

The University of Melbourne WMT 2026 CreoleMT Submission: A Domain-Balanced Approach to Low-Resource Pacific Creole Machine Translation

Raphaël Merx Nick Thieberger Ekaterina Vylomova
The University of Melbourne

Abstract

For our submission to the WMT26 Creole Language Translation Shared Task, we focus on machine translation (MT) models for Pacific creoles: Tok Pisin, Bislama, and Solomon Pijin, with particular attention to broad domain performance. After pre-training on a large collection of domain-imbalanced data, we continue fine-tuning on a diverse mix of domain-balanced data. We rely on a number of data collection and preparation techniques, including LLM-assisted respelling and alignment, back-translation, and distillation from Gemini for domains originally not present in training data. Evaluated on Bouquet and a novel test set made of spoken language transcripts, our models beat open model baselines by 3+ chrF++ points in all directions with human-original references. Looking ahead, we plan to develop human-translated test sets for Solomon Pijin and Bislama, and to distil our best models into much smaller ones that retain broad domain coverage.

1 Introduction

For this second WMT shared task on Creole machine translation (Robinson et al., 2026), we focus on the three main Pacific creoles: Tok Pisin (tpi, Papua New Guinea), Bislama (bis, Vanuatu), and Solomon Pijin (pis, the Solomon Islands). They share common characteristics: they are a lingua franca in their respective countries, with widespread use for daily communication (Eberhard et al., 2025); they are all English-based contact languages from the Melanesian pidgin family (Tryon and Charpentier, 2004); and for Tok Pisin and Bislama, are official languages of their country (Smith and Siegel, 2013; Meyerhoff, 2013).

Together, they are spoken by over 10 million people (Eberhard et al., 2025). However, they have relatively little MT coverage: Google Translate and NLLB only cover Tok Pisin (NLLB Team et al., 2022), with no consumer-facing model deployment

for Bislama or Solomon Pijin. Availability of test sets that cover a wide range of domains (e.g. beyond religion and news) is also limited, with the exception of the recent addition of Tok Pisin to Bouquet (Andrews et al., 2025).

As part of this shared task, we focus on training a single model that supports all three languages, with broad domain coverage. To this end, we collate a training set that includes a mix of parallel data from religious sites, but also synthetic data in the educational and informal domains. We rely on Bouquet tpi as our main test set, with the addition of in-house test sets that cover spoken transcripts for all three languages, provided by the Pacific Creole Project (Passmore et al., 2025).

We provide the following contributions:

- A single joint model per translation direction covering all three languages which we pre-trained on the broad but religion-heavy union of the data that we collected and then continued training on a smaller domain-balanced mix. On Bouquet and the held-out spoken transcripts, it beats every open model baseline by at least 3 chrF++ in all directions with human-original references.
- A demonstration that distilling a strong closed model is a cheap way to better cover some domains that scraped data lacks, and that the gain saturates early: the first 1.5k English→Tok Pisin pairs bring +1.54 chrF++, and the remaining 14.5k add +1.20.
- Evaluation using novel held-out spoken transcripts for all three languages, showing that performance in this domain is poorly represented by other test sets, and that adding just a few thousand rows of this domain to the training mix provides a substantial lift in it, while not moving performance in other test sets.
- Three findings that we believe generalise be-

yond this task: spelling normalisation brings large chrF++ gains, which may not matter as much to end-users as the numbers suggest; domain coverage added in one creole can carry to its neighbours; and joint training helps the two lowest-resource languages while leaving the highest-resource one unchanged.

2 Related Work

Creole MT. Creole languages have long been under-served by NLP, despite communities’ needs that often differ from those assumed by work that simply translates existing English resources (Lent et al., 2022b). Recent efforts have started to close this gap: CreoleVal assembles multitask benchmarks spanning 28 creoles (Lent et al., 2024), and Kreyòl-MT gathers the largest parallel corpus to date across 41 creoles, finding that a genre-diverse model can outperform a genre-specific one (Robinson et al., 2024). The first WMT shared task on creole MT (Robinson et al., 2025) brought these threads into a shared evaluation, of which the present task is the second edition. Past findings include: transfer from English is weaker than the surface similarity between the languages suggests (Lent et al., 2022a); and the available data is dominated by religious text, which transfers poorly to everyday language, while even a few hundred in-domain sentences noticeably lift translation quality (Rowe et al., 2025). This applies to the Pacific creoles we target, with Tok Pisin the only one to have a broad-domain test set.

Cross-domain low-resource MT. Neural MT quality degrades sharply outside its training domain, at times producing fluent but inadequate output (Koehn and Knowles, 2017). The problem is acute for low-resource languages, where the little parallel data that exists is concentrated in a few domains such as religion and news, spurring efforts to build corpora in other domains (Merx et al., 2025). Domain adaptation for NMT (Chu and Wang, 2018) addresses this with data-centric methods, such as selecting or synthesising in-domain data, and model-centric methods, such as continued fine-tuning. A common recipe first trains on all available, largely out-of-domain data and then fine-tunes on in-domain and back-translated data (Imankulova et al., 2019). Back-translating target-side monolingual text (Sennrich et al., 2016) and transfer from massively multilingual models (NLLB Team et al., 2022) are the standard levers

for widening coverage. We combine these: we start from a massively multilingual model, train it on all collected data, and then continue fine-tuning on a domain-rebalanced mix that adds back-translated news and synthetic educational and informal text to counter the religious skew.

3 Data

3.1 Collection

3.1.1 Test sets

We evaluate on three held-out sets, chosen to reach beyond the religious and news domains that dominate the available creole corpora.

Bouquet (854 sentences per language) (Andrews et al., 2025) is a broad-domain benchmark spanning everyday speech, how-to, narration, social-media posts and more. It covers Tok Pisin, but not the other two languages we target, so we build silver references for them using an LLM (Section 3.2).

FLORES-200 (2,009 sentences, Tok Pisin only) (NLLB Team et al., 2022) is the standard multilingual benchmark: general encyclopaedic and news prose with gold human translations, used as a second independent Tok Pisin reference point.

Spoken transcripts (300 sentences per language) (Passmore et al., 2025) are transcribed and translated speech from the Pacific Creole Project, a spoken register not found in other sets and originally produced creole-to-English. We hold out 300 rows per language for testing and fold the rest into training (Tok Pisin 10.0k, Bislama 4.5k, Solomon Pijin 1.6k).

Decontamination. Some of the CreoleMT shared task evaluation set overlaps public text that was in our training data. Before filtering, most English eval sentences already appeared in our data (75% for Tok Pisin, 95% for Bislama, 99% for Solomon Pijin) as did 66% of the Tok Pisin creole side (the Solomon Pijin and Bislama creole sides far less, 3% and 8%). We drop from every training corpus any pair that shares a sentence with an eval source, on either side, after aggressive normalisation; afterwards no eval sentence remains in our training files.

3.1.2 Training data

Table 1 gives the per-source counts. The data falls into three kinds.

Scraped parallel. English $\leftrightarrow$ creole pairs, consolidated, sentence-aligned and deduplicated per language, then decontaminated; this forms the training corpus: 472.7k pairs for Tok Pisin, 42.5k for Solomon Pijin and 80.8k for Bislama. It is heavily religious in domain, with JW300 and the Bible, with, for Bislama, 35.1k back-translated Vanuatu news. Verse-level scripture is joined on the canonical verse id, so those pairs are exact and need no alignment.

Monolingual, back-translated. Creole-only text machine-translated back into English (Gemma-4-31B or Gemini) to add colloquial and news register: Tok Pisin news (34.0k sentences) and web text (FineWeb-2, 18.9k); Vanuatu news and government pages (FineWeb-2, 2.5k documents, 35.2k pairs) for Bislama; and news for Solomon Pijin (33 documents). Most of the FineWeb-2 data however is not used for our best model training mix, after finding it was detrimental to translation quality (favouring English passthrough).

Synthetic (distilled). 6.0k English-to-Tok-Pisin pairs generated by a strong closed model (Gemini 3.5 Flash), sampled evenly from three everyday domains (narration, how-to, social posts) to supply the colloquial registers the scraped data lacks. Figure 1 measures how far this scales.

Usage. We train one joint Tok Pisin + Solomon Pijin + Bislama model per direction. Each first pre-trains on the broad decontaminated union of the data above (596k pairs: 472.7k Tok Pisin, 42.5k Solomon Pijin, 80.8k Bislama), then continues on a smaller on-domain mix: scripture and synthetic data for both directions, plus the Tok Pisin high-quality mix (forward) or held-out spoken transcripts (reverse). Tables 2 and 3 give the full recipe and per-stage scores.

3.2 Data preparation

Three cleaning steps, each applied where a language needed it and each measured as a single-factor ablation (Table 4).

Loanword respelling (Solomon Pijin). The scraped Pijin data, almost all from jw.org, spells English loanwords the English way (“first”, “change”), whereas everyday and test Pijin use local spellings (“fas”, “senis”). We use Gemma-4-31B, guided by a Pijin dictionary, to respell the training targets into the local convention: spelling only, same words

Source	tpi	pis	bis
<i>Pre-train</i>			
JW300 (OPUS)	371,822	–	–
Web scrape	72,319	13,809	17,142
NWT Bible	28,526	28,562	28,568
News, back-transl.	–	147	35,107
Total	472,667	42,518	80,817
<i>Continued fine-tune (eng$\rightarrow$creole)</i>			
NWT Bible	28,526	28,564	28,568
Book of Mormon	5,924	–	5,934
Church study library	26,469	–	27,614
jw.org articles	47,719	–	–
Watchtower magazine	39,210	–	–
News, back-transl.	31,388	–	–
tpi example sentences	3,373	–	–
Scraped parallel (re-used)	–	41,136	52,221
Gemini distill. (eng $\rightarrow$ tpi)	5,991	–	–
Total	188,600	69,700	114,337
<i>Continued fine-tune (creole$\rightarrow$eng)</i>			
NWT Bible	28,526	28,564	28,568
Book of Mormon	5,924	–	5,934
Church study library	26,469	–	27,614
Spoken transcripts	10,035	1,594	4,447
Gemini distill. (eng $\rightarrow$ tpi)	5,991	–	–
Total	76,945	30,158	66,563

Table 1: Training data mix, in number of sentences. Under eng $\rightarrow$ creole, the jw.org, *Watchtower*, back-translated news and example-sentence rows form the Tok Pisin high-quality mix.

and meaning. This was the largest single lever we found (+7.7 chrF++).

Alignment. Extract parallel sentences from parallel documents with an LLM aligner (Gemini). Separately, we pass the Tok Pisin corpus through a Gemma-4-31B judge (using vLLM) that scores each pair and drops the clearly-misaligned ones (490.8k judged, 463.8k kept). Verse-level Bible and Book of Mormon text is joined using verse id.

Silver test sets (Solomon Pijin, Bislama). Neither language has a Bouquet set, so we generate silver references with Gemini from the English source, the Tok Pisin reference, bilingual dictionary entries (SIL Webonary, 3.0k for Pijin; 4.3k for Bislama) and 20 curated example sentences per language. Because the reverse-direction source is then model-generated, we mark those cells in our results tables and do not read them as gold.

4 Models

Baselines. We compare against open multilingual systems that already cover at least one of our three languages: NLLB-200-3.3B (tpi only, NLLB

	English → Tok Pisin			Tok Pisin → English		
	Bouquet	FLORES	Transcript	Bouquet	FLORES	Transcript
NLLB-3.3B, base	46.7	43.1	37.8	39.2	52.4	28.7
Kreyòl-MT	45.6	46.7	36.3	37.7	50.9	29.7
CreoleM2M	40.8	37.9	32.8	32.2	36.1	26.1
Gemini 3.5 Flash [‡]	51.9	50.2	41.0	50.1	52.2	50.0
Gemma-4-31B, zero-shot	45.4	45.5	37.1	46.1	50.7	46.5
NLLB, pre-train (all)	46.9	43.0	38.0	45.4	53.3	36.8
+ continued FT	50.0 ↑3.1	45.7 ↑2.7	40.8 ↑2.8	47.6 ↑2.2	55.5 ↑2.2	36.1 ↓0.7
+ spoken transcr.	49.9 ↓0.1	45.5 ↓0.2	48.7 ↑7.9	47.5 ↓0.1	55.1 ↓0.4	53.9 ↑17.8
+ Gemini distill.*	51.5 ↑1.6	47.6 ↑2.1	47.7 ↓1.0	49.2 ↑1.7	54.6 ↓0.5	53.9 →0.0
Gemma-4-E4B, LoRA pre-train	42.1	38.4	34.1	—	—	—
+ cont FT + transcr. + distill.	47.7	42.4	44.9	—	—	—
Gemma-4-26B, LoRA pre-train	42.1	38.3	34.1	—	—	—
+ cont FT + transcr. + distill.	47.4	42.4	39.9	—	—	—

Table 2: Tok Pisin results, chrF++ on Bouquet (854 human-translated rows), FLORES-200 (2,009 rows) and held-out spoken transcripts (300 rows). **Bold** marks the best score in each column and every score within 1 chrF++ of it (our noise floor); the * row is the system we submit. Pre-training on all data barely moves the forward direction, but continued fine-tuning and Gemini distillation bring us level with Gemini 3.5 Flash, the strongest baseline, everywhere but FLORES. Table 3 reports the same systems and stages for Solomon Pijin and Bislama.

Team et al., 2022), MADLAD-400-3B (bis/pis only, Kudugunta et al., 2023), Kreyòl-MT (tpi only, Robinson et al., 2024) and CreoleM2M (all 3, Lent et al., 2024). We also report Gemma-4-31B zero-shot and, as a closed frontier model, Gemini 3.5 Flash. Among the systems small enough to fine-tune on a single GPU, NLLB-3.3B is the strongest base on Tok Pisin, beating others on Bouquet and on the spoken transcripts in both directions (except Tok Pisin→English transcripts, where it sits 1 chrF++ below Kreyòl-MT); Gemma-4-31B is competitive zero-shot but too large to fine-tune on one GPU. We therefore select NLLB-3.3B for fine-tuning.

Joint multilingual fine-tuning. We add Solomon Pijin and Bislama target tokens to NLLB (initialised from the Tok Pisin token) and fine-tune one joint model per direction on all three languages, using the clean data from Section 3.2. We also add new tokens for punctuation missing in the NLLB tokenizer (e.g. curly quotes, en/em-dashes). Training the related languages together beats per-language fine-tunes on the two minority languages (Table 4).

Continued fine-tuning on diverse mix. Following the joint pre-training above, we continue with fine-tuning on a smaller but more domain-diverse mix. Tables 2 and 3 report the per-stage staircase.

Limited Gemini distillation. Because Gemini 3.5 Flash performs well on tpi zero-shot, and be-

cause we find very limited domain diversity in scraped data, we distill Gemini 3.5 Flash on a couple thousand sentences for each of social media text (TweetEval, Barbieri et al., 2020), how-to articles (WikiLingua, Ladhak et al., 2020) and literature (ROCStories, Mostafazadeh et al., 2016). For each, we split into sentences, and run Gemini forward from English to Tok Pisin. We report on the benefit of this distillation in Figure 1.

Ablations. Table 4 isolates each part of our recipe, varying one factor and holding the rest fixed. Gains are measured on Bouquet chrF++ for the en-xx direction. We also ran a number of data ablations, and ended up using only a limited subset of the data available to us, guided by evaluation on Bouquet chrF++. Data ablations however are mostly not reported in this table.

5 Results

Pretrain. NLLB-3.3B (base) sits 5.2 chrF++ behind Gemini 3.5 Flash on Tok Pisin Bouquet (46.7 vs 51.9). The first stage (broad pre-train on all data) adds almost nothing on Tok Pisin in the forward direction (46.7 to 46.9), but does in the reverse (+6.2, 39.2 to 45.4). For the other two languages, the pre-train is enough to beat every baseline we compare against: 52.8 on Pijin Bouquet, 13.6 above CreoleM2M, and 52.9 on Bislama, 8.4 above it. These Bouquet references are silver (Section 3.2) and should be read with care, but the spoken transcripts are human-original, and they show gains of the

	eng → creole		creole → eng	
	Bouq.	Transc.	Bouq.	Transc.
Solomon Pijin (pis)				
NLLB-3.3B, base	23.9	19.3	40.8	28.3
MADLAD-400-3B	21.4	19.2	40.0	31.6
CreoleM2M	39.2	29.5	41.5	31.1
Gemma-4-31B, zero-shot	38.3	29.2	62.3	51.4
NLLB, pre-train (all)	52.8	36.6	62.5	45.8
+ continued FT	52.8	37.5	61.7	43.1
+ spoken transcr.	52.5	44.2	61.8	55.0
+ Gemini distill.*	53.1	44.0	63.0	55.6
Gemma-4-E4B, LoRA pre-train	45.1	33.2	—	—
+ cont FT + transcr. + distill.	46.3	37.9	—	—
Gemma-4-26B, LoRA pre-train	45.0	33.1	—	—
+ cont FT + transcr. + distill.	46.1	33.9	—	—
Bislama (bis)				
NLLB-3.3B, base	31.7	25.3	38.3	23.6
MADLAD-400-3B	25.7	20.0	40.1	23.3
CreoleM2M	44.5	35.7	40.5	28.4
Gemma-4-31B, zero-shot	51.0	40.3	57.0	46.3
NLLB, pre-train (all)	52.9	44.3	59.0	45.0
+ continued FT	56.4	45.1	59.3	40.3
+ spoken transcr.	56.0	57.8	59.3	54.0
+ Gemini distill.*	57.8	57.5	60.5	53.8
Gemma-4-E4B, LoRA pre-train	48.6	41.3	—	—
+ cont FT + transcr. + distill.	53.0	51.4	—	—
Gemma-4-26B, LoRA pre-train	47.8	39.7	—	—
+ cont FT + transcr. + distill.	52.5	45.0	—	—

Table 3: Solomon Pijin and Bislama results, chrF++; same systems as Table 2. Note Bouquet references are Gemini-generated from the Tok Pisin reference (+ lexicon and example sentences, Section 3.2). Spoken transcripts (300 rows/lang) are human-original in both directions. Unlike Tok Pisin, these two languages gain most from the broad pre-train, which already beats every baseline.

same size over CreoleM2M (+7.1 for Pijin, +8.6 for Bislama), so we read the improvement as real.

Continued fine-tuning. The three continued fine-tuning stages (domain-balanced data, spoken transcripts, and Gemini distillation) close the gap with Gemini-3.5-Flash (the best model we measured for Tok Pisin): our shipped model scores 51.5 against Gemini’s 51.9 on Bouquet, and pulls ahead on the spoken domain, 47.7 against 41.0 on transcripts. FLORES is the exception, where we stay 2.6 points behind (47.6 against 50.2), but FLORES’ age and widespread use may mean contamination in recent models (Tan et al., 2026). In the reverse direction the shipped model is 0.9 below Gemini on Bouquet (49.2 against 50.1) and 3.9 above it on transcripts (53.9 against 50.0).

Ablation	Lang	w/o → w/	Δ
Loanword respelling	pis	45.8 → 53.4	+7.7
Continued FT	tpi	46.9 → 50.0	+3.1
Gemini distil. (6k pairs)	tpi	48.5 → 50.9	+2.4
Joint training	bis	50.7 → 52.5	+1.8
	pis	44.4 → 45.8	+1.4
	tpi	45.6 → 46.3	+0.7
Alignment filtering	tpi	46.0 → 46.3	+0.3

Table 4: Ablations, sorted by effect size; each varies one factor and holds the rest fixed. Scores are Bouquet en-xx chrF++. Respelling and alignment filtering use Gemma-4-31B for cleaning. Continued FT is the two-stage pre-train→continued-FT recipe from Table 2. Gemini distillation adds 6k synthetic eng→tpi pairs (the 0 and 6k points of Figure 1). The joint-training rows compare each language’s single-language NLLB (w/o) against the joint model (w/) on raw-spelling data, so a positive Δ means joint training helps that language.

Fine-tuning LLMs: promising but expensive.

We also try LoRA fine-tuning of Gemma-4 using the same two-stage recipe, at two sizes (E4B and 26B). Both land below the NLLB fine-tunes: 47.7 and 47.4 on Tok Pisin Bouquet against 51.5 for our shipped model, with the same ordering on Pijin (46.3 and 46.1 against 53.1) and Bislama (53.0 and 52.5 against 57.8). Going from E4B to 26B does not improve the scores, which suggests the limit is the LoRA recipe, not model capacity. Gemma-4-31B, on the other hand, is competitive with no fine-tuning at all: zero-shot it reaches 45.4 on Tok Pisin, just below base NLLB, 51.0 on Bislama, above every open baseline we test. Fine-tuning a model of that size therefore looks like a promising direction, but it does not fit our single-GPU budget, so we leave it to future work.

Domain diversity. Breaking Tok Pisin Bouquet down by domain shows where the base model falls short and where fine-tuning helps most. Out of the box, NLLB trails Gemini most on the informal, conversational registers that our scraped corpora (religion and news) barely contain: it is 7.4 chrF++ behind on conversation (48.4 against 55.9), 6.5 on user comments, 5.3 on narration and 5.0 on social posts, against only 2.5 on reflective prose. Our fine-tune recovers most of this gap: it now beats Gemini on how-to (52.9, 2.4 above) and draws level on conversation (56.3 against 55.9) and social posts (51.5 against 52.2). The registers where we still trail are the ones our training data covers least, and only by 1 to 3 points: comments (47.0 against 49.8) and narration (52.1 against 53.6). The remaining




Figure 1: Translation quality gains from adding Gemini distillation in ChrF++ points (top) / ESA through LLM-judge (bottom). Gains are measurable with only 6k sentences (2k each of TweetEval / WikiLingua / ROC-Stories), then plateau.

half-point deficit to Gemini on Bouquet is thus concentrated in a few under-covered registers rather than spread evenly.

5.1 Discussion

Distillation is cheap and beneficial. Figure 1 scales the Gemini distillation set from 0 to 16k pairs, holding the 122.4k human tpi mix and the recipe fixed. The curve rises steeply and then flattens: the first 1.5k pairs are worth +1.54 chrF++, more than half of the +2.73 that 16k pairs buy in total, and the remaining 14.5k pairs add only +1.20 between them. No single increment past the first is individually resolvable (1.5k→3k, 3k→6k, 6k→12k and 12k→16k all have 95% intervals spanning zero), yet the accumulated 1.5k→16k gap is significant, so the plateau slopes gently upward rather than being flat. The takeaway is that distillation from a strong closed model is cheap to exploit and quick to exhaust: a few thousand well-chosen source sentences capture most of the available gain, and buying an order of magnitude more yields roughly one further chrF++ point. Beside ChrF++, we also score with a reference-based LLM judge (GPT-5.6-luna, ESA-style 0–100 with error spans). The LLM-judge also measures translation quality gains: +4.3 ESA from 0 to 6k with a 95% interval clear of zero, and a further +1.0 to 16k that is not resolvable.

Adding transcripts reduces over-translation of spoken text. On the (in-domain) spoken transcript test set, the reverse model trained without transcripts tends to render terse spoken text as more formal, idiomatic English, whereas the transcript-augmented model mirrors the literal register of the references. For example, for *stap klostu wantaim mi olgeta taim* (ref. “stay close to me at all times”), the no-transcript model outputs the idiomatic “always by my side”, while the +transcripts model gives the literal “stay close to me all the time”.

The outsize weight of orthography. Spelling consistency between training data and reference drives chrF++ to an extent that may not reflect end-user priorities. Respelling *pis* to the reference convention gained +7.7 (45.8 to 53.4), while the two languages left unrespelled in the same run barely moved (*bis* −0.2, *tpi* +0.3). Adding more religious-register data then helped *bis*, whose orthography matches (+2.1), but hurt *pis* (−3.6), whose text uses a different yet valid convention (*bi long* rather than *blong*): a 5.7 chrF change from spelling alone, against +2.0 for our best modelling change. As *bi long* is perfectly comprehensible Pijin, and the dominant spelling in the Pijin religious publications we scrape, we question whether chrF++ measures quality or mere conformity to the reference’s spelling, and whether optimising for it trades away subtler gains in terminology or register.

Domain coverage added in one creole may carry to its neighbours. The final stage adds a tpi-only Gemini distillation set. Beyond the +1.6 it adds on Tok Pisin Bouquet (49.9 to 51.5), it also improves the silver **bis** Bouquet set by +1.8 (56.0 to 57.8) and nudges **pis** up +0.6 (52.5 to 53.1), although neither receives new training data at that stage. A plausible explanation is that broadening domain coverage in one language generalises to the other two sharing the joint model. Two caveats here: the added tpi pairs also enlarge the joint mix, so every language sees marginally more optimisation steps, and the silver *bis* references are also Gemini-generated, as is the distillation data, so some of the gain may be convergence toward the reference author’s style rather than improved translation.

Joint training mainly helps the two lowest-resource languages. Fine-tuning the three creoles together rather than separately (holding data and recipe fixed) gains +1.8 for *bis* (50.7 to 52.5) and +1.4 for *pis* (44.4 to 45.8), but only +0.7 for

tpi (45.6 to 46.3), inside our 1 chrF++ noise floor (Table 4). We hold far more tpi data, and across more domains, than pis or bis (463.8k cleaned pairs against 83k for bis and 44.7k for pis), so joint training lets the two data-poor languages draw on the larger and more varied tpi corpus through the shared model, while tpi is already well covered on its own and has little left to gain. This is the cross-creole transfer of the previous finding seen from the data side, and it is why we train one joint model rather than three: the low-resource languages benefit and the high-resource one is not harmed.

6 Conclusion

In this work, we develop a first model that focuses on Pacific creoles (Tok Pisin, Solomon Pijin, Bislama), targeting good performance across domains. We show that pre-training on an unbalanced mix of data, then fine-tuning on a more balanced and diverse mix can bring substantial translation quality gains across domains, and that these cross-domain gains can transfer across languages (e.g. tpi diverse bringing quality gains to pis/bis). We also show the value of forward-distillation from a strong LLM, on domains that are otherwise unrepresented in available training data.

In future work, we plan to develop more reliable test sets for Solomon Pijin and Bislama, as we feel like Gemini-generated test sets, even though they are guided by Tok Pisin ground truth and a lexicon, are unsatisfactory. We also plan to develop a very small (< 50M) model that can translate across 6 directions, and that has good performance across domains.

References

- Pierre Andrews, Mikel Artetxe, Mariano Coria Meglioli, Marta R. Costa-jussà, Joe Chuang, David Dale, Mark Duppenthaler, Nathaniel Paul Ekberg, Cynthia Gao, Daniel Edward Licht, Jean Maillard, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Eduardo Sánchez, Ioannis Tsiamas, Arina Turkatenko, Albert Ventayol-Boada, and Shireen Yates. 2025. [BOUQuET : dataset, benchmark and open initiative for universal quality evaluation in translation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27515–27535, Suzhou, China. Association for Computational Linguistics.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Chenhui Chu and Rui Wang. 2018. [A survey of domain adaptation for neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1304–1319, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2025. *Ethnologue: Languages of the Americas and the Pacific*, 28th edition. SIL International, Dallas, Texas. Online: <https://www.ethnologue.com>.
- Aizhan Imankulova, Raj Dabre, Atsushi Fujita, and Kenji Imamura. 2019. [Exploiting out-of-domain parallel data through multilingual transfer learning for low-resource neural machine translation](#). In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 128–139, Dublin, Ireland. European Association for Machine Translation.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Christopher A. Choquette-Choo, Katherine Lee, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. [MADLAD-400: A multilingual and document-level large audited dataset](#). *Advances in Neural Information Processing Systems (NeurIPS)*.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. [WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048, Online. Association for Computational Linguistics.
- Heather Lent, Emanuele Bugliarello, and Anders Søgaard. 2022a. [Ancestor-to-creole transfer is not a walk in the park](#). In *Proceedings of the Third Workshop on Insights from Negative Results in NLP*, pages 68–74, Dublin, Ireland. Association for Computational Linguistics.
- Heather Lent, Kelechi Ogueji, Miryam de Lhoneux, Orevaoghene Ahia, and Anders Søgaard. 2022b. [What a creole wants, what a creole needs](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6439–6449, Marseille, France. European Language Resources Association.
- Heather Lent, Kushal Tatariya, Raj Dabre, Yiyi Chen, Marcell Fekete, Esther Ploeger, Li Zhou, Ruth-Ann Armstrong, Abee Eijansantos, Catriona Malau, Hans Erik Heje, Ernests Lavrinovics, Diptesh Kanojia, Paul Belony, Marcel Bollmann, Loïc Grobol, Miryam de Lhoneux, Daniel Herscovich, Michel

- DeGraff, and 2 others. 2024. [CreoleVal: Multilingual multitask benchmarks for creoles](#). *Transactions of the Association for Computational Linguistics*, 12:950–978.
- Raphael Merx, Hanna Suominen, Trevor Cohn, and Ekaterina Vylomova. 2025. [OpenWHO: A document-level parallel corpus for health translation in low-resource languages](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 142–160, Suzhou, China. Association for Computational Linguistics.
- Miriam Meyerhoff. 2013. [Bislama](#). In Susanne Maria Michaelis, Philippe Maurer, Martin Haspelmath, and Magnus Huber, editors, *The Survey of Pidgin and Creole Languages, Volume 1: English-based and Dutch-based Languages*. Oxford University Press, Oxford.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- 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, and 20 others. 2022. [No language left behind: Scaling human-centered machine translation](#). *arXiv preprint arXiv:2207.04672*.
- Sam Passmore, Lila San Roque, Kirsty Gillespie, Saurabh Kumar Nath, Kira Davey, Keira Mullan, Tim Cawley, Jennifer Biggs, Rosey Billington, Bethwyn Evans, Nick Thieberger, Nicholas Evans, and Danielle Barth. 2025. [English-based acoustic models perform well in the forced alignment of two English-based pacific creoles](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31172–31183, Vienna, Austria. Association for Computational Linguistics.
- Nathaniel Robinson, Ruth-Ann Armstrong, Claire Bizon Monroc, Rasul Dent, Pranav Gupta, Raj Dabre, Andre Coy, and Kenton Murray. 2026. Findings of the second shared task for creole language machine translation, at WMT 2026. In *Proceedings of the Eleventh Conference on Machine Translation*, Budapest, Hungary. Association for Computational Linguistics.
- Nathaniel Robinson, Claire Bizon Monroc, Rasul Dent, Stefan Watson, Kenton Murray, Raj Dabre, Andre Coy, and Heather Lent. 2025. [Findings of the first shared task for creole language machine translation at WMT25](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 520–531, Suzhou, China. Association for Computational Linguistics.
- Nathaniel R. Robinson, Raj Dabre, Ammon Shurtz, Rasul Dent, Onenamiyi Onesi, Claire Bizon Monroc, Loïc Grobol, Hasan Muhammad, Ashi Garg, Naome A. Etori, Vijay Murari Tiyyala, Olanrewaju Samuel, Matthew Dean Stutzman, Bismarck Bamfo Odoom, Sanjeev Khudanpur, Stephen D. 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.
- Jacqueline Rowe, Edward Gow-Smith, and Mark Hepple. 2025. [Limitations of religious data and the importance of the target domain: Towards machine translation for Guinea-Bissau creole](#). In *Proceedings of the Eighth Workshop on Technologies for Machine Translation of Low-Resource Languages (LoResMT 2025)*, pages 183–200, Albuquerque, New Mexico, U.S.A. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Geoff P. Smith and Jeff Siegel. 2013. [Tok pisin](#). In Susanne Maria Michaelis, Philippe Maurer, Martin Haspelmath, and Magnus Huber, editors, *The Survey of Pidgin and Creole Languages, Volume 1: English-based and Dutch-based Languages*. Oxford University Press, Oxford.
- David Tan, Pinzhen Chen, Josef van Genabith, and Koel Dutta Chowdhury. 2026. [When Flores bloomz wrong: Cross-direction contamination in machine translation evaluation](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 345–358, Rabat, Morocco. Association for Computational Linguistics.
- Darrell T. Tryon and Jean-Michel Charpentier. 2004. *Pacific Pidgins and Creoles: Origins, Growth and Development*. Trends in Linguistics. Studies and Monographs 132. Mouton de Gruyter, Berlin.