

NüshuRescue: Reviving the Endangered Nüshu Language with AI

Ivory Yang¹ Weicheng Ma² Soroush Vosoughi³

^{1,2,3}Department of Computer Science, Dartmouth College

{Ivory.Yang.GR, Weicheng.Ma, Soroush.Vosoughi}@dartmouth.edu

Abstract

The preservation and revitalization of endangered and extinct languages is a meaningful endeavor, conserving cultural heritage while enriching fields like linguistics and anthropology. However, these languages are typically low-resource, making their reconstruction labor-intensive and costly. This challenge is exemplified by Nüshu, a rare script historically used by Yao women in China for self-expression within a patriarchal society. To address this challenge, we introduce NüshuRescue, an AI-driven framework designed to train large language models (LLMs) on endangered languages with minimal data. NüshuRescue automates evaluation and expands target corpora to accelerate linguistic revitalization. As a foundational component, we developed NCGold, a 500-sentence Nüshu-Chinese parallel corpus, the first publicly available dataset of its kind. Leveraging GPT-4-Turbo, with no prior exposure to Nüshu and only 35 short examples from NCGold, NüshuRescue achieved 48.69% translation accuracy on 50 withheld sentences and generated NCSilver, a set of 98 newly translated modern Chinese sentences of varying lengths. In addition, we developed FastText-based and Seq2Seq models to further support research on Nüshu. NüshuRescue provides a versatile and scalable tool for the revitalization of endangered languages, minimizing the need for extensive human input. All datasets and code have been made publicly available at <https://github.com/ivoryyang/NushuRescue>.

1 Introduction

Nüshu is a unique syllabic script originating from the ethnic Yao women in Jiangyong County, Hunan Province, China (Broussard, 2008; Liu, 2014). This rare and historically significant script, often referred to as “women’s script” (Wang, 2020), is the only written language in the world developed by and used exclusively by women (Ann-Gee,

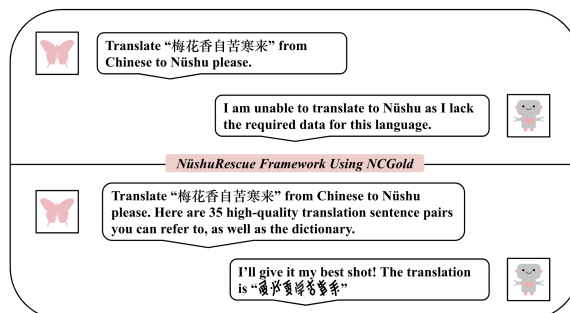


Figure 1: A simplified, stylized rendition of the NüshuRescue generation framework, using samples from NCGold.

2013). In a past era of high female illiteracy, where women’s voices were largely suppressed (Leung, 2003), Nüshu served as a private means of communication and form of female empowerment in a male-dominated society (Liu, 1997). Today, Nüshu is an endangered language with very limited surviving documentation (Luo et al., 2022), making it a critical focus for linguistic preservation and revitalization efforts (Liu, 2018).

In the broader field of Natural Language Processing (NLP), low-resource languages like Nüshu present significant challenges, such as the scarcity of available data (Magueresse et al., 2020) and the absence of established linguistic tools (Joshi et al., 2019; Krasadakis et al., 2024). The preservation of such languages is not only vital for maintaining linguistic diversity (Gibbs, 2002; Brenzinger and de Graaf, 2009), but also for safeguarding the cultural heritage (Romaine, 2007) and historical narratives (Hale, 1992) they encapsulate. To preserve the Nüshu language, this paper provides two main contributions: (1) **NüshuRescue**: A LLM-based data generation framework to construct or scale up text corpora for low-resource languages, using minimal human labor and without presuming any prior knowledge of the target languages. (2.1) **NCGold**: An expert-validated parallel corpus of 500 Nüshu-

Chinese translation sentence pairs. **(2.2) NCSilver:** A parallel corpus of 98 Nüshu-Chinese translation sentence pairs generated through the NüshuRescue framework. Both NCGold and NCSilver are the first publicly available datasets of their kind.

Through our experiments, we find that LLMs like GPT-4-Turbo demonstrate strong language learning capabilities, even for languages with limited prior exposure, allowing the NüshuRescue framework to effectively generate higher-quality text for truly resource-scarce languages. While human expertise remains crucial for annotating seed data and ensuring translation accuracy, NüshuRescue significantly reduces the amount of human labor required compared to fully manual data annotation. As such, we believe that our NüshuRescue framework, as shown in Figure 1, offers a promising solution for the preservation and revitalization of endangered languages.

2 Related Work

Nüshu has garnered renewed cultural interest in recent years (Thurnell-Read et al., 2022), especially after receiving international representation in the Hollywood production “Snow Flower and the Secret Fan” (Dargis, 2011). However, existing linguistic work on Nüshu is limited; “A Compendium of Chinese Nüshu” (Zhao, 1992)¹ remains the only formal print publication with expert-validated translations. Chiu (2012) provided a new perspective in “Heroines of Jiangyong” by compiling Nüshu scripts alongside English translations of their Chinese counterparts. While this is one step removed from a direct Nüshu-to-English translation and lacks formal linguistic mapping, it is nonetheless valuable for introducing Nüshu to the English-speaking domain.

There have been only tentative attempts to explore Nüshu within the NLP community. Tang et al. (2023) and Sun et al. (2023) proposed a two-part project combining research and art, where two AI agents trained on a Chinese-Nüshu dictionary collaborate to create a standard writing system for encoding Chinese, blending linguistic analysis with artistic interpretation. However, no publicly available digital dataset for Nüshu exists, and using NLP methods to analyze its syntax and semantics remains virtually unexplored. The only significant digitalization effort to date has been the creation

¹The very first version of this book was published in 1992; we referenced the most recent e-book released in 2019.




Figure 2: Comparison of mapped Chinese (logographic) and Nüshu (syllabic) characters for the word “ocean”.

of a Nüshu-Chinese Unicode dictionary (Contributors), which remains a work in progress as it continues to be expanded and standardized.

Efforts to document and preserve other under-resourced and endangered languages have been more thoroughly explored. A common challenge across endangered language preservation projects is the need to gather data at scale. One approach is the collection of text corpora through web crawling and data extraction, as demonstrated by de Jesus and Nunes (2024) for Tetun, a language of Timor-Leste. However, this method relies on the presence of online resources, which is not the case for Nüshu. Another approach, transcription, and human annotation is more costly and labor-intensive, as seen in efforts to preserve Yoruba, a West African language (Fagbolu et al., 2015). Given these constraints, it is crucial to combine human annotation efforts with AI and NLP tools, so as to advance Nüshu preservation in a cost-effective and scalable manner.

3 The Nüshu Language

Nüshu is a writing system originally created and exclusively used by ethnic Yao women in rural China (Bissessar, 2013), and therefore possesses internal linguistic and literary rules difficult for those outside the community to fully grasp. Nüshu emerged independently as a phonetic writing system, borrowing sounds from spoken Chinese (particularly the local dialect) and representing them using a syllabic system (Zhao, 2014). Nüshu is *syllabic*, meaning it uses symbols to represent individual syllables in spoken language (Silber, 1995). In contrast, Chinese is *logographic*, meaning it uses characters that represent words or morphemes, with each character typically corresponding to a meaning rather than a specific sound (Tan et al., 2001). For example, Figure 2 illustrates the Nüshu and

Nüshu	Chinese
𠃉	吃, 系, 音, 阴, 因, 姻, 殷, 忆, 与, 裕, 依, 医, 衣, 要, 以, 已, 孕, 荫, 于, 妖, 腰, 语, 药, 叶, 易, 应, 遗, 喻, 瑜, 幼, 意, 任, 如, 儒, 污
𠃉, 𠃉, 𠃉, 𠃉, 𠃉, 𠃉	长

Figure 3: Nüshu-Chinese one-to-many and many-to-one mapping; in extreme cases, one Nüshu word can map up to 35 Chinese words.

Chinese characters for the word “ocean”, shown on the left and right, respectively. In Chinese, the character for “ocean” is logographic and conveys the concept as a whole, with the “water radical” (the three strokes on the left of the word) indicating its connection to water. Removing the “water radical” changes the character’s meaning to “each” or “every”, illustrating how radicals² influence the meaning of Chinese characters. In contrast, the Nüshu symbol for “ocean” is syllabic, representing the phonetic sounds of the spoken word rather than the concept itself.

The phonetic nature of Nüshu creates a multi-mapping challenge with Chinese, as shown in Figure 3. A single Nüshu symbol can correspond to multiple Chinese characters with the same pronunciation but different meanings (one-to-many mapping). In contrast, a single Chinese character may map to several Nüshu symbols representing different phonetic variations (many-to-one mapping) (Moratto, 2022). Although an estimated 600–700 Nüshu characters exist, only 398 have been officially encoded in Unicode Noto Sans Nüshu (Emerson, 2002), leaving many symbols undocumented and creating gaps in translation efforts. Traditional methods like word alignment models (Och, 2003) and statistical machine translation (Sharma et al., 2023), which rely on large, parallel corpora, struggle to handle Nüshu’s complex relationships and limited dataset. Advanced techniques are needed to accommodate Nüshu’s unique structure and support corpus expansion and alignment with other languages.

4 NüshuRescue

This section details our efforts in (1) initializing NLP research on Nüshu via manual translation

²In Chinese, radicals are the building blocks of Chinese characters, often used as a means to organize them in dictionaries. They usually represent common semantic or phonetic components.



Figure 4: Manual Data Collection: Step-by-step transcription process.

(Section 4.1) and ChatGPT-in-the-loop data annotation (Section 4.2), (2) training a base language model for Nüshu (Section 5.1), and (3) enabling machine translation from Nüshu to Chinese (Section 5.2) or from Nüshu to multiple other languages (Section 5.3).

4.1 Manual Data Collection

“A Compendium of Chinese Nüshu” (Zhao, 1992) is the most comprehensive, expert-validated collection of Nüshu calligraphy available (Lo, 2013). It features scanned Nüshu scripts, created with paper and ink, positioned at the top of each page, with the corresponding Chinese translation at the bottom. A sample from the book is shown in Figure A1 in the Appendix. While most works consist of traditional ancient Chinese poems and literature, the collection also includes personal diary entries and an impassioned soliloquy written in response to a mother’s scolding, providing contextual diversity.

The manual data collection process for the 500-sentence NCGold dataset was carried out with meticulous attention to detail. The annotators for the task are college-educated, bilingual in English and Chinese, and have experience in computational linguistics. The step-by-step process is illustrated in Figure 4³. First, each sentence of the Chinese text translations was manually typed out, with the punctuation left out. Next, the Chinese text was copied and pasted into the official online website, Converter of Calligraphy Copybook of Nüshu Standard Characters (Converter), to retrieve all possible

³English Translation: “A noble man empowers the beauty within others”.

Nüshu	Chinese
𠄎	突, 途
𠄏	牌, 裴
𠄐	危
𠄑	萍
𠄒	舟
𠄓	帘, 轮
𠄔	扁
𠄕	雾
𠄖	波
𠄗	晓

Figure 5: A small sample of newly mapped words not in the official Nüshu-Chinese dictionary; an original contribution.

mapped Nüshu words for each Chinese word. Finally, the annotators carefully filtered and matched the single correct Nüshu word to each corresponding Chinese word.

Of the 500 total sentences, 127 had full coverage in the official Nüshu Unicode dictionary, with every Chinese word included. For the remaining 373 sentences, while some Chinese words were covered, others were missing from the limited vocabulary of the dictionary, meaning the online converter could not provide Nüshu mappings for all terms. In these ambiguous cases, we manually consulted the Nüshu column of the official dictionary and visually identified the correct Nüshu characters from the book. As a result, NCGold not only provides the first expert-validated Nüshu-Chinese parallel corpus but also contributes newly mapped Nüshu-Chinese word pairs, a small sample of which is shown in Figure 5 (Figure A2 in the Appendix shows a higher resolution version).

4.2 LLM-Based Data Augmentation

After obtaining NCGold, we established a text translation framework for scaling up the dataset by leveraging LLMs. For clarity, we refer to the manually-translated dataset as **NCGold** and the subsequent LLM-generated dataset as **NCSilver**.

One notable merit of our approach is that the LLM does not require any prior knowledge of Nüshu characters before being introduced to the

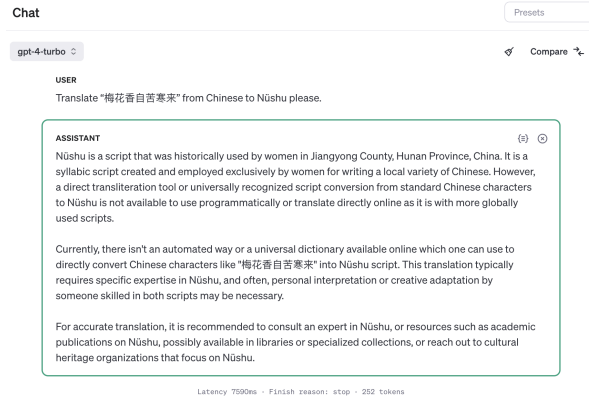


Figure 6: GPT-4-Turbo does not understand Nüshu at the beginning.

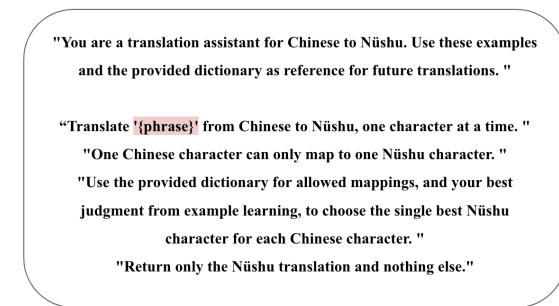


Figure 7: Prompt used on GPT-4-Turbo.

provided examples, as demonstrated in Figure 6⁴. While we use the GPT-4-Turbo model in our experiments, our proposed framework is model-agnostic, allowing for the use of any LLM with comparable language modeling and reasoning capabilities, as illustrated in Figure A3 in the Appendix.

4.2.1 Experimental Settings

In our framework for obtaining Nüshu translations from the LLMs, we first provide translation examples from NCGold and the official Nüshu-Chinese dictionary for GPT-4-Turbo to learn from, uploading them as files to avoid excessively long prompts. The model is then prompted to generate Nüshu translations for Chinese input sentences, as shown in Figure 7.

For quality assurance, 35 NCGold samples were provided for learning, and a round of evaluations was conducted on 50 withheld samples. We manually reviewed the LLM-generated translations for validation, correcting minor character choice errors. The choice of 35 samples was informed by prelim-

⁴English Translation: “A fragrant plum blossom comes from the bitter winter cold”. This is a traditional Chinese idiom that means for good things to happen, one often needs to go through hardship.

inary tests, where fewer examples often led the model to refuse responses. Since there is no strict upper limit for the number of NCGold examples (aside from token limits), incorporating more high-quality examples could further enhance translation performance.

Round	Data Sample	Average Length
1	1-30	10.33
2	31-60	13.63
3	61-90	16.27
4	91-120	19.40
5	121-150	23.10
6	151-180	31.73

Table 1: Average sentence length across data samples for different rounds of annotations.

The UD Chinese GSDSimp dataset (Wang et al., 2021) was used to construct a set of new Chinese sentences for annotation by GPT-4-Turbo. First, we filtered out sentences containing out-of-dictionary words based on our Nüshu-Chinese dictionary, and then categorized the remaining sentences by ascending length. Table 1 details the data samples annotated in each round, along with the average sentence length. Arabic numbers were also removed to avoid biasing the translation evaluations.

Through iterative annotations using ChatGPT and manual corrections, we progressively expanded the translation process to include longer sentences, resulting in a more comprehensive set of Nüshu-Chinese translation instances. This iterative approach to sentence length enabled the generation of progressively longer translations, surpassing the limitations of the existing dataset. However, translation accuracy for longer sentences was slightly lower than for shorter ones, likely due to the fact that all seed sentences were derived from short, ancient Chinese poems.

4.2.2 Generation Pipeline Setup

The pipeline of our LLM-in-the-loop data annotation approach is illustrated in Figure 8, where the LLM agent is queried to translate Chinese sentences into Nüshu using few-shot prompting. In the initial round (R1), seed translation examples are sampled from the NCGold dataset, while subsequent example sets are drawn from a combination of NCGold and NCSilver after the first round of translations. Beginning in the second round (R2), newly generated samples from the previous round are shuffled, with five samples selected to replace




Figure 8: Pipeline for LLM-in-the-loop data annotation approach (Rounds 1-6).

Translating sentences 2 to 50 from 85perfect.csv.
 Translating Chinese sentence 1: 星汉灿烂若出其里
 Translated Nushu: 星汉灿烂若出其里
 Waiting 10 seconds before next translation...

Figure 9: Translation of unseen characters (boxed in pink).

the top 5 samples in the original 35-sample list and the bottom five samples removed. For instance, in R2, the 35 training samples consisted of the top 5 new translations from R1, followed by 30 samples from the original set. By the third round (R3), the 35 samples included five new translations from R2, five from R1, and 25 from the original set. The control test consistently used the original 35 samples across all rounds.

A rule-based translation length validator is integrated into the pipeline to ensure that each Nüshu character maps to a single Chinese character. This validator ensures that the original and translated contents have the same length. If a length mismatch occurs, the generation is rejected, and the LLM is prompted again to resolve the issue. The model is allowed up to 7 retries before moving on and classifying the instance as a failed translation. The enforced one-to-one character translations compensates for limited data, by providing additional linguistic instruction for Nüshu. Our experiments without enforced length validation reported a 31.37% accuracy, versus 48.69% with length validation, showing that data scarcity can be somewhat compensated through incorporating language-specific information.

It is critical to note that while the ChatGPT model is capable of generalizing knowledge from NCGold to unseen characters to produce seemingly valid translations, as exemplified in Figure 9, ex-




Figure 10: Translation success and failure across rounds.

pert knowledge of the Nüshu language is essential to validate these translations and register new characters. Due to the absence of such expert-level knowledge in our current work, we limit our testing to the model’s ability to translate only in-dictionary characters. Nevertheless, these findings suggest that our proposed framework could be highly applicable in future Linguistics research, aiding in the expansion of the Nüshu dictionary and the revival of the language.

4.2.3 Data Annotation Outcomes

Using only 35 sentences from NCGold as seed examples, the GPT-4-Turbo model achieved an accuracy of 48.69% in perfectly translating the 50 left-out NCGold examples on a strict character-by-character match. A manual review of the translations revealed that minor translation errors were predominantly due to ambiguity in Nüshu character selection, as each Nüshu character can correspond to multiple Chinese characters and vice versa.

This complexity in mapping contributed to most of the discrepancies observed. For out-of-domain Chinese sentences, the GPT-4-Turbo model successfully translated 70.00% of sentences with lengths below 16 characters. However, it frequently refused to respond when the sentence lengths exceeded 31 characters. Figure 10 displays the number of successfully translated sentences and failed

Examples	Chinese Input	Nüshu Translated Output
Correct	年时有入口万千人左右	年时有入口万千人左右
Incorrect	但这就是我当时最真实的感觉我不能逃避	但这就是我当时最真实的感觉我不能逃避
Not Translated (Length Mismatch)	调查员亦发现当时负责指示公务机的管制员并没有指示公务机下降至指定高度当时公务机应从尺下降至尺	—

Figure 11: Out-of-domain examples; for the incorrect example, wrongly translated words are in pink.

translations for each round, along with the results from the control group, which is always provided the original 35 learning samples. Figure 11 exemplifies the correct, incorrect, and untranslatable responses from ChatGPT in our experiments.

4.2.4 Other Low-Resource Languages

To evaluate the scalability of our framework, we extended our investigation to additional low-resource languages with somewhat larger datasets than Nüshu but still underrepresented in mainstream language models like GPT-4-Turbo. Cherokee was selected as a representative Native American language, recognized for its moderate usage yet classified as low-resource (Zhang et al., 2022). Despite GPT-4-Turbo’s partial familiarity with Cherokee (Shu et al., 2024), it struggled with accurate

Train Data	BLEU-1	BLEU-2	BLEU-3	METEOR	ROUGE-1	ROUGE-2	ROUGE-L
100	0.0244	0.0074	0.0054	0.0125	0.0245	0.0	0.0245
200	0.0358	0.0147	0.0090	0.0220	0.0396	0.0048	0.0379
300	0.0351	0.0173	0.0133	0.0216	0.0381	0.0094	0.0381
400	0.0705	0.0234	0.0165	0.0398	0.0812	0.0036	0.0767
421	0.0642	0.0231	0.0142	0.0372	0.0748	0.0065	0.0730
432	0.0893	0.0392	0.0313	0.0568	0.1053	0.0182	0.1037
453	0.0758	0.0394	0.0295	0.0542	0.0856	0.0230	0.0825
471	0.0977	0.0429	0.0296	0.0633	0.1174	0.0219	0.1115
487	0.1365	0.0685	0.0528	0.0942	0.1609	0.0408	0.1568
498	0.1093	0.0573	0.0404	0.0760	0.1367	0.0358	0.1308

Table 2: Evaluation results (averaged) for our Seq2Seq MT Model, evaluated on a fixed 100 test sample.

translations. Implementing our framework on a 14k Cherokee-English corpus (Zhang et al., 2020) yielded translation accuracies of 28%, 40%, and 42% with 35, 150, and 300 sample sizes, respectively. These results highlight the persistent challenge of achieving higher accuracy in low-resource language translation without tailored linguistic adaptations.

5 Applications

Leveraging the NCGold and NCSilver datasets, we present both a Nüshu-Chinese language model and a Nüshu-to-Chinese translation model to demonstrate the potential of our data-annotation framework in advancing future research on Nüshu. Our experiments also underscore the critical need for further expanding the Nüshu corpus, enabling researchers to gain deeper insights into the language and the rich cultural heritage it embodies.

5.1 Language Modeling in Nüshu

Having obtained a set of 598 Nüshu-Chinese translation sentence pairs, we employed the FastText library (Athiwaratkun et al., 2018) to train a bilingual skip-gram language model (Mikolov et al., 2013) in both Nüshu and Chinese. Specifically, we separated the Nüshu-Chinese translation pairs from NCGold and NCSilver into individual sentences, shuffled them, and utilized this bilingual corpus to train the FastText model.

For the FastText model, we set the embedding vector size to 300, a window size of 10, and a minimum count of 5. We used the skip-gram model (sg=1), which is well-suited for smaller datasets and rare words, as it predicts surrounding words given a target word—particularly useful for the infrequent characters in the Nüshu corpus.

The model was trained for 20 epochs to ensure convergence without over-fitting, with a negative sampling rate of 10 to speed up training by focusing on relevant examples. The default learning rate of 0.05 is used for training our model.

5.2 Nüshu-to-Chinese Translation

We initially trained a Seq2Seq model for Nüshu-Chinese translation using training samples from NCGold, with a batch size of 3 and for 15 epochs. For the test set, we first shuffled the 500-sentence NCGold and fixed the first 100 samples as the test set. Despite the small dataset size, the model’s performance was promising, comparable to models developed for other truly low-resource languages (Baruah et al., 2021).

Encouraged by these results, we combined NCGold with NCSilver, resulting in a total of 598 sentences, to further improve translation quality. While ChatGPT demonstrated potential in translating Chinese text to Nüshu for NCSilver, it remains a closed-source model with limited interpretability, and its outputs are often noisy, requiring multiple queries for acceptable translations. To mitigate these issues and enhance both translation accuracy and model transparency, we retrained the Seq2Seq model using the combined NCGold and NCSilver datasets. We provide further details on this process in Section 6.

5.3 Nüshu Translates to Other Languages

Given the limited availability of official Nüshu-Chinese translations, data for translating Nüshu into other languages is even more scarce. As mentioned in Section 2, Chiu (2012) published a collection of works featuring Nüshu calligraphy alongside English translations of their corresponding Chinese poems. However, this is not a direct

Nüshu-to-English translation and lacks certain linguistic nuances or fidelity (Rosa et al., 2017). Nevertheless, it introduces the right concept of using Nüshu as a bridge for multilingual translations.

Our framework enables the construction of larger Nüshu datasets than ever before, making it possible to scale up Nüshu-to-other-language translation datasets when combined with such works. This approach opens up new avenues for expanding Nüshu’s reach across languages and enhancing the diversity of translation datasets.

6 Discussion

Given the promising early results of our Seq2Seq model trained on instances from NCGold against the fixed 100 test set, we hypothesized that the model’s performance was limited by the small dataset size. We further proposed that, by continuing to use the data generation framework to incrementally enlarge the dataset, we could achieve even better model performance. To validate this, we trained the Seq2Seq model incrementally with subsets of NCGold, starting with 100, 200, 300 sentences, and so on, up to 400 sentences. The results, as shown in Table 2, confirm our hypothesis, demonstrating that larger dataset sizes consistently improve performance across all evaluated metrics.

To further improve the model, we gradually incorporated the 98 newly generated translations from NCSilver, adding them sequentially in the same quantities as generated during each round of the data generation process. As expected, model performance improved as the dataset expanded, reinforcing the importance of a larger and more diverse training corpus for Nüshu-Chinese translation tasks. This trend is clearly reflected in the steady increase in BLEU, METEOR, and ROUGE scores, demonstrating a positive correlation between dataset size and translation quality.

However, we observed that performance peaked at a training dataset of 487 sentences. This dataset includes 400 original sentences from our book translations and 87 sentences generated through Round 5 (R5). Interestingly, performance declined when data from Round 6 (R6) was added. We hypothesize that this decline is due to the sentences in R6 being significantly longer than those in the rest of the dataset. As shown in Table 1, the average sentence length for R6 is 31.73, compared to 23.1 for R5, while the maximum sentence length in the original 500-sentence dataset is only 18.

We hypothesize that the model struggles with excessively long sentences because it was exposed to primarily shorter training data. Seq2Seq models, particularly those using a fixed-length encoder-decoder architecture, tend to perform well on sentence structures and lengths similar to those seen during training. When presented with longer, more complex sentences, the model may struggle to effectively encode the entire sentence into a fixed representation, leading to loss of important information, truncation, or difficulty decoding the translation. Furthermore, attention mechanisms can become less effective with longer sequences, especially when the model has not been fine-tuned or adequately trained. Consequently, incorporating longer sentences without balancing the model’s exposure to various sentence lengths likely contributed to the performance drop observed after R6.

7 Conclusion and Future Work

In this paper, we introduced NüshuRescue, a novel approach leveraging LLMs to expand text corpora in endangered languages, with Nüshu as a case study. Our method facilitates research on these languages while minimizing reliance on extensive human annotations. Findings show that NüshuRescue can generate high-quality translations from Chinese to Nüshu with only 35 seed pairs, demonstrating its effectiveness in producing accurate translations with minimal input. The model also generalizes well across different sentence lengths and text domains. Notably, NüshuRescue does not require the underlying LLM to have prior knowledge of the target language, making it adaptable to other low-resource languages.

In addition, we digitized 500 Nüshu-Chinese sentence pairs from ancient Chinese poetry, which were used to train both a FastText language model and a Seq2Seq Nüshu-to-Chinese MT model, in combination with the LLM-generated translations. This dataset, encompassing 80% of the only expert-validated Nüshu resource available, is a vital step in preserving Nüshu from extinction. In the absence of other digital resources, our paper serves as a foundational contribution, introducing Nüshu to the NLP community for the first time and providing a data augmentation framework. Looking ahead, NüshuRescue has potential for adaptation for other under-resourced languages, fostering the expansion of corpora and enabling the development of larger, more sophisticated language models.

Limitations

In this paper, we introduced NüshuRescue, a LLM-based framework designed to scale up text corpora in resource-scarce languages. While our experiments demonstrated its effectiveness and generalizability, we encountered limitations in significantly expanding the Nüshu corpus due to the restricted scope of the current Nüshu-Chinese dictionary. The only available expert-validated source of Nüshu content comes from “A Compendium of Chinese Nüshu” (Zhao, 1992), which contains only 374 Nüshu characters and 1453 corresponding Chinese characters.

Although NüshuRescue demonstrated a strong ability to generalize to out-of-vocabulary characters, the absence of expert validation⁵ prevents us from safely releasing the generated data. Consequently, we sampled only 180 sentences from the UD Chinese GSDSimp dataset (Wang et al., 2021) dataset to construct the NCSilver dataset, ensuring that all included Chinese characters were within the existing Nüshu vocabulary. However, as future studies of Nüshu incorporate more characters, NüshuRescue can be easily adapted to further expand the Nüshu corpus.

Ethics Considerations

In this paper, we exclusively utilized publicly available resources, including the “A Compendium of Chinese Nüshu” book (Zhao, 1992) and the UD Chinese GSDSimp dataset (Wang et al., 2021) for our experiments. We relied on off-the-shelf LLMs, such as GPT-4-Turbo, as the backbone model for NüshuRescue. Furthermore, we have thoroughly documented the prompts used to generate Nüshu translations from the LLMs, detailing how these prompts were iteratively refined to improve the quality of the outputs.

Given these considerations, we do not foresee any ethical concerns arising from our work. Our NCGold and NCSilver datasets, as well as related code, have been made publicly available through our GitHub (<https://github.com/ivoryayang/NushuRescue>), ensuring transparency and fostering further research in this area.

⁵Nüshu no longer has any native speakers, as the last known native speaker passed in 2004.

References

- Lee Ann-Gee. 2013. Sustaining women’s culture and preserving a “secret” script: The case of nüshu. *The International Journal of Sustainability in Economic, Social and Cultural Context*, 8(4):89.
- Ben Athiwaratkun, Andrew Gordon Wilson, and Anima Anandkumar. 2018. Probabilistic fasttext for multi-sense word embeddings. *arXiv preprint arXiv:1806.02901*.
- Rupjyoti Baruah, Rajesh Kumar Mundotiya, and Anil Kumar Singh. 2021. Low resource neural machine translation: Assamese to/from other indo-aryan (indic) languages. *Transactions on Asian and Low-Resource Language Information Processing*, 21(1):1–32.
- Charmaine Bissessar. 2013. The movement from secret acts of defiance to manifestation of women’s empowerment. *Advancing Women in Leadership Journal*, 33:69–74.
- Matthias Brenzinger and Tjeerd de Graaf. 2009. Documenting endangered languages and maintaining language diversity. *Linguistic Anthropology*, 238.
- Julia T Broussard. 2008. Nüshu: a curriculum of women’s identity. *TCI (Transnational Curriculum Inquiry)*, 5(2):45–68.
- Elena Suet-Ying Chiu. 2012. Heroines of jiangyong: Chinese narrative ballads in women’s script.
- Nüshu Script Contributors. Nüshu-chinese unicode dictionary. https://github.com/nushu-script/unicode_nushu. Accessed: 2024-09-15.
- Nüshu Script Online Converter. Converter of calligraphy copybook of nüshu standard characters. <https://nushuscript.org/en-US/unicode/>. Accessed: 2024-09-15.
- Manohla Dargis. 2011. Ties that bind: Life-sustaining friendships transcending life’s brutalities. *The New York Times*, pages C12–L.
- Gabriel de Jesus and Sérgio Sobral Nunes. 2024. 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.
- Thomas Emerson. 2002. Chinese script simplification. *Multilingual Computing & Technology*, 52(13):41–46.
- Olutola Fagbolu, Akinwale Ojoawo, Kayode Ajibade, and Boniface Alese. 2015. Digital yoruba corpus. *International Journal of Innovative Science, Engineering and Technology*, pages 2348–7968.
- W Wayt Gibbs. 2002. Saving dying languages. *Scientific American*, 287(2):78–85.

- Ken Hale. 1992. Endangered languages: On endangered languages and the safeguarding of diversity. *language*, 68(1):1–42.
- Pratik Joshi, Christain Barnes, Sebastin Santy, Simran Khanuja, Sanket Shah, Anirudh Srinivasan, Satwik Bhattamishra, Sunayana Sitaram, Monojit Choudhury, and Kalika Bali. 2019. Unsung challenges of building and deploying language technologies for low resource language communities. *arXiv preprint arXiv:1912.03457*.
- Panteleimon Krasadakis, Evangelos Sakkopoulos, and Vassilios S Verykios. 2024. A survey on challenges and advances in natural language processing with a focus on legal informatics and low-resource languages. *Electronics*, 13(3):648.
- Christie KK Leung. 2003. Women who found a way creating a women’s language. *off our backs*, 33(11/12):40–43.
- Fei-Wen Liu. 1997. *Women who de-silence themselves: Male-illegible literature (Nushu) and female-specific songs (Nuge) in Jiangyong County, Hunan Province, China*. Syracuse University.
- Fei-wen Liu. 2014. The biographical writing of nüshu: Its historical implications and cultural practice in changing rural china. *Research on Women in Modern Chinese History/Jindai Zhongguo Funu Shi Yanjiu*, 23.
- Fei-wen Liu. 2018. Practice and cultural politics of “women’s script”: nüshu as an endangered heritage in contemporary china. In *Women Writing Across Cultures*, pages 231–246. Routledge.
- Yuen-yi Lo. 2013. She is a fragment and a whole: A narrative of nüshu artefacts. *Australian Feminist Studies*, 28(78):398–416.
- Wenbin Luo, Yulian Lu, Dallen J Timothy, and Xiaolin Zang. 2022. Tourism and conserving intangible cultural heritage: Residents’ perspectives on protecting the nüshu female script. *Journal of China Tourism Research*, 18(6):1305–1329.
- Alexandre Magueresse, Vincent Carles, and Evan Heetderks. 2020. Low-resource languages: A review of past work and future challenges. *arXiv preprint arXiv:2006.07264*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Riccardo Moratto. 2022. Nüshu in translation: Challenges of translating a unique chinese cultural heritage. In *Chinese literature in the world: Dissemination and translation practices*, pages 195–206. Springer.
- Franz Josef Och. 2003. *Statistical machine translation: From single word models to alignment templates*. Ph.D. thesis, Aachen, Techn. Hochsch., Diss., 2002.
- Suzanne Romaine. 2007. Preserving endangered languages. *Language and linguistics compass*, 1(1-2):115–132.
- Alexandra Assis Rosa, Hanna Pięta, and Rita Bueno Maia. 2017. Theoretical, methodological and terminological issues regarding indirect translation: An overview.
- Sonali Sharma, Manoj Diwakar, Prabhishkek Singh, Vijendra Singh, Seifedine Kadry, and Jungeun Kim. 2023. Machine translation systems based on classical-statistical-deep-learning approaches. *Electronics*, 12(7):1716.
- Peng Shu, Junhao Chen, Zhengliang Liu, Hui Wang, Zihao Wu, Tianyang Zhong, Yiwei Li, Huaqin Zhao, Hanqi Jiang, Yi Pan, et al. 2024. Transcending language boundaries: Harnessing llms for low-resource language translation. *arXiv preprint arXiv:2411.11295*.
- Cathy Lyn Silber. 1995. *Nushu (Chinese women’s script) literacy and literature*. University of Michigan.
- Yuqian Sun, Yuying Tang, Ze Gao, Zhijun Pan, Chuyan Xu, Yurou Chen, Kejiang Qian, Zhigang Wang, Tristan Braud, Chang Hee Lee, et al. 2023. Ai nüshu: An exploration of language emergence in sisterhood through the lens of computational linguistics. In *SIGGRAPH Asia 2023 Art Papers*, pages 1–7.
- Li Hai Tan, Ho-Ling Liu, Charles A Perfetti, John A Spinks, Peter T Fox, and Jia-Hong Gao. 2001. The neural system underlying chinese logograph reading. *Neuroimage*, 13(5):836–846.
- Yuying Tang, Yuqian Sun, Ze Gao, Zhijun Pan, Zhigang Wang, Tristan Braud, Chang Hee Lee, and Ali Asadipour. 2023. Ai nüshu (women’s scripts)-an exploration of language emergence in sisterhood. In *ACM SIGGRAPH Asia 2023 Art Gallery*, pages 1–2.
- Thomas Thurnell-Read, Michael Skey, and Marie Heřmanová. 2022. Introduction: Cultures of authenticity. In *Cultures of authenticity*, pages 1–17. Emerald Publishing Limited.
- Haifeng Wang, Qin Li, Jinho D. Choi, and Nianwen Xue. 2021. [Universal dependencies: Chinese gsd \(simplified, ud version 2.9\)](#). Accessed: 2024-09-15.
- Xiaobo Wang. 2020. Nüshu, the unique female rhetoric in the chinese rhetorical tradition. In *The Routledge Handbook of Comparative World Rhetorics*, pages 297–305. Routledge.
- Shiyue Zhang, Ben Frey, and Mohit Bansal. 2022. How can nlp help revitalize endangered languages? a case study and roadmap for the cherokee language. In

Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1529–1541.

Shiyue Zhang, Benjamin Frey, and Mohit Bansal. 2020. Chren: Cherokee-english machine translation for endangered language revitalization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 577–595.

Liming Zhao. 1992. *A Compendium of Chinese Nüshu (Zhongguo Nüshu Jicheng)*, original print edition, re-released as an ebook in 2019 edition. Tsinghua University Press, Beijing.

Liming Zhao. 2014. Nyushu character unification——theory and rules.

A Appendix Figures

This section presents supplementary figures referenced in the main content, providing additional visual support for illustrating our work.



Figure A1: An example from the book (a public domain page scan included in the product description on Amazon, and therefore not subject to copyright restrictions).

Nüshu	Chinese
𠄎	突, 途
𠄎	牌, 裴
𠄎	危
𠄎	萍
𠄎	舟
𠄎	帘, 轮
𠄎	扁
𠄎	雾
𠄎	波
𠄎	晓

Figure A2: A small sample of newly mapped words not in the official Nüshu-Chinese dictionary; an original contribution.

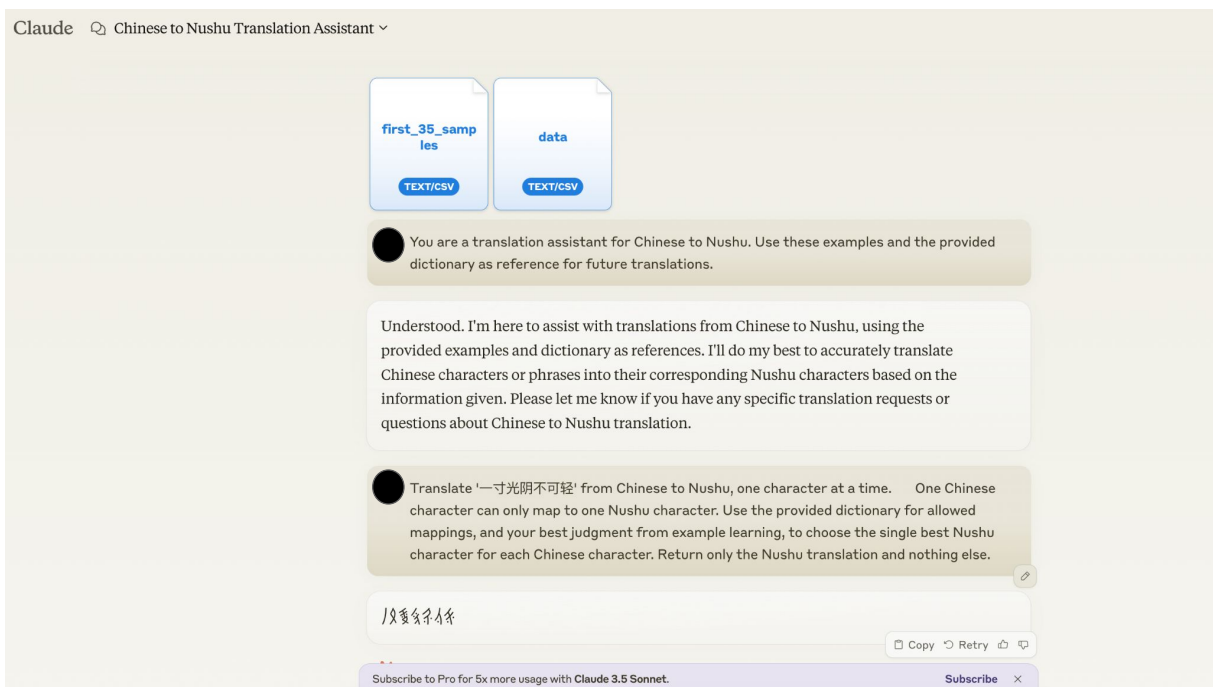


Figure A3: NūshuRescue is model-agnostic; it can work with Claude as well.