

Low-resource Machine Translation: what for? who for?

An observational study on a dedicated Tetun language translation service

Raphael Merx¹, Adérito José Guterres Correia², Hanna Suominen³, Ekaterina Vylomova¹

¹School of Computing and Information Systems, The University of Melbourne

²Instituto Nacional de Linguística, Dili

³The Australian National University

rmerx@student.unimelb.edu.au

Abstract

Low-resource machine translation (MT) presents a diversity of community needs and application challenges that remain poorly understood. To complement surveys and focus groups, which tend to rely on small samples of respondents, we propose an observational study on actual usage patterns of tetun.org, a specialized MT service for the Tetun language, which is the lingua franca in Timor-Leste. Our analysis of 100,000 translation requests reveals patterns that challenge assumptions based on existing corpora. We find that users, many of them students on mobile devices, typically translate text from a high-resource language into Tetun across diverse domains including science, healthcare, and daily life. This contrasts sharply with available Tetun corpora, which are dominated by news articles covering government and social issues. Our results suggest that MT systems for institutionalized minority languages like Tetun should prioritize accuracy on domains relevant to educational contexts, in the high-resource to low-resource direction. More broadly, this study demonstrates how observational analysis can inform low-resource language technology development, by grounding research in practical community needs.

1 Introduction

While machine translation (MT) is often considered a single area for natural language processing (NLP), it covers a large range of applications, end-user requirements, and language directions. MT can be leveraged for everyday communication, specialized sectors (e.g., health or education), learning, or translation of official communication (Dew et al., 2018; Vieira et al., 2021; Herrera-Espejel and Rach, 2023; Paterson, 2023; Merx et al., 2024). It can be used by students in the classroom (Deng and Yu, 2022), workers in an office environment (Brynjolfsson et al., 2019), or by foreign travelers in public

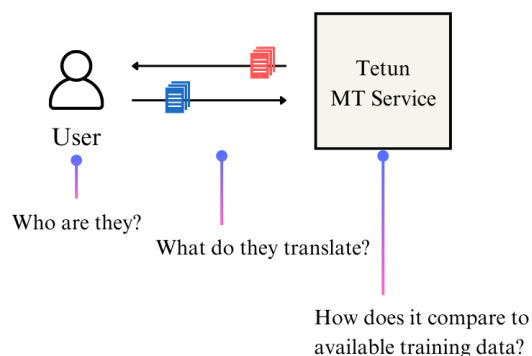


Figure 1: Overview of our approach and research questions

spaces (Carvalho et al., 2023). Its reliability varies significantly depending on the size and domain coverage of available corpora (Dabre et al., 2020; Khiu et al., 2024), and depending on similarity between source and target languages (Xu et al., 2020).

This disparity of potential applications and accuracy across domains or language pairs is particularly salient when translating to or from low-resource languages (Haddow et al., 2022). In low-resource scenarios, the combination of high variability both in MT accuracy across domains (Haddow et al., 2022) and of what native speakers value in NLP applications (Lent et al., 2022), has prompted NLP researchers to better anchor their work in local community perception, including through surveys, focus groups, and product design (Le Ferrand et al., 2022; Mager et al., 2023; Blaschke et al., 2024). These approaches, however, could be complemented by observational studies that rely on inferring needs or lack thereof for MT across domains or demographics from end-user behaviors.

In this paper on low-resource MT, we work with the Tetun language,¹ which is the most spoken

¹Also called Tetun Dili, or Tetum.

language in Timor-Leste (General Directorate of Statistics Timor-Leste, 2015). We rely on a sample of server logs from tetun.org, a volunteer-run MT service for the Tetun language, available in both website and app² format. Until Google Translate added Tetun support in June 2024, this service was the only widely available MT tool for Tetun. Tetun.org has over 70,000 monthly active users, and translates over one million documents per month. From a sample of 100,000 end-user translation text inputs, and their associated machine-translated output, collected between March and August 2024, we aim to respond to the following research questions (Figure 1):

- Who are the main users of tetun.org, what technology do they rely on, and what is the driver of their MT usage?
- What characterizes MT inputs, in terms of domain, length, direction of translation, and language pairs?
- How do MT inputs compare to available corpora for the Tetun language?

Tetun: low-resource, but institutional. Following Nigatu et al. (2024), we define why Tetun is low-resource in our view: (1) Tetun has very limited digital resources, with less than one million monolingual sentences available in open corpora (de Jesus and Nunes, 2024c) (2) Tetun only has around one million native speakers (General Directorate of Statistics Timor-Leste, 2015) and (3) Tetun has very little NLP research coverage, apart from work by Gabriel de Jesus during his PhD candidature at the University of Porto (de Jesus and Nunes, 2024a; de Jesus, 2024; de Jesus and Nunes, 2024b). However, Tetun is an institutional language, with both official status in Timor-Leste, and active usage in education and the media (Grek-sakova, 2018). Our findings, while focused on Tetun, may extend to other institutionalized low-resource languages, but are less likely to apply to non-institutionalized or endangered ones.

2 Related Work

Community-driven approaches to NLP for low-resource languages. Recent years have witnessed a growing focus on developing NLP technologies for low-resource and endangered languages. Researchers have identified that academic

²play.google.com/store/apps/details?id=org.tetun.tetunorg

motivations can be at odds with the actual needs of language communities (Le Ferrand et al., 2022; Liu et al., 2022), which made them advocate for a paradigm shift towards more community-centered approaches in NLP research (Lent et al., 2022; Bird and Yibarbuk, 2024). Studies emphasize the importance of close collaboration with language communities, understanding their specific needs, and involving them in the design and development process of language technologies (Mager et al., 2023; Bird and Yibarbuk, 2024). This approach acknowledges the heterogeneity among low-resource languages and their communities, recognizing that methodologies effective for one may not be universally applicable (Lent et al., 2022; Dolinska et al., 2024).

Linguistic analysis through large-scale data gathering. To our knowledge, this paper is the first large-scale analysis of real-world usage patterns in MT. It takes a similar approach to prior work leveraging large-scale secondary data in linguistics, including those predicting language learning difficulty (Schepens et al., 2020), understanding what language characteristics lead to a language being confused for another (Skirgård et al., 2017), and games that aim to incentivise the development of new language resources (Fiumara et al., 2022).

Domain imbalance and low-resource MT Low-resource language corpora are often concentrated in specific domains, particularly religious texts (Marashian et al., 2025), and lack the broad domain coverage typically found in their high-resource counterparts (Haddow et al., 2022). As a consequence, NLP tools trained on these available corpora tend to lack domain robustness (Ranathunga et al., 2023). This issue is particularly significant for MT systems, as demonstrated by Khiu et al. (2024), who find that domain similarity between training and test data is the largest driver of low-resource MT performance, more so than training data size or language similarity between source and target languages. Building on these findings, our analysis pays particular attention to differences of domain distribution between available corpora and actual translation requests.

3 Methodology

We work with data from tetun.org, a specialized service for automated translation between Tetun Dili (*tdt*) and any of English (*eng*), Indonesian (*ind*), or

Portuguese (*por*). The service also includes bilingual word dictionaries and a spell checker for the Tetun language, but we focus on the MT part.

The MT service uses a transformer model, and inference is run on the server side. The system was initially developed by fine-tuning a M2M100 multilingual model (Fan et al., 2020) on a scraped parallel Tetun-English corpus. To optimize inference speed, this model was then distilled into a smaller transformer-base architecture. For Portuguese and Indonesian translation, the system uses English as a pivot language. The model’s performance was evaluated using BLEU scores on a curated test dataset assembled by tetun.org volunteers.

We get access to the following data sources:

1. A sample of 100,000 server logs from the service, all collected between March and August 2024. Each log contains data on timestamp, text input, MT output, source language code, requested target language code, and device operating system (OS).
2. Summary monthly active users’ data by country and device sourced from Google Analytics (for the website), Google Play Store,³ and Apple App Store.⁴

While data on the user base can be directly extracted from (2), domain analysis requires a choice of methodology that we describe next.

3.1 Domain classification

Components of domain. Recognizing that “domain” is often used interchangeably in NLP research to refer to topic (e.g., “the health domain”) or source (e.g., “the news domain”), we follow the methodology in Saunders (2022), which defines domain as a combination of *topic* (one or more subject matters covered by the text), *provenance* (where the text comes from), and *genre* (language syntax and style).

3.1.1 Topic classification

LDA optimization. We start by identifying topics using Latent Dirichlet Allocation (LDA, Blei et al., 2003) on documents that are more than 5 words long, as very short documents do not offer enough context to identify their topic. We work with the Tetun side (either source Tetun text, or

target machine-translated Tetun text), as this language is always present in any of the language pairs. The text is first lower-cased and punctuation is removed. We run LDA enhanced with bigrams that appear at least 5 times in the corpus for improved context. After running hyperparameter search on number of topics, passes, and alpha, we find a maximum topic coherence of 0.065 using 15 topics, 10 passes, and an asymmetric alpha. While this informs our choice of topic to categorize from, we qualitatively find that LDA is noisy for this task. For example, the health topic also includes words related to project management, such as “objetivu” or “servisu”. It also performs poorly on the categorization of shorter documents, which are likely to contain words that are not part of any identified topics. However, we leverage these identified topics to formulate a list of topics to classify from.

Labels & test set preparation. Through a manual review of each LDA topic, we land on the following list of topics to classify from: *Technology, Daily life & personal experiences, Healthcare & medicine, Law & regulations, Business & work & employment, Government & socio-economic issues, Education, Science & research, Religion & spirituality*. We map each of the 15 LDA topics to one of those topics, and our first author manually annotates a sample of 100 logs against this list of topics for evaluation, where each log can be classified against multiple topics.

Classification using pre-trained models. We evaluate other classification methods that use pre-trained models, either zero-shot transformer-based classifiers or LLMs. Given the lack of pre-trained models that cover Tetun, we rely on the high-resource side of the input-output pair, which can be in English, Indonesian or Portuguese. For zero-shot classifiers, we use the HuggingFace zero shot classification pipeline⁵ with the list of topics passed as candidate labels, in a multi-label mode. We evaluate `bart-large-mnli`⁶ and `deberta-v3-large-zeroshot-v2.0`⁷ (Laurer et al., 2023). For LLMs, we assess Llama 3.1 8B (Dubey et al., 2024), with a prompt that lists the domains to classify from, and that includes examples for in-context learning, shown in Appendix A.

³play.google.com/store/apps/details?id=org.tetun.tetunorg

⁴apps.apple.com/us/app/tetun-org/id1515892208

⁵huggingface.co/tasks/zero-shot-classification

⁶huggingface.co/facebook/bart-large-mnli

⁷huggingface.co/MoritzLaurer/deberta-v3-zeroshot-v2.0

Classification evaluation. We compare micro and weighted average scores on our test set for the different classification approaches in Table 1. We find that Llama 3.1 8B has the highest weighted averaged F1 score on both metrics for this classification task, followed by (in order) deberta-v3-large-zeroshot-v2.0, bart-large-mnli, and last LDA.

We therefore pick Llama 3.1 8B for topic classification, and run topic inference on the whole 100,000 server logs, using a single Nvidia A100 GPU, which takes approximately 30 hours.

Approach	Micro avg F1	Weighted avg F1
LDA	0.21	0.19
bart-mnli	0.40	0.40
deberta-nli	0.66	0.68
Llama 3.1 8B	0.76	0.76

Table 1: Comparison of micro and weighted averaged F1 score for different domain classification approaches. LDA is run on the Tetun side, while other approaches are run on the high-resource side (English/Portuguese/Indonesian).

3.1.2 Provenance classification

After a manual review of the text samples used in topic classification, we propose the following categories for provenance: *Education & research material, News article & press release, Organizational & formal documents, Conversation & correspondence, Literary fiction, Religious text, Other*. We use Llama 3.1 8B here as well, with a prompt that lists the sources to classify from, and that includes examples for in-context learning. On the same sample of 100 logs, which we manually annotate for provenance, we get a weighted F1 score of 0.73, which we deem satisfactory for classifying provenance, given the natural ambiguity of this task, especially for shorter documents.

3.1.3 Genre

After recommendation from the Timor-Leste *Instituto Nacional de Linguística* (INL, National Institute of Linguistics), under genre, we evaluate what fraction of Tetun MT inputs follow official spelling, as defined by the country’s *Matadalan Ortográfiku ba Tetun Nasionál* (Orthographic Guide for National Tetun, 2002).

For calculating the percentage of words that follow INL spelling in each document, we apply the following preprocessing: (1) tokenize Tetun document using an open source Tetun tokenizer (de Je-

sus and Nunes, 2024a), (2) ignore words that start with a capital letter as they are often proper nouns and (3) ignore tokens that start with a number or punctuation character. We further ignore inputs where less than 20% of words were in the official spelling guide, as we empirically find that these documents are almost never in Tetun.

3.1.4 Comparison with available monolingual corpora

For each of the identified domain components (topic, provenance, genre), we provide a comparison between MT inputs and available Tetun corpora. For the latter, we select the Tetun “clean” split of the MADLAD-400 dataset (Kudugunta et al., 2023), as it is the largest Tetun corpus available to date with 40,367 documents. For classifying topic and provenance, given that we work with models that do not support Tetun, we first translate the MADLAD Tetun corpus to English using the MADLAD-3B model. Upon acceptance, we will release our translation of this corpus of 884,000 Tetun sentences and their machine-translated text for future researchers to use.

For all domain analysis, we weight inputs by their number of words. This is to ensure that very short MT inputs do not get as much weight as longer, more informative inputs, or that users who tend to translate several paragraphs at a time are not under-represented compared to users who translate shorter inputs.

4 Results

4.1 User base

Country. Most users are in Timor-Leste: 67,000 of 71,500 (93.4%) Android monthly active users have their device registered in Timor-Leste, and 17,000 of 22,000 (77%) website monthly active users are in Timor-Leste. These results support the hypothesis that the large majority of users are Timorese with English as a second language (ESL), as opposed to foreigners learning Tetun. While usage numbers can be noisy (for example, Timorese users who live abroad might have a mobile device registered outside of Timor-Leste, and foreigners in Timor-Leste using the website cannot be differentiated from local users), given the larger number of Timorese emigrating than people immigrating to Timor (INETL, 2022), they are more likely to under-estimate than to over-estimate the proportion of Timorese users.




Figure 2: Monthly active users by device and service. Note that users who use both the website and mobile app would be double counted.

Device. Most users access tetun.org through their mobile device, either through the Android application (71,500 monthly active users), the iOS application (2,500 monthly active users), or the website in a mobile browser (13,000 monthly active users). Desktop users of the website represent 8,300 monthly active users, with an unknown number of these users being duplicate of mobile users. The distribution of users is illustrated in Figure 2. We discuss the implications of this high representation of mobile devices, particularly Android devices, on NLP research, in Section 5.

Demographics. Server usage patterns seem to indicate that most users are students who use tetun.org for educational purposes. This is supported by (1) the large proportion of requests that occur during evenings before school days (including on Sundays, see Figure 3), at times when office workers are less likely to use the service, and when students are more likely to do their homework; (2) the decrease in service usage during school holidays, with usage going back up the evening before school resumes; (3) domain analysis results presented in Section 4.3.

4.2 Translation direction and text length

Translation direction Using the source and target language codes present in each of the server logs, we extract summary information on the distribution by source and target language. We find that Tetun is more likely to be the target language (over 70% of requests), with English to Tetun translation most in demand (46.7% of requests), followed by Tetun to English (22.3% of requests),

Direction	Requests (% of total)
English to Tetun	46.7 %
Portuguese to Tetun	12.4 %
Indonesian to Tetun	11.4 %
Total Tetun as Target	70.6 %
Tetun to English	22.3 %
Tetun to Portuguese	4.5 %
Tetun to Indonesian	2.5 %
Total Tetun as Source	29.4 %

Table 2: Distribution of translation requests by language pair

Portuguese to Tetun (12.4%), Indonesian to Tetun (11.4%), Tetun to Portuguese (4.5%), and Tetun to Indonesian (2.5%) – see Table 2. For all language pairs, the high-resource to low-resource direction is in higher demand than the low-resource to high-resource direction.

Length of translated text Leveraging an open source Tetun tokenizer (de Jesus and Nunes, 2024a) for Tetun, and the spaCy library⁸ for English/Indonesian/Portuguese, we count the number of words in the input texts. We find that shorter text inputs are much more common, with a sharp decline in frequency as the text length increases. This does not include hits to the tetun.org dictionary, as the dictionary works offline. Overall, inputted text has a median of 8 words, with a median of 12 words when Tetun is the target of translation, and 5 words when Tetun is the source. We note however that this distribution has a long tail, with over 12% of translations having 95 input words or more. A histogram of word count frequency is shown in Figure 4.

4.3 Domain coverage

In this section, we investigate what domains are most or least in demand by tetun.org users, and how they compare to domains present in available Tetun corpora, as represented by the MADLAD corpus. Given that domain similarity between training data and evaluation set is one of the main driver of MT accuracy (Khuu et al., 2024), this investigation can inform future efforts on low-resource corpus gathering.

4.3.1 Topic

Using the input text and its associated MT-translated result from the server logs, we estimate what topics are most or least in-demand for

⁸<https://spacy.io/>, version 3.8.2, model xx_ent_wiki_sm




Figure 3: Number of requests by day / hour, in Asia/Dili (UTC+9). We observe a spike in the evenings before school days (Saturday is a school day in Timor-Leste), when students are more likely to prepare their homework.




Figure 4: Histogram showing the frequency distribution of input text lengths measured by the number of words.

text translation. Our results for topic coverage are shown in Table 3. We include a comparison with topic coverage for the largest monolingual Tetun dataset currently available, MADLAD-400 (Kudugunta et al., 2023).

We find overall that “Science & research” is the topic most in demand, present in 34.2% of MT inputs when weighted by word count, followed by “Healthcare & medicine”, “Education”, and “Business & work & employment”, all present in 20 to 25% of inputs. Overall, “Religion & spirituality” is the least frequent topic in MT inputs, which contrasts with the high representation of the religious domain in most low-resource corpora (Haddow et al., 2022).

We observe a large discrepancy between user-translated text and available corpora, where the main topics covered by MADLAD (Government & socio-economic issues, 83.1%; Law & regula-

Topic	MT text	MADLAD
Science & research	34.2 %	3.2%
Healthcare & medicine	23.9 %	11.9 %
Education	22.8 %	11.3 %
Business & work & empl..	21.2 %	6.2%
Daily life & personal exp.	17.2 %	8.7%
Government & socio. issues	16.2 %	83.1 %
Law & regulations	8.9%	28.2 %
Technology	8.1%	3.9%
Religion & spirituality	6.8%	14.5 %

Table 3: Topic distribution, where each document can cover several topics.

tions, 28.2%) do not match with the main topics in demand by MT users.

We further analyze topic coverage depending on whether the source text is in Tetun or whether Tetun is the target language. When Tetun is the target, topics such as “Science & research” (36.73%), “Healthcare & medicine” (25.52%), and “Education” (23.14%) dominate, reflecting user demand for translating high-resource content into Tetun, likely to support knowledge dissemination and accessibility. In contrast, when Tetun is the source language, translations are predominantly focused on “Daily life & personal experiences” (48.64%), indicating that users frequently translate everyday communication into high-resource languages. This highlights a divergence in translation use cases: translating into Tetun serves educational and professional purposes, whereas translating from Tetun often fulfills personal or practical communication needs.

4.3.2 Provenance

As illustrated in Table 4, we observe a large dominance of “Education & research material”, which alone represents more than half of MT inputs

Provenance	MT text	MADLAD
Education & research material	55.9 %	7.5%
Organizational & formal docs	18.0 %	17.3 %
News article & press release	10.2 %	61.2 %
Conversation & correspond.	8.7%	4.8%
Literary fiction	4.3%	0.9%
Religious text	2.5%	8.2%
Other	0.4%	0.3%

Table 4: Provenance distribution in MT and MADLAD.




Figure 5: Histograms of MT Tetun inputs and MADLAD Tetun texts by fraction of words in the official INL spelling guide.

(55.9%). “Organizational & formal documents” comes as a distant second at 18%. This again stands in contrast with monolingual corpora, where over 61% of documents come from “News articles & press releases”, while only 7.5% come from “Education & research material”.

Our findings on topic and provenance (Tables 3 and 4) are consistent with the hypothesis that most tetun.org users are students (who are more likely than the general population to translate educational material related to science, education or healthcare) and with de Jesus and Nunes (2024c) who found that most Tetun text on the internet comes from news articles.

4.3.3 Genre

We observe that official (INL) spelling is only loosely followed in Tetun MT inputs, with a median of 71% of words spelled correctly, compared to a median of 82% in MADLAD. This is coherent with previous studies that note the high proportion of Tetun texts that do not follow the official spelling, including in the media (Greksakova, 2018). We also observe that variance in the rate of correct spelling is higher in MT inputs (standard deviation of 13%) than in MADLAD (standard deviation of 8%) – see Figure 5.

4.4 Effects of domain mismatch on LM modelling

After uncovering a lack of similarity between MT inputs and available corpora in the previous section, we investigate downstream effects on language modelling, by comparing perplexity on MT Tetun inputs with perplexity on monolingual corpora using Goldfish Tetun⁹ (Chang et al., 2024), a monolingual Tetun language model trained mostly (> 98%) on MADLAD Tetun documents. We use the HuggingFace evaluate¹⁰ library for all perplexity measurements, and only measure on inputs of 5 words or more.

MT inputs vs monolingual corpora We compare perplexity on MT Tetun inputs with perplexity on the Labadain Tetun corpus (de Jesus and Nunes, 2024c), which is the second largest Tetun corpus after MADLAD, and was not used when training the Goldfish model. We also exclude Labadain sentences that are also present in MADLAD. We get an average perplexity of 1,774 on MT Tetun inputs, compared to an average of 153 for Labadain, which illustrates that MT inputs may contain vocabulary, structures, or topics that are underrepresented in scraped monolingual corpora, leading to lower language modelling performance and potentially increased translation errors.

Perplexity per topic Comparing average perplexity per topic, we find that documents in the “Government & socio-economic issues” topic have a much lower perplexity (average of 632) than documents from the “Technology”, “Business & work & employment”, or “Daily life & personal experiences” domains (3,434, 1,871 and 1,864 respectively). This result illustrates that domain imbalance in monolingual corpora has downstream effect on language modelling ability for under-represented domains.

Perplexity per provenance Comparing average perplexity per provenance, we find that documents that come from “News article & press release” have the lowest perplexity at 874, compared to 1,320 for “Education & research material”, 1,662 for “Organizational & formal documents” and 1,865 for “Conversation & correspondence”. Here also, this result shows that document provenances that are under-represented in available corpora impact language modelling performance.

⁹https://huggingface.co/goldfish-models/tet_latn_full

¹⁰<https://github.com/huggingface/evaluate>, version 0.4.3

We report all perplexity results in Appendix B.

5 Recommendations

Anchoring MT evaluation in user demand.

Our results highlight the importance of: (1) broad domain coverage that includes education material, especially about science and health (2) accuracy when translating into, rather than from, the low-resource language, and (3) accuracy when translating very short documents (including individual words). This could start with an evaluation framework that finds a balance between these criteria. In particular, the sharp imbalance between the domain of available Tetun corpora and MT inputs leads us to recommend the systematic evaluation of MT accuracy in a wide variety of domains, and advise against measuring MT accuracy on domains that are over-represented in available corpora (e.g., news or religion), as this would risk giving a poor representation of end-user experience.

Meeting end-users where they are. We recommend a focus on meeting the needs of students who use MT on their mobile device. Given that students seem to represent a large portion of MT users, and that learning outcomes are an important social good, future study of MT perception and improvement for students in low-resource contexts would be beneficial, both to the research community and end users. Similarly, given the high cost and low reliability of internet connections in Timor-Leste ([The Asia Foundation, 2023](#)), we recommend a heightened focus on MT inference on mobile devices, for example through further research into efficient MT, or through the creation of an open source framework for on-device MT that targets mobile devices.¹¹

6 Conclusion

We present an analysis of MT usage for Tetun, a low-resource language, based on server logs from tetun.org, a widely-used MT service. Our study reveals key insights into user demographics, device preferences, translation directions, and domain coverage. We discover that a large share of users are likely students in Timor-Leste using mobile devices, translating education material from high-resource languages into Tetun. Our analysis also uncovers a significant mismatch between the domains of user-translated text and existing Tetun

corpora.

Our recommendations include a stronger emphasis on domains that are not well covered by existing corpora (in particular health and science education material), accuracy for the high-resource \rightarrow low-resource direction, performance on very short texts, and the development of mobile-optimized MT solutions. For human-computer interaction (HCI), we suggest focusing on MT perception and improvement for students in low-resource contexts.

This study provides the first observational analysis of real-world MT usage for a low-resource language. It contributes to a more nuanced understanding of the challenges and opportunities in this field. We hope these insights will guide future research efforts, ultimately leading to more effective and accessible MT solutions for speakers of low-resource languages.

Limitations

System design decisions impacting numbers

Design decisions made by the tetun.org volunteers might have impacted usage numbers. For example, apart from MT, tetun.org includes a multilingual word dictionary (with entry definitions and example sentences), and a Tetun spell checker, which might skew the number of users upward, with potentially some users using the app or website without using the MT functionality. Similarly, the authors chose to only allow translation between Tetun and three higher-resource languages (English, Portuguese, Indonesian), which raises the question of how adding support for more languages would impact usage and balance between languages.

Noise in reported data The methods used to determine user location may not accurately reflect the actual location and language background of the users. For instance, Timorese users living abroad might have devices registered outside Timor-Leste, and some users may access tetun.org through a VPN. These potential inaccuracies may affect our understanding of user demographics. Additionally, server logs only capture successful translation attempts, introducing potential bias: users may avoid tetun.org for domains or language pairs where it performs poorly, skewing the data toward areas where tetun.org is more effective.

Generalizability This paper focuses on Tetun, which might limit the applicability of our findings to other languages, given the variability of

¹¹Similar to the Bergamot / [Browser MT](#) project

use cases, user needs and data availability within low-resource communities (Liu et al., 2022). While some of our findings are applicable to other low-resource languages, in particular those that are also institutionalized lingua franca, they are less likely to be applicable to much lower-resourced languages, in particular oral and endangered ones (Bird and Yibarbuk, 2024).

Ethical considerations

Data privacy While working with real end-user data poses privacy risks, we mitigate them through multiple strategies: (1) Server logs do not include user internet protocol (IP) addresses or personal account information; (2) To protect user confidentiality, we apply Named Entity Recognition using spaCy’s `xx_ent_wiki_sm` model to systematically remove names and addresses from translation inputs and outputs before analysis; (3) Our analysis exclusively reports aggregated, anonymized results, ensuring no individual user data can be traced or identified.

Ethics application and user consent We applied for and received a waiver of consent from our institution’s ethics board, as obtaining explicit user consent was not possible given the anonymous nature of the service. This waiver was granted based on several factors: (1) users agree to research use of translation text under the service’s privacy policy, which states: "We reserve the right to use translated text for research and product improvement purposes. All translated text is processed anonymously, ensuring that no personal identifiers are linked to the content"; (2) the service does not collect user accounts or identifying information; and (3) all analysis is conducted on aggregated data. Data management followed strict institutional protocols - all data was transferred directly from the service’s server to our institution’s high-performance computing environment, with no cross-border data flows, and access was restricted to authorized researchers only.

References

Steven Bird and Dean Yibarbuk. 2024. [Centering the Speech Community](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 826–839, St. Julian’s, Malta. Association for Computational Linguistics.

Verena Blaschke, Christoph Purschke, Hinrich Schuetze, and Barbara Plank. 2024. [What Do Dialect Speakers Want? A Survey of Attitudes Towards Language Technology for German Dialects](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 823–841, Bangkok, Thailand. Association for Computational Linguistics.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.

Erik Brynjolfsson, Xiang Hui, and Meng Liu. 2019. Does machine translation affect international trade? evidence from a large digital platform. *Management Science*, 65(12):5449–5460.

Inês Carvalho, Ana Ramires, and Montserrat Iglesias. 2023. Attitudes towards machine translation and languages among travelers. *Information Technology & Tourism*, 25(2):175–204.

Tyler A. Chang, Catherine Arnett, Zhuowen Tu, and Benjamin K. Bergen. 2024. [Goldfish: Monolingual Language Models for 350 Languages](#). *arXiv preprint*. ArXiv:2408.10441.

Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. A survey of multilingual neural machine translation. *ACM Computing Surveys (CSUR)*, 53(5):1–38.

Gabriel de Jesus. 2024. [Text Information Retrieval in Tetun: A Preliminary Study](#). *arXiv preprint*. ArXiv:2406.07331 [cs].

Gabriel de Jesus and Sérgio Sobral Nunes. 2024a. [Data collection pipeline for low-resource languages: A case study on constructing a tetun text corpus](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4368–4380, Torino, Italia. ELRA and ICCL.

Gabriel de Jesus and Sérgio Nunes. 2024b. [Exploring Large Language Models for Relevance Judgments in Tetun](#). *arXiv preprint*. ArXiv:2406.07299 [cs].

Gabriel de Jesus and Sérgio Nunes. 2024c. [Labadain-30k+: A Monolingual Tetun Document-Level Audited Dataset](#). In *Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024*, pages 177–188, Torino, Italia. ELRA and ICCL.

Xinjie Deng and Zhonggen Yu. 2022. [A Systematic Review of Machine-Translation-Assisted Language Learning for Sustainable Education](#). *Sustainability*, 14(13):7598. Number: 13 Publisher: Multidisciplinary Digital Publishing Institute.

Kristin N. Dew, Anne M. Turner, Yong K. Choi, Alyssa Bosold, and Katrin Kirchoff. 2018. [Development of machine translation technology for assisting health communication: A systematic review](#). *Journal of Biomedical Informatics*, 85:56–67.

- Joanna Dolinska, Shekhar Nayak, and Sumittra Suraratcha. 2024. Akha, Dara-ang, Karen, Khamu, Mlabri and Urak Lawoi’ language minorities’ subjective perception of their languages and the outlook for development of digital tools. In *Proceedings of the Seventh Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 94–99, St. Julians, Malta. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. [Beyond English-Centric Multilingual Machine Translation](#).
- James Fiumara, Christopher Cieri, Mark Liberman, Chris Callison-Burch, Jonathan Wright, and Robert Parker. 2022. [The NIEUW Project: Developing Language Resources through Novel Incentives](#). In *Proceedings of the 2nd Workshop on Novel Incentives in Data Collection from People: models, implementations, challenges and results within LREC 2022*, pages 1–7, Marseille, France. European Language Resources Association.
- General Directorate of Statistics Timor-Leste. 2015. [2015 population and housing census](#).
- Zuzana Greksakova. 2018. *Tetun in Timor-Leste: the role of language contact in its development*. doctoralThesis, 00500::Universidade de Coimbra. Accepted: 2018-09-03T09:37:10Z.
- Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. [Survey of Low-Resource Machine Translation](#). *Computational Linguistics*, 48(3):673–732. Place: Cambridge, MA Publisher: MIT Press.
- Paula Sofia Herrera-Espejel and Stefan Rach. 2023. [The Use of Machine Translation for Outreach and Health Communication in Epidemiology and Public Health: Scoping Review](#). *JMIR Public Health and Surveillance*, 9(1):e50814. Company: JMIR Public Health and Surveillance Distributor: JMIR Public Health and Surveillance Institution: JMIR Public Health and Surveillance Label: JMIR Public Health and Surveillance Publisher: JMIR Publications Inc., Toronto, Canada.
- Instituto Nacional de Linguística. 2002. *Matadalan Ortográfiku ba Tetun Nasionál*. Direção do Instituto Nacional de Linguística, RDTL, Dili.
- Eric Khiu, Hasti Toossi, David Anugraha, Jinyu Liu, Jiayu Li, Juan Flores, Leandro Roman, A. Seza Doğruöz, and En-Shiun Lee. 2024. [Predicting machine translation performance on low-resource languages: The role of domain similarity](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1474–1486, St. Julian’s, Malta. Association for Computational Linguistics.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. [Madlad-400: A multilingual and document-level large audited dataset](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 67284–67296. Curran Associates, Inc.
- Moritz Laurer, Wouter van Atteveldt, Andreu Casas, and Kasper Welbers. 2023. [Building Efficient Universal Classifiers with Natural Language Inference](#). *arXiv preprint*. ArXiv:2312.17543 [cs].
- Éric Le Ferrand, Steven Bird, and Laurent Besacier. 2022. [Fashioning Local Designs from Generic Speech Technologies in an Australian Aboriginal Community](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4274–4285, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Heather Lent, Kelechi Ogueji, Miryam de Lhoneux, Orevaoghene Ahia, and Anders Søgaard. 2022. [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.
- Zoey Liu, Crystal Richardson, Richard Hatcher, and Emily Prud’hommeaux. 2022. [Not always about you: Prioritizing community needs when developing endangered language technology](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3933–3944, Dublin, Ireland. Association for Computational Linguistics.
- Manuel Mager, Elisabeth Mager, Katharina Kann, and Ngoc Thang Vu. 2023. [Ethical Considerations for Machine Translation of Indigenous Languages: Giving a Voice to the Speakers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4871–4897, Toronto, Canada. Association for Computational Linguistics.
- Ali Marashian, Enora Rice, Luke Gessler, Alexis Palmer, and Katharina von der Wense. 2025. [From Priest to Doctor: Domain Adaptation for Low-Resource Neural Machine Translation](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7087–7098, Abu Dhabi, UAE. Association for Computational Linguistics.
- Raphaël Merx, Christine Phillips, and Hanna Suominen. 2024. [Machine translation technology in health: A scoping review](#). *Studies in Health Technology and Informatics*, 318:78–83.

Hellina Hailu Nigatu, Atnafu Tonja, Benjamin Rosman, Tamar Solorio, and Monojit Choudhury. 2024. [The Zeno’s Paradox of ‘Low-Resource’ Languages](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17753–17774, Miami, Florida, USA. Association for Computational Linguistics.

Kate Paterson. 2023. Machine translation in higher education: Perceptions, policy, and pedagogy. *TESOL Journal*, 14(2):e690.

Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2023. [Neural Machine Translation for Low-resource Languages: A Survey](#). *ACM Comput. Surv.*, 55(11):229:1–229:37.

Danielle Saunders. 2022. [Domain Adaptation and Multi-Domain Adaptation for Neural Machine Translation: A Survey](#). *Journal of Artificial Intelligence Research*, 75:351–424.

Job Schepens, Roeland van Hout, and T. Florian Jaeger. 2020. [Big data suggest strong constraints of linguistic similarity on adult language learning](#). *Cognition*, 194:104056.

Hedvig Skirgård, Seán G. Roberts, and Lars Yencken. 2017. [Why are some languages confused for others? Investigating data from the Great Language Game](#). *PLOS ONE*, 12(4):e0165934. Publisher: Public Library of Science.

The Asia Foundation. 2023. [Digital Youth in Timor-Leste](#).

Timor-Leste National Institute of Statistics (INETL). 2022. [Timor-Leste Population and Housing Census 2022 - Thematic Report Migration](#). Technical report, INETL.

Lucas Nunes Vieira, Minako O’Hagan, and Carol O’Sullivan. 2021. [Understanding the societal impacts of machine translation: a critical review of the literature on medical and legal use cases](#). *Information, Communication & Society*, 24(11):1515–1532. Publisher: Routledge [eprint: https://doi.org/10.1080/1369118X.2020.1776370](https://doi.org/10.1080/1369118X.2020.1776370).

Jitao Xu, Josep-Maria Crego, and Jean Senellart. 2020. Boosting neural machine translation with similar translations. In *Annual Meeting of the Association for Computational Linguistics*, pages 1570–1579. Association for Computational Linguistics.

A Prompt for domain classification

```
TEXT_PROMPT = """<text to classify>
{text}
</text to classify>"""
```

```
EXAMPLE_USER_1 = TEXT_PROMPT.format(text="Please
meet_me_at_the_park_at_3pm.")
```

```
EXAMPLE_RESPONSE_1 = """<rationale>
```

```
This text is about going to the park with a
friend, a situation that is part of daily life
experiences.
</rationale>
```

```
<domain>
daily life & personal experiences
</domain>"""
```

```
EXAMPLE_USER_2 = TEXT_PROMPT.format(text="
Mindfulness-based_interventions_(MBIs)_are_
increasingly_being_integrated_into_oncological_
treatment_to_mitigate_psychological_distress_and_
_promote_emotional_and_physical_well-being.")
EXAMPLE_RESPONSE_2 = """<rationale>
This text is about the use of mindfulness-based
interventions in oncological treatment, a health
topic. It seems to be from a scientific or
research context.
</rationale>
```

```
<domain>
healthcare & medicine; science & research
</domain>"""
```

```
EXAMPLE_USER_3 = TEXT_PROMPT.format(text="The_
government_has_announced_new_policies_to_boost_
employment_opportunities_in_rural_areas.")
EXAMPLE_RESPONSE_3 = """<rationale>
This text discusses government policies aimed at
improving employment, which relates to both
governance and socio-economic issues, as well as
business and employment.
</rationale>
```

```
<domain>
government & socio-economic issues; business &
work & employment
</domain>"""
```

```
EXAMPLE_USER_4 = TEXT_PROMPT.format(text="Neraca
adalah_suatu_laporan_keuangan_jadi_di_dalamnya_
pasti_terdapat_tiga_bagian_terpenting_yaitu_aset_
_atau_harta,_liabilitas_atau_utang,_dan_ekuitas.
")
EXAMPLE_RESPONSE_4 = """<rationale>
This text describes a financial report,
specifically the balance sheet, which is an
essential component of accounting and finance.
It belongs to the business and finance domain.
</rationale>
```

```
<domain>
business & work & employment
</domain>"""
```

```
messages = [
    {"role": "system", "content": f"You_are_a_
linguist_assistant,_helping_classify_user-
inputted_text_into_a_list_of_domains:<list_of_
domains>{';'.join(TOPICS)}</list_of_domains>._
You_will_be_given_text_to_classify,_and_asked_
_the_domains_that_this_text_belongs_to._Only_ever_
_use_the_domains_listed_in<list_of_domains>"},
    {"role": "user", "content": EXAMPLE_USER_1},
    {"role": "system", "content":
EXAMPLE_RESPONSE_1},
    {"role": "user", "content": EXAMPLE_USER_2},
    {"role": "system", "content":
EXAMPLE_RESPONSE_2},
```

```

    {"role": "user", "content": EXAMPLE_USER_3},
    {"role": "system", "content":
EXAMPLE_RESPONSE_3},
    {"role": "user", "content": EXAMPLE_USER_4},
    {"role": "system", "content":
EXAMPLE_RESPONSE_4},
]

```

B Perplexity results

Topic	Average perplexity
Science & research	1,208
Healthcare & medicine	1,198
Education	1,372
Business & work & empl.	1,871
Daily life & personal exp.	1,864
Government & socio. issues	632
Law & regulations	1,274
Technology	3,434
Religion & spirituality	1,127

Table 5: Perplexity per topic, as measured on Tetun MT input documents of 5 words or more.

Provenance	Average perplexity
Education & research material	1,321
Organizational & formal docs	1,662
News article & press release	874
Conversation & correspond.	1,866
Literary fiction	2,119
Religious text	1,473
Other	3,876

Table 6: Perplexity per provenance, as measured on Tetun MT input documents of 5 words or more.