

VietMix: A Naturally-Occurring Parallel Corpus and Augmentation Framework for Vietnamese-English Code-Mixed Machine Translation

Hieu Tran^{*,1}, Phuong-Anh Nguyen-Le^{*,1}, Huy Nghiem¹, Quang-Nhan Nguyen²,
Wei Ai¹, Marine Carpuat¹

¹ University of Maryland, College Park, ² Harvard University

¹{hieutt, nlpa, nghiemh, aiwei, marine}@umd.edu

²quangnhan_nguyen@fas.harvard.edu

Abstract

Machine translation (MT) systems universally degrade when faced with code-mixed text. This problem is more acute for low-resource languages that lack dedicated parallel corpora. This work directly addresses this gap for Vietnamese-English, a language context characterized by challenges including orthographic ambiguity and the frequent omission of diacritics in informal text. We introduce VIETMIX, the first expert-translated, naturally occurring parallel corpus of Vietnamese-English code-mixed text. We establish VIETMIX's utility by developing a data augmentation pipeline that leverages iterative fine-tuning and targeted filtering. Experiments show that models augmented with our data outperform strong back-translation baselines by up to +3.5 xCOMET points and improve zero-shot models by up to +11.9 points. Our work delivers a foundational resource for a challenging language pair and provides a validated, transferable framework for building and augmenting corpora in other low-resource settings.

1 Introduction

The proliferation of code-mixing, the blending of multiple languages within a single utterance (Poplack, 1980), is dominant in global digital communication. While this reflects the realities of multilingualism, it severely degrades the performance of state-of-the-art machine translation (MT) systems, especially those overwhelmingly trained on non-code-mixed formal text (Zhang et al., 2023). In low-resource language pairs, the absence of large-scale code-mixed parallel corpora creates a cycle of poor model performance and data scarcity. Benchmark datasets have emerged for medium-resource pairs, e.g., Sheth et al. (2025), but low-resource communities remain underserved.

Throughout this paper, we adopt the sociolinguistic term "code-mixing" to refer specifically to intra-

sentential blending of linguistic units (Bokamba, 1989). While NLP works often use "code-mixing" as an umbrella term to cover all forms of language alternation (Doğruöz et al., 2021; Winata et al., 2023), our choice of "code-mixing" more precisely describes the hybridization phenomena that characterize our corpus.

This paper focuses on Vietnamese-English, where a confluence of unique linguistic properties amplifies the code-mixing challenge. First, the two share significant orthographic ambiguity. For instance, the token `ban` is the English word meaning ("prohibition") but corresponds to at least five distinct Vietnamese words with different tones (bạn (friend), bản (village), bàn (table), bán (to sell), ban (to grant)), rendering context-free language identification nearly impossible. Second, informal digital text frequently omits the diacritics essential for disambiguating meaning in Vietnamese. Third, the prevalence of English-origin proper nouns and named entities further confounds this task of identifying true code-mixing. Automatic data curation, therefore, becomes exceptionally difficult to achieve and, until now, has hindered the development of a dedicated natural corpus.

Addressing this, we introduce VIETMIX, the first expert-translated parallel corpus of naturally occurring Vietnamese-English code-mixing (§ 3). Our core contribution lies in both this resource and the validated methodology for its creation. VIETMIX contains over 10K code-mixed samples collected from Threads.com, with a held-out test set of 1K sentences translated and adjudicated by linguistic experts to serve as a gold-standard evaluation.

Accompanying VIETMIX, we developed a three-stage data augmentation pipeline to establish strong baselines and provide a transferable framework for the community (§ 4). Critically, we show that augmenting models with our filtered, synthetic data provides consistent and significant gains over zero-shot models by up to +11.9 xCOMET points and

*These authors contributed equally to this work.

strong back-translation baselines by up to +3.5 points. This proves the value of our quality-focused curation approach. Experimental validation across multiple model architectures (§ 5) confirms that this methodology provides more faithful translations, especially in capturing pragmatic and idiomatic nuances of code-mixed language (§ 6).

We release the held-out test set under a Data Use Agreement (non-commercial, research-only). Code and preprocessing recipes are publicly accessible.¹

2 Related Work

Our work intersects three research areas: code-mixed MT, resource creation for low-resource languages, and synthetic data generation.

Code-mixed MT. The challenge that mixed-language text poses to NLP systems is well-documented, with initial works focusing on core tasks like LID and part-of-speech tagging (Solorio and Liu, 2008; Li and Fung, 2012). Centralized benchmarks like GLUECoS (Khanuja et al., 2020) and LinCE (Aguilar et al., 2020) enabled systematic evaluation of model degradation on code-mixed inputs. These efforts have shown that even large multilingual models struggle to transfer non-code-mixed capabilities to code-mixed scenarios.

Within this landscape, MT presents distinct challenges, primarily in the scarcity of parallel data. Consequently, research has concentrated on a few relatively high-resource language pairs, with Hindi-English being the most prominent with annotated datasets (Sheth et al., 2025) and specialized augmentation techniques (Kartik et al., 2024). While work has expanded to low-resource pairs like Kazakh-Russian (Borisov et al., 2025) and dialects of Mandarin (Lu et al., 2022), these efforts remain more limited in scope and do not address unique challenges in Southeast Asian tonal languages. Our work addresses this linguistic and geographic gap.

Parallel Corpora for Vietnamese MT. The ecosystem of NLP resources for Vietnamese has matured significantly from foundational work on tokenization and syntactic parsing (Vu et al., 2003; Nguyen et al., 2006) to the comprehensive Vn-CoreNLP toolkit (Vu et al., 2018). In MT, large-scale multi-domain corpora like PhoMT (Doan et al., 2021) and MTet (Ngo et al., 2022) have been instrumental in advancing translation quality for non-code-mixed Vietnamese-English. Re-

cent efforts focused on domain specialization, namely, medical texts (Vo et al., 2024) and other language pairs like Japanese-Vietnamese (Do, 2021) or Korean-Vietnamese (Nguyen et al., 2019).

Despite the pervasiveness of code-mixing in Vietnamese social media, prior work on Vietnamese code-mixing has been limited to small-scale and non-parallel speech corpora (Nguyen and Bryant, 2020) or monolingual text mining (Nguyen et al., 2024a; Hoang et al., 2023). VIETMIX is the first resource to fill this gap and provides a gold-standard, expert-translated benchmark to catalyze research on this phenomenon.

Synthetic Data for Low-Resource MT. Given high costs in creating parallel corpora, researchers have long turned to monolingual and comparable web corpora as a source of training data. Early statistical MT methods mined noisy parallels or lexicons to bootstrap low-fidelity translation systems before iteratively adding comparable monolingual corpora for retraining (Abdul-Rauf and Schwenk, 2009; Do et al., 2010). More recent neural methods have refined this process using multilingual embeddings and self-supervised retraining to extract parallel data with minimal supervision (Keung et al., 2020; Shi et al., 2022).

LLMs introduce a new paradigm for synthetic data generation as they can output vast amounts of parallel data through both forward- and back-translation (Nguyen et al., 2024b; Lim et al., 2024), with back-translation often yielding higher adequacy improvements (Moslem et al., 2022). However, naive back-translation can distort code-mixing patterns or "correct" them into non-code-mixed text and compound errors at language-switching boundaries. This failure highlights the need for more robust, quality-aware augmentation protocols. Our work builds on this precedent by proposing a forward-translation pipeline coupled with a multi-stage filtering mechanism, combining heuristic constraints and a neural quality classifier, to ensure syntactic plausibility and pragmatic appropriateness of the generated synthetic code-mixed.

3 VIETMIX Curation and Annotation

This section details the multi-stage process for creating VIETMIX. Our methodology emphasizes rigorous data filtering, expert human translation, and transparent translation quality control.

¹<https://github.com/umd-vietnlp/code-switching>

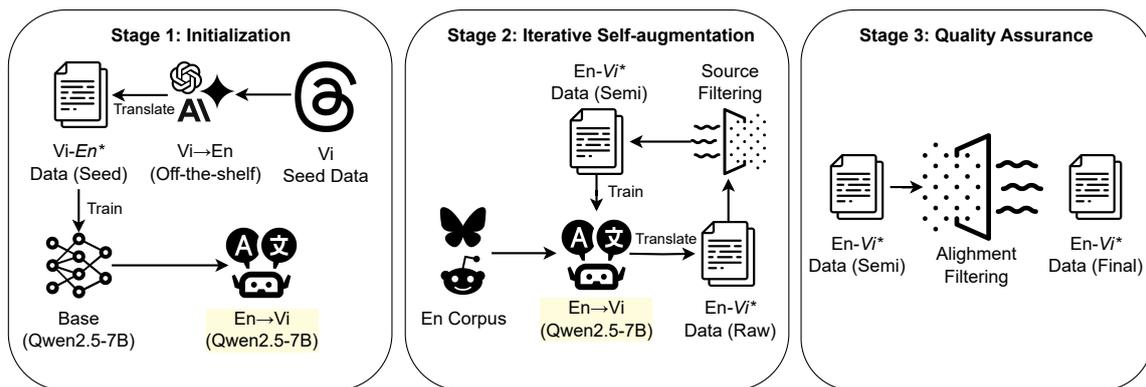


Figure 1: Overview of our three-stage iterative data augmentation pipeline converting naturally occurring code-mixed Vi texts into synthetic parallel Vi-En corpora. **Stage 1** focuses on seed data generation from naturally occurring Vi text and base model (En→Vi translation) bootstrapping. **Stage 2** involves iterative synthetic Vi generation from En corpora and progressive fine-tuning with filtering. **Stage 3** implements final quality assurance to produce the augmented parallel dataset. *Italicized language code* denotes synthetic/generated data.

3.1 Data Sourcing

VIETMIX consists of public posts from the social media Threads.com between October 2024 and February 2025. Adhering to strict data minimization and ethical research principles, we scraped publicly accessible content without user authentication using public APIs.

3.2 Data Preprocessing and Filtering

To ensure the quality of the corpus, the raw data underwent a comprehensive filtering pipeline to isolate relevant code-mixed content. We first applied several heuristic filters, removing posts that were overly short (≤ 10 words), contained links, consisted solely of emojis or special symbols.

Following this initial cleanup, a crucial step was to perform sentence-level Language Identification (LID) to accurately identify true code-mixing. This is a non-trivial task, as off-the-shelf LID tools are often unreliable for Vi-En due to high orthographic overlap (e.g., the word "ban") and the common omission of diacritics. To address this, we developed a robust two-stage LID process:

1. **Initial Filtering.** We used the lingua library for a high-recall pass to identify all Vietnamese posts containing English tokens.
2. **Ensemble LLM Validation.** This initial subset was then processed by a majority-vote ensemble of three LLMs (GPT-4o (OpenAI Team, 2024), Claude-3.5 (Anthropic, 2024), and Gemini-1.5 (Gemini Team, 2024)) to accurately classify samples, resolving ambiguities in orthography and named entities.

We validated this pipeline on a manually annotated 400-sample set, where it achieved a 87.98% F1-score and 89.75% accuracy, significantly outperforming baseline methods (see Appendix A.4).

The final stage of our filtering pipeline is the removal of personally identifiable information (PII) to ensure user privacy. All data undergoes an automated process to detect and redact sensitive spans like phone numbers, email addresses, and similar artifacts. To guarantee the integrity of our benchmark, the test set undergoes an additional, rigorous manual review to identify and handle any complex edge cases missed by the automated system. Across all splits, posts containing PII that cannot be cleanly redacted are discarded entirely.

3.2.1 Final Corpus Statistics and Splits

Applying our validated pipeline to the sourced data yields the final VIETMIX corpus of 10,254 high-confidence code-mixed posts. We then perform a random split into a training set (8,791 posts; seed data for synthetic generation), a development set (462 posts; *source-only* with no target references, used for model selection and ablations via reference-free models), and a held-out test set (1,002 posts; reserved for expert translation). Table 1 reports corpus statistics, including the Code-Mixing Index (CMI; Srivastava and Singh 2021) and Switching Point Frequency (SPF, Pratapa et al. 2018).

Statistic	Seed			Aug.
	Train	Dev	Test	
CMI	21.66	21.36	20.66	18.20
SPF	0.49	0.49	0.49	0.48
# tokens (src)	372,415	17,179	31,801	3,460,247
# tokens (tgt)	369,429	0	31,952	3,347,241
# samples	8,790	462	1,002	81,768

Table 1: Corpus statistics for the naturally occurring Vi-En code-mixed seed data splits (Train, Dev, Test) and the synthetically generated augmented (Aug.) data. Metrics include Code-Mixing Index (CMI), Switching Point Frequency (SPF), token counts for source (src) and target (tgt), and number of samples. The ‘Test’ set of the seed data forms the basis of the VietMix corpus.

3.3 Expert Translation of the Gold-Standard Test Set

To create a gold-standard benchmark, the 1,002 samples in the held-out test set were translated into English. This was performed by two professional translators with expertise in social media language, following a rigorous quality control process:

- **Phase 1 (Alignment):** Translators first independently translated 70 samples deemed as most challenging. They then cross-evaluated each other’s work, rating on an adapted version of the MQM-Core framework², which focuses on six dimensions, including accuracy, style, and locale conventions (Appendix B). Disagreements were mediated by a Vietnamese language instructor to establish clear conventions for handling challenging cases (e.g., puns, cultural references, neologisms).
- **Phase 2 (Translation & Validation):** Translators then completed the remaining of the data. All translations underwent a final verification pass by our language expert to ensure adherence to established guidelines and refine for maximum semantic fidelity.

This process ensures VIETMIX becomes a high-quality, ecologically valid benchmark suitable for evaluations of code-mixed MT systems.

4 Quality-Aware Data Augmentation

The VIETMIX corpus provides a high-quality foundation, but its size is insufficient for training robust systems. Standard augmentation techniques like back-translation often fail on code-mixed text,

²<https://themqm.org/the-mqm-typology/>

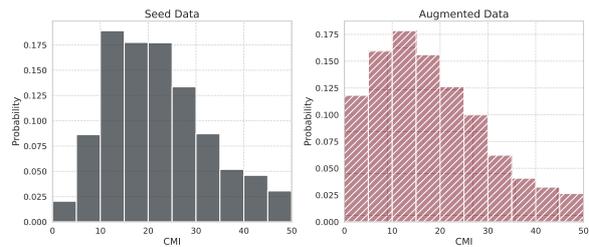


Figure 2: Comparative distribution of the Code-Mixing Index (CMI) in the naturally occurring VietMix seed corpus (Left/Real) versus the synthetically generated augmented corpus (Right/Augmented). The CMI quantifies the extent of language mixing within text samples.

correct it into a non-code-mixed form, or produce unnatural translations at language boundaries. Addressing this, we developed a three-stage data augmentation pipeline designed to generate a large-scale, high-quality synthetic parallel Vietnamese-English code-mixing corpus. Our pipeline’s architecture, visualized in Figure 1, prioritizes quality through iterative retraining and filtering.

Stage 1: Seed Data Generation and Base Model Bootstrapping.

The pipeline begins with naturally-occurring code-mixed Vietnamese (Vi) from our training split (n = 8,791).

- **Initial Translation.** We employ GPT-4o to generate initial English translations (En) for this Vi sample. This creates the natural Vi ↔ synthetic En seed parallel data.
- **Base Model Fine-tuning.** We then fine-tune a pre-trained multilingual model, Qwen2.5-7B, on this seed data for the En→Vi translation task. Qwen2.5-7B has strong documented performance on East Asian languages, including Vietnamese (Yang et al., 2024). This model is now "bootstrapped" with a foundational understanding of Vi and can serve as the core engine for subsequent augmentation.

Stage 2: Iterative Synthetic Data Generation.

This stage generates large volumes of synthetic Vi that mirror real Vi in style and domain.

- **Source Data Selection.** We select large monolingual English corpora from sources stylistically similar to our seed data (Bluesky³ and Reddit⁴) to minimize domain mismatch

³<https://huggingface.co/datasets/Roronotalt/bluesky>

⁴<https://huggingface.co/datasets/sentence-transformers/reddit-title-body>

and ensure the generated text reflects the informal, conversational register of social media.

- **Iterative Translation and Retraining.** Using the bootstrapped $E_n \rightarrow V_i$ model from Stage 1 to translate the English source text, we produce a large corpus of synthetic V_i . Iteration is key to this stage. After generating an initial batch of V_i , we apply quality filters as described in Appendix A, combine the filtered synthetic data with our original seed, and retrain the $E_n \rightarrow V_i$ translation model. Iteration nudges the model to progressively generate more natural code-mixed patterns.

Stage 3: Final Quality Assurance. Several cycles of iterative generation and filtering aggregate to a final corpus of over 200K pairs. As a final quality check, we use a reference-free quality estimation metric, xCOMET (Guerreiro et al., 2024), to score the semantic alignment of each E_n - V_i pair. Through manual inspection, we established 0.9 as a threshold and discarded all pairs below this score to prune examples with significant meaning shifts.

Our pipeline yielded a final augmented corpus of over 80K parallel examples. We use this corpus in our experiments to measure the impact of quality-aware synthetic data generation. The CMI distribution of this final corpus compared against the original seed is reported in Figure 2.

5 Experimental Validation

We conduct a series of experiments to empirically validate the utility of our contributions. Our main objective is to determine the extent fine-tuning on the VIETMIX corpus and our quality-aware augmented data improve code-mixed MT performance.

5.1 Experimental Setup

We evaluate three diverse open-source multilingual backbones to reduce architecture-specific effects: GemmaX2-28-9B, a sparse Mixture-of-Experts model specialized for multilingual translation with Vietnamese in pre-training (Cui et al., 2025); Qwen2.5-7B, a dense model with strong performance in East Asian languages (Yang et al., 2024); and Llama3.1-8B, a general-purpose dense model with predominantly English pre-training to serve as contrast to multilingual models (Team, 2024). Unless otherwise noted, we fine-tune the *base* variants using Low-Rank Adaptation (LoRA). Full hyperparameters are in Appendix C.

We assess each backbone under four training conditions to isolate the contribution of our data:

Instruction-tuned (zero-shot). Off-the-shelf instruction variants without additional fine-tuning; used as lower-bound baselines.

Back-Translation (BT). Trained on 80K+ synthetic pairs via standard back-translation, using the English side of our augmented corpus (§ 4) and regenerating Vietnamese with an $E_n \rightarrow V_i$ model fine-tuned on the VIETMIX seed data.

Seed-only. Base variants fine-tuned exclusively on the VIETMIX seed set (8,791 pairs).

Augmented (Ours). Base variants fine-tuned on our quality-filtered synthetic corpus (80K+ pairs).

5.2 Evaluation Protocol

Our primary semantic metrics are reference-based xCOMET⁵ and reference-free COMETKiwi⁶. Both were chosen for their support for Vietnamese and their demonstrated strong correlation with human judgments in recent WMT evaluations across multiple domains, especially for domains like social media chat (Freitag et al., 2022; Rei et al., 2023; Guerreiro et al., 2024). We avoid BLEU as a headline metric as prior work shows that typical test-set references exhibit limited diversity and tend to concentrate around translationese, which depresses BLEU and weakens its alignment with human preferences (Freitag et al., 2020; Wu et al., 2024). For broader comparability with other studies, we also report scores for BLEU in Appendix C.

To complement the automatic scores, we employ an LLM-as-a-Judge framework using GPT-4.1 for pairwise preference judgments between outputs from our *Seed-only* and *Augmented* models. This provides a qualitative validation of the improvements indicated by the metrics (Kocmi and Federmann, 2023). For structured and reliable outputs, prompts were managed using DSPy (Khattab et al., 2024), with full specifications reported in Appendix D.2.

To ensure the reliability of our findings, we employ paired permutation tests to determine if the observed improvements across system pairs are statistically significant.⁷

⁵<https://huggingface.co/Unbabel/XCOMET-XXL>

⁶<https://huggingface.co/Unbabel/wmt23-cometkiwi-da-xxl>

⁷https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.permutation_test.html

Model	xCOMET	Δ	COMETKiwi	Δ
Instruction models				
Qwen2.5-7B-Inst	77.72	-	67.16	-
Llama3.1-8B-Inst	66.96	-	57.96	-
GemmaX2-28-9B-v0.1	77.83	-	65.85	-
Back-Translation (BT)				
Qwen2.5-7B	78.27	-	65.94	-
Llama3.1-8B	77.57	-	64.02	-
GemmaX-9B	78.26	-	67.51	-
Ours				
with Seed only				
Qwen2.5-7B	78.74	-	69.72	-
Llama3.1-8B	78.19	-	68.79	-
GemmaX-9B	80.31	-	71.11	-
with Augmented				
Qwen2.5-7B	79.28	+0.54	69.99	+0.26
Llama3.1-8B	78.90	+0.71	69.44	+0.65
GemmaX-9B	81.73	+1.41	71.60	+0.49

Table 2: Performance on Vi→En (VietMix test). We report reference-based xCOMET and reference-free COMETKiwi for instruction-tuned baselines, back-translation (BT), our models fine-tuned only on the seed data, and our models fine-tuned with the augmented dataset. Higher is better. The Δ columns show *Augmented* minus *Seed-only*. Colored boxes highlight performance differences compared to *Seed-only*: blue = strong ($p < .001$), green = significant ($p < .05$), yellow = non-significant ($p > .05$).

6 Results

We analyze empirical results of our experiments using quantitative metrics, followed by a qualitative error analysis to provide a holistic view.

6.1 Main Quantitative Results

Table 2 presents our primary results. All models improve significantly when fine-tuned on the VIETMIX seed corpus compared to their zero-shot performance, confirming that even a small amount of high-quality, in-domain data is crucial. Building on this, our quality-aware augmentation consistently and substantially outperforms the BT baseline across all models. For instance, GemmaX-9B improves to 81.73/71.60 (+3.47/+4.09 over BT), with Llama3.1-8B and Qwen2.5-7B showing similar strong gains. This validates that a principled, quality-centered strategy is superior to naive synthetic data generation.

Our augmentation also provides a consistent boost over the *Seed-only* models, though the gains vary in an insightful way. While GemmaX-9B (+1.41/+0.49) and Llama3.1-8B (+0.71/+0.65) benefit significantly, the COMETKiwi gain for Qwen2.5-7B is minimal (+0.26). This is expected, as Qwen2.5-7B was the generator for the augmented data. Fine-tuning a model on its own out-

Metric	# Samples	Agreement w/ GPT-4.1
COMETKiwi	493	80.1%
xCOMET	555	81.3%

Table 3: Agreement between COMET metrics and GPT-4.1 pairwise preferences on En-Vi code-mixed test items.

puts offers little new signal and can reinforce existing biases. In contrast, the other models benefit from the augmented pairs as teacher-style supervision from a different model, leading to more effective knowledge transfer. Overall, the largest incremental gains appear with GemmaX-9B, while all backbones benefit from our quality-aware augmentation.

6.2 Validating the Evaluation Protocol

Automated MT metrics can be unreliable on code-mixed text. We therefore validated our protocol for Vi-En code-mixing. We selected xCOMET and COMETKiwi because they (i) support Vietnamese and (ii) are widely used in recent MT work (Xu et al., 2024; Cui et al., 2025), while acknowledging that their calibration on Vi-En code-mixed inputs is not well established.

We conducted a focused agreement study between COMETKiwi, xCOMET, and GPT-4.1 acting as a pseudo-human judge (LLM-as-a-judge), which prior work reports can reach >80% agreement with human judgments (Zheng et al., 2023). Concretely, we compared GPT-4.1 pairwise preferences between *Seed-only* and *Augmented* system outputs against the “wins” predicted by each COMET metric on the same pairs. A metric declared a win when one translation’s score exceeded the other by more than 2%; otherwise we treated the case as a tie and excluded it from agreement calculations.⁸

As shown in Table 3, both metrics correlate strongly with the LLM judge’s preferences on our code-mixed test set. With agreement rates of 81.3% for xCOMET and 80.1% for COMETKiwi, our study provides considerable confidence that these metrics are reliable proxies for human judgment in this domain. Furthermore, the LLM judge preferred outputs from the *Augmented* model in over 55% of non-tie cases, supporting our earlier evidence of perceptually improved translation quality.

⁸We adopt a 2% margin to reduce spurious wins due to small score fluctuations.

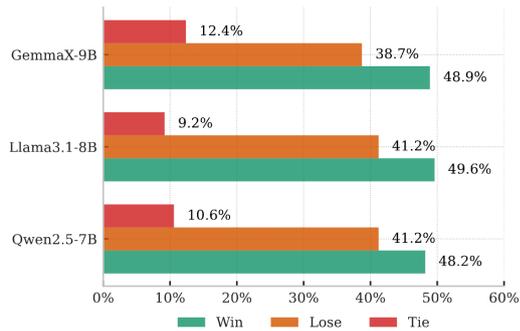


Figure 3: LLM-as-a-Judge assessment comparing En translations from models fine-tuned on augmented data vs. seed data only. Bars represent the percentage of judgments where GPT-4.1 favored the augmented model’s translation (Win), the seed-only model’s translation (Lose for augmented), or found them equivalent (Tie) across three base models (GemmaX2, Llama3.1, Qwen2.5). Semantic faithfulness against expert Viet-Mix references was the primary criterion.

We therefore report COMET scores as primary automatic metrics, while recommending continued use of human or LLM-as-a-judge audits for calibration on code-mixed evaluation.

6.3 Qualitative Error Analysis

To understand *how* our approach improves translation beyond aggregate scores, we conducted a manual error analysis over representative test items. We observe consistent gains from the augmentation pipeline in four areas that are central to natural code-mixed communication: pragmatic intent around proper nouns, idiomatic phrasing, transitions at switch boundaries, and acronym handling with contextual disambiguation.

Pragmatics near proper nouns. Exposure to monolingual En helps models preserve speaker intent when English entities anchor the clause. For “Threads cho lên xu hướng...”, the *Seed-only* output “**Threads for trending...**” misses the call-to-action, whereas the *Augmented* output “**Let’s get Threads trending...**” captures the pragmatic force by dropping a literal filler and adopting an imperative framing.

Idiomatic phrasing. Augmented models learn more faithful colloquial paraphrases, reducing literal calques. For “Loewe puzzle... không bao giờ sợ lỗi thời”, *Seed-only* yields “**never afraid of being outdated**”, while *Augmented* produces “**never go out of style**”, a conventional idiom that better matches intent and register.

Smoother switch boundaries. At complex

Model	w/ Ref	w/o Ref
Llama3.1-8B (Augmented)	78.90	80.87
Llama3.1-8B (9K Seed)	78.19	80.34
Llama3.1-8B-Inst	66.96	67.65

Table 4: Comparison of xCOMET scores for Llama3.1-8B translations: the reference-based variant (w/ Ref) versus the reference-free variant (w/o Ref)

code-mixed junctures, *Seed-only* systems tend to mash word-for-word fragments across languages (e.g., “**apply where pass that**”), whereas *Augmented* systems produce coherent clauses that respect switch-sensitive structure (e.g., “**you will pass everywhere**”). This suggests improved modeling of clause-level alignment rather than token-level overlap.

Acronyms and contextual disambiguation. Augmented models also resolve acronym-induced ambiguities and country/nationality confluences. In “Du học sinh Mỹ... SSN/ITIN...”, *Seed-only* misreads both actor and identifiers (“**American international students... No social SSN/ITIN**”), while *Augmented* correctly renders “**international students in the US... no social security number (SSN/ITIN)**”, leveraging redundant cues in the source to bridge concepts across languages.

Overall, these patterns indicate benefits beyond lexical accuracy: *Augmented* models better preserve pragmatic intent, select idiomatic equivalents, maintain fluency across switch points, and use acronym context to disambiguate meaning. Because lexically “correct” translations are not always the most useful or natural, the pipeline encourages learning latent rules of code-mixing and offers a practical recipe for other low-resource pairs where realism and quality control are essential.

7 Discussion

Our work presents a validated data augmentation methodology and a foundational resource for a challenging, low-resource code-mixed translation task. This section synthesizes key findings and their implications for multilingual NLP.

7.1 Quality Trumps Quantity: Filtering is Crucial for Code-Mixed Augmentation

For complex linguistic phenomena such as code-mixing, data quality is paramount. Fine-tuning on our VIETMIX seed corpus yields substantial gains over zero-shot, indicating that state-of-the-art multilingual models, as-is, underperform on natural,

Source	Model Hypothesis	Reference	w/ Ref	w/o Ref
Khám 120 brands xong vẫn chưa hiểu tại sao chiếc áo bà ba mà bà tư bỏ mặc	After inspecting 120 brands, I still don't understand why the shirt grandma wears is like that.	After an inspection of 120 brands, I still don't get why the fourth lady is wearing the third's shirt.	71.43	91.31
Bé Ánh zia tới q2 rồi muzik mai trời bé zia nhà thiú	Little Ánh is coming to q2 tomorrow, and the family will celebrate her birthday	Little Anh's back in District 2, Muzik! Let's carry her back home tomorrow!	57.53	80.30

Table 5: Qualitative examples of translation errors by Llama3.1-8B (Augmented) that are overlooked by reference-free xCOMET but identified when using human references. The table shows the source text, model hypothesis, human reference, and the corresponding xCOMET scores with and without references.

in-domain code-mixing. This establishes that a small amount of high-quality, in-domain data is an essential first step toward meaningful improvement.

Beyond the seed effect, our augmentation experiments provide strong evidence for the quantity-quality trade-off. As shown in Figure 4, adding unfiltered synthetic data depressed performance below the *Seed-only* baseline, whereas quality-aware filtering restored and exceeded baseline. Thus, pruning low-quality synthetic examples is not optional but necessary. For specialized, low-resource MT, a smaller high-quality set is more valuable than a much larger noisy corpus.

7.2 Human References Are Essential for Code-Mixed Evaluation

Code-mixing challenges evaluations. While convenient, reference-free metrics that rely on source-hypothesis similarity can be misleading. To substantiate this claim, we re-evaluated Llama3.1-8B using both reference-based and reference-free variants of xCOMET. As shown in Table 4, reference-free scores were consistently artificially inflated, producing overconfidence that masks critical errors. Table 5 provides qualitative snapshots of these failures. In the first example, the reference-free metric sees a superficial word overlap and gives this semantically unfaithful translation a score of 91.31, while the reference-based metric assigns a much lower score of 71.43. In the second, the model completely inverts the meaning of the source, yet the reference-free metric still assigns a higher score, 80.30, due to surface-level alignment.

These observations underscore that evaluations of complex linguistic phenomena—such as code-mixing—must be anchored in human-crafted references. In this regard, VIETMIX supplies expert-curated, switch-sensitive references for reliable evaluation, enabling more faithful measurement of

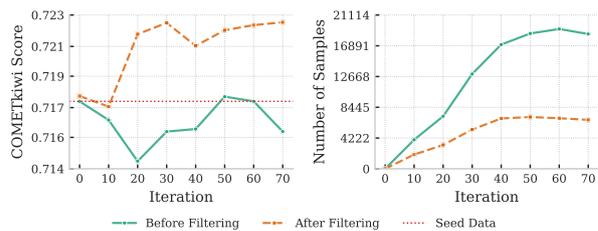


Figure 4: Impact of the Quality Assurance (QA) phase on synthetic data utility and yield across augmentation iterations. **(Left)** COMETkiwi scores on a development set for data before and after QA filtering, benchmarked against the seed-only data baseline (horizontal line). **(Right)** Number of synthetic samples generated per iteration before and after QA filtering, illustrating the yield of high-quality data.

progress in real-world code-mixed usage.

8 Conclusion

This work introduces VIETMIX, an expert-translated benchmark for Vietnamese–English code-mixing that fills a key evaluation gap. We show that state-of-the-art models underperform on this task but improve markedly with the high-quality, in-domain data VIETMIX provides. We also present a quality-aware augmentation pipeline and show that principled filtering is essential, since unfiltered synthetic data degrades performance. Our method yields significant gains over strong baselines and produces translations that better capture pragmatic and idiomatic nuance. By releasing the benchmark and the pipeline, we offer a practical blueprint for addressing data scarcity and evaluation in other low-resource settings and enable more robust and culturally competent MT systems.

9 Limitations

We acknowledge several limitations. First, VIETMIX is sourced from Threads.com, and our augmen-

tation data from Reddit and Bluesky. While this promotes stylistic consistency, linguistic patterns here may not generalize to other platforms, genres, or modalities. Future work should aim to create resources and evaluations that cover a broader range of platforms and genres.

Second, our data pipeline relies on proprietary LLMs. This introduces limitations in reproducibility, cost, and the possibility of inheriting opaque biases. Therefore, a key avenue for future research is to use VIETMIX to distill the capabilities of these large models to smaller and open specialist models.

Third, our experiments exclusively evaluated the code-mixed to non-code-mixed $V_i \rightarrow E_n$ direction. This deliberate choice targets our primary real-world application: enabling understanding of code-mixed text by monolingual speakers (e.g., in chat translation or content moderation). The inverse task of generating natural code-mixed from a non-code-mixed source ($E_n \rightarrow V_i$) is a fundamentally different challenge in stylization that goes beyond translation. Critically, this direction also presents immense evaluation challenges because a single non-code-mixed source can have an indefinite number of "correct" code-mixed targets, and this renders reference-based metrics inadequate. While linguistically fascinating, we have scoped this more ill-defined problem as future work.

Fourth, while our expert annotation process was rigorous and accounted for dialectal diversity, its scale is necessarily limited. Expanding expert-curated corpora requires significant resources. Capturing the full spectrum of sociolinguistic variation remains an open challenge.

Finally, our evaluation centers on intrinsic translation quality. While we validated our metrics, we have not yet measured the impact of improved translation on more downstream, extrinsic tasks. This is an important next step for future work to assess the real-world utility of this methodology.

10 Ethical Considerations

Data making up the VIETMIX corpus was collected via a public API in accordance with the platform's terms of service. To protect user privacy, we implemented rigorous PII removal. We first applied an automated script to detect and delete posts with common PII patterns (i.e., phone numbers, email addresses). Second, every sample in the final data underwent manual review and anonymization to remove any remaining potentially sensitive infor-

mation. As our work involved only publicly available data and professional translation services (not human subjects research), it did not require institutional review board (IRB) approval. VIETMIX (test set), all training and augmentation pipeline code will be publicly released under a CC BY-NC-SA 4.0 license to restrict use to non-commercial, research-only purposes.

To ensure linguistic inclusivity, translators and quality evaluators represent the three major dialectal variations of Vietnamese (Northern, Central, and Southern). Our annotation guidelines, developed collaboratively and refined through an adjudication process, will be released with the corpus to ensure transparency. We acknowledge that data sourced from social media platforms can reflect existing societal biases. While our sourcing from multiple platforms and quality filtering aim to mitigate the impact of potentially harmful content, models may still learn and reproduce subtle, harmful associations present in the training data. We consider the auditing of our models and corpus for social bias a critical area for future work.

11 Acknowledgements

We thank Zoey Ki, Calvin Bao, and the members of the CLIP lab at University of Maryland for their constructive feedback. We also thank the annotators who contributed to building the test set. Their careful work was essential to this project.

References

- Sadaf Abdul-Rauf and Holger Schwenk. 2009. Exploiting comparable corpora with ter and terp. In *Proceedings of the 2nd Workshop on Building and Using Comparable Corpora: From Parallel to Non-Parallel Corpora*, BUCC '09, page 46–54, USA. Association for Computational Linguistics.
- Gustavo Aguilar, Sudipta Kar, and Thamar Solorio. 2020. [LinCE: A centralized benchmark for linguistic code-switching evaluation](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1803–1813, Marseille, France. European Language Resources Association.
- Anthropic. 2024. [The claude 3 model family: Opus, sonnet, haiku](#). Technical report, Anthropic.
- Eyamba G Bokamba. 1989. Are there syntactic constraints on code-mixing? *World Englishes*, 8(3):277–292.
- Maksim Borisov, Zhanibek Kozhirbayev, and Valentin Malykh. 2025. [Low-resource machine translation](#)

- for code-switched Kazakh-Russian language pair. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pages 66–76, Albuquerque, USA. Association for Computational Linguistics.
- Menglong Cui, Pengzhi Gao, Wei Liu, Jian Luan, and Bin Wang. 2025. **Multilingual machine translation with open large language models at practical scale: An empirical study.** In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5420–5443, Albuquerque, New Mexico. Association for Computational Linguistics.
- Thi-Ngoc-Diep Do. 2021. **Mining japanese-vietnamese multi-level parallel text corpus from wikipedia data resource.** In *2021 RIVF International Conference on Computing and Communication Technologies (RIVF)*, pages 1–6.
- Thi Ngoc Diep Do, Laurent Besacier, and Eric Castelli. 2010. **A fully unsupervised approach for mining parallel data from comparable corpora.** In *Proceedings of the 14th Annual Conference of the European Association for Machine Translation*, Saint Raphaël, France. European Association for Machine Translation.
- Long Doan, Linh The Nguyen, Nguyen Luong Tran, Thai Hoang, and Dat Quoc Nguyen. 2021. **PhoMT: A high-quality and large-scale benchmark dataset for Vietnamese-English machine translation.** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4495–4503, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- A. Seza Doğruöz, Sunayana Sitaram, Barbara E. Bullock, and Almeida Jacqueline Toribio. 2021. **A survey of code-switching: Linguistic and social perspectives for language technologies.** In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1654–1666, Online. Association for Computational Linguistics.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. **BLEU might be guilty but references are not innocent.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 61–71, Online. Association for Computational Linguistics.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Koçmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. **Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust.** In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Google Gemini Team. 2024. **Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context.**
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. **xcomet: Transparent machine translation evaluation through fine-grained error detection.** *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Phu Gia Hoang, Canh Duc Luu, Khanh Quoc Tran, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2023. **ViHOS: Hate speech spans detection for Vietnamese.** In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 652–669, Dubrovnik, Croatia. 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.**
- Kartik Kartik, Sanjana Soni, Anoop Kunchukuttan, Tanmoy Chakraborty, and Md. Shad Akhtar. 2024. **Synthetic data generation and joint learning for robust code-mixed translation.** In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15480–15492, Torino, Italia. ELRA and ICCL.
- Phillip Keung, Julian Salazar, Yichao Lu, and Noah A. Smith. 2020. **Unsupervised bitext mining and translation via self-trained contextual embeddings.** *Transactions of the Association for Computational Linguistics*, 8:828–841.
- Simran Khanuja, Sandipan Dandapat, Anirudh Srinivasan, Sunayana Sitaram, and Monojit Choudhury. 2020. **GLUECoS: An evaluation benchmark for code-switched NLP.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3575–3585, Online. Association for Computational Linguistics.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan A, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. **DSPy: Compiling declarative language model calls into state-of-the-art pipelines.** In *The Twelfth International Conference on Learning Representations*.
- Tom Koçmi and Christian Federmann. 2023. **Large language models are state-of-the-art evaluators of translation quality.** In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.

- Ying Li and Pascale Fung. 2012. [Code-switch language model with inversion constraints for mixed language speech recognition](#). In *Proceedings of COLING 2012*, pages 1671–1680, Mumbai, India. The COLING 2012 Organizing Committee.
- Zheng Wei Lim, Nitish Gupta, Honglin Yu, and Trevor Cohn. 2024. Mufu: Multilingual fused learning for low-resource translation with llm. *arXiv preprint arXiv:2409.13949*.
- Sin-En Lu, Bo-Han Lu, Chao-Yi Lu, and Richard Tzong-Han Tsai. 2022. [Exploring methods for building dialects-Mandarin code-mixing corpora: A case study in Taiwanese hokkien](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6287–6305, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, John Kelleher, and Andy Way. 2022. [Domain-specific text generation for machine translation](#). In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 14–30, Orlando, USA. Association for Machine Translation in the Americas.
- Chinh Ngo, Trieu H Trinh, Long Phan, Hieu Tran, Tai Dang, Hieu Nguyen, Minh Nguyen, and Minh-Thang Luong. 2022. Mtet: Multi-domain translation for english and vietnamese. *arXiv preprint arXiv:2210.05610*.
- Cam-Tu Nguyen, Trung-Kien Nguyen, Xuan-Hieu Phan, Le-Minh Nguyen, and Quang-Thuy Ha. 2006. [Vietnamese word segmentation with CRFs and SVMs: An investigation](#). In *Proceedings of the 20th Pacific Asia Conference on Language, Information and Computation*, pages 215–222, Huazhong Normal University, Wuhan, China. Tsinghua University Press.
- Li Nguyen and Christopher Bryant. 2020. [CanVEC - the canberra Vietnamese-English code-switching natural speech corpus](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4121–4129, Marseille, France. European Language Resources Association.
- Quang-Phuoc Nguyen, Anh-Dung Vo, Joon-Choul Shin, Phuoc Tran, and Cheol-Young Ock. 2019. [Korean-vietnamese neural machine translation system with korean morphological analysis and word sense disambiguation](#). *IEEE Access*, 7:32602–32616.
- Thanh-Nhi Nguyen, Thanh-Phong Le, and Kiet Nguyen. 2024a. [ViLexNorm: A lexical normalization corpus for Vietnamese social media text](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1421–1437, St. Julian’s, Malta. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Mahani Aljunied, Shafiq Joty, and Lidong Bing. 2024b. [Democratizing LLMs for low-resource languages by leveraging their English dominant abilities with linguistically-diverse prompts](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3501–3516, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI OpenAI Team. 2024. [Gpt-4 technical report](#).
- Shana Poplack. 1980. Sometimes i’ll start a sentence in spanish y termino en espanol: toward a typology of code-switching.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Adithya Pratapa, Gayatri Bhat, Monojit Choudhury, Sunayana Sitaram, Sandipan Dandapat, and Kalika Bali. 2018. [Language modeling for code-mixing: The role of linguistic theory based synthetic data](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1543–1553, Melbourne, Australia. Association for Computational Linguistics.
- Ricardo Rei, Nuno M. Guerreiro, Jos  Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, Jos  G. C. de Souza, and Andr  Martins. 2023. [Scaling up CometKiwI: Unbabel-IST 2023 submission for the quality estimation shared task](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 841–848, Singapore. Association for Computational Linguistics.
- Rajvee Sheth, Himanshu Beniwal, and Mayank Singh. 2025. Comi-lingua: Expert annotated large-scale dataset for multitask nlp in hindi-english code-mixing. *arXiv preprint arXiv:2503.21670*.
- X. Shi, P. Yue, X. Liu, C. Xu, and L. Xu. 2022. [Obtaining parallel sentences in low-resource language pairs with minimal supervision](#). *Computational Intelligence and Neuroscience*, 2022(1):5296946.
- Thamar Solorio and Yang Liu. 2008. [Part-of-Speech tagging for English-Spanish code-switched text](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 1051–1060, Honolulu, Hawaii. Association for Computational Linguistics.
- Vivek Srivastava and Mayank Singh. 2021. [Challenges and limitations with the metrics measuring the complexity of code-mixed text](#). In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 6–14, Online. Association for Computational Linguistics.
- Llama Team. 2024. [The llama 3 herd of models](#).
- Nhu Vo, Dat Quoc Nguyen, Dung D. Le, Massimo Piccardi, and Wray Buntine. 2024. [Improving Vietnamese-English medical machine translation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language*

- Resources and Evaluation (LREC-COLING 2024)*, pages 8955–8962, Torino, Italia. ELRA and ICCL.
- Hai Quan Vu, Nam Trung Pham, Duc Hoang Ha Nguyen, Bao Toan Huynh, Hoai Bac Le, and Kiem Hoang. 2003. [Towards a multi-objective corpus for Vietnamese language](#). In *Proceedings of the 17th Pacific Asia Conference on Language, Information and Computation*, pages 416–422, Sentosa, Singapore. COLIPS PUBLICATIONS.
- Thanh Vu, Dat Quoc Nguyen, Dai Quoc Nguyen, Mark Dras, and Mark Johnson. 2018. [VnCoreNLP: A Vietnamese natural language processing toolkit](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 56–60, New Orleans, Louisiana. Association for Computational Linguistics.
- Genta Winata, Alham Fikri Aji, Zheng Xin Yong, and Thamar Solorio. 2023. [The decades progress on code-switching research in NLP: A systematic survey on trends and challenges](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2936–2978, Toronto, Canada. Association for Computational Linguistics.
- Minghao Wu, Jiahao Xu, and Longyue Wang. 2024. [TransAgents: Build your translation company with language agents](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 131–141, Miami, Florida, USA. Association for Computational Linguistics.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024. [Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation](#).
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Ruochen Zhang, Samuel Cahyawijaya, Jan Christian Blaise Cruz, Genta Winata, and Alham Fikri Aji. 2023. [Multilingual large language models are not \(yet\) code-switchers](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12567–12582, Singapore. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin,

Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).

A Data Augmentation Filtering

Our data augmentation consists of two stages with different mechanisms. This appendix section elucidates the methods employed in each.

A.1 Heuristics

In the first filtration stage, we implement heuristic constraints at the lexical and sub-lexical levels. These constraints eliminate pathological artifacts to make sentences more syntactically coherent.

A.1.1 Proportional Length Equilibrium

We compute the target-to-source token ratio and establish a permissible interval of $[0.5, 1.5]$. Candidate pairs exhibiting ratios outside this interval are immediately eliminated. This constraint mitigates catastrophic under-translation and verbose over-translation.

A.1.2 Lexical Repetition

To intercept repetition artifacts, we extract 5-word n-grams from the generated Vi translations and compute a weighted repetition coefficient:

$$xR_{lex} = \frac{\sum_{g:count(g)>1} count(g)}{\sum_g count(g)} \quad (1)$$

Candidate pairs with $R_{lex} \geq 0.3$ are eliminated because of lexical repetition redundancy.

A.1.3 Sub-lexical Repetition

To intercept sub-lexical repetition artifacts, we extract 10-character n-grams, calculate their frequency distribution, and isolate the top $k = \min(\lfloor \sqrt{N} \rfloor, N - U)$ most frequent sequences, where N is the total n-gram count and U is the unique n-gram count. Then, we compute:

$$R_{char} = \frac{\sum_{i=1}^k f_i}{\sum_{i=1}^n f_i} \quad (2)$$

Candidate pairs with $R_{lex} \geq 0.2$ are eliminated. This threshold and our threshold above for lexical repetition are manually determined based on qualitative observation of generated artifacts.

A.1.4 Code-mixing Equilibrium

To preserve a relatively authentic code-mixing ratio, we implement a constraint that English tokens within the generated Vietnamese must not exceed 30% of total tokens. We derive this parameter from an empirical analysis of natural code-mixing distributions in our real code-mixed

corpus. This proportion represents the median value observed across our real Threads.net corpus, though variations exist depending on topic.

All heuristic constraints function as first-stage filtration and eliminate approximately 37.8% of generated synthetic pairs.

A.2 Neural Classifier

Our second filtration stage seeks to approximate a naturalness assessment of generated texts. To approximate generated pairs' semantic-pragmatic naturalness, we implement a neural classification architecture described in the following subsections.

A.2.1 Model Selection and Classification

We fine-tune XLM-RoBERTa,⁹ given its effectiveness in cross-lingual tasks, to distinguish authentic Vietnamese code-mixed text from synthetic generations. The classifier received positive exemplars from realistic human-written text in our Vi Threads.net corpus and negative exemplars we generated from two sources:

- Low-quality generations with lots of visible, unrealistic artifacts from Llama-3.2-3B,¹⁰ selected for demonstrable deficiencies in Vi language modeling.
- Moderate-quality but still containing many synthetic artifacts from Llama-3.1-8B.¹¹

Both Llama models underwent initial fine-tuning on our seed dataset to translate En to Vi. We use these models to generate synthetic Vietnamese data when prompted with samples from En datasets described in § 4. This stratified negative sampling helps the classifier develop sensitivity and recognize more patterns in synthetic text.

A.2.2 Implementation in Filtering

After candidate Vi samples pass through the heuristic filters described in Section A.1, samples that satisfy heuristic constraints undergo neural classification. A straightforward decision rule is applied: samples receiving a "synthetic" classification confidence score of ≥ 0.5 are discarded. This threshold is chosen to balance maintaining corpus size and ensuring some naturalness in generated texts. It is very challenging for synthetically generated code-mixed texts to achieve high levels of resemblance

⁹<https://huggingface.co/FacebookAI/xlm-roberta-base>

¹⁰<https://huggingface.co/meta-llama/Llama-3.2-3B>

¹¹<https://huggingface.co/meta-llama/Llama-3.1-8B>

Table 6: Adapted MQM-Core evaluation criteria and rating scale (1-5, where 5 is best) used for human translator assessment and alignment discussions during the VietMix corpus creation. The table details the operationalization of six quality dimensions for evaluating code-mixed social media translations.

Criterion	5	4	3	2	1
Terminology	Excellent handling of all terminology with appropriate decisions on cultural elements	Good handling of most terminology with minor inconsistencies	Some terms translated correctly, others inconsistently handled	Frequent terminology errors affecting understanding	Inappropriate handling of most cultural terms and stylistic choices
Accuracy	Perfect semantic correspondence	Very good semantic alignment with minimal reorganization	Core meaning preserved but with some omissions/additions	Notable meaning loss or additions	Significant distortion of meaning
Linguistic Conventions	Flawless grammar despite any source text errors	Good grammatical construction with very few errors	Acceptable grammar with some minor issues	Multiple grammatical issues affecting readability	Major grammatical errors making text difficult to understand
Style	Excellent balance of faithfulness and natural expression	Natural-sounding with good interpretative choices	Balanced translation with some natural elements	Stiff translation with little interpretive flexibility	Extremely rigid, word-for-word translation
Locale Conventions	Perfect handling of all locale-specific content and conventions	Good handling of locale conventions with appropriate notes	Adequate maintenance of most locale conventions	Poor handling of locale-specific content	Failed to maintain or explain any cultural elements
Audience Appropriateness	Perfect social media tone maintained throughout	Good social media tone with minor variations	Generally appropriate with occasional inconsistencies	Frequently mismatched tone for the platform	Completely inappropriate tone for social media

with human-written texts. Because of that, it is not our goal to invent a methodology for synthetic data to become as close as possible to the naturalness of human texts. This classification component only seeks to eliminate most semantically and pragmatically implausible code-mixed constructions that might otherwise contaminate our training data.

A.3 Filtering Approaches Evaluated but Not Incorporated in our Final Pipeline Implementation

A.3.1 Back-translation

Back-translation, despite being a well-established technique for monolingual data augmentation, presents challenges for our scenario. Preliminary experiments with back-translation reveal significant distortions in code-mixing ratios. The back-translation model frequently "corrected" code-mixing by normalizing to either predominantly Vi or En. Back-translation also compounds er-

rors at language switching points, per our qualitative examination. We believe that evaluating back-translated outputs in code-mixed scenarios, especially Vi, is especially challenging due to the variability in acceptable code-mixing patterns. This variability also creates challenges in establishing quality thresholds.

A.3.2 Direct Preference Optimization (DPO)

Members who are native Vietnamese speakers in the research team together annotated for human preference. We performed comparative evaluations of multiple translation outputs for the same source text and assessed naturalness on a calibrated scale. Our results rendered this approach impractical for our current work due to subjective preference inconsistencies arising from regional dialectal variation. Annotators may disagree on "cleaned-up" code-mixing or the retention of messy code-mixing in real-world social media data. Preference-based op-

Table 7: LID Evaluation Results

Approach	Accuracy	Precision	Recall	F1 Score
Lingua	73.25	98.65	40.78	57.71
GPT-4o (single)	88.00	83.25	91.62	87.23
Ensemble LLMs	89.75	92.59	83.80	87.98

timization, therefore, was not performed to allocate efforts to developing our core filtering methodology. Future research could explore how to resolve inter-annotator disagreements in code-mixed language annotations arising from stylistic variations.

A.4 Language Identification Evaluation

In the absence of a standard En–Vi code-mixed LID benchmark, we manually annotated a 400-sample set drawn from Bluesky to stress challenging phenomena (e.g., named entities, dates, orthographic overlap). Each sample was labeled as *monolingual Vi* or *code-mixed*. We evaluated three approaches on this set: (i) *lingua* (dictionary/statistical), (ii) a single GPT-4o classifier, and (iii) a majority-vote *Ensemble LLMs* (GPT-4o, Claude 3.5, Gemini 1.5). For reporting, we use standard binary metrics—Accuracy, Precision, Recall, and F1—treating *code-mixed* as the positive class. Table 7 summarizes results.

The two-stage filter used in the main pipeline aligns with these results: an initial *lingua* pass efficiently removes clear monolingual cases, while the LLMs leverage context to disambiguate orthographically identical tokens and named entities. In the 400-sample set, error analysis indicates fewer than 5% of cases involve named-entity misclassification (e.g., *Vietnam* in an English context), supporting the reliability of code-mixed labels for downstream curation.

B MQM-Core Dimensions Contextual Adaptions for Human Translators’ Ratings and Alignment Discussions

Table 6 displays our adapted MQM-Core quality rating metrics descriptions. We operationalized six out of seven evaluative dimensions and created a Likert scale measurement framework for each dimension to facilitate translators’ discussions throughout their annotation work and structure the team’s collaborative discussions. Productive discussions on dimensions in which professional human translators most disagree with each other inform us where personal style and interpretive expe-

rience most influence translations. In this context, since innovative language use is common, maximizing naturalness in target En translations can mean sacrificing semantic equivalence.

Contemporary neural evaluation systems predominantly align with technical translators’ strict accuracy preferences. Hence, they largely overlook fluid interpretative ranges that non-technical translators rely upon to carry out and evaluate their work, especially in contexts abundant with cultural meanings. Future research may empirically examine how heterogeneity in translators’ profiles (e.g., technical translators vs. literary translators) can inform the creation of evaluation metrics that dynamically adapt their criterion weights based on stylistic and domain heterogeneity.

Figure 5 snapshots the interface we show translators each session before they produce translations to ground their task interpretations on specific instructions.

C Experimental Details

All models were fine-tuned using Low-Rank Adaptation (LoRA) (Hu et al., 2021) to efficiently adapt pre-trained LMs to our Vi–En translation task. The fine-tuning configuration was:

Training Hyperparameters

- Learning rate: 2.5×10^{-4}
- Epochs: 3
- Batch size: 128
- Warmup steps: 10
- Precision: BF16
- LoRA rank: 256
- LoRA modules: [q_proj, k_proj, v_proj, o_proj]

All experiments and model evaluations reported in this paper were conducted on the held-out test set described in § 3, after model development and hyperparameter tuning on the development set.

BLEU Results BLEU results are provided for comparability and are *not* used as headline metrics (see main text). We report **detokenized BLEU** computed with SACREBLEU (Post, 2018) using the default 13a tokenizer and case-sensitive (mixed) scoring, with exponential smoothing. Exact scores appear in Table 8.

Setting	Qwen2.5-7B	Llama3.1-8B	GemmaX-9B
Back-Translation (BT)	20.48	19.97	20.44
Seed (9K)	23.02	22.98	24.23
Augmented	22.40	22.14	22.71

Table 8: Detokenized BLEU on Vi→En VIETMIX test.

D Prompts

D.1 Seed data generation

Vietnamese code-mixed seed data underwent a two-stage refinement. Initially, we send samples from the seed data to GPT-4o, Gemini-1.5, and Claude-3.5-sonnet for detecting code-mixing, identifying Personally Identifiable Information (PII), and performing an initial translation; samples identified with PII were discarded. Subsequently, PII-cleared code-mixed samples were translated, primarily by GPT-4o. In cases of GPT-4o failure or poor output, translations were selected from Gemini or Claude. The prompt for this multi-model ensemble (also described in § 4) is in Figure 6.

D.2 LLM-as-a-Judge translation assessments

Figure 7 presents the DSPy signature (CompareEnglishTranslations) defining the prompt used for the LLM-as-a-Judge evaluation of translation quality, as referenced in § 5 and § 6.

Your Main Task - Công việc chính:

- You will see a series of Vietnamese sentences from social media - Anh/chị sẽ thấy một tập hợp câu được viết bằng Tiếng Việt từ mạng xã hội
- You will need to translate each sentence into natural, fluent English, or rate the quality of provided translations - Anh/chị sẽ cần dịch câu qua tiếng Anh với giọng văn chuẩn và tự nhiên, hoặc đánh giá chất lượng bản dịch được cho.
- We will ask you to capture both the meaning and tone of the original - Anh/chị sẽ cần giữ bản dịch gần nghĩa và hành văn của câu gốc nhất có thể.

Additional Guidelines - Chỉ dẫn thêm:

- Depending on the question, you may type directly into the text box provided or select multiple-choice responses - Dựa trên yêu cầu của câu hỏi, anh/chị có thể gõ bản dịch vào ô 'text' của bản khảo sát hoặc chọn phương án trắc nghiệm.
- You must translate the complete sentence and not skip any parts - Anh/chị sẽ cần dịch toàn bộ câu và không bỏ qua thành phần câu nào.
- You must maintain the original meaning and context as accurately as possible - Anh/chị sẽ cần giữ nghĩa và ngữ cảnh câu gần nhất có thể.
- You must use natural English expressions rather than word-for-word translations - Anh/chị sẽ cần tạo bản dịch Tiếng Anh tự nhiên nhất có thể, tránh dịch từng từ 'word-for-word' kém tự nhiên.
- If there are slang terms or cultural references you recognize, and no appropriate English equivalents exist, you should simply describe it with a phrase - Trong trường hợp anh/chị gặp tiếng lóng 'slang' hoặc dẫn chứng văn hoá Việt Nam khó dịch, và anh/chị không thể tìm bản dịch tiếng Anh phù hợp, ta sẽ dùng một cụm từ để mô tả từ đó. *For example - Ví dụ:* “Nón quai thao = Vietnamese traditional headwear with a flat palm and fringes”; “Đàn bầu = traditional Vietnamese instrument with one string and a cylindrical bamboo resonator”
- If there is Vietnamese-English code-mixing, you should retain English word forms in your translation if it does not alter meanings. - Trong trường hợp anh/chị gặp câu lẫn tiếng Anh-Việt, ta cần giữ cấu trúc của từ hoặc cụm từ tiếng Anh trong bản dịch nếu ngữ nghĩa không thay đổi.
- If there is dialectal Vietnamese word or grammatical structures you don't understand, you can skip that translation. - Trong trường hợp anh/chị gặp tiếng Việt địa phương mình không hiểu, vui lòng bỏ qua không dịch câu đó.

Important Notes - Yêu cầu quan trọng:

- You can NOT to use online translation tools (like Google Translate) or AI assistants (like ChatGPT) - Anh/chị sẽ KHÔNG được sử dụng các công cụ dịch thuật online (như Google Translate) hoặc công cụ trí tuệ nhân tạo (như ChatGPT)
- You can NOT leave the task interface to look up concepts, as your authentic understanding suffices - Anh/chị sẽ KHÔNG được rời khỏi trang thực hành tác vụ để tra cứu từ, bản dịch chỉ cần anh/chị đưa ra ý hiểu tự nhiên nhất của bản thân.
- We value your authentic translation skills and cultural knowledge - Chúng tôi trân trọng công sức, kỹ năng dịch thuật tự nhiên, cũng như vốn hiểu biết văn hoá của anh/chị.

Figure 5: Bilingual (Vietnamese and English) instructions provided to professional translators. Key guidelines include maintaining meaning and tone, handling slang and cultural references, and retaining English word forms in code-mixed segments if meanings can be preserved.

Seed Data Detection and Translation

Detect if the given Vietnamese text contains any English words. If English words are present, detect if the given text contains any personally identifiable information (PII), then translate the entire text into English.

Steps

0. ***Language Detection***: Detect whether the given text contains English words.
1. ***PII Detection***: Detect whether the text contains any PII (e.g. name, address, email, phone number, bank account).
2. ***Translation***: Translate the entire text into English accurately.

Output Format

- ***Detection Result***: Output "True" if English words are present, otherwise "False".
- ***PII Result***: "True" if any PII (e.g., name, address, email, phone number, bank account, etc.) is present, otherwise "False".
- ***Translation***: If applicable, provide the English translation.

Examples

Example 1:

Input: "Tôi thích đi shopping vào cuối tuần."
- Detection Result: True
- PII Result: False
- Translation: "I like to go shopping on weekends."

Example 2:

Input: "Tôi thích đi mua sắm vào cuối tuần."
- Detection Result: False

Example 3:

Input: "Tôi tên là John, số điện thoại của tôi là 0987654321 và tôi thích đi shopping vào cuối tuần."
- Detection Result: True
- PII Result: True
- Translation: "My name is John, my phone number is 0987654321, and I like to go shopping on weekends."

Notes

- Only proceed to PII detection and translation if English words are detected.
- PII detection must detect correctly common identifiers (names, addresses, emails, phone numbers, bank accounts, etc.).

Input: {input}

Figure 6: LLM prompt structure for the added data processing of naturally occurring Vietnamese-English code-mixed seed data. The prompt guides the model through a multi-step task: (1) detecting the presence of English words, (2) identifying any left-over personally identifiable information (PII) undetected in initial processing, and (3) translating the entire text into English if code-mixing is detected.

```

1 class CompareEnglishTranslations(dspy.Signature):
2     """You are an expert bilingual linguist in Vietnamese and English. Please
3     choose the better English translation for the given Vietnamese source, using
4     the provided reference translation to understand the intended meaning of the
5     source. First, provide a rationale explaining your thought process, then
6     determine which translation is better. Output exactly one of: A, B, Tie, or
7     BothBad"""
8
9     source = dspy.InputField(
10         desc="The Vietnamese source text."
11     )
12     translation_a = dspy.InputField(
13         desc="The English translated version from model A."
14     )
15     translation_b = dspy.InputField(
16         desc="The English translated version from model B."
17     )
18     reference = dspy.InputField(
19         desc="A human written English reference translation."
20     )
21     rationale = dspy.OutputField(
22         desc="A brief explanation of the thought process."
23     )
24     better_translation: Literal['A', 'B', 'Tie', 'BothBad'] = dspy.OutputField(
25         desc="Which translation is better? Output exactly one of: A, B, Tie, or
26         BothBad."
27     )

```

Figure 7: DSPy signature (CompareEnglishTranslations) defining the LLM-as-a-Judge prompt for evaluating translation quality. This prompt instructs an expert bilingual LLM to compare two English translations (A and B) of a Vietnamese source text against a human reference, provide a rationale, and determine the better translation (A, B, Tie, or BothBad).