.3
lij_Latn (Ligurian)	10.9	+11.1	+11.0	+14.8
Average	27.1	+4.6	-2.4	+1.5
Baseline < 25 Average	10.4	+10.7	+9.4	+13.5
Turkic Language Family				
tur_Latn (Turkish)	33.3	-0.6	-	-
uzn_Latn (Northern Uzbek)	1.9	+0.4	+4.7	+5.4
tuk_Latn (Turkmen)	2.4	+1.4	+4.6	+5.7
azj_Latn (North Azerbaijani)	7.6	+4.0	-0.1	+2.2
crh_Latn (Crimean Tatar)	8.9	+3.2	+2.1	+3.0
Average	5.2	+2.3	+2.8	+4.0
Baseline < 25 Average	5.2	+2.3	+2.8	+4.0
Indic Language Family				
hin_Deva (Hindi)	33.4	-0.4	-	-
hne_Deva (Chattisgarhi)	15.5	+5.0	+13.6	+14.0
bho_Deva (Bhojpuri)	10.4	+3.3	+6.1	+7.8
mag_Deva (Magahi)	18.0	+4.5	+12.0	+13.2
mai_Deva (Maithili)	8.1	+7.1	+16.2	+16.4
Average	13.0	+5.0	+12.0	+12.8
Baseline < 25 Average	13.0	+5.0	+12.0	+12.8
Creole Language Family				
hat (Haitian)	40.6	-0.3	-	-
gcf (Guadeloupean)	5.7	+1.7	-	-
mart1259 (Martinican)	4.7	-0.4	-	-
acf (Saint Lucian Patois)	4.4	+1.3	+7.1	+8.4
gcr (French Guianese Creole)	2.5	+2.8	-	-
lou (Louisiana Creole)	5.7	+3.1	-	-
mfc (Mauritian)	5.2	+1.5	+3.1	+3.5
ref (Réunion Creole)	16.3	+1.1	-	-
crs (Seychellois)	11.4	+2.5	+3.9	+4.3
Average	7.0	+1.7	+4.7	+5.4
Baseline < 25 Average	7.0	+1.7	+4.7	+5.4
General Average	17.3	+2.6	+2.2	+4.0
General Baseline < 25 Average	11.5	+3.3	+5.1	+6.5

Table 13: M2M BLEU for best performing M→D, D→M, M↔D by language. Performance gains/loss are relative to Best Baseline.

Language	Best Baseline	Best M→D	Best D→M	Best M↔D
Austronesian Language Family				
ind_Latn (Indonesian)	41.5	-1.0	-	-
jav_Latn (Javanese)	15.2	+4.9	+9.6	+11.2
sun_Latn (Sundanese)	12.2	+5.4	+10.2	+12.1
smo_Latn (Samoan)	3.2	+1.1	+1.1	+2.2
mri_Latn (Maori)	4.8	+0.1	+0.4	+2.0
ceb_Latn (Cebuano)	14.0	-0.6	-0.0	-0.3
zsm_Latn (Standard Malay)	36.6	+0.4	-1.3	-0.6
tgl_Latn (Tagalog)	25.6	-6.4	-4.7	-8.8
ilo_Latn (Ilocano)	7.5	+1.5	+2.4	+4.0
fij_Latn (Fijian)	3.3	+0.6	+0.7	+2.2
plt_Latn (Plateau Malagasy)	3.3	+1.0	+3.9	+4.6
pag_Latn (Pangasinan)	8.2	+3.5	+4.2	+5.7
Average	12.2	+1.0	+2.4	+3.1
Baseline < 25 Average	8.0	+1.9	+3.6	+4.8
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	38.2	-0.8	-	-
acm_Arab (Mesopotamian Arabic)	32.4	+0.0	-1.0	-0.1
acq_Arab (Ta'izzi-Adeni Arabic)	34.1	-0.1	-0.2	-0.2
aeb_Arab (Tunisian Arabic)	26.7	+0.5	-0.1	+0.4
ajp_Arab (South Levantine Arabic)	37.4	-2.0	-4.0	-4.7
apc_Arab (North Levantine Arabic)	33.2	-0.2	-1.2	-1.8
ars_Arab (Najdi Arabic)	37.6	-0.7	-0.5	-0.5
ary_Arab (Moroccan Arabic)	23.5	+0.4	-1.8	-0.8
arz_Arab (Egyptian Arabic)	29.2	+0.1	-1.7	-1.0
Average	31.8	-0.3	-1.3	-1.1
Baseline < 25 Average	23.5	+0.4	-1.8	-0.8
Romance Language Family				
ita_Latn (Italian)	33.8	-0.7	-	-
spa_Latn (Spanish)	31.0	-0.7	-6.6	-3.6
fra_Latn (French)	44.0	-3.3	-9.8	-9.7
por_Latn (Portuguese)	46.9	-3.2	-7.7	-9.1
ron_Latn (Romanian)	42.0	-3.8	-10.2	-11.7
glg_Latn (Galician)	35.5	-0.2	-5.7	-4.5
cat_Latn (Catalan)	36.5	-1.5	-5.7	-5.5
oci_Latn (Occitan)	35.1	+0.9	-2.2	-0.8
ast_Latn (Asturian)	31.8	+0.1	-5.1	-4.6
lmo_Latn (Lombard)	19.8	+5.2	+0.9	+6.4
vec_Latn (Venetian)	30.2	+3.2	-0.8	+2.5
scn_Latn (Sicilian)	17.7	+6.4	+2.7	+6.8
srd_Latn (Sardinian)	16.6	+6.7	+5.9	+9.8
fur_Latn (Friulian)	20.3	+5.1	+1.9	+6.5
lij_Latn (Ligurian)	20.7	+6.7	+2.8	+9.9
Average	30.6	+1.5	-2.8	-0.5
Baseline < 25 Average	19.0	+6.0	+2.8	+7.9
Turkic Language Family				
tur_Latn (Turkish)	31.7	-1.7	-	-
uzn_Latn (Northern Uzbek)	4.8	+2.3	+1.1	+4.2
tuk_Latn (Turkmen)	8.4	+1.5	-0.3	+1.7
azj_Latn (North Azerbaijani)	10.5	+1.6	-1.0	+0.6
crh_Latn (Crimean Tatar)	13.3	+1.9	-1.4	+1.9
Average	9.3	+1.8	-0.4	+2.1
Baseline < 25 Average	9.3	+1.8	-0.4	+2.1
Indic Language Family				
hin_Deva (Hindi)	34.8	-1.7	-	-
hne_Deva (Chattisgarhi)	23.3	+2.2	+6.0	+6.3
bho_Deva (Bhojpuri)	16.2	+2.4	+3.2	+4.2
mag_Deva (Magahi)	24.9	+2.5	+5.5	+5.8
mai_Deva (Maithili)	17.2	+4.2	+7.6	+7.6
Average	20.4	+2.8	+5.6	+6.0
Baseline < 25 Average	20.4	+2.8	+5.6	+6.0
Creole Language Family				
hat (Haitian)	29.7	-1.6	-	-
gcf (Guadeloupean)	7.0	+3.8	-	-
mart1259 (Martinican)	5.8	+2.7	-	-
acf (Saint Lucian Patois)	5.6	+0.5	-2.6	+3.9
gcr (French Guianese Creole)	4.7	+3.1	-	-
lou (Louisiana Creole)	12.1	-0.1	-	-
mfc (Mauritian)	8.5	-0.6	-5.3	-0.4
ref (Réunion Creole)	20.9	-2.7	-	-
crs (Seychellois)	15.0	+0.2	-8.7	+0.3
Average	10.0	+0.9	-5.8	+1.0
Baseline < 25 Average	10.0	+0.9	-5.8	+1.0
General Average	20.7	+1.2	+0.8	+2.4
General Baseline < 25 Average	12.5	+2.4	+2.3	+5.0

Table 14: Aya-23 BLEU for best performing M→D, D→M, M↔D by language. Performance gains/loss are relative to Best Baseline.

Language	Best Baseline	Best M→D	Best D→M	Best M↔D
Austronesian Language Family				
ind_Latn (Indonesian)	87.3	+0.0	-	-
jav_Latn (Javanese)	70.9	+1.2	-1.5	-0.4
sun_Latn (Sundanese)	72.3	+1.3	-3.4	-2.4
smo_Latn (Samoan)	27.6	+2.8	+13.3	+12.2
mrj_Latn (Maori)	25.2	+6.5	+17.6	+17.2
ceb_Latn (Cebuano)	64.7	+1.4	-10.3	-2.5
zsm_Latn (Standard Malay)	86.5	-0.0	-8.0	-6.3
tgl_Latn (Tagalog)	73.8	+1.0	-18.6	-6.6
ilo_Latn (Ilocano)	57.0	+2.0	-4.2	+1.6
fij_Latn (Fijian)	28.0	+7.5	+14.6	+17.0
plt_Latn (Plateau Malagasy)	47.0	+3.0	-3.5	+4.1
pag_Latn (Pangasinan)	36.3	+13.2	+10.5	+11.0
Average	53.6	+3.6	+0.6	+4.1
Baseline < 25 Average	47.7	+4.3	+3.7	+6.4
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	80.6	+0.0	-	-
acm_Arab (Mesopotamian Arabic)	76.0	+0.7	+0.3	+0.6
acq_Arab (Ta'izzi-Adeni Arabic)	77.6	+0.3	+0.0	+0.0
acb_Arab (Tunisian Arabic)	68.8	+1.7	+0.8	+1.8
ajp_Arab (South Levantine Arabic)	73.2	+1.5	-0.4	+0.2
apc_Arab (North Levantine Arabic)	72.8	+1.2	+0.6	+1.1
ars_Arab (Najdi Arabic)	80.0	+0.0	+0.0	-0.3
ary_Arab (Moroccan Arabic)	60.8	+3.7	+2.4	+3.7
arz_Arab (Egyptian Arabic)	71.4	+1.4	+0.6	+1.1
Average	72.6	+1.3	+0.5	+1.0
Baseline < 25 Average	71.5	+1.5	+0.6	+1.2
Romance Language Family				
ita_Latn (Italian)	86.3	+0.0	-	-
spa_Latn (Spanish)	85.2	-0.2	-8.1	-4.4
fra_Latn (French)	87.8	-0.2	-10.6	-5.2
por_Latn (Portuguese)	88.0	-0.1	-10.2	-4.1
ron_Latn (Romanian)	88.2	-0.1	-12.9	-6.5
glg_Latn (Galician)	86.6	-0.0	-9.2	-3.6
cat_Latn (Catalan)	87.1	-0.1	-9.9	-5.5
oci_Latn (Occitan)	80.0	+0.2	-2.8	-2.3
ast_Latn (Asturian)	79.8	+0.8	-7.3	-3.9
lmo_Latn (Lombard)	44.5	+19.3	+6.6	+18.5
vec_Latn (Venetian)	54.0	+17.9	+8.3	+16.5
scn_Latn (Sicilian)	44.1	+18.9	+8.7	+18.8
srđ_Latn (Sardinian)	37.9	+21.4	+11.3	+23.4
fur_Latn (Friulian)	42.5	+19.9	+8.2	+20.7
lij_Latn (Ligurian)	45.0	+19.1	+11.0	+20.7
Average	67.9	+8.3	-1.2	+5.9
Baseline < 25 Average	44.7	+19.4	+9.0	+19.8
Turkic Language Family				
tur_Latn (Turkish)	86.7	-0.2	-	-
uzn_Latn (Northern Uzbek)	39.5	+5.0	+16.0	+17.0
tuk_Latn (Turkmen)	34.4	+10.1	+12.2	+15.5
azj_Latn (North Azerbaijani)	60.8	+9.4	-1.1	+1.4
crh_Latn (Crimean Tatar)	55.7	+7.7	+3.5	+5.2
Average	47.6	+8.0	+7.6	+9.8
Baseline < 25 Average	47.6	+8.0	+7.6	+9.8
Indic Language Family				
hin_Deva (Hindi)	86.0	+0.2	-	-
hne_Deva (Chattisgarhi)	66.2	+6.9	+10.2	+11.9
bho_Deva (Bhojpur)	63.6	+7.2	+7.3	+9.8
mag_Deva (Magahi)	69.2	+6.7	+9.5	+11.1
mai_Deva (Maithili)	58.1	+13.8	+16.2	+18.2
Average	64.3	+8.7	+10.8	+12.8
Baseline < 25 Average	64.3	+8.7	+10.8	+12.8
French-Creole Language Family				
hat (Haitian)	73.6	+0.7	-	-
gcf (Guadeloupean)	52.4	+1.6	-	-
mart1259 (Martinican)	50.7	+1.0	-	-
acf (Saint Lucian Patois)	43.4	+3.3	-	-
ger (French Guianese Creole)	43.9	+3.5	-	-
lou (Louisiana Creole)	43.4	+6.0	-	-
mfe (Mauritian)	45.0	+4.8	-	-
ref (Réunion Creole)	47.1	+2.7	-	-
crs (Seychellois)	54.6	+6.2	-	-
Average	47.6	+3.6	-	-
Baseline < 25 Average	47.6	+3.6	-	-
General Average	60.0	+5.5	+3.3	+7.2
General Baseline < 25 Average	53.1	+7.0	+6.9	+9.6

Table 15: m2mL’s COMET for best performing M→D, D→M, M↔D by language. Performance gains/loss are relative to Best Baseline.

Language	Best Baseline	Best M→D	Best D→M	Best M↔D
Austronesian Language Family				
ind_Latn (Indonesian)	88.5	-0.1	-	-
jav_Latn (Javanese)	67.5	+6.2	+1.1	+6.6
sun_Latn (Sundanese)	63.7	+9.0	+1.7	+8.9
smo_Latn (Samoan)	35.4	+6.2	+3.9	+11.4
mrj_Latn (Maori)	40.5	+2.2	+3.6	+9.0
ceb_Latn (Cebuano)	61.1	+0.3	-2.2	+0.3
zsm_Latn (Standard Malay)	85.7	+0.2	-4.3	-3.2
tgl_Latn (Tagalog)	76.4	-5.4	-7.3	-9.4
ilo_Latn (Ilocano)	48.2	+7.0	+2.6	+9.1
fij_Latn (Fijian)	40.1	+3.7	+3.5	+9.7
plt_Latn (Plateau Malagasy)	46.1	+1.6	+2.2	+8.0
pag_Latn (Pangasinan)	43.1	+13.6	+0.4	+12.6
Average	55.2	+4.1	+0.5	+5.7
Baseline < 25 Average	49.5	+5.5	+1.9	+8.4
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	86.5	+0.0	-	-
acm_Arab (Mesopotamian Arabic)	83.6	+0.7	-0.6	+0.5
acq_Arab (Ta'izzi-Adeni Arabic)	84.7	+0.4	-0.5	+0.1
acb_Arab (Tunisian Arabic)	78.8	+1.4	-0.4	+1.1
ajp_Arab (South Levantine Arabic)	83.4	+0.5	-1.5	-0.6
apc_Arab (North Levantine Arabic)	82.7	+0.8	-1.0	+0.0
ars_Arab (Najdi Arabic)	86.2	-0.1	-0.7	-0.1
ary_Arab (Moroccan Arabic)	74.6	+2.3	-1.8	+1.1
arz_Arab (Egyptian Arabic)	82.5	+0.3	-1.5	-0.4
Average	82.0	+0.8	-1.0	+0.2
Baseline < 25 Average	74.6	+2.3	-1.8	+1.1
Romance Language Family				
ita_Latn (Italian)	87.2	-0.2	-	-
spa_Latn (Spanish)	86.0	-0.2	-4.0	-1.9
fra_Latn (French)	87.8	-0.3	-5.4	-3.4
por_Latn (Portuguese)	88.2	-0.4	-3.8	-2.7
ron_Latn (Romanian)	87.9	-0.6	-3.5	-4.4
glg_Latn (Galician)	85.0	+0.4	-4.7	-1.6
cat_Latn (Catalan)	83.2	+0.7	-4.3	-1.7
oci_Latn (Occitan)	75.1	+3.3	-2.1	+2.0
ast_Latn (Asturian)	77.1	+2.4	-3.8	-0.7
lmo_Latn (Lombard)	63.4	+7.5	-1.4	+8.0
vec_Latn (Venetian)	74.7	+3.6	-2.8	+2.7
scn_Latn (Sicilian)	62.5	+8.9	-0.3	+8.8
srđ_Latn (Sardinian)	58.3	+10.8	+4.1	+12.2
fur_Latn (Friulian)	64.7	+7.2	-0.4	+7.6
lij_Latn (Ligurian)	63.1	+8.7	+0.3	+10.4
Average	75.5	+3.7	-2.4	+2.5
Baseline < 25 Average	62.4	+8.6	+0.5	+9.4
Turkic Language Family				
tur_Latn (Turkish)	87.3	-0.1	-	-
uzn_Latn (Northern Uzbek)	58.0	+6.2	-2.1	+6.2
tuk_Latn (Turkmen)	59.9	+3.9	-6.6	+3.8
azj_Latn (North Azerbaijani)	73.1	+2.7	-4.1	+1.4
crh_Latn (Crimean Tatar)	69.1	+3.4	-3.9	+3.3
Average	65.0	+4.0	-4.2	+3.7
Baseline < 25 Average	65.0	+4.0	-4.2	+3.7
Indic Language Family				
hin_Deva (Hindi)	88.3	-1.0	-	-
hne_Deva (Chattisgarhi)	77.2	+2.8	+3.2	+4.3
bho_Deva (Bhojpur)	75.2	+2.1	+1.6	+3.2
mag_Deva (Magahi)	78.9	+2.9	+3.0	+3.7
mai_Deva (Maithili)	74.6	+4.2	+4.6	+5.2
Average	76.5	+3.0	+3.1	+4.1
Baseline < 25 Average	76.5	+3.0	+3.1	+4.1
French-Creole Language Family				
hat (Haitian)	72.2	-0.6	-	-
gcf (Guadeloupean)	54.9	+2.5	-	-
mart1259 (Martinican)	54.0	+3.0	-	-
acf (Saint Lucian Patois)	48.7	+0.9	-	-
ger (French Guianese Creole)	52.1	+0.7	-	-
lou (Louisiana Creole)	53.6	+0.7	-	-
mfe (Mauritian)	54.7	+0.7	-	-
ref (Réunion Creole)	51.1	+3.4	-	-
crs (Seychellois)	61.5	+1.4	-	-
Average	53.8	+1.7	-	-
Baseline < 25 Average	53.8	+1.7	-	-
General Average	63.5	+4.5	+2.6	+6.6
General Baseline < 25 Average	55.3	+6.2	+4.6	+9.0

Table 16: aya-23-8b’s COMET for best performing M→D, D→M, M↔D by language. Performance gains/loss are relative to Best Baseline.

Language	off-the-shelf	fthrln	ftrandaug	ftrandaug-cloud	-shell	-cloud
Austronesian Language Family						
ind_Latn (Indonesian)	38.9	-0.1	-0.4	-0.4	<u>-0.5</u>	-0.7
jav_Latn (Javanese)	22.9	<u>+0.4</u>	+0.0	-0.1	<u>+1.1</u>	+1.0
sun_Latn (Sundanese)	22.9	<u>+0.6</u>	+0.1	+0.1	<u>+1.4</u>	+0.8
sno_Latn (Samoan)	2.3	+0.1	+0.1	-0.4	-0.3	-0.1
mri_Latn (Maori)	2.3	<u>+0.2</u>	<u>+0.2</u>	+0.2	<u>+0.6</u>	+0.3
ceb_Latn (Cebuano)	22.2	<u>+0.4</u>	+0.2	+0.3	<u>+1.5</u>	+0.8
zsm_Latn (Standard Malay)	38.7	<u>+0.3</u>	-0.3	-0.2	<u>+0.3</u>	-0.2
tgl_Latn (Tagalog)	27.8	<u>+0.5</u>	+0.1	+0.1	<u>+1.2</u>	+0.7
ilo_Latn (Ilocano)	13.4	+0.6	<u>+0.9</u>	+0.5	<u>+2.4</u>	+2.0
fij_Latn (Fijian)	2.3	<u>+0.2</u>	<u>+0.2</u>	+0.2	+0.6	<u>+0.7</u>
plt_Latn (Plateau Malagasy)	2.5	+0.0	+0.0	+0.0	<u>+5.1</u>	+3.7
pag_Latn (Pangasinan)	6.4	+0.2	<u>+0.3</u>	<u>+0.3</u>	<u>+4.3</u>	+2.9
Average	14.9	<u>+0.3</u>	+0.2	+0.1	<u>+1.7</u>	+1.1
Baseline < 25 Average	10.8	<u>+0.3</u>	+0.2	+0.2	<u>+1.9</u>	+1.3
Arabic Language Family						
arb_Arab (Modern Standard Arabic)	29.6	<u>+0.3</u>	-0.1	-0.2	<u>+0.2</u>	-0.3
acm_Arab (Mesopotamian Arabic)	22.4	<u>+0.5</u>	<u>+0.5</u>	+0.5	<u>+0.9</u>	+0.8
acq_Arab (Ta'izzi-Adeni Arabic)	24.7	<u>+0.5</u>	+0.2	+0.1	<u>+0.8</u>	+0.4
aeb_Arab (Tunisian Arabic)	17.7	<u>+0.3</u>	+0.2	+0.3	+0.8	<u>+1.1</u>
ajp_Arab (South Levantine Arabic)	24.5	<u>+0.1</u>	-0.2	-0.1	<u>+0.6</u>	<u>+0.6</u>
apc_Arab (North Levantine Arabic)	21.7	-0.5	-0.4	-0.5	-0.1	-0.2
ars_Arab (Najdi Arabic)	28.5	<u>+0.2</u>	-0.1	-0.1	<u>+0.4</u>	-0.5
ary_Arab (Moroccan Arabic)	11.9	<u>+1.2</u>	+1.0	+1.1	<u>+3.1</u>	+2.8
arz_Arab (Egyptian Arabic)	18.2	<u>+0.1</u>	+0.0	+0.0	<u>+1.1</u>	+0.9
Average	21.2	<u>+0.3</u>	+0.2	+0.2	<u>+1.0</u>	+0.7
Baseline < 25 Average	20.2	<u>+0.3</u>	+0.2	+0.2	<u>+1.0</u>	+0.9
Romance Language Family						
ita_Latn (Italian)	30.8	+0.9	+0.4	+0.4	<u>+0.9</u>	+0.5
spa_Latn (Spanish)	26.8	<u>+0.3</u>	-0.2	-0.2	<u>+0.5</u>	+0.1
fra_Latn (French)	41.6	<u>+0.2</u>	-0.4	-0.4	<u>+0.4</u>	-0.3
por_Latn (Portuguese)	45.8	+0.2	-0.2	-0.2	-0.1	-0.6
ron_Latn (Romanian)	40.8	<u>+0.7</u>	+0.2	+0.2	<u>+0.4</u>	+0.1
glg_Latn (Galician)	37.5	<u>+0.7</u>	+0.2	+0.1	<u>+0.7</u>	+0.5
cat_Latn (Catalan)	42.7	<u>+0.2</u>	+0.1	+0.0	<u>+0.2</u>	-0.1
oci_Latn (Occitan)	44.6	+0.0	+0.1	+0.1	+0.0	+0.0
ast_Latn (Asturian)	34.3	+0.4	-0.2	-0.3	<u>+1.1</u>	+0.1
lmo_Latn (Lombard)	9.6	+0.2	<u>+0.5</u>	+0.5	<u>+12.1</u>	+9.9
vec_Latn (Venetian)	17.2	-0.1	+0.4	+0.5	<u>+10.5</u>	+8.4
scn_Latn (Sicilian)	7.3	<u>+2.1</u>	+1.2	+1.1	<u>+11.0</u>	+8.7
srd_Latn (Sardinian)	5.3	+0.3	<u>+0.5</u>	+0.5	<u>+12.0</u>	+9.8
fur_Latn (Friulian)	8.6	-0.4	-0.1	+0.0	<u>+10.9</u>	+8.2
lij_Latn (Ligurian)	9.8	<u>+1.1</u>	+0.8	+0.7	<u>+12.2</u>	+9.7
Average	26.6	<u>+0.4</u>	+0.2	+0.2	<u>+5.1</u>	+3.9
Baseline < 25 Average	9.6	+0.5	<u>+0.6</u>	<u>+0.6</u>	<u>+11.5</u>	+9.1
Turkic Language Family						
tur_Latn (Turkish)	33.3	-0.4	-0.5	-0.4	<u>-0.6</u>	-0.8
uzn_Latn (Northern Uzbek)	1.6	+0.2	+0.2	<u>+0.3</u>	<u>+0.7</u>	<u>+0.7</u>
tuk_Latn (Turkmen)	<u>2.4</u>	-0.2	-0.2	-0.1	<u>+1.4</u>	+0.7
azj_Latn (North Azerbaijani)	7.4	-0.3	+0.2	<u>+0.2</u>	<u>+4.2</u>	+2.7
crh_Latn (Crimean Tatar)	8.5	-0.2	<u>+0.4</u>	<u>+0.4</u>	<u>+3.6</u>	+2.3
Average	5.0	-0.1	<u>+0.2</u>	<u>+0.2</u>	<u>+2.5</u>	+1.6
Baseline < 25 Average	5.0	-0.1	<u>+0.2</u>	<u>+0.2</u>	<u>+2.5</u>	+1.6
Indic Language Family						
hin_Deva (Hindi)	33.4	-0.5	-0.6	-0.6	<u>-0.4</u>	<u>-0.4</u>
hne_Deva (Chattisgarhi)	15.2	+0.2	+0.3	+0.3	+4.8	<u>+5.3</u>
bho_Deva (Bhojpuri)	10.1	<u>+0.3</u>	+0.2	+0.3	+3.3	<u>+3.6</u>
mag_Deva (Magahi)	17.2	+0.7	+0.7	+0.8	+4.8	<u>+5.3</u>
mai_Deva (Maithili)	7.7	<u>+0.4</u>	+0.2	+0.2	+6.9	<u>+7.5</u>
Average	12.6	<u>+0.4</u>	<u>+0.4</u>	<u>+0.4</u>	+5.0	<u>+5.4</u>
Baseline < 25 Average	12.6	<u>+0.4</u>	<u>+0.4</u>	<u>+0.4</u>	+5.0	<u>+5.4</u>
Creole Language Family						
hat (Haitian)	39.7	+0.9	+0.6	+0.7	+0.5	<u>+0.6</u>
gcf (Guadeloupean)	4.4	+1.2	+1.2	<u>+1.3</u>	+2.8	<u>+3.0</u>
mart1259 (Martinican)	0.0	<u>+4.7</u>	<u>+4.7</u>	+4.1	+3.6	<u>+4.3</u>
acf (Saint Lucian Patois)	3.7	+0.7	+0.4	+0.7	<u>+2.0</u>	+1.7
gcr (French Guianese Creole)	2.5	-0.8	-0.8	-0.7	<u>+2.8</u>	+2.5
lou (Louisiana Creole)	5.3	+0.2	<u>+0.4</u>	+0.4	<u>+3.5</u>	+3.4
mfe (Mauritian)	5.1	+0.0	+0.0	+0.1	<u>+1.6</u>	+1.4
ref (Réunion Creole)	14.8	<u>+1.5</u>	+1.4	<u>+1.5</u>	<u>+2.6</u>	<u>+2.6</u>
crs (Seychellois)	11.2	-0.4	-0.1	<u>+0.2</u>	<u>+2.7</u>	+2.6
Average	5.9	+0.9	+0.9	<u>+1.0</u>	<u>+2.7</u>	<u>+2.7</u>
Baseline < 25 Average	5.9	+0.9	+0.9	<u>+1.0</u>	<u>+2.7</u>	<u>+2.7</u>
General Average	16.8	<u>+0.4</u>	+0.3	+0.3	+3.0	+2.5
General Baseline < 25 Average	10.9	<u>+0.4</u>	<u>+0.4</u>	<u>+0.4</u>	<u>+3.8</u>	+3.2

Table 17: M→D BLEU scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each M→D approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The overall best score for each language is **bolded** and the best score in each paradigm is underlined.

Language	off-the-shelf	fthrIn	ftrandaug	ftrandaug-cloud	-shell	-cloud
Austronesian Language Family						
ind_Latn (Indonesian)	41.5	-0.4	-0.3	-0.4	<u>-1.0</u>	-1.1
jav_Latn (Javanese)	14.2	<u>+1.0</u>	+0.4	+0.8	+5.9	+5.5
sun_Latn (Sundanese)	11.0	+0.8	<u>+1.2</u>	+0.9	+6.6	+6.6
sno_Latn (Samoan)	<u>3.2</u>	-0.1	-0.4	-0.3	+1.1	+1.1
mri_Latn (Maori)	4.8	-1.5	-1.2	-1.4	+0.0	+0.1
ceb_Latn (Cebuano)	14.0	-2.3	-2.1	-2.2	-1.1	<u>-0.6</u>
zsm_Latn (Standard Malay)	36.5	<u>+0.1</u>	-0.1	<u>+0.1</u>	+0.5	+0.2
tgl_Latn (Tagalog)	25.6	-6.5	-5.6	-6.2	-7.4	<u>-6.4</u>
ilo_Latn (Ilocano)	7.5	-1.9	-1.6	-1.8	+1.1	+1.5
fij_Latn (Fijian)	3.3	-0.2	-0.1	-0.4	+0.6	+0.5
plt_Latn (Plateau Malagasy)	3.3	-1.0	-1.0	-0.9	+0.6	+1.0
pag_Latn (Pangasinan)	7.3	+0.7	<u>+0.9</u>	+0.8	+4.4	+4.2
Average	<u>11.9</u>	-1.0	-0.9	-1.0	+1.1	+1.2
Baseline < 25 Average	<u>7.6</u>	-0.6	-0.4	-0.5	+2.1	+2.2
Arabic Language Family						
arb_Arab (Modern Standard Arabic)	38.1	+0.0	+0.1	+0.1	<u>-0.7</u>	<u>-0.7</u>
acm_Arab (Mesopotamian Arabic)	31.6	+0.3	+0.3	+0.8	+0.5	+0.8
acq_Arab (Ta'izzi-Adeni Arabic)	33.7	+0.2	+0.3	+0.4	+0.3	+0.3
aeb_Arab (Tunisian Arabic)	<u>26.7</u>	-0.6	-0.5	-1.0	+0.2	+0.5
ajp_Arab (South Levantine Arabic)	37.4	-2.2	-1.7	-2.1	-2.2	<u>-2.0</u>
apc_Arab (North Levantine Arabic)	<u>33.2</u>	-1.5	-1.0	-1.3	-0.3	<u>-0.2</u>
ars_Arab (Najdi Arabic)	37.1	+0.0	+0.5	+0.2	-0.4	<u>-0.2</u>
ary_Arab (Moroccan Arabic)	<u>23.5</u>	-1.8	-2.3	-1.8	+0.4	+0.1
arz_Arab (Egyptian Arabic)	28.9	<u>+0.3</u>	<u>+0.3</u>	+0.3	+0.0	+0.4
Average	31.5	-0.7	-0.5	-0.6	-0.2	<u>-0.0</u>
Baseline < 25 Average	<u>23.5</u>	-1.8	-2.3	-1.8	+0.4	+0.1
Romance Language Family						
ita_Latn (Italian)	32.6	+1.2	+1.0	+0.7	<u>+0.5</u>	<u>+0.5</u>
spa_Latn (Spanish)	29.0	+1.7	+1.9	+2.0	+1.1	+1.3
fra_Latn (French)	44.0	-2.1	-2.4	-2.1	-3.5	<u>-3.3</u>
por_Latn (Portuguese)	46.9	-1.9	-1.7	-1.6	-3.7	<u>-3.2</u>
ron_Latn (Romanian)	42.0	-2.0	-1.9	-1.8	-4.8	<u>-3.8</u>
glg_Latn (Galician)	35.5	-0.4	-0.2	-0.6	-0.5	<u>-0.2</u>
cat_Latn (Catalan)	36.5	-2.6	-2.3	-2.3	-2.1	<u>-1.5</u>
oci_Latn (Occitan)	<u>35.1</u>	-2.6	-2.6	-2.1	+0.5	+0.9
ast_Latn (Asturian)	<u>31.8</u>	-1.5	-1.5	-2.0	-0.3	+0.1
lmo_Latn (Lombard)	17.4	+1.4	<u>+2.4</u>	+1.7	+7.6	+7.4
vec_Latn (Venetian)	26.7	+2.7	<u>+3.5</u>	+2.9	+6.6	+6.7
scn_Latn (Sicilian)	16.7	+0.6	<u>+1.0</u>	+0.3	+7.4	+7.1
srd_Latn (Sardinian)	16.4	-0.5	<u>+0.2</u>	-0.2	+6.9	+6.8
fur_Latn (Friulian)	17.7	+2.4	<u>+2.6</u>	<u>+2.6</u>	+7.7	+7.6
lij_Latn (Ligurian)	17.4	+2.9	<u>+3.2</u>	<u>+3.3</u>	+10.0	+9.6
Average	29.5	-0.1	<u>+0.2</u>	+0.0	+2.4	+2.5
Baseline < 25 Average	17.1	+1.4	<u>+1.9</u>	+1.5	<u>+7.9</u>	+7.7
Turkic Language Family						
tur_Latn (Turkish)	31.7	-0.8	-1.4	-1.4	-2.5	<u>-1.7</u>
uzn_Latn (Northern Uzbek)	3.8	+0.6	+0.8	<u>+1.0</u>	+3.2	+3.3
tuk_Latn (Turkmen)	6.1	+2.1	<u>+2.3</u>	+2.0	+3.5	+3.8
azj_Latn (North Azerbaijani)	9.8	<u>+0.7</u>	+0.4	+0.6	+2.0	+2.3
crh_Latn (Crimean Tatar)	11.9	+1.3	+1.0	<u>+1.4</u>	+2.9	+3.3
Average	7.9	+1.2	+1.1	<u>+1.3</u>	+2.9	+3.2
Baseline < 25 Average	7.9	+1.2	+1.1	<u>+1.3</u>	+2.9	+3.2
Indic Language Family						
hin_Deva (Hindi)	34.8	-1.5	-1.5	-1.5	-2.5	<u>-1.7</u>
hne_Deva (Chattisgarhi)	23.3	-2.2	-2.2	-2.9	+2.2	+2.0
bho_Deva (Bhojpuri)	<u>16.2</u>	-1.2	-1.3	-1.2	+2.4	+2.0
mag_Deva (Magahi)	24.9	-1.7	-1.9	-1.7	+2.5	+2.5
mai_Deva (Maithili)	<u>17.2</u>	-2.2	-1.9	-1.9	+4.1	+4.2
Average	<u>20.4</u>	-1.8	-1.8	-1.9	+2.8	+2.7
Baseline < 25 Average	<u>20.4</u>	-1.8	-1.8	-1.9	+2.8	+2.7
Creole Language Family						
hat (Haitian)	8.9	+20.5	+20.8	+20.6	+18.7	+19.2
gcf (Guadeloupean)	6.4	+0.3	+0.5	<u>+0.6</u>	+3.0	+4.4
mart1259 (Martinican)	4.7	+0.9	+1.0	<u>+1.1</u>	+2.8	+3.8
acf (Saint Lucian Patois)	2.6	+2.9	+2.8	<u>+3.0</u>	+3.5	+3.5
gcr (French Guianese Creole)	4.7	-0.4	-0.4	-4.7	+1.6	+3.1
lou (Louisiana Creole)	9.6	+1.8	+1.6	+2.5	+2.4	+2.2
mfe (Mauritian)	3.4	+4.7	+4.6	+5.1	+4.4	+4.5
ref (Réunion Creole)	20.9	-0.5	-3.0	-3.0	-2.7	-2.9
crs (Seychellois)	5.7	+8.8	<u>+9.3</u>	+8.4	+9.3	+9.5
Average	7.3	<u>+2.3</u>	+2.1	+1.6	+3.0	+3.5
Baseline < 25 Average	7.3	<u>+2.3</u>	+2.1	+1.6	+3.0	+3.5
General Average	19.7	-0.1	+0.0	-0.1	+1.9	+2.0
General Baseline < 25 Average	11.3	<u>+0.5</u>	<u>+0.5</u>	+0.4	+3.4	+3.5

Table 18: M→D BLEU scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each M→D approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The overall best score for each language is **bolded** and the best score in each paradigm is underlined.

Language	off-the-shelf	fthr1	randaug-shell	randaug-cloud	-shell	-cloud
Austronesian Language Family						
ind_Latn (Indonesian)	87.4	-0.1	-0.2	-0.2	-0.1	-0.1
jav_Latn (Javanese)	70.1	+0.9	-0.0	-0.1	+2.1	+1.6
sun_Latn (Sundanese)	71.6	+0.7	-0.2	-0.3	+2.0	+1.6
smo_Latn (Samoan)	27.4	+0.2	+0.0	+0.0	+3.1	+1.9
mri_Latn (Maori)	25.7	-0.6	-0.6	-0.6	+5.9	+3.2
ceb_Latn (Cebuano)	64.2	+0.5	-0.4	-0.4	+1.8	+1.1
zsm_Latn (Standard Malay)	86.5	-0.0	-0.1	-0.1	-0.0	-0.1
tgl_Latn (Tagalog)	73.2	+0.6	-0.4	-0.4	+1.6	+0.7
ilo_Latn (Ilocano)	56.1	+1.0	+0.2	-0.1	+3.0	+2.6
fj_Latn (Fijian)	27.7	+0.3	-0.1	-0.0	+7.8	+5.4
plt_Latn (Plateau Malagasy)	44.3	+2.8	+1.6	+1.3	+5.8	+3.1
pag_Latn (Pangasinan)	35.5	+0.8	+1.0	+0.9	+14.0	+10.4
Average	52.9	+0.7	+0.1	+0.0	+4.3	+2.9
Baseline < 25 Average	46.9	+0.7	+0.2	+0.1	+5.1	+3.4
Arabic Language Family						
arb_Arab (Modern Standard Arabic)	80.6	-0.2	-0.4	-0.4	+0.0	-0.3
acm_Arab (Mesopotamian Arabic)	76.0	-0.2	-0.4	-0.3	+0.7	+0.5
acq_Arab (Ta'izzi-Adeni Arabic)	77.6	-0.2	-0.5	-0.5	+0.3	-0.1
aeb_Arab (Tunisian Arabic)	68.8	-0.2	-0.5	-0.5	+1.7	+1.2
ajp_Arab (South Levantine Arabic)	73.2	+0.0	-0.5	-0.5	+1.5	+1.0
apc_Arab (North Levantine Arabic)	72.8	-0.2	-0.6	-0.7	+1.2	+0.8
ars_Arab (Najdi Arabic)	80.0	-0.3	-0.5	-0.5	+0.0	-0.4
ary_Arab (Moroccan Arabic)	60.8	+0.1	-0.3	-0.3	+3.7	+2.8
arz_Arab (Egyptian Arabic)	71.4	+0.0	-0.5	-0.6	+1.4	+1.0
Average	72.6	-0.1	-0.5	-0.5	+1.3	+0.8
Baseline < 25 Average	71.5	-0.1	-0.5	-0.5	+1.5	+1.0
Romance Language Family						
ita_Latn (Italian)	86.3	-0.1	-0.2	-0.2	+0.0	-0.2
spa_Latn (Spanish)	85.2	-0.3	-0.4	-0.4	-0.2	-0.2
fra_Latn (French)	87.8	-0.2	-0.3	-0.3	-0.2	-0.3
por_Latn (Portuguese)	88.0	-0.1	-0.1	-0.1	-0.1	-0.1
ron_Latn (Romanian)	88.2	-0.1	-0.1	-0.1	-0.1	-0.2
glg_Latn (Galician)	86.6	+0.0	-0.0	-0.1	-0.0	-0.0
cat_Latn (Catalan)	87.1	-0.1	-0.1	-0.1	-0.1	-0.0
oci_Latn (Occitan)	80.0	+0.0	-0.4	-0.4	+0.2	+0.0
ast_Latn (Asturian)	79.8	+0.2	-0.2	-0.2	+0.8	+0.4
lmo_Latn (Lombard)	44.5	-2.3	-1.0	-1.0	+19.3	+17.1
vec_Latn (Venetian)	54.0	-1.2	+0.0	+0.0	+17.9	+16.1
scn_Latn (Sicilian)	44.1	-0.8	+0.2	+0.1	+18.9	+15.9
srd_Latn (Sardinian)	37.9	-0.9	-0.2	-0.3	+21.4	+17.5
fur_Latn (Friulian)	42.5	-1.6	-0.7	-0.7	+19.9	+16.4
lij_Latn (Ligurian)	45.0	-0.8	-0.1	-0.0	+19.1	+15.9
Average	67.9	-0.6	-0.2	-0.3	+8.3	+7.0
Baseline < 25 Average	44.7	-1.3	-0.3	-0.3	+19.4	+16.5
Turkic Language Family						
tur_Latn (Turkish)	86.7	-0.1	-0.3	-0.3	-0.2	-0.2
uzn_Latn (Northern Uzbek)	39.5	+0.8	-0.1	-0.0	+5.0	+3.0
tuk_Latn (Turkmen)	34.4	-1.8	-1.3	-1.1	+10.1	+4.2
azj_Latn (North Azerbaijani)	60.8	-0.9	+0.3	+0.2	+9.4	+5.4
crh_Latn (Crimean Tatar)	55.7	-0.3	-0.2	-0.3	+7.7	+4.9
Average	47.6	-0.6	-0.3	-0.3	+8.0	+4.4
Baseline < 25 Average	47.6	-0.6	-0.3	-0.3	+8.0	+4.4
Indic Language Family						
hin_Deva (Hindi)	86.2	-0.2	-0.2	-0.2	-0.0	-0.0
hne_Deva (Chhattisgarhi)	66.4	-0.2	-0.2	-0.3	+6.2	+6.7
bho_Deva (Bhojpuri)	63.6	-0.0	-0.1	-0.1	+6.5	+7.2
mag_Deva (Magahi)	69.0	+0.2	+0.2	+0.2	+6.3	+6.9
mai_Deva (Maithili)	57.9	+0.2	+0.0	-0.1	+13.1	+14.0
Average	64.2	+0.0	-0.0	-0.1	+8.0	+8.7
Baseline < 25 Average	64.2	+0.0	-0.0	-0.1	+8.0	+8.7
French-Creole Language Family						
hat (Haitian)	73.6	+0.4	+0.4	+0.4	+0.7	+0.6
gcf (Guadeloupean)	52.4	-0.9	-1.0	-1.0	+1.6	+1.3
mart1259 (Martinican)	50.7	+0.1	+0.0	-0.8	+1.0	+1.5
acf (Saint Lucian Patois)	43.4	+0.7	+0.7	+0.7	+3.3	+3.3
ger (French Guianese Creole)	43.9	-1.0	-1.1	-1.2	+3.5	+3.3
lou (Louisiana Creole)	43.4	+0.9	+0.9	+1.0	+6.0	+5.7
mfe (Mauritian)	45.0	-0.1	-0.2	-0.0	+4.8	+4.2
ref (Réunion Creole)	47.1	-1.4	-1.2	-1.2	+2.7	+1.8
crs (Seychellois)	54.6	-0.4	-0.3	-0.2	+6.2	+5.6
Average	47.6	-0.3	-0.3	-0.3	+3.6	+3.3
Baseline < 25 Average	47.6	-0.3	-0.3	-0.3	+3.6	+3.3
General Average	60.0	-0.1	-0.2	-0.2	+5.5	+4.4
General Baseline < 25 Average	53.1	-0.2	-0.2	-0.2	+7.0	+5.7

Table 19: M→D COMET scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each M→D approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.

Language	off-the-shelf	fthr1	randaug-shell	randaug-cloud	-shell	-cloud
Austronesian Language Family						
ind_Latn (Indonesian)	88.5	+0.1	+0.1	+0.0	-0.0	-0.1
jav_Latn (Javanese)	67.5	+0.8	+0.8	+1.0	+6.2	+6.2
sun_Latn (Sundanese)	63.7	+0.9	+1.3	+1.2	+8.7	+9.0
smo_Latn (Samoan)	35.4	-2.6	-2.0	-2.5	+6.3	+6.2
mri_Latn (Maori)	40.5	-6.7	-5.6	-6.1	+1.7	+2.2
ceb_Latn (Cebuano)	61.1	-3.7	-3.5	-3.7	+0.1	+0.3
zsm_Latn (Standard Malay)	85.7	+0.1	+0.1	+0.1	+0.3	+0.2
tgl_Latn (Tagalog)	76.4	-4.1	-3.8	-4.3	-6.5	-5.4
ilo_Latn (Ilocano)	48.2	-4.0	-3.5	-3.9	+6.6	+7.0
fij_Latn (Fijian)	40.1	-4.4	-3.7	-4.2	+3.5	+3.7
plt_Latn (Plateau Malagasy)	46.1	-1.0	-1.1	-0.9	+0.9	+1.6
pag_Latn (Pangasinan)	43.1	+2.3	+3.1	+3.1	+13.8	+13.6
Average	55.2	-2.0	-1.6	-1.8	+3.8	+4.1
Baseline < 25 Average	49.5	-2.0	-1.6	-1.8	+5.3	+5.5
Arabic Language Family						
arb_Arab (Modern Standard Arabic)	86.1	+0.5	+0.4	+0.3	+0.3	+0.4
acm_Arab (Mesopotamian Arabic)	83.2	+0.3	+0.4	+0.5	+1.0	+1.1
acq_Arab (Ta'izzi-Adeni Arabic)	84.3	+0.4	+0.4	+0.5	+0.7	+0.8
aeb_Arab (Tunisian Arabic)	78.4	+0.4	+0.4	+0.4	+1.6	+1.8
ajp_Arab (South Levantine Arabic)	83.7	-0.3	-0.3	-0.4	+0.2	+0.3
apc_Arab (North Levantine Arabic)	82.7	-0.1	+0.0	-0.1	+0.8	+0.8
ars_Arab (Najdi Arabic)	85.4	+0.7	+0.8	+0.7	+0.7	+0.7
ary_Arab (Moroccan Arabic)	74.0	+0.2	+0.6	+0.2	+2.7	+2.9
arz_Arab (Egyptian Arabic)	82.4	-0.0	+0.1	-0.1	+0.3	+0.4
Average	81.7	+0.2	+0.3	+0.2	+1.0	+1.1
Baseline < 25 Average	74.0	+0.2	+0.6	+0.2	+2.7	+2.9
Romance Language Family						
ita_Latn (Italian)	87.1	+0.2	+0.1	+0.1	-0.1	-0.1
spa_Latn (Spanish)	86.1	-0.3	-0.2	-0.2	-0.4	-0.3
fra_Latn (French)	88.6	-0.7	-0.8	-0.9	-1.2	-1.1
por_Latn (Portuguese)	88.4	-0.3	-0.3	-0.3	-0.8	-0.7
ron_Latn (Romanian)	88.6	-0.7	-0.7	-0.7	-1.4	-1.3
glg_Latn (Galician)	85.2	-0.5	-0.2	-0.5	+0.1	+0.3
cat_Latn (Catalan)	84.0	-0.9	-0.8	-0.8	-0.4	-0.1
oci_Latn (Occitan)	74.9	-0.3	+0.2	+0.1	+3.3	+3.5
ast_Latn (Asturian)	77.7	-1.1	-0.6	-1.0	+1.5	+1.8
lmo_Latn (Lombard)	58.8	+3.3	+4.6	+3.9	+12.2	+12.1
vec_Latn (Venetian)	70.3	+3.9	+4.3	+4.0	+7.6	+7.9
scn_Latn (Sicilian)	59.5	+1.9	+3.0	+2.5	+12.0	+11.9
srd_Latn (Sardinian)	56.9	-0.1	+1.4	+0.9	+12.0	+12.2
fur_Latn (Friulian)	59.7	+4.1	+5.0	+4.8	+11.9	+12.2
lij_Latn (Ligurian)	56.9	+4.8	+6.2	+5.6	+15.3	+14.9
Average	74.0	+0.9	+1.5	+1.3	+5.1	+5.2
Baseline < 25 Average	58.4	+2.8	+4.0	+3.5	+12.7	+12.7
Turkic Language Family						
tur_Latn (Turkish)	87.1	+0.2	+0.3	-0.0	+0.1	-0.0
uzn_Latn (Northern Uzbek)	56.5	+1.4	+1.7	+1.3	+7.6	+7.7
tuk_Latn (Turkmen)	51.2	+8.8	+8.7	+8.5	+12.6	+12.2
azj_Latn (North Azerbaijani)	70.5	+2.6	+2.5	+2.4	+5.3	+5.2
crh_Latn (Crimean Tatar)	64.8	+4.4	+4.1	+4.1	+7.8	+7.6
Average	60.7	+4.3	+4.3	+4.1	+8.3	+8.2
Baseline < 25 Average	60.7	+4.3	+4.3	+4.1	+8.3	+8.2
Indic Language Family						
hin_Deva (Hindi)	88.3	-1.0	-1.0	-1.0	-1.1	-1.0
hne_Deva (Chattisgarhi)	77.2	-1.4	-1.5	-1.4	+2.9	+2.8
bho_Deva (Bhojpur)	75.2	-2.5	-2.8	-2.5	+1.9	+2.1
mag_Deva (Magahi)	78.9	-1.2	-1.1	-1.0	+2.7	+2.9
mai_Deva (Maithili)	74.6	-1.8	-1.6	-2.0	+3.9	+4.2
Average	76.5	-1.7	-1.8	-1.7	+2.8	+3.0
Baseline < 25 Average	76.5	-1.7	-1.8	-1.7	+2.8	+3.0
French-Creole Language Family						
hat (Haitian)	51.1	+21.1	+21.2	+21.3	+20.3	+20.5
gcf (Guadeloupean)	53.7	+1.1	+1.0	+0.6	+2.5	+3.6
mart1259 (Martinican)	49.8	+4.2	+3.2	+4.1	+6.7	+7.2
acf (Saint Lucian Patois)	42.4	+6.4	+5.8	+6.4	+6.3	+7.2
gcr (French Guianese Creole)	44.6	+7.5	+4.4	+3.8	+7.5	+8.2
lou (Louisiana Creole)	49.9	+3.7	+3.2	+3.9	+4.7	+4.4
mfe (Mauritian)	45.1	+9.7	+9.8	+9.8	+10.1	+10.4
rcf (Réunion Creole)	48.3	+2.8	+4.5	+3.6	+6.3	+6.2
crs (Seychellois)	47.4	+14.0	+13.9	+13.6	+15.7	+15.4
Average	47.6	+6.2	+5.7	+5.7	+7.5	+7.8
Baseline < 25 Average	48.0	+7.8	+7.4	+7.4	+8.9	+9.2
General Average	62.9	+0.5	+0.5	+0.4	+5.0	+4.6
General Baseline < 25 Average	54.5	+1.0	+1.1	+1.0	+7.2	+6.6

Table 20: M→D COMET scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each M→D approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.

	-ccloud	-ccloud-multsrcs
Arabic Language Family (FloRes Dataset)		
acm_Arab	30.4	30.5
acq_Arab	31.4	31.8
aeb_Arab	25.4	26.0
ajp_Arab	33.5	35.2
apc_Arab	30.3	32.4
arb_Arab	37.1	36.0
ars_Arab	36.1	35.2
ary_Arab	19.7	23.8
arz_Arab	27.9	29.4
Average	30.2	31.1
Romance Language Family		
ast_Latn	31.6	33.4
cat_Latn	33.4	35.8
fra_Latn	40.9	41.4
fur_Latn	24.4	24.1
glg_Latn	34.1	36.4
ita_Latn	32.0	33.5
lij_Latn	25.6	25.0
lmo_Latn	26.2	25.5
oci_Latn	36.5	37.3
por_Latn	43.8	45.2
ron_Latn	38.2	40.4
scn_Latn	26.1	25.5
spa_Latn	30.7	30.9
srd_Latn	24.1	23.5
vec_Latn	32.7	31.0
Average	32.0	32.6
Turkic Language Family		
azj_Latn	11.3	11.6
crh_Latn	17.0	16.4
tuk_Latn	9.4	9.7
tur_Latn	29.2	27.7
uzn_Latn	7.6	9.2
Average	14.9	14.9

Table 21: Summary of adopting multiple HRLs as sources for the M→D approach on the model Aya-23. These results are computed over a 300 sample-subset of the FloRes-200 test set, for Aya-23

Language	off=the-shelf	fthr1	-shell
bho_Deva	10.1	10.0	13.3
hin_Deva	34.8	33.4	33.3
hne_Deva	14.5	13.2	19.5
mag_Deva	17.4	16.0	21.3
mai_Deva	9.6	7.1	16.1
Average	17.3	15.9	20.7

Table 22: Using indicorp to noise for languages in the Indic family. The M2M model was tested. These results are computed over a 300 sample-subset of the FloRes-200 test set.

²⁵https://web.archive.org/web/20091027151616/http://geocities.com/fijidictionary/eng_fijian_

Language	off-the-shelf	fthr1	-shell
bho_Deva	10.1	10.0	13.3
hin_Deva	34.8	33.4	33.3
hne_Deva	14.5	13.2	19.5
mag_Deva	17.4	16.0	21.3
mai_Deva	9.6	7.1	16.1
Average	17.3	15.9	20.7

Table 23: Using the NLLB dataset to noise the languages in the Indic family. The M2M model was tested. These results are computed over a 300 sample-subset of the FloRes-200 test set.

Language Family	-ccloud	-ccloud-sem
Austronesian	17.9	18.0
Arabic	22.8	22.6
Romance	30.5	30.5
Turkic	11.8	12.0
Indic	21.0	21.0
French-Creole	12.1	-

Table 24: Additional semantic noisers as described in § E.1 for M→D are applied to the model M2M. Language family means were taken.

dict.html

²⁶<https://toidp.com/search/>

²⁷<https://ilocano.pinoydictionary.com/>

²⁸<https://education.nsw.gov.au/content/dam/main-education/teaching-and-learning/curriculum/multicultural-education/eald/eald-bilingual-dictionary-samoan.pdf>

²⁹<https://www.webonary.org/turkmen/en/>

³⁰<https://www.scribd.com/doc/37996625/Kirim-Tatarca-dan-Turkce-ye-SozluK>

³¹<https://derja.ninja/>

³²<https://www.learnmoroccan.com/dictionary>

³³<https://tajinequiparle.com/en/english-moroccan-arabic-dictionary/>

³⁴<https://www.thoughtco.com/sicilian-english-dictionary-2011651>

³⁵<https://arlef.it/en/>

³⁶<https://ligu.re/en/>

Language	off-the-shelf	-cont	-func	-all
Austronesian Language Family				
ind_Latn (Indonesian)	38.9	-	-	-
jav_Latn (Javanese)	22.9	-0.2	+0.2	+0.2
sun_Latn (Sundanese)	22.9	-2.6	+1.0	-1.5
smo_Latn (Samoan)	2.3	+0.2	+1.0	+2.1
mri_Latn (Maori)	2.3	+0.5	+1.2	+2.3
ceb_Latn (Cebuano)	22.2	-2.7	-9.5	-10.4
zsm_Latn (Standard Malay)	38.7	-6.0	-1.5	-7.5
tgl_Latn (Tagalog)	27.8	-6.3	-14.3	-16.3
ilo_Latn (Ilocano)	13.4	+0.5	-4.0	-3.3
fij_Latn (Fijian)	2.3	+0.3	+1.3	+1.9
plt_Latn (Plateau Malagasy)	2.5	+0.7	+2.4	+3.8
pag_Latn (Pangasinan)	6.4	-0.6	+5.5	+5.1
Average	14.9	-1.5	-1.5	-2.1
Baseline < 25 Average	10.8	-0.4	-0.1	+0.0
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	29.6	-	-	-
acm_Arab (Mesopotamian Arabic)	22.4	-0.5	+1.0	+0.4
acq_Arab (Ta'izzi-Adeni Arabic)	24.7	-0.6	+0.7	-0.2
aeb_Arab (Tunisian Arabic)	17.7	-1.3	+1.3	+0.3
ajp_Arab (South Levantine Arabic)	24.5	-5.3	-0.8	-5.7
apc_Arab (North Levantine Arabic)	21.7	-3.6	+1.0	-2.8
ars_Arab (Najdi Arabic)	28.5	+0.0	+0.1	+0.0
ary_Arab (Moroccan Arabic)	11.9	-1.3	+2.0	+0.1
arz_Arab (Egyptian Arabic)	18.2	-3.1	+0.7	-2.2
Average	21.2	-2.0	+0.7	-1.2
Baseline < 25 Average	20.2	-2.2	+0.8	-1.4
Romance Language Family				
ita_Latn (Italian)	30.8	-	-	-
spa_Latn (Spanish)	26.8	-10.0	-7.2	-16.9
fra_Latn (French)	41.6	-16.5	-13.0	-28.3
por_Latn (Portuguese)	45.8	-16.2	-14.9	-30.0
ron_Latn (Romanian)	40.8	-16.1	-15.6	-28.7
glg_Latn (Galician)	37.5	-14.3	-11.1	-23.5
cat_Latn (Catalan)	42.7	-14.9	-12.8	-26.1
oci_Latn (Occitan)	44.6	-10.7	-5.1	-15.7
ast_Latn (Asturian)	34.3	-10.9	-7.6	-18.4
lmo_Latn (Lombard)	9.6	+1.6	+5.2	+6.7
vec_Latn (Venetian)	17.2	-1.7	+8.0	+5.4
scn_Latn (Sicilian)	7.3	+1.6	+8.5	+10.0
srd_Latn (Sardinian)	5.3	+1.8	+10.5	+13.8
fur_Latn (Friulian)	8.6	+4.3	+6.0	+10.8
lij_Latn (Ligurian)	9.8	+2.2	+9.8	+12.1
Average	26.6	-7.1	-2.8	-9.2
Baseline < 25 Average	9.6	+1.6	+8.0	+9.8
Turkic Language Family				
tur_Latn (Turkish)	33.3	-	-	-
uzn_Latn (Northern Uzbek)	1.6	+2.2	+1.0	+5.0
tuk_Latn (Turkmen)	2.4	+2.3	+1.6	+4.6
azj_Latn (North Azerbaijani)	7.4	-0.2	+0.1	-0.2
crh_Latn (Crimean Tatar)	8.5	+1.3	+1.3	+2.5
Average	5.0	+1.4	+1.0	+3.0
Baseline < 25 Average	5.0	+1.4	+1.0	+3.0
Indic Language Family				
hin_Deva (Hindi)	33.4	-	-	-
hne_Deva (Chattisgarhi)	15.2	-1.3	+13.9	+9.9
bho_Deva (Bhojpur)	10.1	+0.3	+6.1	+6.4
mag_Deva (Magahi)	17.2	+0.8	+12.8	+12.8
mai_Deva (Maithili)	7.7	+1.4	+15.3	+16.6
Average	12.6	+0.3	+12.0	+11.4
Baseline < 25 Average	12.6	+0.3	+12.0	+11.4
Creole Language Family				
hat (Haitian)	39.7	-	-	-
acf (Saint Lucian Patois)	3.7	+2.9	+3.7	+7.8
mfe (Mauritian)	5.1	+1.6	+1.1	+3.2
crs (Seychellois)	11.2	+0.8	+4.1	+3.6
Average	5.9	+2.5	+3.7	+5.6
Baseline < 25 Average	5.9	+2.5	+3.7	+5.6
General Average	16.8	-1.4	+1.5	-0.8
General Baseline < 25 Average	10.9	+0.9	+4.3	+4.5

Table 25: D→M BLEU scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each D→M approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The best score for each language in the D→M paradigm is **bolded** and underlined.

Language	off-the-shelf	-cont	-func	-all
Austronesian Language Family				
ind_Latn (Indonesian)	41.5	-	-	-
jav_Latn (Javanese)	14.2	+2.9	+7.8	+10.6
sun_Latn (Sundanese)	11.0	+3.6	+10.1	+11.4
smo_Latn (Samoan)	3.2	-0.1	+0.4	+1.1
mri_Latn (Maori)	4.8	+0.4	-0.8	+0.4
ceb_Latn (Cebuano)	14.0	-0.0	-3.0	-1.7
zsm_Latn (Standard Malay)	36.5	-3.7	-1.2	-4.6
tgl_Latn (Tagalog)	25.6	-4.7	-12.5	-12.4
ilo_Latn (Ilocano)	7.5	+0.3	+0.8	+2.4
fij_Latn (Fijian)	3.3	+0.4	-0.1	+0.7
plt_Latn (Plateau Malagasy)	3.3	+1.1	+1.3	+3.9
pag_Latn (Pangasinan)	7.3	-1.0	+5.1	+4.9
Average	11.9	-0.1	+0.7	+1.5
Baseline < 25 Average	7.6	+0.8	+2.4	+3.7
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	38.1	-	-	-
acm_Arab (Mesopotamian Arabic)	31.6	-1.7	-0.2	-1.9
acq_Arab (Ta'izzi-Adeni Arabic)	33.7	-1.3	+0.2	-1.1
aeb_Arab (Tunisian Arabic)	26.7	-2.8	-0.1	-2.6
ajp_Arab (South Levantine Arabic)	37.4	-9.1	-4.0	-12.5
apc_Arab (North Levantine Arabic)	33.2	-6.7	-1.2	-8.2
ars_Arab (Najdi Arabic)	37.1	-0.1	+0.0	-0.3
ary_Arab (Moroccan Arabic)	23.5	-6.2	-1.8	-6.9
arz_Arab (Egyptian Arabic)	28.9	-6.4	-1.4	-6.8
Average	31.5	-4.3	-1.1	-5.0
Baseline < 25 Average	23.5	-6.2	-1.8	-6.9
Romance Language Family				
ita_Latn (Italian)	32.6	-	-	-
spa_Latn (Spanish)	29.0	-10.1	-4.6	-13.4
fra_Latn (French)	44.0	-15.5	-9.8	-21.9
por_Latn (Portuguese)	46.9	-15.8	-7.7	-20.7
ron_Latn (Romanian)	42.0	-15.3	-10.2	-22.4
glg_Latn (Galician)	35.5	-12.4	-5.7	-15.3
cat_Latn (Catalan)	36.5	-10.6	-5.7	-14.7
oci_Latn (Occitan)	35.1	-8.1	-2.2	-9.1
ast_Latn (Asturian)	31.8	-8.3	-5.1	-11.8
lmo_Latn (Lombard)	17.4	-2.0	+3.3	+0.2
vec_Latn (Venetian)	26.7	-4.7	+2.7	-3.7
scn_Latn (Sicilian)	16.7	-0.6	+3.7	+1.9
srd_Latn (Sardinian)	16.4	-1.4	+6.1	+4.4
fur_Latn (Friulian)	17.7	-0.4	+4.5	+4.4
lij_Latn (Ligurian)	17.4	-0.5	+6.1	+4.8
Average	29.5	-7.6	-1.8	-8.4
Baseline < 25 Average	17.1	-1.0	+4.7	+3.1
Turkic Language Family				
tur_Latn (Turkish)	31.7	-	-	-
uzn_Latn (Northern Uzbek)	3.8	+1.3	+0.6	+2.1
tuk_Latn (Turkmen)	6.1	+0.7	+0.9	+2.0
azj_Latn (North Azerbaijani)	9.8	-2.7	-0.3	-3.1
crh_Latn (Crimean Tatar)	11.9	-1.4	+0.0	-0.8
Average	7.9	-0.6	+0.3	+0.0
Baseline < 25 Average	7.9	-0.6	+0.3	+0.0
Indic Language Family				
hin_Deva (Hindi)	34.8	-	-	-
hne_Deva (Chattisgarhi)	23.3	-3.3	+6.0	+2.0
bho_Deva (Bhojपुरी)	16.2	-0.2	+3.2	+2.8
mag_Deva (Magahi)	24.9	-0.3	+5.5	+5.3
mai_Deva (Maithili)	17.2	-0.4	+7.6	+7.4
Average	20.4	-1.1	+5.6	+4.4
Baseline < 25 Average	20.4	-1.1	+5.6	+4.4
Creole Language Family				
hat (Haitian)	8.9	-	-	-
acf (Saint Lucian Patois)	2.6	+0.2	-0.1	+0.5
mfe (Mauritian)	3.4	-0.3	-0.2	-0.3
crs (Seychellois)	5.7	-1.1	+0.6	-0.5
Average	7.3	-3.8	-3.3	-3.5
Baseline < 25 Average	7.3	-3.8	-3.3	-3.5
General Average	19.7	-2.2	+1.2	-1.6
General Baseline < 25 Average	11.3	-0.0	+3.0	+2.7

Table 26: D→M BLEU scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each D→M approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The best score for each language in the D→M paradigm is **bolded** and underlined.

Language	off-the-shelf	-cont	-func	-all
Austronesian Language Family				
ind_Latn (Indonesian)	87.4	-	-	-
jav_Latn (Javanese)	70.1	-0.6	-0.0	-0.6
sun_Latn (Sundanese)	71.6	-2.6	+0.9	-2.7
smo_Latn (Samoan)	27.4	+7.5	+3.9	+13.6
mri_Latn (Maori)	25.7	+7.4	+7.9	+17.0
ceb_Latn (Cebuano)	64.2	-3.3	-8.4	-9.8
zsm_Latn (Standard Malay)	86.5	-6.7	-1.5	-8.0
tgl_Latn (Tagalog)	73.2	-7.9	-14.2	-18.0
ilo_Latn (Ilocano)	56.1	+0.8	-4.3	-3.2
fij_Latn (Fijian)	27.7	+7.9	+7.4	+14.9
plt_Latn (Plateau Malagasy)	44.3	-1.3	-7.3	-0.7
pag_Latn (Pangasinan)	35.5	+2.4	+8.6	+11.3
Average	52.9	+0.3	-0.6	+1.3
Baseline < 25 Average	46.9	+2.0	+1.0	+4.4
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	80.6	-	-	-
acm_Arab (Mesopotamian Arabic)	76.0	-1.1	+0.3	-0.9
acq_Arab (Ta'izzi-Adeni Arabic)	77.6	-1.3	+0.0	-1.2
aeb_Arab (Tunisian Arabic)	68.8	-1.6	+0.8	-0.8
ajp_Arab (South Levantine Arabic)	73.2	-6.3	-0.4	-6.5
apc_Arab (North Levantine Arabic)	72.8	-4.8	+0.6	-4.3
ars_Arab (Najdi Arabic)	80.0	+0.0	+0.0	+0.1
ary_Arab (Moroccan Arabic)	60.8	-2.3	+2.4	-0.8
arz_Arab (Egyptian Arabic)	71.4	-4.4	+0.6	-3.9
Average	72.6	-2.7	+0.5	-2.3
Baseline < 25 Average	71.5	-3.1	+0.6	-2.6
Romance Language Family				
ita_Latn (Italian)	86.3	-	-	-
spa_Latn (Spanish)	85.2	-18.6	-8.1	-28.7
fra_Latn (French)	87.8	-20.8	-10.6	-30.6
por_Latn (Portuguese)	88.0	-18.9	-10.2	-28.6
ron_Latn (Romanian)	88.2	-20.8	-12.9	-32.3
glg_Latn (Galician)	86.6	-18.6	-9.2	-28.0
cat_Latn (Catalan)	87.1	-16.7	-9.9	-27.7
oci_Latn (Occitan)	80.0	-10.2	-2.8	-13.3
ast_Latn (Asturian)	79.8	-14.0	-7.3	-21.8
lmo_Latn (Lombard)	44.5	+2.3	+6.6	+8.7
vec_Latn (Venetian)	54.0	-1.7	+8.3	+5.6
scn_Latn (Sicilian)	44.1	+2.6	+8.7	+11.2
srd_Latn (Sardinian)	37.9	+3.6	+11.3	+16.1
fur_Latn (Friulian)	42.5	+6.4	+8.2	+14.6
lij_Latn (Ligurian)	45.0	+3.3	+11.0	+13.0
Average	67.9	-8.7	-1.2	-10.1
Baseline < 25 Average	44.7	+2.8	+9.0	+11.5
Turkic Language Family				
tur_Latn (Turkish)	86.7	-	-	-
uzn_Latn (Northern Uzbek)	39.5	+11.4	+4.9	+16.0
tuk_Latn (Turkmen)	34.4	+8.2	+3.3	+12.2
azj_Latn (North Azerbaijani)	60.8	-0.7	+0.9	-1.1
crh_Latn (Crimean Tatar)	55.7	+1.9	+2.4	+3.5
Average	47.6	+5.2	+2.9	+7.6
Baseline < 25 Average	47.6	+5.2	+2.9	+7.6
Indic Language Family				
hin_Deva (Hindi)	86.2	-	-	-
hne_Deva (Chattisgarhi)	66.4	-2.2	+10.0	+6.3
bho_Deva (Bhojpuri)	63.6	+0.9	+7.2	+7.0
mag_Deva (Magahi)	69.0	+0.7	+9.7	+9.3
mai_Deva (Maithili)	57.9	+3.9	+16.4	+17.3
Average	64.2	+0.8	+10.8	+10.0
Baseline < 25 Average	64.2	+0.8	+10.8	+10.0
General Average	62.5	-2.8	+0.9	-1.9
General Baseline < 25 Average	54.6	+1.2	+4.1	+5.4

Table 27: D→M COMET scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each D→M approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.

Language	off-the-shelf	-cont	-func	-all
Austronesian Language Family				
ind_Latn (Indonesian)	88.5	-	-	-
jav_Latn (Javanese)	67.5	+1.1	+3.5	+3.4
sun_Latn (Sundanese)	63.7	+1.7	+7.6	+6.1
smo_Latn (Samoan)	35.4	+3.9	+0.5	+5.3
mri_Latn (Maori)	40.5	+3.6	-2.8	+3.4
ceb_Latn (Cebuano)	61.1	-2.2	-7.7	-7.7
zsm_Latn (Standard Malay)	85.7	-4.3	-1.0	-5.1
tgl_Latn (Tagalog)	76.4	-7.3	-17.0	-19.2
ilo_Latn (Ilocano)	48.2	+2.6	-2.4	+0.4
fji_Latn (Fijian)	40.1	+3.5	-3.2	+2.2
plt_Latn (Plateau Malagasy)	46.1	+2.2	-4.5	+1.0
pag_Latn (Pangasinan)	43.1	+0.4	+7.7	+7.4
Average	55.2	+0.5	-1.7	-0.3
Baseline < 25 Average	49.5	+1.9	-0.1	+2.4
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	86.1	-	-	-
acm_Arab (Mesopotamian Arabic)	83.2	-1.5	-0.2	-1.7
acq_Arab (Ta'izzi-Adeni Arabic)	84.3	-1.2	-0.1	-1.3
aeb_Arab (Tunisian Arabic)	78.4	-2.8	-0.0	-3.3
ajp_Arab (South Levantine Arabic)	83.7	-8.0	-1.8	-10.1
apc_Arab (North Levantine Arabic)	82.7	-6.2	-1.0	-7.7
ars_Arab (Najdi Arabic)	85.4	-0.1	+0.1	-0.0
ary_Arab (Moroccan Arabic)	74.0	-6.5	-1.3	-7.7
arz_Arab (Egyptian Arabic)	82.4	-6.2	-1.4	-7.9
Average	81.7	-4.1	-0.7	-5.0
Baseline < 25 Average	74.0	-6.5	-1.3	-7.7
Romance Language Family				
ita_Latn (Italian)	87.1	-	-	-
spa_Latn (Spanish)	86.1	-15.1	-4.1	-19.6
fra_Latn (French)	88.6	-14.4	-6.1	-20.9
por_Latn (Portuguese)	88.4	-13.7	-4.1	-17.4
ron_Latn (Romanian)	88.6	-14.5	-6.0	-21.9
glg_Latn (Galician)	85.2	-13.7	-4.8	-19.0
cat_Latn (Catalan)	84.0	-11.3	-5.0	-16.7
oci_Latn (Occitan)	74.9	-9.0	-1.9	-10.5
ast_Latn (Asturian)	77.7	-10.7	-4.4	-15.0
lmo_Latn (Lombard)	58.8	-3.7	+3.2	-1.3
vec_Latn (Venetian)	70.3	-7.0	+1.5	-7.0
scn_Latn (Sicilian)	59.5	-2.1	+2.7	-1.2
srd_Latn (Sardinian)	56.9	-2.6	+5.5	+1.7
fur_Latn (Friulian)	59.7	-1.0	+4.5	+3.2
lij_Latn (Ligurian)	56.9	-1.1	+6.4	+4.1
Average	74.0	-8.6	-0.9	-10.1
Baseline < 25 Average	58.4	-2.1	+4.5	+1.3
Turkic Language Family				
tur_Latn (Turkish)	87.1	-	-	-
uzn_Latn (Northern Uzbek)	56.5	-2.2	-0.6	-2.3
tuk_Latn (Turkmen)	51.2	+1.2	+2.2	+2.9
azj_Latn (North Azerbaijani)	70.5	-9.8	-1.5	-11.0
crh_Latn (Crimean Tatar)	64.8	-3.8	+0.5	-3.2
Average	60.7	-3.7	+0.1	-3.4
Baseline < 25 Average	60.7	-3.7	+0.1	-3.4
Indic Language Family				
hin_Deva (Hindi)	88.3	-	-	-
hne_Deva (Chattisgarhi)	77.2	-2.9	+3.2	-0.4
bho_Deva (Bhojpuri)	75.2	-1.6	+1.6	+0.5
mag_Deva (Magahi)	78.9	-0.2	+3.0	+2.8
mai_Deva (Maithili)	74.6	-0.6	+4.6	+4.0
Average	76.5	-1.3	+3.1	+1.7
Baseline < 25 Average	76.5	-1.3	+3.1	+1.7
General Average	65.9	-3.5	+0.1	-3.3
General Baseline < 25 Average	56.6	+0.3	+2.9	+3.3

Table 28: D→M COMET scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each D→M approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.

Dictionary	Medium	Length (words)	→ English	→ Related Language	Inflections	Part-of-Speech
Fijian						
Ronald Gatty's Fijian-English Dictionary (Gatty, 2009)	PDF	~ 7585	✓			✓
Geocities Fijian-English Dictionary ²⁵	table	~ 867	✓			✓
Ilocano						
The Online Ilokano Dictionary Project ²⁶	MySQL database	~ 2060	✓			
The Pinoy Dictionary ²⁷	web app	unknown	✓			✓
Samoan						
University of Hawai'i at Manoa (Kobayashi et al., 2016)	web app	13703	✓		✓	✓
New South Wales Department of Education and Training ²⁸	PDF	~ 3160	✓		✓	✓
Maori						
Te Aka (Moorfield, 2024)	web app	unknown	✓			✓
Turkmen						
Peace Corps Turkmenistan (Garrett et al., 1996)	PDF	~ 10000	✓		✓	✓
Türkmençe-iñlisçe sözlük ²⁹	web app	~ 7179	✓		✓	✓
Crimean Tatar						
Kırım Tatarca'dan Türkçe'ye Sozluk ³⁰	PDF	~ 1440		✓		
Tunisian Arabic						
Derja Ninja ³¹	web app	19,231	✓			
Peace Corps (Abdelkader et al., 1977)	PDF	~ 6000	✓		✓	✓
Moroccan Arabic						
Traductor Darija (El Ouamari, 2024)	web app	unknown	✓			
learnmoroccan ³²	web app	unknown	✓			
Online English-Moroccan Arabic (Darija) Dictionary ³³	web app	unknown	✓		✓	
Egyptian Arabic						
Lisaan Masry (Green, 2020)	web app	unknown	✓		✓	✓
Sicilian						
Michael San Filippo's English-Sicilian Dictionary ³⁴	table	678	✓			
Dieli (Sicilian, English) (Dieli, 2011)	table	~ 19152	✓	✓	✓	✓
Friulian						
Agjenzie regional pe lenghe furlane ³⁵	web app	unknown		✓		
Ligurian						
ligu.re ³⁶	web app	unknown	✓	✓		
Bhojpuri						
(Bafna et al., 2024a)	JSON	21983		✓		
Magahi						
(Bafna et al., 2024a)	JSON	30784		✓		
Maithili						
(Bafna et al., 2024a)	JSON	12069		✓		

Table 29: Lexicons specific to a language

Dictionary Source	Total Words	Functional Words	Content Words	Coverage (%)
French - Italian				
Combined	238677	643	238034	89.3
PanLex (Kamholz et al., 2014)	234647	370	234277	38.2
Swadesh (Swadesh, 2015)	217	23	194	1.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	6636	538	6098	84.1
Spanish - Italian				
Combined	209274	538	208736	90.7
PanLex (Kamholz et al., 2014)	205057	349	204708	38.5
Swadesh (Swadesh, 2015)	242	29	213	1.3
Statistically Aligned FloRes Data (Dyer et al., 2013)	6880	413	6467	85.1
Portuguese - Italian				
Combined	162212	507	161705	92.1
PanLex (Kamholz et al., 2014)	157815	257	157558	35.3
Swadesh (Swadesh, 2015)	277	30	247	1.6
Statistically Aligned FloRes Data (Dyer et al., 2013)	6846	422	6424	88.3
Catalan - Italian				
Combined	75100	441	74659	91.0
PanLex (Kamholz et al., 2014)	70435	230	70205	31.3
Swadesh (Swadesh, 2015)	253	23	230	1.5
Statistically Aligned FloRes Data (Dyer et al., 2013)	6851	398	6453	87.8
Romanian - Italian				
Combined	65578	529	65049	89.3
PanLex (Kamholz et al., 2014)	59734	250	59484	21.7
Swadesh (Swadesh, 2015)	241	22	219	0.9
Statistically Aligned FloRes Data (Dyer et al., 2013)	7493	505	6988	86.8
Standard Malay - Indonesian				
Combined	43854	354	43500	95.8
PanLex (Kamholz et al., 2014)	40501	250	40251	44.1
Statistically Aligned FloRes Data (Dyer et al., 2013)	5994	271	5723	92.5
Galician - Italian				
Combined	39031	414	38617	91.2
PanLex (Kamholz et al., 2014)	33989	88	33901	26.8
Swadesh (Swadesh, 2015)	290	24	266	1.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	6883	489	6394	88.5
Magahi - Hindi				
Combined	33143	3274	29869	90.3
PanLex (Kamholz et al., 2014)	10	0	10	0.0
(Bafna et al., 2024a)	30784	3171	27613	51.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	4966	427	4539	82.6
North Azerbaijani - Turkish				
Combined	26518	324	26194	93.5
PanLex (Kamholz et al., 2014)	19772	95	19677	17.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	8134	331	7803	92.0
Asturian - Italian				
Combined	25267	554	24713	91.3
PanLex (Kamholz et al., 2014)	19343	94	19249	17.1
Swadesh (Swadesh, 2015)	222	27	195	1.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	7140	543	6597	89.9
Bhojpuri - Hindi				
Combined	24453	1262	23191	87.1
PanLex (Kamholz et al., 2014)	38	0	38	0.1
Swadesh (Swadesh, 2015)	235	26	209	1.4
(Bafna et al., 2024a)	21983	1088	20895	50.3
Statistically Aligned FloRes Data (Dyer et al., 2013)	5044	515	4529	78.2

Table 30: Part 1 —Total, functional, and content word counts across all lexicons. We also report the word type coverage, a measure of how extensively the lexicon documents a LRL. To calculate coverage percentage, we compared the CRL lexicon against the vocabulary size of the FloRes dev set.

Dictionary Source	Total Words	Functional Words	Content Words	Coverage (%)
Sicilian - Italian				
Combined	23755	575	23180	90.2
Dieli (Dieli, 2011)	10380	143	10237	11.8
PanLex (Kamholz et al., 2014)	17080	206	16874	15.9
Swadesh (Swadesh, 2015)	219	22	197	1.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	7213	528	6685	88.8
Northern Uzbek - Turkish				
Combined	23307	378	22929	94.5
PanLex (Kamholz et al., 2014)	16288	116	16172	14.3
Statistically Aligned FloRes Data (Dyer et al., 2013)	8198	350	7848	93.7
Occitan (post 1500) - Italian				
Combined	21464	401	21063	88.3
PanLex (Kamholz et al., 2014)	15708	76	15632	16.5
Swadesh (Swadesh, 2015)	291	34	257	1.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	6877	409	6468	86.5
Turkmen - Turkish				
Combined	15425	323	15102	94.4
PanLex (Kamholz et al., 2014)	8520	64	8456	12.7
Swadesh (Swadesh, 2015)	259	21	238	1.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	7935	339	7596	93.9
Maithili - Hindi				
Combined	15154	676	14478	89.1
PanLex (Kamholz et al., 2014)	16	0	16	0.0
(Bafna et al., 2024a)	12069	471	11598	39.7
Statistically Aligned FloRes Data (Dyer et al., 2013)	5211	467	4744	83.8
Venetian - Italian				
Combined	14531	396	14135	89.3
PanLex (Kamholz et al., 2014)	8317	109	8208	6.8
Statistically Aligned FloRes Data (Dyer et al., 2013)	6688	355	6333	88.8
Lombard - Italian				
Combined	13603	432	13171	81.2
PanLex (Kamholz et al., 2014)	6565	25	6540	1.3
Swadesh (Swadesh, 2015)	336	44	292	0.8
Statistically Aligned FloRes Data (Dyer et al., 2013)	6875	380	6495	81.0
Ligurian - Italian				
Combined	12591	363	12228	86.7
PanLex (Kamholz et al., 2014)	6191	21	6170	6.3
Swadesh (Swadesh, 2015)	216	24	192	0.7
Statistically Aligned FloRes Data (Dyer et al., 2013)	6703	356	6347	85.9
Crimean Tatar - Turkish				
Combined	11213	325	10888	93.6
PanLex (Kamholz et al., 2014)	3470	51	3419	6.3
Swadesh (Swadesh, 2015)	242	18	224	1.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	8260	305	7955	93.3
Sardinian - Italian				
Combined	10720	352	10368	87.0
PanLex (Kamholz et al., 2014)	4850	52	4798	7.6
Statistically Aligned FloRes Data (Dyer et al., 2013)	6356	342	6014	86.0
Friulian - Italian				
Combined	10596	282	10314	87.5
PanLex (Kamholz et al., 2014)	5166	75	5091	11.4
Swadesh (Swadesh, 2015)	215	23	192	1.5
Statistically Aligned FloRes Data (Dyer et al., 2013)	6160	297	5863	86.5

Table 31: Part 2 —Total, functional, and content word counts across all lexicons. We also report the word type coverage, a measure of how extensively the lexicon documents a LRL. To calculate coverage percentage, we compared the CRL lexicon against the vocabulary size of the FloRes dev set.

Dictionary Source	Total Words	Functional Words	Content Words	Coverage (%)
Awadhi - Hindi				
Combined	10470	1356	9114	31.8
PanLex (Kamholz et al., 2014)	12	0	12	0.0
(Bafna et al., 2024a)	10462	1355	9107	31.8
Tagalog - Indonesian				
Combined	10176	284	9892	79.5
PanLex (Kamholz et al., 2014)	4916	76	4840	8.9
Swadesh (Swadesh, 2015)	211	20	191	1.5
Statistically Aligned FloRes Data (Dyer et al., 2013)	5801	253	5548	78.1
Egyptian Arabic - Standard Arabic				
Combined	10077	284	9793	93.4
PanLex (Kamholz et al., 2014)	2244	26	2218	4.5
Swadesh (Swadesh, 2015)	299	0	299	1.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	8203	317	7886	93.2
Maori - Indonesian				
Combined	9908	208	9700	78.0
PanLex (Kamholz et al., 2014)	6700	86	6614	16.5
Swadesh (Swadesh, 2015)	306	35	271	3.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	3851	172	3679	75.1
Moroccan Arabic - Standard Arabic				
Combined	9199	408	8791	93.1
PanLex (Kamholz et al., 2014)	464	2	462	0.7
Swadesh (Swadesh, 2015)	308	0	308	1.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	8590	445	8145	93.0
Najdi Arabic - Standard Arabic				
Combined	9147	117	9030	98.5
PanLex (Kamholz et al., 2014)	1	0	1	0.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	9146	115	9031	98.5
Ta'izzi-Adeni Arabic - Standard Arabic				
Statistically Aligned FloRes Data (Dyer et al., 2013)	8915	232	8683	96.1
Tunisian Arabic - Standard Arabic				
Combined	8813	325	8488	94.1
PanLex (Kamholz et al., 2014)	237	20	217	0.2
Swadesh (Swadesh, 2015)	205	0	205	0.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	8579	343	8236	94.1
Mesopotamian Arabic - Standard Arabic				
Combined	8705	275	8430	95.1
PanLex (Kamholz et al., 2014)	6	0	6	0.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	8702	282	8420	95.1
South Levantine Arabic - Standard Arabic				
Combined	8480	336	8144	93.8
PanLex (Kamholz et al., 2014)	300	24	276	1.0
Swadesh (Swadesh, 2015)	242	0	242	1.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	8196	343	7853	93.8
Javanese - Indonesian				
Combined	8439	384	8055	92.3
PanLex (Kamholz et al., 2014)	2382	79	2303	5.5
Swadesh (Swadesh, 2015)	219	24	195	0.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	6245	319	5926	92.0

Table 32: Part 3 —Total, functional, and content word counts across all lexicons. We also report the word type coverage, a measure of how extensively the lexicon documents a LRL. To calculate coverage percentage, we compared the CRL lexicon against the vocabulary size of the FloRes dev set.

Dictionary Source	Total Words	Functional Words	Content Words	Coverage (%)
North Levantine Arabic - Standard Arabic				
Statistically Aligned FloRes Data (Dyer et al., 2013)	8411	416	7995	93.6
Italian - Italian				
Statistically Aligned FloRes Data (Dyer et al., 2013)	8118	226	7892	98.3
Plateau Malagasy - Indonesian				
Combined	8106	433	7673	85.9
PanLex (Kamholz et al., 2014)	2420	85	2335	3.5
Statistically Aligned FloRes Data (Dyer et al., 2013)	5908	383	5525	85.7
Sundanese - Indonesian				
Combined	7640	418	7222	91.7
PanLex (Kamholz et al., 2014)	1245	51	1194	4.3
Swadesh (Swadesh, 2015)	301	40	261	1.7
Statistically Aligned FloRes Data (Dyer et al., 2013)	6633	415	6218	91.4
Cebuano - Indonesian				
Combined	6969	263	6706	77.6
PanLex (Kamholz et al., 2014)	1441	37	1404	3.3
Swadesh (Swadesh, 2015)	297	27	270	1.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	5582	238	5344	77.1
Morisyen - Haitian				
Combined	6513	754	5759	36.0
PanLex (Kamholz et al., 2014)	90	2	88	0.8
Swadesh (Swadesh, 2015)	218	28	190	3.3
Statistically Aligned Bible Data (McCarthy et al., 2020)	6384	766	5618	35.5
Pangasinan - Indonesian				
Combined	6473	416	6057	86.6
PanLex (Kamholz et al., 2014)	546	49	497	2.2
Statistically Aligned FloRes Data (Dyer et al., 2013)	6067	402	5665	86.4
Iloko - Indonesian				
Combined	6329	238	6091	76.6
PanLex (Kamholz et al., 2014)	743	30	713	2.9
Swadesh (Swadesh, 2015)	219	23	196	1.4
Statistically Aligned FloRes Data (Dyer et al., 2013)	5696	240	5456	76.3
Seselwa Creole French - Haitian				
Combined	6051	635	5416	29.1
PanLex (Kamholz et al., 2014)	103	2	101	0.7
Statistically Aligned Bible Data (McCarthy et al., 2020)	5983	643	5340	28.9
Chhattisgarhi - Hindi				
Combined	5293	476	4817	80.8
PanLex (Kamholz et al., 2014)	11	0	11	0.0
Statistically Aligned FloRes Data (Dyer et al., 2013)	5282	467	4815	80.8
Saint Lucian Creole French - Haitian				
Combined	4867	464	4403	50.6
PanLex (Kamholz et al., 2014)	85	3	82	1.0
Statistically Aligned Bible Data (McCarthy et al., 2020)	4796	462	4334	50.6
Samoan - Indonesian				
Combined	4471	155	4316	72.4
PanLex (Kamholz et al., 2014)	872	39	833	4.1
Statistically Aligned FloRes Data (Dyer et al., 2013)	3785	138	3647	71.8

Table 33: Part 4 —Total, functional, and content word counts across all lexicons. We also report the word type coverage, a measure of how extensively the lexicon documents a LRL. To calculate coverage percentage, we compared the CRL lexicon against the vocabulary size of the FloRes dev set.

Dictionary Source	Total Words	Functional Words	Content Words	Coverage (%)
Fijian - Indonesian				
Combined	4280	205	4075	77.5
PanLex (Kamholz et al., 2014)	630	46	584	4.1
Swadesh (Swadesh, 2015)	222	33	189	1.3
Statistically Aligned FloRes Data (Dyer et al., 2013)	3705	150	3555	77.0

Table 34: Part 5 —Total, functional, and content word counts across all lexicons. We also report the word type coverage, a measure of how extensively the lexicon documents a LRL. To calculate coverage percentage, we compared the CRL lexicon against the vocabulary size of the FloRes dev set.

feature	Aya		M2M	
	lrc	fi	lrc	fi
lrl_wiki_count	0.0	0.07	0.0	0.02
hrl_wiki_count	0.0	0.11	0.0	0.2
chrf_sim	-0.03	0.12	-0.07	0.26
baseline_performance	-0.3	0.48	-0.12	0.45
in_pretraining	-1.72	0.0	-1.67	0.0
token_fertility_ratio	-6.3	0.21	-3.66	0.07

Table 35: Random forest feature importance values and linear regression coefficients for different language features. lrc = linear regression coefficient; fi = feature importance value.

HRL	Natural %	Mean % over all CRLs
ind	24.5	43.1
arb	27.7	16.3
ita	40.3	38.5
tur	16.8	15.0
hin	39.7	33.4

Table 36: Comparing the natural percentage of functional words in an HRL against the mean % of words identified as function words over all its CRLs for D→M.

Language	% in -cont	% in -func	% in -all
Austronesian Language Family			
jav_Latn (Javanese)	42.9	32.2	75.1
sun_Latn (Sundanese)	40.2	31.7	71.9
smo_Latn (Samoan)	33.2	52.2	85.4
mrj_Latn (Maori)	41.2	46.8	87.9
ceb_Latn (Cebuano)	26.0	49.7	75.7
zsm_Latn (Standard Malay)	61.8	27.5	89.3
tgl_Latn (Tagalog)	31.1	46.5	77.6
ilo_Latn (Ilocano)	22.6	51.0	73.6
fij_Latn (Fijian)	35.5	50.4	85.9
plt_Latn (Plateau Malagasy)	36.4	39.3	75.7
pag_Latn (Pangasinan)	22.7	47.3	70.0
Average	35.8	43.1	78.9
Arabic Language Family			
acm_Arab (Mesopotamian Arabic)	40.2	17.9	58.1
acq_Arab (Ta'izzi-Adeni Arabic)	39.4	19.0	58.4
aeb_Arab (Tunisian Arabic)	43.9	13.4	57.3
ajp_Arab (South Levantine Arabic)	42.8	17.4	60.2
apc_Arab (North Levantine Arabic)	38.4	18.5	56.9
ars_Arab (Najdi Arabic)	39.6	20.4	60.0
ary_Arab (Moroccan Arabic)	42.4	15.9	58.4
arz_Arab (Egyptian Arabic)	54.9	7.6	62.6
Average	42.7	16.3	59.0
Romance Language Family			
spa_Latn (Spanish)	46.5	40.4	86.9
fra_Latn (French)	51.8	33.8	85.6
por_Latn (Portuguese)	52.2	32.1	84.3
ron_Latn (Romanian)	42.4	33.4	75.8
glg_Latn (Galician)	47.0	33.5	80.5
cat_Latn (Catalan)	43.2	39.3	82.5
oci_Latn (Occitan)	36.4	40.5	76.9
ast_Latn (Asturian)	37.9	37.9	75.8
lmo_Latn (Lombard)	24.6	45.7	70.4
vec_Latn (Venetian)	35.5	38.5	74.0
scn_Latn (Sicilian)	39.4	34.6	74.0
srd_Latn (Sardinian)	28.5	48.8	77.3
fur_Latn (Friulian)	41.2	38.8	80.0
lij_Latn (Ligurian)	34.5	41.1	75.6
Average	40.1	38.5	78.5
Turkic Language Family			
uzn_Latn (Northern Uzbek)	49.2	13.7	62.9
tuk_Latn (Turkmen)	45.7	14.8	60.5
azj_Latn (North Azerbaijani)	49.4	15.4	64.8
crh_Latn (Crimean Tatar)	41.9	16.1	58.1
Average	46.6	15.0	61.6
Indic Language Family			
hne_Deva (Chattisgarhi)	34.4	42.9	77.3
bho_Deva (Bhojpuri)	58.5	30.0	88.4
mag_Deva (Magahi)	53.8	35.4	89.2
mai_Deva (Maithili)	60.9	25.2	86.1
Average	51.9	33.4	85.2
Creole Language Family			
acf (Saint Lucian Patois)	43.0	39.3	82.3
mfe (Mauritian)	34.5	37.9	72.4
crs (Seychellois)	24.7	40.6	65.3
Average	34.1	39.3	73.3

Table 37: Percentage of words switched out in D→M approaches for each language

	off-the-shelf	-all-eng	-cont-eng	-func-eng
ind (Indonesian)	15.9	10.1	10.2	13.4
arb (Standard Arabic)	29.5	10.3	10.2	28.8
ita (Italian)	29.3	14.4	16.6	22.8
tur (Turkish)	12.5	3.7	4.1	7.6
hin (Hindi)	22.9	4.9	4.8	19.5

Table 38: D→M results for Aya-23 where low-resource words are swapped with English words. Experiment was performed on 300 FLORES examples.

	off-the-shelf	-all-eng	-cont-eng	-func-eng
ind (Indonesian)	13.5	11.1	11.0	11.0
arb (Standard Arabic)	20.9	9.3	9.3	19.9
ita (Italian)	19.6	14.3	14.3	16.8
tur (Turkish)	10.9	4.6	4.7	5.5
hin (Hindi)	17.3	5.4	5.4	12.9

Table 39: D→M results for M2M where low-resource words are swapped with English words. Experiment was performed on 300 FLORES examples.

Language	off-the-shelf	-cloud-cont	-cloud-func	-cloud-all
Austronesian Language Family				
ind_Latn (Indonesian)	38.9	-	-	-
jav_Latn (Javanese)	22.9	+0.4	+1.1	+0.7
sun_Latn (Sundanese)	22.9	-1.6	+1.5	-1.2
smo_Latn (Samoan)	2.3	+1.1	+1.6	+2.8
mri_Latn (Maori)	2.3	+2.4	+1.4	+2.8
ceb_Latn (Cebuano)	22.2	-1.9	-7.3	-9.1
zsm_Latn (Standard Malay)	38.7	-6.4	-2.2	-7.7
tgl_Latn (Tagalog)	27.8	-5.1	-10.8	-14.0
ilo_Latn (Ilocano)	13.4	+2.3	-1.9	-1.8
fij_Latn (Fijian)	2.3	+2.6	+1.4	+2.1
plt_Latn (Plateau Malagasy)	2.5	+5.2	+4.3	+5.1
pag_Latn (Pangasinan)	6.4	+2.5	+7.7	+6.5
Average	14.9	+0.1	-0.3	-1.3
Baseline < 25 Average	10.8	+1.4	+1.1	+0.9
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	29.6	-	-	-
acm_Arab (Mesopotamian Arabic)	22.4	-0.1	+1.5	+0.7
acq_Arab (Ta'izzi-Adeni Arabic)	24.7	-0.4	+0.7	+0.0
aeb_Arab (Tunisian Arabic)	17.7	-0.3	+2.0	+0.7
ajp_Arab (South Levantine Arabic)	24.5	-4.3	-1.1	-5.7
apc_Arab (North Levantine Arabic)	21.7	-4.0	+0.8	-3.0
ars_Arab (Najdi Arabic)	28.5	-0.3	-0.3	-0.2
ary_Arab (Moroccan Arabic)	11.9	+0.4	+3.6	+0.8
arz_Arab (Egyptian Arabic)	18.2	-1.9	+1.7	-1.8
Average	21.2	-1.4	+1.1	-1.1
Baseline < 25 Average	20.2	-1.5	+1.3	-1.2
Romance Language Family				
ita_Latn (Italian)	30.8	-	-	-
spa_Latn (Spanish)	26.8	-8.5	-3.9	-11.5
fra_Latn (French)	41.6	-13.4	-8.5	-20.1
por_Latn (Portuguese)	45.8	-13.5	-8.6	-20.9
ron_Latn (Romanian)	40.8	-13.9	-11.1	-22.0
glg_Latn (Galician)	37.5	-11.8	-5.8	-16.7
cat_Latn (Catalan)	42.7	-12.5	-8.3	-19.5
oci_Latn (Occitan)	44.6	-10.3	-4.6	-13.8
ast_Latn (Asturian)	34.3	-8.6	-5.7	-14.1
lmo_Latn (Lombard)	9.6	+9.2	+12.9	+11.1
vec_Latn (Venetian)	17.2	+5.2	+11.4	+7.9
scn_Latn (Sicilian)	7.3	+8.7	+13.1	+12.1
srd_Latn (Sardinian)	5.3	+10.2	+17.8	+18.0
fur_Latn (Friulian)	8.6	+10.9	+13.2	+14.3
lij_Latn (Ligurian)	9.8	+9.9	+15.9	+15.4
Average	26.6	-2.7	+2.0	-4.3
Baseline < 25 Average	9.6	+9.0	+14.1	+13.1
Turkic Language Family				
tur_Latn (Turkish)	33.3	-	-	-
uzn_Latn (Northern Uzbek)	1.6	+3.2	+1.8	+5.7
tuk_Latn (Turkmen)	2.4	+3.3	+2.2	+5.7
azj_Latn (North Azerbaijani)	7.4	+1.4	+2.4	+0.8
crh_Latn (Crimean Tatar)	8.5	+2.4	+2.9	+3.4
Average	5.0	+2.6	+2.3	+3.9
Baseline < 25 Average	5.0	+2.6	+2.3	+3.9
Indic Language Family				
hin_Deva (Hindi)	33.4	-	-	-
hne_Deva (Chattisgarhi)	15.2	+2.6	+14.3	+10.9
bho_Deva (Bhojpuri)	10.1	+3.9	+7.7	+8.1
mag_Deva (Magahi)	17.2	+5.3	+14.0	+13.2
mai_Deva (Maithili)	7.7	+8.4	+16.0	+16.8
Average	12.6	+5.1	+13.0	+12.2
Baseline < 25 Average	12.6	+5.1	+13.0	+12.2
Creole Language Family				
hat (Haitian)	39.7	-	-	-
acf (Saint Lucian Patois)	3.7	+4.0	+4.7	+9.1
mfe (Mauritian)	5.1	+2.8	+2.2	+3.6
crs (Seychellois)	11.2	+2.1	+4.5	+3.1
Average	5.9	+3.7	+4.6	+6.1
Baseline < 25 Average	5.9	+3.7	+4.6	+6.1
General Average	16.8	+1.1	+3.7	+1.3
General Baseline < 25 Average	10.9	+3.7	+6.2	+5.6

Table 40: $M \leftrightarrow D$ BLEU scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each $M \leftrightarrow D$ approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The best score for each language in the $M \leftrightarrow D$ paradigm is **bolded** and underlined.

Language	off-the-shelf	-cloud-cont	-cloud-func	-cloud-all
Austronesian Language Family				
ind_Latn (Indonesian)	41.5	-	-	-
jav_Latn (Javanese)	14.2	+7.6	+9.9	+12.2
sun_Latn (Sundanese)	11.0	+8.5	+12.3	+13.3
smo_Latn (Samoan)	3.2	+1.8	+1.9	+2.2
mri_Latn (Maori)	4.8	+2.0	-0.0	+1.8
ceb_Latn (Cebuano)	14.0	-0.3	-2.1	-0.7
zsm_Latn (Standard Malay)	36.5	-3.4	-0.5	-3.4
tgl_Latn (Tagalog)	25.6	-8.8	-13.0	-11.5
ilo_Latn (Ilocano)	7.5	+2.3	+2.4	+4.0
fij_Latn (Fijian)	3.3	+1.8	+0.9	+2.2
plt_Latn (Plateau Malagasy)	3.3	+2.5	+2.7	+4.6
pag_Latn (Pangasinan)	7.3	+3.8	+6.6	+6.4
Average	11.9	+1.6	+1.9	+2.8
Baseline < 25 Average	7.6	+3.3	+3.8	+5.1
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	38.1	-	-	-
acm_Arab (Mesopotamian Arabic)	31.6	-1.0	+0.7	-0.8
acq_Arab (Ta'izzi-Adeni Arabic)	33.7	-1.0	+0.2	-0.5
aeb_Arab (Tunisian Arabic)	26.7	-2.2	+0.4	-1.9
ajp_Arab (South Levantine Arabic)	37.4	-9.2	-4.7	-11.7
apc_Arab (North Levantine Arabic)	33.2	-6.2	-1.8	-7.5
ars_Arab (Najdi Arabic)	37.1	-0.2	+0.0	+0.0
ary_Arab (Moroccan Arabic)	23.5	-4.3	-0.8	-4.5
arz_Arab (Egyptian Arabic)	28.9	-4.5	-0.7	-5.5
Average	31.5	-3.6	-0.8	-4.0
Baseline < 25 Average	23.5	-4.3	-0.8	-4.5
Romance Language Family				
ita_Latn (Italian)	32.6	-	-	-
spa_Latn (Spanish)	29.0	-6.7	-1.6	-9.3
fra_Latn (French)	44.0	-14.4	-9.7	-18.6
por_Latn (Portuguese)	46.9	-14.1	-9.1	-18.0
ron_Latn (Romanian)	42.0	-14.7	-11.7	-20.3
glg_Latn (Galician)	35.5	-10.4	-4.5	-12.0
cat_Latn (Catalan)	36.5	-10.0	-5.5	-11.8
oci_Latn (Occitan)	35.1	-5.4	-0.8	-5.8
ast_Latn (Asturian)	31.8	-6.9	-4.6	-10.3
lmo_Latn (Lombard)	17.4	+3.9	+8.8	+5.1
vec_Latn (Venetian)	26.7	+0.9	+6.0	-0.2
scn_Latn (Sicilian)	16.7	+4.5	+7.8	+5.3
srd_Latn (Sardinian)	16.4	+4.6	+10.0	+8.2
fur_Latn (Friulian)	17.7	+6.9	+9.1	+8.6
lij_Latn (Ligurian)	17.4	+6.4	+13.2	+10.1
Average	29.5	-4.0	+0.5	-4.9
Baseline < 25 Average	17.1	+5.3	+9.8	+7.4
Turkic Language Family				
tur_Latn (Turkish)	31.7	-	-	-
uzn_Latn (Northern Uzbek)	3.8	+4.1	+3.9	+5.2
tuk_Latn (Turkmen)	6.1	+3.8	+3.9	+4.0
azj_Latn (North Azerbaijani)	9.8	-1.4	+1.3	-1.6
crh_Latn (Crimean Tatar)	11.9	+0.8	+3.3	+1.0
Average	7.9	+1.8	+3.1	+2.1
Baseline < 25 Average	7.9	+1.8	+3.1	+2.1
Indic Language Family				
hin_Deva (Hindi)	34.8	-	-	-
hne_Deva (Chattisgarhi)	23.3	-0.7	+6.3	+2.4
bho_Deva (Bhojpuri)	16.2	+1.0	+4.2	+3.4
mag_Deva (Magahi)	24.9	+1.9	+5.8	+5.2
mai_Deva (Maithili)	17.2	+4.3	+7.6	+7.3
Average	20.4	+1.6	+6.0	+4.6
Baseline < 25 Average	20.4	+1.6	+6.0	+4.6
Creole Language Family				
hat (Haitian)	8.9	-	-	-
acf (Saint Lucian Patois)	2.6	+6.0	+5.2	+6.9
mfe (Mauritian)	3.4	+4.7	+3.5	+3.4
crs (Seychellois)	5.7	+7.2	+9.6	+7.3
Average	7.3	+2.6	+2.8	+2.5
Baseline < 25 Average	7.3	+2.6	+2.8	+2.5
General Average	19.7	+0.4	+2.9	+0.6
General Baseline < 25 Average	11.3	+3.6	+5.7	+5.1

Table 41: $M \leftrightarrow D$ BLEU scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each $M \leftrightarrow D$ approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25. The best score for each language in the $M \leftrightarrow D$ paradigm is **bolded** and underlined.

Language	off-the-shelf	-c1oud-cont	-c1oud-func	-c1oud-all
Austronesian Language Family				
ind_Latn (Indonesian)	87.4	-	-	-
jav_Latn (Javanese)	70.1	+0.4	+1.3	+0.1
sun_Latn (Sundanese)	71.6	-1.7	+1.5	-1.9
smo_Latn (Samoan)	27.4	+12.4	+7.5	+15.5
mri_Latn (Maori)	25.7	+16.6	+10.4	+18.7
ceb_Latn (Cebuano)	64.2	-2.0	-5.6	-7.9
zsm_Latn (Standard Malay)	86.5	-6.3	-1.4	-7.6
tgl_Latn (Tagalog)	73.2	-6.0	-9.0	-13.5
ilo_Latn (Ilocano)	56.1	+2.6	-1.4	-1.2
fij_Latn (Fijian)	27.7	+17.3	+10.5	+16.5
plt_Latn (Plateau Malagasy)	44.3	+6.9	-0.4	+4.3
pag_Latn (Pangasinan)	35.5	+11.9	+15.9	+15.3
Average	52.9	+4.7	+2.6	+3.5
Baseline < 25 Average	46.9	+7.2	+4.4	+6.6
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	80.6	-	-	-
acm_Arab (Mesopotamian Arabic)	76.0	-0.5	+0.6	-0.4
acq_Arab (Ta'izzi-Adeni Arabic)	77.6	-1.1	+0.0	-1.1
aeb_Arab (Tunisian Arabic)	68.8	-0.8	+1.8	-0.3
ajp_Arab (South Levantine Arabic)	73.2	-5.1	+0.2	-5.7
apc_Arab (North Levantine Arabic)	72.8	-3.7	+1.1	-3.8
ars_Arab (Najdi Arabic)	80.0	-0.3	-0.3	-0.3
ary_Arab (Moroccan Arabic)	60.8	-0.4	+3.7	+0.1
arz_Arab (Egyptian Arabic)	71.4	-3.5	+1.1	-3.2
Average	72.6	-1.9	+1.0	-1.8
Baseline < 25 Average	71.5	-2.2	+1.2	-2.0
Romance Language Family				
ita_Latn (Italian)	86.3	-	-	-
spa_Latn (Spanish)	85.2	-14.0	-4.4	-18.0
fra_Latn (French)	87.8	-14.6	-5.2	-19.5
por_Latn (Portuguese)	88.0	-13.5	-4.1	-17.6
ron_Latn (Romanian)	88.2	-15.8	-6.5	-21.0
glg_Latn (Galician)	86.6	-14.0	-3.6	-18.2
cat_Latn (Catalan)	87.1	-13.2	-5.5	-18.9
oci_Latn (Occitan)	80.0	-9.2	-2.3	-11.4
ast_Latn (Asturian)	79.8	-11.6	-3.9	-15.4
lmo_Latn (Lombard)	44.5	+15.0	+18.5	+15.8
vec_Latn (Venetian)	54.0	+10.9	+16.5	+10.5
scn_Latn (Sicilian)	44.1	+15.3	+18.8	+16.5
srd_Latn (Sardinian)	37.9	+17.8	+23.4	+22.5
fur_Latn (Friulian)	42.5	+18.3	+20.7	+20.5
lij_Latn (Ligurian)	45.0	+14.5	+20.7	+18.3
Average	67.9	-1.0	+5.9	-2.6
Baseline < 25 Average	44.7	+15.3	+19.8	+17.3
Turkic Language Family				
tur_Latn (Turkish)	86.7	-	-	-
uzn_Latn (Northern Uzbek)	39.5	+13.7	+7.4	+17.0
tuk_Latn (Turkmen)	34.4	+12.6	+7.5	+15.5
azj_Latn (North Azerbaijani)	60.8	+2.6	+5.1	+1.4
crh_Latn (Crimean Tatar)	55.7	+4.3	+5.9	+5.2
Average	47.6	+8.3	+6.5	+9.8
Baseline < 25 Average	47.6	+8.3	+6.5	+9.8
Indic Language Family				
hin_Deva (Hindi)	86.2	-	-	-
hne_Deva (Chattisgarhi)	66.4	+3.7	+11.7	+8.1
bho_Deva (Bhojpuri)	63.6	+6.8	+9.8	+8.8
mag_Deva (Magahi)	69.0	+6.3	+11.3	+10.3
mai_Deva (Maithili)	57.9	+14.2	+18.3	+18.2
Average	64.2	+7.8	+12.8	+11.4
Baseline < 25 Average	64.2	+7.8	+12.8	+11.4
General Average	62.5	+2.1	+4.8	+1.8
General Baseline < 25 Average	54.6	+6.8	+8.1	+7.8

Table 42: $M \leftrightarrow D$ COMET scores by language for the model M2M. Performance gains/losses are relative to off-the-shelf. Averages for each $M \leftrightarrow D$ approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.

Language	off-the-shelf	-cloud-cont	-cloud-func	-cloud-all
Austronesian Language Family				
ind_Latn (Indonesian)	88.5	-	-	-
jav_Latn (Javanese)	67.5	+6.6	+8.2	+8.0
sun_Latn (Sundanese)	63.7	+8.9	+12.1	+11.1
smo_Latn (Samoan)	35.4	+11.4	+6.8	+10.5
mri_Latn (Maori)	40.5	+9.0	+0.9	+8.0
ceb_Latn (Cebuano)	61.1	+0.3	-3.8	-3.0
zsm_Latn (Standard Malay)	85.7	-3.2	-0.4	-3.5
tgl_Latn (Tagalog)	76.4	-9.4	-15.6	-14.7
ilo_Latn (Ilocano)	48.2	+9.1	+5.3	+7.1
fij_Latn (Fijian)	40.1	+9.7	+1.9	+7.5
plt_Latn (Plateau Malagasy)	46.1	+8.0	+2.9	+8.1
pag_Latn (Pangasinan)	43.1	+12.6	+14.2	+13.3
Average	55.2	+5.7	+3.0	+4.8
Baseline < 25 Average	49.5	+8.4	+5.4	+7.8
Arabic Language Family				
arb_Arab (Modern Standard Arabic)	86.1	-	-	-
acm_Arab (Mesopotamian Arabic)	83.2	+0.1	+0.9	-0.0
acq_Arab (Ta'izzi-Adeni Arabic)	84.3	+0.0	+0.6	-0.2
aeb_Arab (Tunisian Arabic)	78.4	-0.7	+1.4	-0.8
ajp_Arab (South Levantine Arabic)	83.7	-5.1	-0.8	-6.1
apc_Arab (North Levantine Arabic)	82.7	-3.4	+0.0	-4.4
ars_Arab (Najdi Arabic)	85.4	+0.8	+0.7	+0.6
ary_Arab (Moroccan Arabic)	74.0	-1.7	+1.6	-2.8
arz_Arab (Egyptian Arabic)	82.4	-3.6	-0.3	-4.6
Average	81.7	-1.7	+0.5	-2.3
Baseline < 25 Average	74.0	-1.7	+1.6	-2.8
Romance Language Family				
ita_Latn (Italian)	87.1	-	-	-
spa_Latn (Spanish)	86.1	-9.4	-2.1	-11.6
fra_Latn (French)	88.6	-9.4	-4.2	-13.3
por_Latn (Portuguese)	88.4	-8.2	-3.0	-11.3
ron_Latn (Romanian)	88.6	-10.6	-5.1	-15.0
glg_Latn (Galician)	85.2	-8.4	-1.8	-10.2
cat_Latn (Catalan)	84.0	-7.5	-2.5	-9.7
oci_Latn (Occitan)	74.9	-2.8	+2.2	-3.7
ast_Latn (Asturian)	77.7	-5.4	-1.3	-8.7
lmo_Latn (Lombard)	58.8	+7.5	+12.6	+8.1
vec_Latn (Venetian)	70.3	+1.4	+7.1	+0.1
scn_Latn (Sicilian)	59.5	+8.3	+11.8	+8.3
srd_Latn (Sardinian)	56.9	+8.7	+13.6	+10.3
fur_Latn (Friulian)	59.7	+9.8	+12.6	+10.5
lij_Latn (Ligurian)	56.9	+11.7	+16.5	+12.7
Average	74.0	-1.0	+4.0	-2.4
Baseline < 25 Average	58.4	+9.2	+13.4	+10.0
Turkic Language Family				
tur_Latn (Turkish)	87.1	-	-	-
uzn_Latn (Northern Uzbek)	56.5	+8.5	+7.7	+8.7
tuk_Latn (Turkmen)	51.2	+11.9	+12.5	+11.8
azj_Latn (North Azerbaijani)	70.5	-0.5	+4.0	-1.7
crh_Latn (Crimean Tatar)	64.8	+3.8	+7.7	+3.7
Average	60.7	+5.9	+8.0	+5.6
Baseline < 25 Average	60.7	+5.9	+8.0	+5.6
Indic Language Family				
hin_Deva (Hindi)	88.3	-	-	-
hne_Deva (Chattisgarhi)	77.2	-0.1	+4.3	+1.3
bho_Deva (Bhojpuri)	75.2	+0.9	+3.2	+2.0
mag_Deva (Magahi)	78.9	+2.1	+3.7	+3.0
mai_Deva (Maithili)	74.6	+3.7	+5.2	+4.7
Average	76.5	+1.6	+4.1	+2.8
Baseline < 25 Average	76.5	+1.6	+4.1	+2.8
General Average	65.9	+1.9	+4.1	+1.3
General Baseline < 25 Average	56.6	+6.7	+7.7	+7.3

Table 43: $M \leftrightarrow D$ COMET scores by language for the model Aya-23. Performance gains/losses are relative to off-the-shelf. Averages for each $M \leftrightarrow D$ approach are computed for each language family as well as for the general body of languages studied. These averages are recomputed for languages whose off-the-shelf BLEU score < 25.