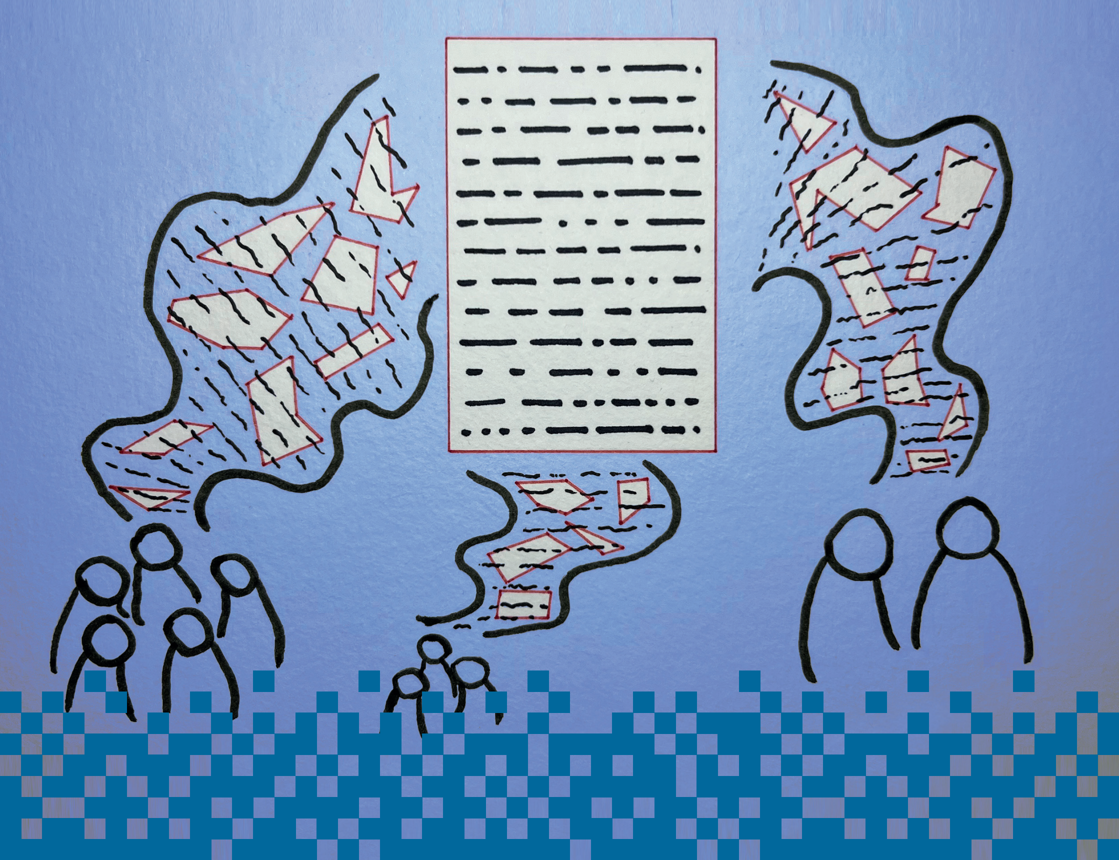

Utilisation of Artificial Intelligence in Accurate Translation of Peace Agreements: A Practical Assessment

Adam Farquhar

PeaceRep
Peace and Conflict Resolution Evidence Platform

PEACE ANALYTICS SERIES

THE UNIVERSITY of EDINBURGH

PeaceRep
Peace and Conflict Resolution Evidence Platform

Author: Adam Farquhar

PeaceRep: The Peace and Conflict Resolution Evidence Platform
School of Law, Old College, The University of Edinburgh
South Bridge, Edinburgh EH8 9YL

Tel. +44 (0)131 651 4566
Fax. +44 (0)131 650 2005
E-mail: peacerep@ed.ac.uk
PeaceRep.org
in https://www.linkedin.com/company/peacerep/

Developed as part of the Peace and Conflict Resolution Evidence Platform, this report intends to support the research programme's work and provide insights into the use of Artificial Intelligence for translation purposes.

This research is supported by the Peace and Conflict Resolution Evidence Platform (PeaceRep), funded by UK International Development from the UK government. However, the views expressed are those of the authors and do not necessarily reflect the UK government's official policies. Any use of this work should acknowledge the authors and the Peace and Conflict Resolution Evidence Platform.

About the author:
Adam Farquhar is a Research Associate and Data Officer at PeaceRep. He supports the management, development, and coding of the PA-X Peace Agreement Database and its sub-databases. His research interests include the application of geocoding and AI in peacebuilding. Adam holds an MA in Comparative Ethnic Conflict from Queens University Belfast and a BA in International Studies from Huntingdon College in Montgomery, Alabama, USA.

Acknowledgements:
I would like to express my sincere gratitude to Dr Sanja Badanjak, Tim Epple, Dr Roy Gardner, Dr Johanna Amaya-Panche, Alice Raymond and Dr Winnie Xia for their valuable contributions to this research report. Their thorough review and insightful suggestions have significantly enhanced the quality of this work. Their expertise and time are greatly appreciated.

Cover images: Yasmine Boudiaf & LOTI
Design: Smith Design Agency
DOI: http://dx.doi.org/10.7488/era/5883

Contents

Executive Summary

This report explores how Artificial Intelligence (AI) technologies can enhance the translation of peace agreements for the PA-X Peace Agreements Database. Peace agreements are vital instruments in conflict resolution, making accurate translations essential. This report addresses the challenges of translating these complex documents, which has traditionally relied on academic researchers, translation professionals, and domain experts. The report examines AI-driven approaches to improve sustainability, efficiency, and linguistic precision in this critical work.

Key Findings

- AI offers valuable support for human translation of peace agreements, providing a scalable, efficient solution to a traditionally labour-intensive process. Recent advances in Natural Language Processing (NLP) and Generative AI models potentially deliver faster translation times, handle greater document volumes, and significantly reduce costs associated with manual translations.

- However, challenges persist, including English language bias in training data, lack of domain-specific expertise for legal and diplomatic terminology, and AI models' limitations in fully grasping cultural nuances and multilingual contexts. These factors present significant barriers to the reliability and accuracy of AI translations for peace agreements. By carefully navigating technological, ethical, and environmental challenges associated with AI, translation of specialised documents such as peace agreements can be achieved with lower costs, which can in turn enhance the contributions.

Recommendations

While this assessment focuses specifically on PA-X and peace agreements, the findings generate recommendations that apply broadly to the translation of most specialized, technical documents. The following recommendations provide a framework that can be generalized beyond PA-X to other contexts requiring accurate translation of complex technical content:

- **Adopt a Semi-automated Approach**
  Combining AI's efficiency with human expertise mitigates translation inaccuracies, particularly for languages underrepresented in AI training datasets. This hybrid model ensures translation fidelity while leveraging AI to streamline the process.

- **Address Ethical and Environmental Considerations**
  As AI models require substantial energy consumption, evaluating the environmental impact of these technologies is crucial. Additionally, addressing inherent biases in AI training data is essential to ensure equitable and accurate translations across languages and cultures.

- **Optimise Workflows**
  Developing efficient workflows that integrate AI translation tools can help in preliminary analysis and extraction of key metadata from peace agreements, enhancing the database's comprehensiveness and accessibility.

- **Foster Collaborations**
  Strengthening collaboration with translators, academic institutions and leveraging global networks can enrich the translation process by incorporating diverse linguistic expertise and contextual understanding.

- **Invest in Continuous Research and Development**
  Further research is needed to address current limitations of AI in translation, such as improving models' capacity to understand cultural contexts and technical jargon, while keeping abreast of advances in Optical Character Recognition (OCR) technology to support AI translations.

Introduction

Context of Peace Agreements and Translation Challenges

Translating peace agreement texts accurately is critical to conflict resolution research and practice. In her book On the Law of Peace, Christine Bell describes peace agreements as "delicately crafted, legal-looking texts, with preambles and articles all speaking the language of legal obligation. In substance they link commitments to ceasefires to new constitutional arrangements for how power will be held and exercised. They are usually formally signed by both domestic and international actors. Both form and substance are intrinsically "legal'" (Bell, 2008, p. 5). Accurate translations ensure that the true meaning and nuances of these documents are preserved as far as possible, mitigating any variation in interpretation that could arise from linguistic inaccuracies or ambiguities. A single misinterpreted word in translation can distort our understanding and analysis of peace agreement provisions, potentially undermining research credibility. Furthermore, the context in which a word is used is crucial, as misinterpretation in translation can affect global understanding of peace agreement practices and, in rare cases, could influence how stakeholders interpret implementation requirements if they rely on translated versions rather than originals. This presents both research accuracy and ethical concerns, as poor translations could lead to flawed analysis and potentially impact peace processes if widely disseminated.

Agreement documents often come in various languages and formats, presenting a unique set of challenges for translators. PeaceRep's current practice of employing post-graduate students who are both native speakers and legal scholars as translators, although beneficial, faces limitations due to overall availability, specifically in certain languages. Moreover, ethical considerations regarding the reliance on temporary academic labour and the sustainability of such an approach in the long term need to be addressed.

This report will look at the application of AI in peace agreement translation. First, key terminology will be clarified in the AI and translation fields to establish a shared understanding of concepts. The report will then explore the current state of AI translation technology, identifying both advancements and limitations, with particular attention to training issues in AI models and their implications for multilingual peace agreement work. It will then examine various translation modalities—manual, semi-automated, and fully automated approaches—evaluating their respective strengths and challenges in the peace agreement context. Through case studies and comparative analysis, the report will demonstrate practical applications and limitations of AI translation in real-world peace agreement scenarios.

Ethical considerations will then be addressed, focusing on environmental impacts, bias in AI training data, and data ownership concerns. The report concludes with recommendations for implementing AI translation in peace processes, mitigating risks, and suggestions for further research and development that could enhance translation practices in conflict resolution contexts.

Purpose of Exploring AI in this context

The advent of different forms of AI technologies has facilitated a significant leap forward in the field of language translation. These advancements promise to enhance the scale, efficiency, and accuracy of translation processes, potentially improving the efficiency of how peace agreements are translated for the PA-X Peace Agreements Database and similar projects. The deployment of AI in this environment could lead to faster translation times, reduced costs, and the ability to handle a larger volume of documents. Most importantly, leveraging AI for translations could address the sustainability challenges faced by this and similar projects (due to short-term funding horizons), ensuring their long-term effectiveness in supporting peace and transition processes. This report provides an overview assessment of the feasibility, benefits, and limitations of integrating AI technologies into PeaceRep's translation workflow, to enhance the programme's capacity to contribute to the study and advancement of peace and transition processes worldwide, and to provide key lessons learned to similar projects in the peacemaking and peacebuilding space.

Key Terms

- **Artificial Intelligence (AI)**
  Refers to the technology that takes in information and makes predictions, creates content, gives recommendations, or makes decisions that affect real worlds or digital environments. (Organisation for Economic Co-operation and Development [OECD], 2019).

- **Generative AI**
  Encompasses a range of computer-based methods designed to create content that appears original and significant. This includes producing new pieces of text, visuals, or sound based information it has learned from existing data (Feuerriegel et al., 2024).

- **Large Language Model (LLM)**
  A large language model is a high-capacity neural network trained on diverse and extensive datasets, capable of performing a wide range of natural language processing tasks without explicit supervision (Radford et al., 2019).

- **Natural Language Processing (NLP)**
  This is a way for computers to interact with human languages in order to perform various tasks such as language translation (Henry, 2024, p. 3)

- **Tokenisation**
  In the context of large language models, tokenisation is the processes of breaking down text into smaller pieces, called 'tokens'. These tokens can represent various parts of a sentence, phrase or word. An LLM then uses this information to analyse the meaning of individual pieces of text, as well as its meaning in relation to other parts of the text (Singh & Strouse, 2024). This process is a critical component in AI translation.

- **Transformer Model**
  A transformer model is a computer-based method that analyses large amounts of data such as text at once instead of one by one. Transformers excel at understanding context and relationships within text by simultaneously considering the significance of each word in relation to others, enhancing tasks such as translation, text generation, and content summarisation (Vaswani et al., 2023). This is an example of a technology that is currently relevant to the field, though it should be recognise that this is subject to continuous new developments.

Overview of AI Translation Technology

The utilization of artificial intelligence (AI) in translation most commonly takes the form of machine learning models, a development with deeper historical roots than many realize. Although recent advances have generated significant excitement, machine translation technologies have been evolving for several decades. Roberts (2018) demonstrates that machine translation and interpretation technologies have long been considered potential solutions for addressing language barriers in democratic systems, with online translation becoming "ubiquitous globally" well before the current wave of generative AI. Similarly, Lucas et al. (2015) established methodologies for processing, managing, and analysing multilingual textual datasets using machine translation tools, indicating the maturity of these approaches in comparative research. De Vries et al. (2018) further validated the reliability of machine translation for bag-of-words models by comparing human-translated and machine-translated texts, finding "highly similar" results with only "minor differences across languages." These models are trained on extensive datasets and are capable of making predictions or decisions with little to no human intervention.

At the core of these technologies is Natural Language Processing (NLP), which involves using computers to support human understanding of language presented in the form of text. Machine learning is one of the primary tools employed by NLP to process and analyse text data. Deep learning, a subset of machine learning that utilizes artificial neural networks, has significantly advanced the field of natural language processing and translation (LeCun et al., 2015).

The tokenisation of text in natural language processing is critical in how AI models are able to be accurately translate complex text. The term ‘tokenisation' simply refers to taking a text and breaking it down into ‘tokens', or smaller more manageable sections of the text, that can be read by the algorithm. This is key to the advantage NLP can provide for translation. Text can be broken down into phrases, words, or even parts of words. This allows the algorithm to compare each token to the tokens around it to map the context and how it is being used. This is crucial in translation, as the same word or phrase could have different meanings or context depending on where it sits among other parts of the text (Rust et al., 2021). Current issues with tokenisation and text will be covered in the next section.

Neural machine translation (NMT), a deep learning-based approach, has greatly improved the quality of AI-generated translations by learning and exploiting complex patterns in language data (Bahdanau et al., 2014; Vaswani et al., 2017). Generative AI, which can create new text or content based on vast amounts of training data, is a form of NMT and has driven much of the recent excitement in AI fields, including translation. It often takes the tokens discussed above and learns the different patterns within this data to generate new text data. Underpinning this is a recent innovation called the Transformer Model. Simply put, this model allows the machine learning algorithm to process large quantities of the tokenised text at the same time, instead of one by one (Vaswani et al., 2023).

Training Issues in AI Models

As with any model used in the world of artificial intelligence, the data that it is trained on is critical to its accuracy and relevance to the subject matter it is being used for. Large Language Models (LLMs) are some of the most exciting technologies in AI to be able to tackle language translation, and understanding the data that they are trained on can help understand what their limitations and biases are.

LLMs such as ChatGPT 4 have been trained on data that is predominantly in the English language. This makes the tool less effective in multilingual environments. While information is scarce on ChatGPT 4's training data, 93% of ChatGPT 3's training data was in English ("Why AI Needs to Learn New Languages," n.d.). Tests done on these models also have shown that queries around cultural issues such as gun control and refugee policy yield responses that closer align with American sentiment.

A good example of how this becomes a problem is in countries such as India, where multiple languages are spoken. Because of the lack of training data in a lot of the languages spoken in India, translation accuracy can be much lower. GPT 4 fails to understand the Indian lexicon around Lesbian, Gay, Bisexual, Transgender, Queer and others (LGBTI+) issues (Joshi et al., 2023), which is one of the categories captured in the PA-X database. Understanding how natural language is used in different languages spoken in India, as well as other languages, is critical to how linguistic complexity in a single setting can have significant consequences. The research community should be aware of how such complexity multiplies globally.

The tokenisation of these languages themselves is also part of the problem, as tokenisation optimisation, or the way the algorithm splits up the language for processing, when based on an English model, is not optimised for other languages. If tokenisation was optimised for another specific model, it would take a fraction of the computing power and potentially require less data. It would be beneficial to keep abreast of efforts to optimise the tokenisation of data of languages other than English, as this would enhance the technology's accuracy and reduce its computing power requirements. Recent developments in language-specific AI models demonstrate the effectiveness of this approach. For example, 'Jais' or 'Jais chat' is a generative AI model specialized in Arabic language processing. Despite being trained on a smaller dataset, Jais has demonstrated performance that matches or exceeds ChatGPT in various benchmark tests for Arabic text translation and comprehension (Sengupta et al., 2023).

Various initiatives worldwide are developing LLMs trained on indigenous and local language corpora. The KenCorpus project at Maseno University in Kenya collects and digitizes text and audio from African languages to train LLMs (Wanjawa et al., 2023). In Southeast Asia, the CENDOL project focuses on training LLMs across Indonesian languages while incorporating an understanding of local Indonesian customs and culture (Cahyawijaya et al., 2024). These examples represent a small portion of emerging global projects developing LLMs for indigenous and local languages. Any programme interested in AI-assisted translation must keep up-to-date with these rapidly evolving multilingual generative AI models.

Modalities of Translating Peace Agreement Texts with AI

In the context of translating peace agreements for the PeaceRep programme, and for similar research projects, the basic modalities of translation available are manual, semi-automated, and fully automated processes. Each option has distinct advantages and drawbacks, particularly concerning accuracy, resource intensity, and ethical considerations. It is important to note that PA-X presents a particularly challenging case for automated translation because it functions as an archive of legal documents that constitute a form of "lex pacificatoria" (Bell, 2008). As legal texts, peace agreements require an exceptionally high standard of precision in translation. This creates a dual consideration: first, organisations dealing with less technically complex or legally binding documents may face fewer challenges with AI translation; and second, PA-X represents something of a stress test for AI translation capabilities – if automated systems can adequately handle the complexity of peace agreements, they would likely perform well across most other translation contexts.

Manual Translation

Manual translation, traditionally the gold standard for accuracy and nuance in language, involves human translators interpreting and converting text from one language to another. This method ensures a high level of contextual understanding, which is essential for the nuanced language of peace agreements. The words we speak and write contain information that can be analysed, and the nuance and flexibility of language used represents valuable information that is important to capture when studying peace agreements (Mayo, 2021). When PeaceRep projects require translations, they are frequently commissioned from native speakers, who often come from the regions affected by conflict, and who thus have a better contextual understanding of the agreement.

While PA-X benefits from a unique pool of legally-trained native speakers for languages like French, Spanish, and Portuguese through Edinburgh University, manual translation remains resource-intensive for other languages such as Russian, Pashto, and Dari, requiring complex and costly external contracting. This creates uneven coverage and significant financial barriers, particularly when considering potential expansions to include related resources like news reports or external documents that would require extensive translation to avoid English-language bias. When faced with this issue in the past, PeaceRep has sourced external companies who have had far greater access to different languages when required. However, this comes at a large increase in cost, which may not be sustainable in the long term.

An additional challenge for PA-X and similar projects is the inherently unpredictable nature of global peace processes. Unlike many translation projects with steady, predictable content sources, peace agreements emerge from dynamic political contexts that can shift dramatically without warning. For example, initiatives like President Petro's 'total peace' approach in Colombia have suddenly generated large volumes of Spanish-language texts requiring translation. Such unexpected surges in specific languages make advance planning difficult, requiring significant flexibility in translation budgets and approaches. This unpredictability further complicates reliance on traditional translation methods alone, as maintaining a stable network of human translators across all potentially needed languages becomes impractical. Instead, it necessitates the ability to rapidly scale translation capabilities in response to emerging situations, potentially through a combination of approaches rather than a single fixed solution.

It is important to note that while this report focuses on ethical issues around AI, there are ethical issues that arise from the use of manual translators. While PA-X only collects documents that are publicly available, there is still information that can appear in these agreements, such as contact details, that requires a certain degree of sensitivity. Some brief ethical training or guidelines should be made available for translators such as post-graduate students, who are not professionally trained in translation and who might not be aware of some of these concerns.

The options available to PeaceRep if only manual methods are used are a combination of:

- Continuing to use post-graduate students from the Law School or wider University, with increased guidance on practical and ethical considerations specific to our project.
- Using professional translation companies for languages not covered by the University or when we receive many agreements that require translation.
- Only using professional translation companies going forward.
- Relying on PA-X coders and other members of the consortium to respond to translation needs.

These options are also reliant on continual funding as commissioning translation is costly.

Semi-Automated Translation

Semi-automated translation represents a middle ground, combining human oversight with automated tools to enhance efficiency and manage resource demands better. In this model, initial translations are generated by AI-based tools, and then human translators review, correct, and refine the output. This approach leverages the speed and scalability of AI whilst maintaining a level of accuracy and contextual sensitivity only humans can provide. However, semi-automated processes still require significant human intervention, particularly for languages or dialects with less representation in training data, which can limit the efficiency gains from automation.

In PA-X, the semi-automated approach is already partially being enacted, as different AI tools have made the search for peace agreements easier by giving researchers the ability to translate agreements to an acceptable level of accuracy to make a judgement as to whether the document warrants inclusion into the database. Applications that assist in this task include:

- **Google Translate**

  Google translate is a widely-used, free web-based translation service developed by Google. It supports translation between 133 languages for text, documents, and websites. The service offers multiple modalities including text translation, speech translation, instant camera translation, and website translation. Google Translate has evolved to incorporate neural machine translation technology, significantly improving its accuracy and fluency. While useful for quick translations, it may lack the nuanced understanding required for complex peace agreement texts.

- **Copy Fish**

  Copyfish is a browser extension designed to extract text from various media formats including images, videos, and PDFs. It supports multiple languages and can recognize text in complex documents. This tool is useful to help identify the contents of a document in image form in order to assess its suitability for PA-X.

Claude AI

Claude AI is a large language model developed by Anthropic, trained on diverse text data to understand language, generate responses, analyse information, and assist with various tasks through natural conversation. It processes language through deep learning techniques to provide contextually relevant and coherent text across a wide range of tasks including writing, analysis, coding, and information retrieval. While capable of handling multiple languages and domains, Claude has limitations regarding knowledge beyond its October 2024 cut-off date and potential biases in its training data that may affect responses in specialized contexts like legal, cultural, or technical domains.

ChatGPT 4

ChatGPT 4, developed by OpenAI, is an advanced large language model capable of generating human-like text based on input prompts. It leverages deep learning techniques to understand and generate contextually relevant and coherent text. While not specifically designed for translation, ChatGPT 4 can be used for various language-related tasks including drafting, summarization, and potentially assisting in translation efforts. It is another tool to get a first assessment of peace agreements that are already in machine readable format. The model's use in full translation is not advised as it is trained on biased data and does not understand certain legal and cultural contexts.

Generative AI models, such as ChatGPT 4, offer a new way using a semi-automated approaches to translation. Models such as these, while lacking some of the nuance of human translated material, can extract key metadata from agreements in different languages that would be otherwise time consuming for a translator or PA-X coder. Once a translation is determined to be needed for an agreement, this could mean that key information such as dates, locations, conflict levels, etc. can be extracted and entered into the database before the document has been translated by a human coder, thus saving time. There are drawbacks to this approach. Workflows will need to be tested for efficiency, as hybrid approaches to combining human and AI resources can often increase the initial time needed for tasks. This is particularly true for 'rare' languages, for which current AI models have a lower accuracy rate when translating. This means that work of an individual simply 'verifying' an AI-translated text could be potentially more time consuming and increases the likelihood that key terminology could be mistranslated or misunderstood.

Therefore, options available to the programme using a semi-automated approach are:

- Using AI translation solutions to assist in the discovery and approval of potential peace agreement documents.
- Using AI, particularly generative AI, to extract key metadata from peace agreements in different languages to decrease administration time.
- Using AI to translate an agreement, and a native speaker/language expert to verify the output against the original text.
- Developing efficient hybrid translation workflows.

Fully-Automated Translation

Fully automated translation, powered by advances in generative AI and neural machine translation technologies, offers the potential for rapid translation at scale. These systems can translate vast amounts of text quickly, with the overall cost of translation technology decreasing year over year as the technology matures. If PeaceRep needs to translate a large number of documents in a short amount of time, one of these solutions could help reduce the amount of time it takes to translate. This ought to be the case for other organisations with similar translation needs.

As PeaceRep progresses from collecting peace agreement data to developing a suite of tools that analyse peace processes holistically and incorporate multiple text sources, fully automated translation solutions could be more appropriate. PeaceRep already uses Universal Sentence Encoder in some work, which can analyse texts and their similarities in multiple languages without having to initially use a translation programme (Cer et al., 2018, Gardner, 2023, pp. 11-12). However, while this is exceptionally useful for the development of research data, it is not certain whether the translation would match the quality of expert human translators.

Despite significant advancements, fully automated systems often struggle with the subtleties of legal language, cultural nuances, and technical jargon specific to peace agreements. These challenges not only highlight the current limitations of AI translation models but also emphasize the crucial role of extensive research and contextual knowledge in accurately translating these documents. The way language is used in peace agreements can, in fact, inform the classification of the agreement itself, which will be further explored in the case studies section of this report. The ethical concerns discussed later in this report also become more prevalent when using this approach. For the PeaceRep programme, the choice between these translation options must be informed by the specific needs of accurate, ethically sound translations that respect the programme's resource constraints. An effective translation for PeaceRep must accurately convey the original text's intent, legal implications, and cultural context, while remaining cost-effective and ethically sound. It would involve a careful balancing of accuracy, efficiency, and the programme's broader goals, including building sustainable partnerships and respecting the communities involved in peace processes.

Case Studies and Comparative Analysis

The following section will briefly discuss the current workflow involved when receiving an agreement which needs translation, followed by a discussion of three different case studies which are illustrative of some of the challenges of using automated translation for peace agreements.

Current Workflow

Agreements in PA-X are sourced from various mediums. Sources in languages other than English are rarely received as machine-readable text that can be entered into the PA-X database for 'coding' (the process of categorising peace agreement segments according to content categories, such as women, girls and gender). Instead, texts often arrive as images, or in other formats where the copy function is not adequate.

This means the text needs to be extracted from the source material manually or with Optical Character Recognition (OCR) software. Both methods can be mistake-prone, but OCR is the most fraught, as up to this point the best software available in this field still struggles to accurately recognise different characters, line spacing and languages other than English accurately. As the programme deals with language in signed legal documents, complete accuracy is critical. This accuracy problem creates a roadblock in the workflow of automated translation: before a document can be translated using machine learning, it needs to be converted to plain text.

Currently the agreements are translated by paid translators, so the agreement can be provided to them in the form of the original source of the agreement. Unlike computers, humans are very good at reading text which is not typed, or which is distorted – for example on pages of printed documents that have been bent or folded. From there, the translated agreements are sent back to PeaceRep in plain text format that can then be added to the PA-X database.

Analysis of Case Studies (See Appendices)

Appendix 1 shows a test where three peace agreements were selected for their representation of the issues that arise when trying to translate agreements using AI resources. The first translation shown is provided by a human translator. The second translation is provided by the popular generative AI programme, ChatGPT.

The first agreement is entitled 'The Peace Agreement District Kurram executed in Para Chanar', signed on 14 July 2023. This local inter-tribal six-point ceasefire agreement outlines the areas of enforcement in the town, provides for elements of demobilization, and sets out terms of violation. As this agreement was identified as being in Urdu, it was sent to Umar Shehzad, a native speaker, for translation. Through his own expertise, he identified that the agreement was mainly in Urdu, except for one word, " تیگہ " or "taega", which occurs in the third clause of the peace agreement and is an important provision of the agreement. It states, '3. Starting with immediate effect, the "taega" has been placed between the fighting parties for one year.' (Bell et al., 2024). This nuance is important because a literal translation would interpret this word as 'stone,' whereas in reality, it refers to a traditional ceremony meant to enact a ceasefire for up to one year. Clarifying this distinction is crucial for understanding the complexities of this local peacemaking process, which is conveyed through a Pashto term within an Urdu text. The ChatGPT translation does not provide a full translation of the agreement. Instead, it recognizes the presence of both Urdu and Pashto in the text but is unable to translate it.

The second example in the Appendices is the agreement entitled 'Acte de Non Agression' (Bell et al., 2024). This agreement was highlighted as it illustrates a common issue in the workflow of adding peace agreements to the database: the need to extract the text from the image of the agreement. When using human translators, this is not a concern, since humans are able to read the text in the image. If AI translation techniques are to be used in the translation of peace agreements, the plain text will have to be extracted from the image first. The agreement example demonstrates the typical condition of source materials. The image is blurry and tilted at an angle. This makes it difficult for optical character recognition software to extract text accurately. The appendices show that when it is extracted using OCR, the text is not accurate and therefore not as accurate as the human translated example.

The third example uses the agreement 'Acuerdo sobre Cese al Fuego y de Hostilidades Bilateral y Definitivo y Dejación de las Armas entre el Gobierno Nacional y las FARC-EP' (Bell et al., 2024). This example highlights another problem inherent in using AI for translation, which is length. The PeaceRep programme frequently encounters agreements that are ten pages or longer, and this can present problems for automated translation. In this case, the agreement is 15 pages long. ChatGPT 4 would not translate the document as its length violated its use case policy. As mentioned previously, a document needs to be machine readable before it can be translated in a machine learning programme. The current generative AI programmes have length restrictions on documents.

To understand how a programme using AI would work with the larger document in the third example, Google Translate was used. This programme did a better job translating the document, mainly due to the source document being clear with little noise. It did, however, present formatting issues and omitted the first two paragraphs of the agreement. Again, as the accuracy and formatting of the text can affect how it is coded for PA-X, this does not provide a suitable option for full automation.

Ethical Concerns

Ethical considerations underpin PeaceRep's approach to research and the application and advancement of PeaceTech. Ethical considerations are especially important when working with AI, particularly Natural Language Models, which use machine learning. Ethical concerns around the use of AI in the translation of peace agreements would be best categorised as 'Good Practice and Process Concerns' (p. 182), as described in Christine Bell's 2024 book, *PeaceTech: Digital Transformations to End Wars*. This categorisation of ethical concerns around PeaceTech focuses on commitments to forms of ethical practice in the design of research in both the tech and peacebuilding worlds to produce positive outcomes. Three types of ethical concerns stand out relating to the translation of peace agreements: environmental concerns, bias in training of AI models, and the problem of offering peace agreements as additional source of data for commercial AI models to ingest. As an organisation engaged in researching peace and transition processes and their outcomes, PeaceRep is aware of the environmental factors contributing to conflict. Therefore, it is imperative for to consider the substantial environmental impact associated with deploying large-scale NLP models. These models demand extensive computational resources, and accurately assessing their carbon footprint presents considerable challenges. A pertinent study by Luccioni et al. (2022) delves into the complexities of quantifying such impacts, underscoring the necessity for PeaceRep to remain at the forefront of discussions on sustainable AI utilisation in peace research and beyond.

In pursuit of minimizing environmental harm, understanding these complexities of the significant carbon emissions associated with NLP models becomes crucial. Research by Strubell et al. (2019, p. 1) reveals that the emissions from training a single NLP model, inclusive of the tuning and experimental phases, can rival the lifetime emissions of an automobile. This stark comparison emphasizes the need for PeaceRep to rigorously evaluate and adopt environmentally responsible frameworks for assessing the impact of the programme's PeaceTech and NLP initiatives. Moreover, maintaining transparency about PeaceRep's engagement with NLP technologies aligns with the programme's values and strengthens PeaceRep's commitment to sustainability.

A significant ethical concern with NLP models is training bias. As stated earlier in the report, the lack of multilingual understanding could lead to biased translations and incorrect analysis. This is shown to be prevalent in the ChatGPT models (Zhuo et al., 2023).

The cultural context and composition of the training data used for the translation model could end up prioritising one set of interpretations over another. *In The Ghost in the Machine has an American Accent: Value Conflict in GPT-3*, Johnson and colleagues (2021) demonstrate that sometimes issues more important to US culture could affect interpretation of a document. Understanding the unique cultural contexts of the peace agreement translation will be further analysed in a case study in this report.

While it can be challenging to address the ethical concerns surrounding language bias in translations holistically, there are certain measures that can be adopted. One crucial step would be to continue to engage a diverse group of researchers, board members, and partners from all over the world on these ethical challenges. A diverse group can help the programme stay in touch with technological needs and developments worldwide. With this approach, the programme can better understand how new technology addresses ethical concerns identified in the report and improve its translation capabilities.

Finally, this report will address the ethics around data ownership of AI translation models. While this report will not address the legal implications of this subject, the overall ethical concerns that appear when using data from gained from other countries is important to understand. In line with the 'do no harm' approach in PeaceRep's data-driven work, we often handle data that could be considered 'property' of the people and areas we research (Kamocki & Witt, 2022). This is an important consideration when evaluating translation programmes that process large amounts of data from various sources and the methods by which they obtain it. Establishing a system within PeaceRep to evaluate NLP models' adherence to different standards of ethical ownership practises will be important going forward. This could complement the existing principles applied to the PA-X database, known as the FAIR principles. These principles require that datasets be findable, accessible, interoperable and reusable (Wilkinson et al., 2016).

Recommendations

This section proposes strategic guidelines and workflow optimisations for the use of AI translation in peace agreement processing. These recommendations aim to enhance the effectiveness of translating peace agreements by leveraging artificial intelligence, while considering the specific scale of this work and the invaluable role of human experts and international partners.

Strategies for Implementing AI Translation in Peace Processes

The use of generative AI, including machine translation technologies, offers a promising avenue for drafting initial translations of peace agreements. These AI-driven tools can act as efficient exploratory mechanisms, providing a valuable first pass at translating complex texts. However, it is important to recognize that peace agreement translation operates on a smaller scale, where the extensive use of generative AI might not be practical or desirable due to the nuanced nature of these texts. To ensure the long-term viability and accuracy of translated agreements, this report makes the following recommendations:

- **Adopt a Semi-automated Approach**
  Combining AI's efficiency with human expertise could mitigate translation inaccuracies, particularly for languages underrepresented in AI training datasets. This hybrid model ensures the fidelity of translations while leveraging AI to streamline the process.

- **Address Ethical and Environmental Considerations**
  As AI models require substantial energy consumption, evaluating the environmental impact of deploying these technologies is crucial. Furthermore, addressing inherent biases in AI training data is essential to ensure equitable and accurate translations across languages and cultures.

- **Optimise Workflows**
  Developing efficient workflows that integrate AI translation tools can help in the preliminary analysis and extraction of key metadata from peace agreements, enhancing database comprehensiveness and accessibility.

- **Foster Collaborations**
  Strengthening partnerships with academic institutions, professional translators, and leveraging global networks can enrich the translation process, incorporating diverse linguistic expertise and contextual understanding.

- **Invest in Continuous Research and Development**
  Further research is needed to identify and address the current limitations of AI in translation, such as improving models' capacity to understand cultural contexts and technical jargon, as well as keeping abreast of advances in Optical Character Recognition (OCR) technology to support AI translations.

Mitigation of Risks and Addressing Limitations

To mitigate the risks associated with AI translation and address its limitations, it is imperative for stakeholders, particularly those in conflict-affected countries, and practitioners to enhance their understanding of AI and its ethical challenges. Improving knowledge around AI technologies will enable strategic decision-making in a rapidly evolving environment. It is crucial that these decisions are co-created with actors from conflict-affected countries and other underrepresented groups to ensure inclusivity and relevance in the translations of peace agreements.

Suggestions for Further Research and Development

Looking ahead, it is essential to:

- Stay informed about advancements in OCR and AI Translation technologies. Innovations in these areas can significantly improve the accuracy and efficiency of translating peace agreements.

- Monitor efforts to establish Large Language Models (LLMs) in various countries, particularly those trained on local languages, as they could provide valuable insights and tools for peace agreement work. Exploring the development of small-scale, local LLMs could offer tailored solutions for translating peace agreements, ensuring they are culturally and linguistically aligned with the regions they pertain to.

- Be aware of and remain transparent about the impact of these models on climate change. As the effects of climate change become more pronounced, particularly in countries disproportionately exposed to both climate extremes and conflict, ethical considerations of using NLP models will become increasingly important.

- Consider who owns the data involved in these models, and how the models are trained.

In conclusion, embracing these strategies and recommendations can improve translation workflows for peace agreements, mitigate risks associated with AI technologies, and pave the way for further research and development in this field. The collaborative and informed approach outlined here ensures that AI translation serves as a complementary tool to invaluable human expertise, ultimately contributing to more effective and impactful peace processes research.

References

Bahdanau, D., Cho, K., & Bengio, Y. (2016). Neural Machine Translation by Jointly Learning to Align and Translate. arXiv.Org. https://doi.org/10.48550/arxiv.1409.0473

Bell, C. (2008). *On the law of peace : peace agreements and the lex pacificatoria.* Oxford University Press.

Bell, C., & Badanjak, S. (2019). Introducing PA-X: A new peace agreement database and dataset. Journal of Peace Research, 56(3), 452-466. Available at https://pax.peaceagreements.org/

Bell, C. (2024). PeaceTech : Digital Transformation to End Wars. (1st ed.). Springer International Publishing AG. https://doi.org/10.1007/978-3-031-38894-1.

Bell, Christine, Sanja Badanjak, Adam Farquhar, Juline Beaujouan, Tim Epple, Robert Forster, Astrid Jamar, Kevin McNicholl, Sean Molloy, Kathryn Nash, Jan Pospisil, Robert Wilson, Laura Wise. (2024). PA-X Codebook, Version 8. Peace and Conflict Resolution Evidence Platform (PeaceRep), University of Edinburgh.

Cabrera, L. (2024). Babel Fish Democracy? Prospects for Addressing Democratic Language Barriers through Machine Translation and Interpretation. American Journal of Political Science, 68(2), 767–782. https://doi.org/10.1111/ajps.12728

Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., Sung, Y.-H., Strope, B., & Kurzweil, R. (2018). *Universal Sentence Encoder.* https://doi.org/10.48550/arxiv.1803.11175

Dařena, F., & Süss, M. (2020). Quality of Word Vectors and its Impact on Named Entity Recognition in Czech. European Journal of Business Science and Technology, 6(2), 154–169. https://doi.org/10.11118/ejobsat.2020.010

de Vries, E., Schoonvelde, M., & Schumacher, G. (2018). No Longer Lost in Translation: Evidence that Google Translate Works for Comparative Bag-of-Words Text Applications. *Political Analysis,* 26(4), 417–430. doi:10.1017/pan.2018.26

Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering,* 66(1), 111–126. https://doi.org/10.1007/s12599-023-00834-7

Gardner, R. (2023). Semantic Analysis to Support Peace Analytics [Peace Analytics Series]. PeaceRep: The Peace and Conflict Resolution Evidence Platform, University of Edinburgh. https://peaceagreements.org/.

Generative Models. (n.d.) Retrieved March 7, 2024, from https://openai.com/research/generativeimodels.

Henry, N. (2024). Extracting named actors from text: Using named entity recognition (NER) in peace and conflict studies. PeaceRep: The Peace and Conflict Resolution Evidence Platform. https://peaceagreements.org/.

Johnson, R. L., Panai, E., Pistilli, G., Kalpokiene, J., Menédez-González, N., Bertulfo, D. J., & Duran, L. D. D. (n.d.). The Ghost in the Machine has an American accent: Value conflict in GPT-3.

Joshi, A., Rawat, S., & Dange, A. (2023). Evaluation of large language models using an Indian language LGBTI+ lexicon. The AI Ethics Journal, 3(1), Article 1. https://doi.org/10.47289/AIEJ20231109

Kamocki, P., & Witt, A. (2022). Ethical Issues in Language Resources and Language Technology – Tentative Categorisation. In N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, J. Odijk, & S. Piperidis (Eds.), *Proceedings of the Thirteenth Language Resources and Evaluation Conference* (pp. 559–563). European Language Resources Association. https://aclanthology.org/2022.lrec-1.59

LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. Nature 521, 436–444 (2015). https://doi.org/10.1038/nature14539

Lucas, C., Nielsen, R. A., Roberts, M. E., Stewart, B. M., Storer, A., & Tingley, D. (2015). Computer-Assisted Text Analysis for Comparative Politics. *Political Analysis,* 23(2), 254–277. doi:10.1093/pan/mpu019

Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2022). Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model (arXiv:2211.02001). arXiv. http://arxiv.org/abs/2211.02001

Mayo, W. (2021). Language Is Data. *Expert.Ai.* https://www.expert.ai/blog/language-is-data/

Organisation for Economic Co-operation and Development. (2019). Recommendation of the Council on Artificial Intelligence. OECD Legal Instruments. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

Park, H. H., Zhang, K. J., Haley, C., Steimel, K., Liu, H., & Schwartz, L. (2021). Morphology Matters: A Multilingual Language Modeling Analysis. Transactions of the Association for Computational Linguistics, 9, 261–276. https://doi.org/10.1162/tacl_a_00365

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI. Retrieved from https://openai.com/research/language-models-are-unsupervised-multitask-learners.

Rust, P., Pfeiffer, J., Vulić, I., Ruder, S., & Gurevych, I. (2021). How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models. In C. Zong, F. Xia, W. Li, & R. Navigli (Eds.), Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers) (pp. 3118–3135). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.243

Sengupta, N., Sahu, S. K., Jia, B., Katipomu, S., Li, H., Koto, F., ... & Xing, E. (2023). Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models. arXiv preprint arXiv:2308.16149.

Singh, A. K., & Strouse, D. J. (2024). Tokenization counts: The impact of tokenization on arithmetic in frontier LLMs (arXiv:2402.14903). arXiv. http://arxiv.org/abs/2402.14903

S. M., Mahendra, M. I., Putri, D. A., Wilie, B., Winata, G. I., Aji, A. F., Purwarianti, A., & Fung, P. (2024). *Cendol: Open Instruction-tuned Generative Large Language Models for Indonesian Languages* (arXiv:2404.06138). arXiv. https://doi.org/10.48550/arXiv.2404.06138

Son, J., & Kim, B. (2023). Translation Performance from the User's Perspective of Large Language Models and Neural Machine Translation Systems. Information, 14(10), Article 10. https://doi.org/10.3390/info14100574

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP (arXiv:1906.02243). arXiv. http://arxiv.org/abs/1906.02243

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). Attention Is All You Need (arXiv:1706.03762). arXiv. http://arxiv.org/abs/1706.03762

Wanjawa, B., Wanzare, L., Indede, F., McOnyango, O., Ombui, E., & Muchemi, L. (2023). Kencorpus: A Kenyan Language Corpus of Swahili, Dholuo and Luhya for Natural Language Processing Tasks. *Journal for Language Technology and Computational Linguistics,* 36(2), 1–27. https://doi.org/10.21248/jlcl.36.2023.243

Why AI needs to learn new languages. (n.d.). The Economist. Retrieved January 27, 2024, from https://www.economist.com/science-and-technology/2024/01/24/why-ai-needs-to-learn-new-languages.

Wilkinson, M. D., Dumontier, M., Aalbersberg, Ij. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J. W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., … Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3(1), 160018–160018. https://doi.org/10.1038/sdata.2016.18

Windsor, L. C., Cupit, J. G., & Windsor, A. J. (2019). Automated content analysis across six languages. PLOS ONE, 14(11), e0224425. https://doi.org/10.1371/journal.pone.0224425

Zhuo, T. Y., Huang, Y., Chen, C., & Xing, Z. (2023). Red teaming ChatGPT via Jailbreaking: Bias, Robustness, Reliability and Toxicity (arXiv:2301.12867). arXiv. http://arxiv.org/abs/2301.12867

Appendices: Tests on the Translation of the Text

For full appendices and supporting materials on tests on the translation of example texts, please scan the QR code below or visit:

https://edin.ac/41WtEwv

View Our Key Findings

peacerep@ed.ac.uk
PeaceRep

www.peacerep.org
pax.peaceagreements.org

About Us

PeaceRep: The Peace and Conflict Resolution Evidence Platform is a research consortium based at Edinburgh Law School. Our research is rethinking peace and transition processes in the light of changing conflict dynamics, changing demands of inclusion, and changes in patterns of global intervention in conflict and peace/mediation/transition management processes.

Consortium members include: Conciliation Resources, Centre for Trust, Peace and Social Relations (CTPSR) at Coventry University, Edinburgh Law School, International IDEA, LSE Conflict and Civicness Research Group, LSE Middle East Centre, Queens University Belfast, University of St Andrews, University of Stirling, and the World Peace Foundation at Tufts University.

PeaceRep is funded by UK International Development from the UK government.

PeaceRep: The Peace and Conflict Resolution Evidence Platform
peacerep@ed.ac.uk | https://peacerep.org

University of Edinburgh, School of Law
Old College, South Bridge EH8 9YL