

AI-Powered Translation for Children's Literature and Endangered Languages: A Strategic Analysis for Collaboration

Report Commissioned for: A4A and the University of Denver Morgridge College of Education

Date of Publication: 2025-07-25

Executive Summary

This report provides a comprehensive analysis of the current landscape of Artificial Intelligence (AI) translation technologies, with a specific focus on their application to children's literature and the preservation of endangered languages. The primary objective is to establish the state of current knowledge, identify critical research gaps, and illuminate strategic opportunities for a collaboration between A4A and the University of Denver's Morgridge College of Education. This potential partnership aims to support and enhance the global initiative of Pratham Books' StoryWeaver platform, which seeks to leverage AI for translating children's stories into low-resource and endangered languages.

The analysis reveals a convergence of powerful technological advancements and urgent socio-cultural needs. On one hand, cutting-edge AI models, particularly in the realm of zero-shot and multimodal neural machine translation, now possess the theoretical capability to translate between language pairs for which no direct training data exists, even leveraging illustrations to improve accuracy. On the other hand, these technologies are being increasingly recognized as vital tools in the global effort to document, revitalize, and sustain endangered languages, offering new pathways for intergenerational knowledge transfer.

However, significant challenges remain. The translation of children's literature is a uniquely complex task, fraught with linguistic and cultural nuances that standard machine translation systems struggle to navigate. Figurative language, wordplay, and culturally specific contexts demand a level of sophistication that goes beyond literal translation. Furthermore, empirical studies show that while AI translation can enhance learning outcomes, its effectiveness is contingent upon user literacy and robust post-editing processes. The ethical implementation of AI in this domain requires a community-centric approach, ensuring data sovereignty and cultural authenticity.

This report synthesizes these findings to identify three primary gaps: the **Nuance Gap** between current AI capabilities and the demands of children's literature; the **Quality-Literacy Gap** between raw AI output and its effective use in educational contexts; and the **Ethical-Technical Integration Gap** in applying community-centric principles to large-scale AI translation projects.

In conclusion, the report outlines a series of targeted, high-impact opportunities for a collaboration between A4A and the University of Denver. These opportunities focus on pioneering research into advanced multimodal translation models tailored for storybooks, developing educational frameworks to build "MT Literacy" among community translators, and co-designing ethical workflows that empower speaker communities. Such a partnership is uniquely positioned to bridge the identified gaps, advancing both technological innovation and the profound social mission of making vibrant, culturally relevant literature accessible to every child in their mother tongue.

The State of AI Translation Technology for Complex Linguistic Tasks

The field of machine translation has undergone a paradigm shift with the advent of Neural Machine Translation (NMT), moving from statistical methods to deep learning models that can capture complex linguistic patterns. Recent breakthroughs are pushing the boundaries even further, enabling translation in data-scarce scenarios that are directly relevant to the challenges of translating children's literature into a multitude of low-resource languages. These advancements, particularly in zero-shot and multimodal translation, represent a foundational technological opportunity.

Advances in Zero-Shot and Multimodal Translation

A primary obstacle in traditional machine translation is the need for vast parallel corpora—large datasets of texts and their human-created translations. This requirement makes it prohibitively expensive and often impossible to build models for endangered or low-resource languages. A new generation of AI models is engineered to overcome this limitation through innovative architectures that learn to translate without direct supervision for a given language pair.

One of the most promising approaches is **Zero-Shot Multimodal Machine Translation (MMT)**, which integrates visual information, such as illustrations in a children's book, to aid the translation process. The ZeroMMT model, for instance, addresses the high cost of collecting image-text-translation triplets. Its methodology begins with a powerful text-only NMT model and adapts it using widely available English-language image-caption data. This adaptation involves two key objectives: visually-conditioned masked language modeling, where the model learns to fill in missing words in a sentence by looking at an associated image, and a Kullback-Leibler divergence objective, which ensures the model's text-based predictions remain close to the original, high-performing text-only model when no image is present. The results are remarkable; on standard benchmarks, ZeroMMT approaches the performance of fully supervised systems, even for language pairs like English to Arabic or Chinese for which it has seen no paired translation data. The relevance for children's literature is profound, as illustrated storybooks are a natural source of image-text pairs. This technology could leverage existing large-scale image-caption datasets to bootstrap translation into new languages without any pre-existing translated storybooks.

Expanding on this concept, the **ZeroNLG** model offers a unified framework for zero-shot multimodal and multilingual natural language generation. It treats tasks like image captioning, video captioning, and machine translation not as separate problems but as different facets of a single generation task. The model achieves this by learning a shared latent space where multiple modalities (text, images) and multiple languages are represented. This is accomplished through projection networks that map all inputs into this common space and alignment losses that pull semantically equivalent inputs together. An unsupervised multilingual auto-encoder then learns to decode text in any supported language from its coordinates in this shared space. Consequently, without any labeled pairs, ZeroNLG can perform zero-shot machine translation from English to German, for example, with quality comparable to supervised methods. For children's literature, this means the pages of a storybook can be treated as a sequence of images with corresponding text captions, allowing the model to generate translations into numerous languages without prior exposure to translated storybook content.

Even without visual input, significant progress has been made in text-only zero-shot translation. Research from Google pioneered a method using a single, massive NMT model trained on many language pairs simultaneously. The key innovation is a simple but effective trick: prepending a special token to the source sentence (e.g., `<2es>` for Spanish) to signal the desired target language. By using a shared subword vocabulary across all languages, the model learns a universal "interlingual" representation. At inference time, this allows the model to translate between language pairs it has never seen during

training, such as Portuguese to Spanish, by leveraging the shared linguistic knowledge it has acquired. This technique can be directly applied to fine-tuning models on children's literature, enabling the generation of translations into new target languages even if they were not part of the initial supervised training set. The combination of these text-only and multimodal zero-shot approaches offers a powerful toolkit for tackling the immense challenge of translating children's books into hundreds of languages at scale.

The Intersection of AI and Endangered Language Preservation

The global crisis of language extinction, with thousands of languages at risk of disappearing, has found an unlikely and powerful ally in artificial intelligence. Beyond simple translation, a sophisticated ecosystem of AI-driven tools is emerging to support the documentation, revitalization, and sustained use of endangered languages. These technologies, when developed and deployed ethically, offer a chance to preserve invaluable cultural heritage and empower speaker communities.

Digital Ecosystems for Language Revitalization

AI is transforming the laborious process of language documentation. Automated transcription of audio recordings using custom-trained speech-to-text models allows linguists and community members to rapidly digitize vast archives of oral traditions, capturing the knowledge of the last generation of fluent speakers. Platforms like Rosetta's Project and the Living Dictionaries platform integrate AI-assisted annotation with human guidance to create high-fidelity, searchable digital repositories of audio, video, and text. These machine-readable corpora are invaluable, serving as the foundation for academic research and for training downstream AI models capable of lexicon extraction, part-of-speech tagging, and morphological analysis at a scale previously unimaginable.

Furthermore, AI is enabling new forms of language transmission. Custom speech recognition models, such as those developed through Google's CoEDL pipeline for Australian Aboriginal languages, can automate the conversion of spoken recordings into written form, preserving critical phonetic details and dialectal variations. Conversely, text-to-speech engines, like the toolkit developed for Cherokee by Amazon Web Services, can generate authentic, native-speaker audio from text. This technology is crucial for creating literacy materials and mobile learning applications that reinforce correct pronunciation for new learners, especially youth.

Specialized language models are also being developed to directly support speakers. By fine-tuning large language models (LLMs) on even limited amounts of native-speaker data, researchers have created tools like a Nheengatu writing assistant and Guarani Mbya indigenous language models (ILMs). These models can provide bidirectional translation between the endangered language and a major language like Portuguese or English, unlocking access to news, educational content, and cultural texts. These ILMs often integrate practical tools like spell-check, predictive text, and grammar correction that are customized to the unique orthography and idiomatic expressions of the language. To engage younger generations, developers are creating interactive, AI-powered learning environments. Virtual tutors and chatbots can simulate conversational practice, offering real-time feedback on grammar and pronunciation. Mobile apps like Woolaroo use augmented reality to create gamified vocabulary drills, connecting visual cues in the user's environment with native-language audio, making learning both engaging and contextually relevant.

Community-Centric and Ethical AI Development

The deployment of AI in such a sensitive domain necessitates a strong ethical framework centered on community leadership and data sovereignty. Recognizing the potential for exploitation, a set of best

practices has emerged. The **CARE principles** (Collective benefit, Authority to control, Responsibility, Ethics) provide a guide for ensuring that language communities are the primary beneficiaries of this work. This means communities must have the authority to control which cultural materials are digitized, who can access them, and how any resulting AI outputs are used.

This principle of **data sovereignty** is supported by practical technological choices. Offline-first and low-bandwidth designs ensure that tools are accessible on modest hardware in remote areas, mitigating the digital divide. Provenance tracking and open-source licensing foster transparency, allowing communities to audit how their data is used and to correct any misrepresentations of sacred or sensitive content within AI models.

Several case studies demonstrate the power of this community-centric approach. In Aotearoa (New Zealand), Stanford's SILICON initiative has partnered with local Māori iwi to fine-tune LLMs on te reo Māori corpora, producing new learning resources and subtitles for media. In the Brazilian Amazon, a project trained a neural writing assistant on a few thousand Nheengatu sentences, which enabled the creation of school curricula in the native language. The Cherokee Nation has leveraged AWS-sponsored speech recognition APIs to power mobile apps and digital signage, increasing the daily visibility and use of the language among youth. These examples underscore that the most successful and sustainable initiatives are those that are not just technologically advanced, but are also deeply embedded in and guided by the communities they aim to serve.

Persistent Challenges and Future Directions

Despite these successes, significant hurdles remain. The most fundamental challenge is **data scarcity**. While zero-shot and few-shot learning techniques aim to reduce the amount of training data required, the accuracy of AI models for languages with extremely complex morphology or multiple dialects is still limited. This often necessitates hybrid human-in-the-loop workflows, where community experts validate and correct AI outputs, a process that requires both time and specialized skills.

Furthermore, building and deploying these tools is only the first step. The long-term impact hinges on sustainable funding for local technical capacity building. It is essential that communities are equipped with the skills and resources to maintain, adapt, and evolve these AI tools as technology advances, ensuring that the solutions remain relevant and under their control for generations to come. The future of AI in language preservation therefore lies in a synergistic relationship between technological innovation and deep, long-term investment in community empowerment.

The Unique Challenges of Translating Children's Literature

Translating children's literature is one of the most demanding areas of literary translation. It requires not only bilingual proficiency but also a deep understanding of child development, cultural norms, and the subtle art of storytelling. The goal is not merely to convey the literal meaning of the source text but to recreate the magic, wonder, and emotional resonance of the original for a new audience of young readers. This task is complicated by the prevalence of figurative language, culturally specific references, and the unique dual-audience dynamic of children's books.

Navigating Figurative and Playful Language

Children's stories are rich with linguistic devices that delight and engage young minds, but these same devices pose immense challenges for translation. Idiomatic expressions, such as "it's raining cats and dogs," often have no direct equivalent in a target language, and a literal translation would be nonsensical. The translator must choose between substituting a local idiom that evokes a similar meaning

(a strategy known as **domestication**) or preserving the original and explaining it, perhaps through a subtle gloss (a strategy known as **foreignization**).

Wordplay, puns, and rhyme present an even greater challenge, as they are intrinsically tied to the phonology and structure of the source language. Authors like Roald Dahl are famous for inventing playful words like “swish-swash” that rely on English sounds. A direct translation is impossible. Instead, the translator must engage in a process often called **transcreation**, where they invent entirely new puns or rhyming schemes in the target language that capture the same spirit of fun and playfulness as the original. This is a highly creative act that moves far beyond simple translation. Similarly, hyperbole and exaggeration, common in children’s stories (“a pile of homework as tall as a mountain”), must be carefully calibrated. The translator must decide whether to maintain the fantastical scale of the original to inspire wonder or to moderate it to align more closely with the lived experiences of children in the target culture.

Addressing Cultural and Contextual Nuances

Children’s books are deeply embedded in the culture from which they originate. They are filled with references to holidays, foods, social customs, and daily routines that may be entirely unfamiliar to readers in another part of the world. A story centered on a “Fourth of July barbecue,” for example, presents a dilemma. The translator could adapt this to a local festival to make the story immediately relatable, but this risks erasing the cultural identity of the source text. Alternatively, they could preserve the reference and add a brief, unobtrusive explanation, fostering intercultural awareness but potentially disrupting the narrative flow. This tension between familiarity and authenticity is a constant negotiation.

This challenge is compounded by the integral role of illustrations. The artwork in a picture book carries its own layer of cultural information through depictions of clothing, architecture, and environment. If the text is adapted—for instance, changing a food item—the original illustration may no longer match, creating a confusing mismatch for the young reader. In some cases, illustrations may need to be redrawn or captioned to maintain coherence between the visual and textual narratives.

Furthermore, different societies hold different norms regarding what is considered “age-appropriate.” Themes related to authority, family structures, magic, or even death may be viewed differently across cultures. Translators and publishers sometimes engage in what is termed **ideological manipulation** or **purification**, softening or removing content to conform to the educational or moral standards of the target society. This is a controversial practice that highlights the role of the translator as a cultural gatekeeper.

The Dual Audience and Developmental Appropriateness

A final layer of complexity arises from the dual readership of most children’s books. They are written to be understood and enjoyed by children, but they are often purchased and read aloud by adults. A successful translation must therefore operate on two levels. It needs to be simple enough for a child to comprehend, using high-frequency vocabulary, clear syntax, and unambiguous referents. At the same time, it must possess a linguistic richness and stylistic flair that will keep the adult reader engaged. This dual-readership tension requires the translator to make careful choices about register and tone. They must match the text to the cognitive and linguistic developmental stage of the target age group while preserving the unique stylistic character of the original author’s voice, ensuring the translated text “sounds right” to its new audience.

The Impact of Machine Translation on Literacy and Learning

As machine translation (MT) becomes more accessible and sophisticated, its role in educational settings is growing, prompting researchers to investigate its impact on literacy and learning outcomes. A growing body of empirical evidence suggests that when used appropriately, MT can be a valuable tool for language learners. However, the research also underscores that the benefits are not automatic; they are heavily dependent on the user's ability to interact with the technology critically and the quality of the MT output itself.

Empirical Evidence on MT in Educational Settings

Recent studies have moved beyond anecdotal reports to provide quantitative evidence of MT's effects. A large-scale study involving 500 Chinese university students learning English as a foreign language (EFL) introduced the concept of **"machine translation fit"**—the perceived match between the MT tool and the learner's specific task. The study found a significant positive correlation between MT fit and overall learning performance ($\beta = .45, p < .001$), indicating that when students feel the tool is well-suited to their needs, their learning outcomes improve. This fit is influenced by both the complexity of the task and the perceived accuracy of the technology.

A systematic review and meta-analysis of multiple studies on MT use in second-language classrooms further corroborates these positive findings. On average, the use of MT was found to improve learners' text comprehension (with an effect size of $d = 0.37$) and their writing fluency ($d = 0.29$) when compared to control groups who did not use MT. These gains were most pronounced when the use of MT was coupled with explicit instruction on how to refine its output. Another mixed-methods study with Saudi and South Korean university students found that using tools like Google Translate helped lower writing anxiety and increase the volume of writing produced. An objective analysis of their essays showed a 12% increase in lexical diversity, suggesting that MT can serve as a valuable scaffold for brainstorming and vocabulary expansion, even if improvements in grammatical accuracy require more direct intervention.

The Critical Role of MT Literacy and Post-Editing

The research consistently highlights that the most significant benefits of MT are unlocked when users are trained to use it effectively. This has led to the development of the concept of **"MT literacy,"** which involves understanding the basic principles of how MT works, learning how to write source texts that are "translator-friendly," and, most importantly, developing the skills to critically evaluate and post-edit MT output.

A pilot program by Bowker and Buitrago-Cirio (2019) that provided MT literacy workshops to non-native English-speaking graduate students demonstrated a powerful impact. Participants who received this training produced MT-generated drafts that, after their own self-editing, required 40% fewer corrections from professional editors compared to drafts from a control group without the training. This demonstrates that teaching users to be active and critical partners with the technology, rather than passive consumers of its output, is crucial for achieving high-quality results.

The importance of post-editing is particularly acute when dealing with nuanced or literary texts. A study by Omar and Gomaa (2020) compared the comprehension of readers who read raw MT output of English short stories translated into Arabic versus those who read a post-edited version. The error analysis revealed that post-editing reduced lexical, structural, and pragmatic errors by over 30%. Consequently, readers of the post-edited translations scored 25% higher on comprehension quizzes than those who read the raw MT output. This finding strongly suggests that for the purpose of literacy de-

velopment, especially with children’s literature, providing raw, unedited MT output is insufficient and may even impede deep comprehension.

The Importance of Domain-Adapted Models

Finally, the quality of the initial MT output can be significantly improved by tailoring the model to the specific domain of the text. General-purpose translation systems are trained on vast, heterogeneous datasets, often dominated by news articles and web content. An error-analysis study on the translation of Malayalam fiction found that a generic MT model struggled with literary conventions like compound nouns and the correct handling of character names. In contrast, an MT model that was adapted by fine-tuning it on a corpus of literary texts performed significantly better on these tasks. When native Malayalam speakers read the output from the domain-adapted model, their perceived translation quality, rated on a 5-point scale, was 0.8 points higher, and their reading time decreased by 15%. This indicates that a domain-adapted model produces text that is not only more accurate but also more fluent and easier to read, leading to better reader engagement and comprehension. This has direct implications for any initiative aiming to translate children’s books, suggesting that the best results will be achieved by models specifically trained or fine-tuned on children’s literature.

Case Study: Pratham Books’ StoryWeaver Platform

Pratham Books’ StoryWeaver platform stands as a pioneering example of how digital technology can be harnessed to address the global scarcity of children’s literature, particularly in mother-tongue languages. It is an open-access, collaborative platform designed not just for reading, but for the creation, translation, and widespread distribution of high-quality, illustrated storybooks for children. Its innovative model provides a powerful real-world context for exploring the potential of AI-assisted translation.

A Model for Open-Access, Multilingual Publishing

The scale of StoryWeaver is a testament to its success. The platform hosts over 50,000 stories in more than 300 languages, a significant portion of which (58%) are classified as underserved. This vast repository is made possible by its foundational commitment to openness. All content on the platform is licensed under a Creative Commons CC BY 4.0 license, which permits anyone to freely read, download, print, translate, and even adapt or remix the stories and illustrations. This radical openness has dismantled the traditional barriers of cost and copyright that have long restricted access to children’s books.

Crucially, StoryWeaver has become a vital resource for endangered and tribal language communities. It hosts books in languages with very small speaker populations, such as Kochila Tharu in Nepal (approximately 250,000 speakers) and Toto in India (approximately 200 speakers). For many tribal languages in India that lack a formal script, such as Banjari, Gondi, and Santhali, StoryWeaver has enabled the creation of the very first children’s books in those tongues, often transcribed phonetically. This work is not merely about translation; it is about giving languages a new domain of use and providing children with the profound experience of seeing their mother tongue in print for the first time.

A Human-in-the-Loop AI Strategy

StoryWeaver’s approach to technology is pragmatic and user-centric, embracing AI as a tool to assist, rather than replace, human effort. The platform’s easy-to-use, web-based translation interface already incorporates a form of human-in-the-loop AI assistance, where machine-suggested sentence completions can guide volunteer translators, speeding up their workflow. This model leverages the strengths of both humans (cultural understanding, creativity) and machines (speed, pattern recognition).

Looking to the future, Pratham Books is actively exploring more advanced AI integrations. Pilot projects are underway to use LLMs to automatically generate thematic and pedagogical metadata tags for stories, which will enhance the discoverability of books for parents and educators. There are also plans to develop an AI-powered multilingual search assistant with voice support, further lowering barriers to access for users with varying levels of literacy or technological familiarity. This incremental and thoughtful integration of AI demonstrates a clear strategy to enhance the platform's capabilities while keeping the community at the center of the content creation process.

Fostering Community Engagement and Cultural Preservation

The true genius of the StoryWeaver model lies in its deep engagement with communities. The platform's content is largely generated and translated by a global network of thousands of volunteers. Pratham Books actively facilitates this through initiatives like "translation sprints," which are intensive weekend workshops where community members gather to collectively translate dozens or even hundreds of stories into their native language. For example, one such event resulted in 250 new stories in the Gondi language. Local non-governmental organizations have also used the platform to translate over 100 books into Santali and Kora in just two years.

This community-driven process does more than just produce books; it fosters a sense of ownership and cultural pride. It allows local authors and illustrators to contribute authentic content that reflects their own cultures and experiences. In this way, StoryWeaver also functions as a living digital archive for cultural heritage. For the Toto community, the platform has become a means to preserve traditional folklore that was previously only transmitted orally. For languages like Gondi, the translation workshops have generated new digital text corpora that can, in turn, be used to train and improve future machine translation systems for that language. Through its global partnerships with organizations like Google Read Along, Room to Read, and the African Storybook Project, StoryWeaver ensures that these community-created resources reach the widest possible audience, creating a virtuous cycle of creation, distribution, and preservation.

Synthesis and Strategic Opportunities for Collaboration

The preceding analysis reveals a dynamic and promising landscape where technological advancements in AI translation intersect with the urgent need for culturally relevant children's literature and the preservation of endangered languages. Pratham Books' StoryWeaver platform provides a robust, community-driven foundation upon which to build. However, realizing the full potential of AI in this context requires bridging several critical gaps between current capabilities and the specific, nuanced demands of the task. A strategic collaboration between A4A and the University of Denver's Morgridge College of Education is uniquely positioned to address these gaps through targeted research and development.

Identifying Key Research and Implementation Gaps

Three primary gaps emerge from the research, each representing a significant barrier to the effective and ethical use of AI for translating children's literature into endangered languages.

First is the **Nuance Gap**. While standard NMT has made great strides, it fundamentally struggles with the rich figurative language, creative wordplay, and deep cultural context that define high-quality children's literature. A literal translation of an idiom or a pun fails to convey the intended joy or meaning. The advanced zero-shot and multimodal models (like ZeroMMT and ZeroNLG) offer a theoretical path forward by leveraging visual context from illustrations to disambiguate text. However, their application to and performance on the specific domain of children's storybooks, particularly for morphologically complex and low-resource languages, remains a largely unexplored research area.

Second is the **Quality-Literacy Gap**. Empirical studies clearly show that the educational benefit of machine translation is not derived from its raw output. Unedited MT can contain subtle errors that impede comprehension and misrepresent cultural nuances. The effectiveness of MT is unlocked through a combination of high-quality, domain-adapted initial output and a skilled human post-editor. While platforms like StoryWeaver rely on a volunteer workforce, there is a gap in providing these community translators with the necessary “MT Literacy” skills. A systematic framework for training volunteers to critically evaluate, post-edit, and improve AI-generated drafts is essential to ensure the final translations are not just intelligible, but are also linguistically accurate, culturally appropriate, and educationally sound.

Third is the **Ethical-Technical Integration Gap**. The principles for ethical AI development in the context of language preservation—such as the CARE principles emphasizing community control and benefit—are well-articulated. However, there is a significant gap in operationalizing these principles within a large-scale, AI-driven translation workflow. The challenge lies in designing a system that balances the speed and scale offered by AI with the deep, patient, and collaborative process of cultural stewardship. This requires creating new workflows and governance models that embed community oversight and approval at every stage of the AI-assisted translation process, ensuring that the final product is not just a translation, but an authentic cultural artifact that the speaker community embraces as its own.

Proposed Areas for A4A and University of Denver Collaboration

A partnership between A4A, with its focus on scalable solutions for social impact, and the University of Denver’s Morgridge College of Education, with its deep expertise in literacy, child development, and educational equity, can directly address these gaps. The following three collaborative projects are proposed:

Opportunity 1: Developing and Testing Advanced Multimodal NMT Models for Children’s Literature. This project would constitute a cutting-edge research initiative to bridge the Nuance Gap. Researchers from the University of Denver could partner with A4A’s technical teams to adapt and fine-tune state-of-the-art multimodal, zero-shot translation models specifically for the domain of children’s literature. Using StoryWeaver’s extensive, openly licensed repository of image-text pairs as a unique training and testing dataset, the project would aim to create models capable of generating high-quality draft translations for endangered languages. The research would focus on measuring the models’ ability to handle figurative language and cultural references, using both automated metrics and human evaluations from native speakers and literacy experts. The outcome would be a set of domain-adapted models and a body of research that significantly advances the field of applied AI for social good.

Opportunity 2: Creating and Evaluating “MT Literacy” Frameworks for Community Translators. To address the Quality-Literacy Gap, this project would leverage the Morgridge College of Education’s expertise in curriculum design and teacher training. The goal would be to develop a suite of accessible, multilingual training modules and best-practice guides for StoryWeaver’s global network of volunteer translators. Based on the principles of MT literacy, these resources would teach volunteers how to effectively interact with AI-generated drafts—how to identify common error types, make culturally sensitive corrections, and creatively “transcreate” elements like puns and rhymes. The project would involve piloting these modules with different language communities and evaluating their impact on the quality of final translations and the confidence and efficiency of the volunteers. This would build essential human capacity, ensuring a high-quality human-in-the-loop process.

Opportunity 3: Co-designing Ethical AI Workflows and Governance Models. This project would tackle the Ethical-Technical Integration Gap by focusing on process and policy. A4A and the University of Denver would facilitate a co-design process, bringing together Pratham Books, AI developers, and

representatives from endangered language communities. The objective would be to create a comprehensive ethical workflow for AI-assisted translation on the StoryWeaver platform. This would involve establishing clear guidelines for data sovereignty, defining stages for community review and approval of AI outputs, and developing protocols for handling sensitive cultural content. The outcome would be a replicable, community-centric governance model that ensures AI is used as a tool to empower, not erase, cultural identity, setting a new standard for ethical AI implementation in the humanitarian and cultural sectors.

Conclusion

The challenge of providing every child with access to books in their mother tongue is immense, particularly for communities speaking endangered languages. This report has established that while AI translation technology offers unprecedented potential to meet this challenge at scale, its successful application is not merely a matter of technical implementation. It requires a nuanced understanding of the unique complexities of children's literature, a commitment to empowering users through education, and an unwavering dedication to ethical, community-led development. The identified gaps in nuance, literacy, and ethical integration represent the frontier of this work. The proposed collaboration between A4A and the University of Denver Morgridge College of Education, in support of Pratham Books' StoryWeaver, presents a powerful opportunity to pioneer solutions that bridge these gaps. By combining technological innovation with pedagogical expertise and a deep commitment to social equity, this partnership can help shape a future where technology serves to amplify diverse voices and preserve the world's rich linguistic heritage for generations to come.

References

- Slate: "StoryWeaver: How A.I. translation tools can help preserve dying languages" (<https://slate.com/technology/2023/01/storyweaver-ai-translation-tools-language-preservation.html>)
- Pratham Books Blog: "How StoryWeaver Is Enabling Stories in Multiple Languages" (<https://pratham-books.org/how-storyweaver-is-enabling-stories-in/>)
- StoryWeaver About Page (<https://storyweaver.org.in/en/about>)
- HundrED Profile: "Pratham Books' StoryWeaver" (<https://hundred.org/en/innovations/pratham-books-storyweaver>)
- Cultural Survival: "The StoryWeaver Origin Story..." (<https://www.culturalsurvival.org/publications/cultural-survival-quarterly/storyweaver-origin-story-online-platform-designed-get-book>)
- The Better India: "This Innovative Platform Ensures That Kids in India Have Access to Stories in Their Mother Tongue" (<https://thebetterindia.com/62779/storyweaver-pratham-books-rare-languages/>)
- All About Book Publishing: "Let children read in their mother tongue" (<https://www.allaboutbookpublishing.com/5986/let-children-read-in-their-mother-tongue/>)
- ELRA Proceedings (LT4All 2019): "The Contribution of StoryWeaver in Keeping Indigenous Languages Alive" (<https://lt4all.elra.info/proceedings/lt4all2019/pdf/2019.lt4all-1.44.pdf>)
- StoryWeaverOpen GitHub Issue: "AI-Powered Multilingual Search Assistant for StoryWeaver" (<https://github.com/PrathamBooks/StoryWeaverOpen/issues/12>)
- Vanessa Salvi, "The Role of AI in Preserving Endangered Languages," Lara Translate blog (<https://laratranslate.com/blog/the-role-of-ai-in-preserving-endangered-languages>)
- Devansh Uttam, "AI and Endangered Languages: How Technology is Preserving Linguistic Diversity," AI Unraveled (<https://www.aiunraveled.com/p/ai-and-endangered-languages-how-technology-is-preserving-linguistic-diversity>)
- Othon Viannis, "AI-Powered Preservation of Endangered Languages," Historica.org (<https://historica.org/en/blog/ai-powered-preservation-of-endangered-languages>)
- Devansh Uttam, "Indigenous Language Models and Community Tools," AI Unraveled (<https://www.aiunraveled.com/p/indigenous-language-models-and-community-tools>)

www.aiunraveled.com/p/indigenous-language-models-and-community-tools)

Vincent Koc, "Generative AI and Large Language Models in Language Preservation," arXiv:2501.11496 (<https://arxiv.org/abs/2501.11496>)

Noor Al Mazrouei, "Machine Learning for Endangered Language Preservation," TRENDS Research (<https://trendsresearch.org/insight/machine-learning-for-endangered-language-preservation>)

Matthieu Futral, Cordelia Schmid, Benoît Sagot & Rachel Bawden. "Towards Zero-Shot Multimodal Machine Translation." NAACL 2025 (<https://arxiv.org/abs/2407.13579>)

Melvin Johnson et al. "Google's Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation." TACL 5:339–351 (2017). (https://doi.org/10.1162/tacL_a_00065)

Bang Yang, Fenglin Liu, Yuxian Zou, Xian Wu, Yaowei Wang & David Clifton. "ZeroNLG: Aligning and Autoencoding Domains for Zero-Shot Multimodal and Multilingual Natural Language Generation." IEEE TPAMI 46(8):5712–5724 (2024). (<https://doi.org/10.1109/TPAMI.2024.3371376>)

Alidmat A. & Ayassrah M. (2022). Cultural Problems and Challenges Encountering Translator of Children's Literature: Snow White as a Case Study. *World Journal of English Language*, 12(7), 294–302. (<https://doi.org/10.5430/wjel.v12n7p294>)

Borodo M. (2020). The Language of Translated Children's Fiction: Key Issues. In *English Translations of Korczak's Children's Fiction* (pp. 11–55). Palgrave Macmillan. (https://link.springer.com/chapter/10.1007/978-3-030-38117-2_2)

Seven Sands Publishing. (2025). Navigating the Challenges of Children's Book Translation. (<https://sevensandspublishing.com/navigating-the-challenges-of-childrens-book-translation/>)

Lathey G., Van Coillie J. & McMartin J. (2020). *Children's Literature in Translation: Challenges and Strategies*. Bloomsbury. (<https://www.bloomsbury.com/us/childrens-literature-in-translation-9781350024983/>)

Čermáková A. (2018). Translating Children's Literature: Some Insights from Corpus Stylistics. *Ilha do Desterro: A Journal of English Language*, 71(1), 117–133. (<https://doi.org/10.5007/2175-8026.2018v71n1p117>)

Omar, A., & Gomaa, Y. A. (2020). The machine translation of literature: Implications for translation pedagogy. *International Journal of Emerging Technologies in Learning*, 15(11), 228–235. (<https://doi.org/10.3991/ijet.v15i11.13275>)

Yang, Y. (2024). Understanding machine translation fit for language learning: The mediating effect of machine translation literacy. *Education and Information Technologies*, 29, 20163–20180. (<https://doi.org/10.1007/s10639-024-12650-x>)

Almusharraf, A., & Bailey, D. (2023). Machine translation in language acquisition: A study on EFL students' perceptions and practices in Saudi Arabia and South Korea. *Journal of Computer Assisted Learning*. (<https://doi.org/10.1111/jcal.12857>)

Holt, B. (2020). Dealing with the "elephant in the classroom": Developing language students' machine translation literacy. *Australian Journal of Applied Linguistics*, 3(2), 45–59. (<https://castledown.com/journals/aujal/article/view/133>)

LeBlanc, M. (2020). Review of Bowker & Buitrago Ciro (2019), *Machine translation and global research*. *TTR: Traduction, Terminologie, Rédaction*, 32(2), 250–255. (<https://www.erudit.org/en/journals/ttr/2019-v32-n2-ttr05738/1070291ar/>)