Extracting Named Actors from Text: Using Named Entity Recognition (NER) in Peace and Conflict Studies

Niamh Henry

PeaceRep
Peace and Conflict Resolution Evidence Platform

PEACE ANALYTICS SERIES

Author: Niamh Henry

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
@Peace_Rep_
https://www.facebook.com/PeaceRepResearch
https://www.linkedin.com/company/peacerep/

This research is supported by the Peace and Conflict Resolution Evidence Platform (PeaceRep), funded by the UK Foreign, Commonwealth & Development Office (FCDO) for the benefit of developing countries. The information and views set out in this publication are those of the authors. Nothing herein constitutes the views of FCDO. Any use of this work should acknowledge the authors and the Peace and Conflict Resolution Evidence Platform.

About the author: Niamh Henry is a Research Fellow in Data Engineering with the Peace and Conflict Resolution Evidence Programme (PeaceRep) at the University of Edinburgh. She works on the organisation and extension of peace and conflict data and develops innovative PeaceTech tools to support better understanding peace and transition processes. Niamh is a co-author of the new Peace Agreements Actor Dataset (PAA-X), in addition to being a co-creator of the PA-X Tracker, a new peace and transition process tracker from PeaceRep (https://pax.peaceagreements.org/tracker/). She holds an MS in Information Science from the University of Amsterdam, and an MA in Digital Media and Information Studies from the University of Glasgow.

Thanks to Allyson Doby, Alice Raymond, and Rick Smith of Smith Design Agency for production work.

Design: Smith Design Agency

Contents

Peace Analytics Series

PeaceRep's Peace Analytics Series features the research methodology underlying the PeaceTech innovations of the PeaceRep programme.

The series includes: data scoping research; 'how to' discussions relating to particular challenges in the field of visualisations and geocoding; and other proof-of-concept tech-based innovations, such as the use of natural language processing. It is intended to present the methodologies and decisions behind our PeaceTech digital research, to make it transparent, and to contribute to establishing a new research digital infrastructure in the field of peace and conflict studies, by supporting others to reuse and repurpose our methodologies and findings.

Executive Summary

This report explores the potential role of Named Entity Recognition (NER) in peace and conflict studies, drawing insights from PeaceRep's use of the technology to extract signatory organisations from a large-scale database containing peace agreements – the PA-X Peace Agreement Database. PeaceRep's successful application of NER in extracting signatory data from peace agreement texts provides an example of how to navigate complex textual data relevant to ending conflict, in a time where we increasingly require data as evidence for decision making. As text remains the primary medium in this domain, NER emerges as a crucial tool in identifying, extracting, and structuring key data, facilitating streamlined research and information extraction.

Highlighting NER fundamentals and methodological approaches, this report advocates for transparent artificial intelligence, human oversight, and a functional approach tailored to peace and conflict studies. Accompanied by practical demonstrations in Python Jupyter Notebooks, the report showcases NER's applications—from document summarisation to geospatial and temporal analysis. Ultimately, it emphasises NER's potential in deciphering complex textual data while emphasising, and recommending, the need for nuanced approaches that balance NER's strengths with its limitations in this critical field.

Definitions

Term	Definition
Peace Analytics	This use of data pertaining to peace and conflict related issues, for use in discovery, communication, and interpretation of patterns in the data, so as to inform conflict resolution and peacebuilding practice. (Bell, 2024)
Natural Language Processing (NLP)	Natural Language Processing (NLP) involves the interaction between computers and human languages, enabling machines to understand, interpret, process and generate human language.
Named Entity Recognition (NER)	Named Entity Recognition (NER) is a subtask of NLP that identifies and categorises named entities (such as names of persons, organisations, locations) within textual data.
Structured Data	Data that is organised and formatted in a predefined manner, typically residing in databases with a well-defined schema, making it easily searchable and analysable.
Unstructured Data	Data that lacks a predefined structure or format, such as text documents, audio recordings, or images, making it more challenging to analyse and compare using traditional methods.
Slightly Structured Data	Data that falls between structured and unstructured data, having some organisation or loose formatting but not as rigid as structured data. For example, Factivia or BBC Monitoring structure sources with a range of metadata. However, not all the content is structured.

Term	Definition
Data “in-the-wild”	Refers to real-world data that is typically unfiltered, unprocessed, and obtained directly from primary sources, reflecting the actual context or environment. For example, from social media, news outlets and field reports.
String Matching Algorithms	Algorithms designed to identify similarities or patterns between strings of characters, commonly used in text processing or data cleaning tasks. Similar to executing ‘Find’ (ctrl+F) or ‘Search’ on a document/webpage as it returns exact matches for the exact characters queried (the *string* of characters).
Training a Model	In machine learning, "training a model" refers to the process of feeding labelled/annotated data into an algorithm to enable it to learn patterns and make predictions or classifications.
Training Data	The dataset used to train machine learning models, containing examples or instances with associated labels or target outcomes used for model learning and validation.
Annotated Documents	Documents that have been annotated or labelled with additional information, such as named entities, sentiments, or classifications, aiding in the training of machine learning models.
Supervised Learning	A type of machine learning where the algorithm learns from labelled/annotated training data, making predictions or classifications based on patterns identified during training.

Term	Definition
Unsupervised Learning	Machine learning where algorithms explore data without labelled outcomes, finding patterns or structures autonomously, often used for clustering or dimensionality reduction.
Segmenting Documents	The process of dividing larger documents or texts into smaller, more manageable sections or segments, such as sentences or paragraphs, often to facilitate analysis or extraction of specific information.
Stop Words	Commonly used words in a language (such as "the," "and," "is") that are often filtered out in NLP or text analysis tasks due to the noise from their high frequency and low semantic value.
Metadata	Data that provides information about other data, describing various attributes such as source, format, authorship, or creation date, aiding in the organisation and retrieval of data.
Overfitting	Occurs in machine learning when a model could learn to perform extremely well on the training data (so have a high accuracy) but struggles to generalise to new, unseen data due to training the model on specific data and categories, rather than a diverse set of categories. This can also occur if noise or irrelevant patterns are captured during training, or if there is not a sufficient amount of training data.
Web-crawler	An automated tool designed to systematically browse and index content on the internet. They can be used for gathering information and monitoring particular sources.

1 Introduction to NER in Peace and Conflict Studies

This report provides an overview of how named entity recognition (NER) can be used in peace and conflict studies. As a form of natural language processing (NLP), NER can be used to extract information from text to produce structured text and datasets for use in many fields, including peace and conflict research and practice.

The report is informed by the work of the Peace and Conflict Resolution Evidence Platform (PeaceRep). PeaceRep is pioneering Peace Analytics – the use of data pertaining to peace and conflict related issues, for use in discovery, communication, and interpretation of patterns in the data, so as to inform conflict resolution and peacebuilding practice (Bell, 2024).

As part of our Peace Analytics work, PeaceRep is experimenting with the use of NER models to extract information from global peace agreements in our PA-X Peace Agreements Database,[1] through categorising unstructured text into structured text and datasets to provide evidence-based data on peace and conflict resolution processes.

The rapid development of the Internet and diversification of information carriers have exponentially increased the amount of information people encounter daily. Finding and extracting key information in this massive pool of data has become an urgent problem.[2] This is particularly prevalent within the field of peace and conflict studies, where text is an important medium for recording information across peace and conflict resolution institutions, governments, and organisations. Agreements to end conflict through negotiation and compromise are often formally documented through text-based documents like peace agreements, United Nations resolutions, national constitutions, and other texts. Then copious documentary sources often document attempts to implement the commitments and injunctions made in these documents. There is therefore a need for key data to be identified, extracted, and structured across these text sources to enable information extraction for research and practice. NLP techniques can help sift through and make sense of this vast amount of text related to peace and conflict in complex contexts.

This report will outline the basics of NER and the primary purpose it was used for with reference to PA-X, to illustrate the vast applications of this technology and potential uses within the peace and conflict resolution sphere. The report will evaluate a range of standard NER methods for differing levels of technical and coding skillsets. To provide readers with the ability to execute NER on any set of relevant texts, this report is accompanied by two Python Jupyter Notebooks to provide a straightforward example code base for running NER on different types of text (Word documents and CSV files), and to identify, extract and categorise entities or actors relevant to peace and transition processes under a unique identifier, using our manually generated actor dictionaries. This practice can be fruitful for a range of uses relevant to peace and conflict studies, such as summarising documents, extracting actors' names from texts, tracking events and incidents, geospatial analysis via extracting place names, temporal analysis by extracting dates, and monitoring media and social media, amongst many other powerful applications.

2 Named Entity Recognition

What is Named Entity Recognition?

Named Entity Recognition (NER) is an entity extraction technology that uses natural language processing (NLP) techniques to identify, extract and categorise meaningful entities in text. NLP enables machines to understand, interpret, and generate human language in a 'machine-readable' manner. A 'named entity' is the most basic information unit in natural language, with NER being a prerequisite and foundation for researchers to conduct research on NLP and text mining. A named entity can be a word or a group of words or phrases that refers to that entity – usually, but not limited to, proper nouns such as place, location or organisation names, or other key phrases such as dates and people's names. NER is the basis for many advanced tasks in NLP, such as syntax analysis, text summary, information retrieval and automatic Q&A, to name a few.[3]

NER models are highly useful and powerful applications across fields. NER can take any unstructured text – such as reports, documents, articles, blogs, emails – and effectively and efficiently recognise any mentions of people, organisations, geopolitical entities, locations, products, events, works of art, dates, quantities (such as money, percentages, or other cardinal values), and many more entity types, depending on what the model has been trained on (what data/information has been previously supplied to the model in terms of entity types and how they are mentioned in text to enable the algorithm to identify patterns, and apply them to new unseen text). There are a range of approaches to NER, differing in levels of complexity, accuracy, time, and resources, that variously suit a range of purposes, needs and abilities.

How is NER Relevant to Peace and Conflict Studies?

NER is particularly useful in the realm of peace and conflict studies, which relies on data to inform evidence-based decision-making and practice. However, the nature of information storage norms in the field can make it difficult to access information in a way that can be efficiently extracted and aggregated. Text-based files have historically been the preferred medium for recording information, and for good reason – textual reports can be easily disseminated, and they can capture the complexity of issues such as narratives, nuances, and concerns, and can therefore be more subjective and inclusive than structured data collections. This is particularly important when documenting individual stories surrounding human rights violations and human suffering, as the complexity and severity of the issues cannot be communicated in the same way through structured data collections.

Although these are valid reasons for storing information in unstructured text files, there are downsides to this approach. Textual reports can be biased and subjective, they can be difficult to preserve and archive, and most importantly, in terms of Peace Analytics, they can be difficult to analyse and compare. As a result, there has been a growing interest in collecting and structuring data in peace and conflict settings, but there are still a number of challenges to doing this: the lack of reliable and accessible data; the need for customised tools and methods to enable practitioners to analyse and visualise structured data in a meaningful way; and the need to build capacity and awareness on the benefits of using structured information for analysis. However, NER can bridge this gap by extracting structured data from unstructured text for a wide range of valuable and required applications – some of which are outlined in the section below. NER is currently being used in data gathering and information extraction in peace and conflict studies, by data collectors as outlined in the Section below, but its full potential has not yet been unlocked.

What Text and Data can NER be Performed On?

A wide variety of well-known and trusted sources of information provide vast amounts of potentially valuable data relevant to peace and conflict resolution. However, there is not enough time and resources for one person, or a team of people, to read, understand, extract, and categorise information from all documents in even one of these sources, especially in such a time-sensitive field. NER models can be leveraged to easily extract information from a variety of sources for a range of applications. This method can be conducted on varying levels, depending on the users' needs, technical ability, resources, and requirements for the desired outcome.

NER is currently used in data creation pipelines to create slightly structured text sources which do some pre-processing to aggregate media and news reports from a range of outlets across the globe, with varying contexts and languages, to make them easily accessible to users via topics, countries, themes, specific events, and many more factors. This includes popular tools such as Factiva, which aggregates content from a range of sources; the GDELT Project (Global Database of Events, Language and Tone) which monitors the world's news media and traverses it into the CAMEO (Conflict and Mediation Event Observations) framework,[4] which categorises events into different event codes; or BBC Monitoring, which tracks, translates, analyses, and summarises global media to help users make sense of world events.

Although NER is used to categorise content into particular events and topics, these sources also provide organised text sources that NER can then be performed on. For instance, a user can take all aggregated content from a specific news source or surrounding a specific topic or event – such as reports on armed conflict in Democratic Republic of the Congo – and run NER to identify mentions of entities such as armed actors, states, locations, or key dates.

Although these slightly structured and aggregated text collections are useful resources, they are not easily accessible to everyone in the field. There are usually financial burdens – such as subscriptions, which are often offered on a quote-only rate so it is unknown how expensive the venture will be – which is important for planning, budgeting, and executing reliable methodologies for researchers or other organisations within the field, or there are often pay-per-query rates which make the end-price of the data collection unknown. Furthermore, running queries to extract information from these sources is not straightforward: high levels of technical abilities, capacities and resources are required to access the raw content through sophisticated queries, and it often returns 'big data' which requires extensive reliable storage of raw data, and mechanisms to easily extract and interpret the data.

Nevertheless, despite the hurdles in the use and accessibility of these slightly structured content aggregators for data collection and extraction, there are also a range of other relevant text-based content sources within the peace and conflict domain that are open and free, easily accessed, and hold key information. For instance, the PA-X Peace Agreements Database and corpus[5] hosts, records, and archives over 2,000 peace agreements signed since 1990, which equates to over 12.5 thousand pages of textual provisions that have been formally agreed by conflict parties (roughly 1.8 million words). The UN Digital Library has over one million records of documents, publications, votes and speeches that have been transcribed into text, in addition to UN Security Council documents such as resolutions, presidential statements, annual reports and Secretary General reports, mission reports (and many more), accessible in a range of languages. ReliefWeb hosts and regularly updates reports (~900,000 reports at time of writing) from a range of organisations which can be categorised by country, theme, date and content, which may hold crucial information embedded within the assessment or situation report texts. Human Rights Watch records human rights abuses in textual reports on their website. Additionally, news outlets and other media content providers report on events that are of great interest to the field. These are only a few examples of readily available texts and content to extract relevant data for peace and conflict studies.

How can NER be Utilised in Peace and Conflict Studies?

NER can be applied to this plethora of domain-specific (and non-domain specific) text for a range of applications. For instance, it can take unstructured text – any text from word documents, reports, or news articles from any of the above-mentioned sources (and many more) – and identify within the texts key entities that have been mentioned, extract them, and categorise them into different entity types. This can be used for a wide variety of purposes; for instance, NER models can extract place names and coordinates from text, to allow textual data to be mapped for geospatial analysis that can aid in visualising spatial patterns of violence and peace activities. These methods can be applied to identify organisations, locations and types of assistance in reports and news articles, which can provide the possibility of humanitarian aid tracking, or alternatively, to identify hotspots of conflict and track and monitor movement as part of security and risk assessments. For instance, 'liveuamap' uses web-crawlers to find news-worthy stories which are then fact checked by experts. These crawlers use NER in their pipelines to identify, extract and categorise named entities and thus relevant news stories to them.

Additionally, NER can be used to extract events and code conflict events by identifying, extracting and categorising key entities such as names of conflict parties, locations where event(s) occurred, dates, casualties and other key and relevant information to record individual events with standardised properties, such as in the Uppsala Conflict Data Program (UCDP)[6] and the Armed Conflict Location Event Dataset (ACLED).[7] This type of data allows users to keep up to date on the number of events happening within a specific area, compare levels of conflict by actors, location, date, or event type, in addition to providing raw data that enables a wide variety of analysis that cannot be done directly on the original content (raw text/news articles/reports), but is a growing requirement for empirical and evidence-based research within the field. Alternatively, it can be used to analyse the role of the media itself by comparing how different entities involved in peace and conflicts are portrayed in the media, and gain insights into the role of the media in shaping public perceptions of conflicts. On more 'in the wild' data such as social media, NER can be utilised for monitoring to identify and track discussions of key entities, which can provide real time insights into public perceptions and reactions, in addition to humanitarian aid tracking and security and risk assessment.

3 Approaches to NER

This section will outline different approaches to NER; from simple, basic and traditional approaches, no- or low-code approaches to mixed methods, or modern high-tech methods using machine learning/deep-learning techniques. There is no 'best' approach for using NER across all use cases, with the effectiveness of each depending on the output requirements.

Rule-based NER

Commonly used in NER is the 'rule-based' approach, which follows pattern-based rules on how words are used together in speech and context-based rules on how a word is used within the text, and then determines if the entity should be recognised as an entity, and if so, what type. Certain entities can be discerned through distinct patterns, such as capitalisation, prefixes, suffixes, or specific character sequences. For instance, strings of capitalised words or terms preceded by "Mrs," "Dr," or "Mr" often indicate a person's name. Organisations frequently employ abbreviations or punctuation marks; thus, a criterion can be established to recognise words comprising multiple capitalised or non-capitalised terms separated by full stops as an 'organisation'. Contextual cues play a pivotal role in guiding rules for entity identification; for instance, locations may be identified by proximity to words like "in," "at," or "near," while temporal references often encompass words like "on," "at," or include month designations/names such as “April”. Sequences of rules can be amalgamated to efficiently capture intricate entities, such as recognising a person's name followed by a title and subsequently an organisation. Exploiting grammatical associations between words facilitates dependency parsing by the algorithm to ascertain entities based on their syntactical positions as subjects or objects within sentences.

Manually forming rules for the models can be a time-consuming process, but fortunately many existing models are standardised and available in 'out-of-the-box' libraries. Selecting a model to execute NER depends on the use case, as each works best in different contexts, or recognises different entity types. For example, the Natural Language Toolkit (NLTK)[8] has been longstanding in NLP since 2001 and has a vast array of libraries for research. Stanford's Stanza library has a wide range of languages included in its NER models, with up to 34 languages and dialects tested.[9] If a specific language model is required, such as Ukrainian, one of Stanza's models may work best out of those two options.

However, if running on English, their most reliable English model only recognises four different entity types: PER (Person), LOC (Location), ORG (Organisation) and MISC (Miscellaneous), which may not be sufficient for all use cases. NLTK can recognise eight different entity types; however, it only provides a basic NER functionality compared to newer and more powerful models. First launched in 2015, later than its counterparts listed above, spaCy[10] has become the go-to library for NLP due to the speed and efficiency of the models' performance. Additionally, there is the powerful unsupervised BERT (Bidirectional Encoder Representations from Transformers)[11] from Google, although it may require high levels of technical knowledge and ability to execute effectively and efficiently.

The key advantages of standardised libraries using a rule-based approach are a high level of precision, which is a key requirement for information extraction in the field, and there is no need for training data which is time-consuming and resource-heavy to manually annotate.[12] These models work particularly well if the user wants to extract things with a clear and structured pattern – for example, country names and dates are well recognised with rule-based models, but names of less-known organisations may not be identified. Additionally, a key benefit is the standardisation of results – if the model was executed on the same text by another user, it would return the same results. This is crucial to show reproducibility in recording the provenance of information and data.

However, there are limitations to this approach, as when entities are mentioned and recognised by the model in the text, they are standalone entities denoted exactly how they are mentioned in that piece of text, when there may be other differently abbreviated mentions of the same entity throughout the text, and additional information on the entity, that cannot be denoted under a single identifier from this method alone. This limitation may be important for some use cases; for instance, if one wants to find all mentions of the United Nations within the text, they are not just looking for mentions of the words 'United' + 'Nations' as a phrase, but also 'UN', 'UNSG, 'UNSC', 'UNICEF', 'UNHCR', 'UNMISS', 'DPPA', 'OHCHR' – and many other ways to refer to the UN Secretariat, missions, and agencies, funds and programmes).

Mixed Methods NER

The mixed methods approach can combat these limitations of a rule-based approach. For instance, a 'dictionary-based' approach is a simple method whereby string-matching algorithms check whether an entity is present in the vocabulary (denoted exactly as it is in the text), and if it is correct, it is labelled with the corresponding entity type. This approach addresses some key issues using out-of-the-box-models: alternate naming conventions and spellings that refer to the same actor can be included in the dictionary, information on the actor can be stored (for example, Actor ID in other datasets, country/location based), and new actor types rather than those basic categories provided by the model can be labelled depending on the actor recognised.

The dictionary-based method has a simple and effective fool-proof approach, by being flexible to different naming conventions, and providing the model with a list of entities that should be recognised. Nevertheless, these dictionaries are hard to keep up to date, and only provide string matching. Fortunately, there is already vast information readily available on relevant entities within particular contexts, and datasets that contain information to form these dictionaries. For example, the Integrated Crisis Early Warning System (ICEWS) Actor Dictionaries,[13] last updated in 2015, contains large amount of information on more than 10,000 individuals and group actors globally, including information on what the type of actor they are, the country in which they are based, alternate naming references, and their time-dependent affiliations, which remains an extremely useful resource for this approach, despite not having been updated in eight years. Other reliable datasets such as UCDP release information on actors they have recorded in their collections, with the UCDP Actor Dataset[14] containing information on the original name (both in English and the original language), whether there have been any name changes and if so, what; whether there have been any splinters in the group, and if so, the name and ID of the new group, and more useful information to enhance the knowledge base on conflict actors. Correlates of War, additionally, hosts information on actors such as Intergovernmental Organisations (IGOs),[15] which denotes the members and date of origin, in addition to previous names of the IGOs, the short names (acronyms) and long organisation names. These examples of existing datasets are not wholly comprehensive, and have not been recently updated; however, they provide good starting points for building dictionaries and Gazetteers on entities that are involved in peace and conflict resolution activities.

Machine Learning NER

Over the last decade there have been huge advances in NER machine learning based models, both supervised and un-supervised. Supervised machine learning is when algorithms are fed with features to return corresponding predetermined ('known') labels or outputs – for example, text is labelled with PERSON when a person is mentioned, ORGANISATION when an organisation is mentioned and so on. The model then learns from these examples to predict these labels on other examples.[16] A supervised NER model is trained on annotated documents, which denotes entities in the text with their corresponding entity types. This approach is particularly useful if the text is domain-specific and requires customisation of the entity types to be detected (in healthcare, for example, there is often a need to train their own models as terms for genomes and proteins are very healthcare-specific), or when there is a large amount of reliable training data available. This approach can be beneficial, as it opens the door for customisation of models to incorporate different entity types and recognise entities that the rule-based ones may ignore. However, the annotated training data needs to be large, diverse, and accurate enough for the model to learn from – which can be boring and time intensive for the annotator and will likely be then biased toward the training data if there is no 'gold standard' to learn from.

Rather than training on annotated text, there is the possibility of the Gazetteer-based approach, which does not entail learning specific rules for the system. Rather, it involves the development of a Gazetteer (a list/dictionary/dataset) of frequent and regular named entities that need to be recognised within the text and maximises the retrieval of them throughout the text (similar to the dictionary approach discussed above). Gazetteers can be especially useful when dealing with entities that may not follow typical linguistic patterns or when there is a need to capture specific and domain-specific information accurately. One notable application of this methodology in the field is exemplified by the PeaceTech Lab, which has employed Gazetteers and lexicons of hate speech across various languages to effectively identify instances of hate speech within conflict-affected regions.[17]

Machine learning approaches, whereby users train their own models based on their own annotated training data, can be highly useful and powerful as they can improve accuracy levels of the standard models, and can be customised for different entity types and adaptable to new types. Additionally, building a custom model means heightened privacy – if the data is sensitive or proprietary it does not need to be shared with any external annotators or third-party services and the user can take complete control over the annotations.

Useful tools can be used along with spaCy to conduct the annotations process to ensure they are in the correct format for training NER models. Prodigy is a scriptable annotation tool for users to make annotations on the text to quickly train models with a few examples. It can assist with a range of functions for improving, training, and correcting spaCy NER models to be more customisable to the project requirements. Users can accept or reject pre-tagged entities by the model, add new entities or note that there are no entities present in the text, which all can assist in improving the model's accuracy. Users can put their own annotations in the Prodigy format and train a custom model from scratch, with different entity types. This requires a few hundred annotations at minimum before running the first iteration of testing to see how accurate and reliable the training is, and things can be tweaked, corrected and new annotations added.

It is essential to bear in mind the possible implications of training your own custom models when deciding if this is the best approach for your project. For instance, if there is not enough training data in terms of quantity, the model performance may struggle with over-fitting as it may not generalise well to new and unseen examples. As large amounts of annotated data are required, which usually requires domain expertise and human annotators, it can be time consuming and expensive to hire these annotators. Additionally, there may be biases from the annotations – for instance, if there are any mistakes in the training data, they can be inherited by the model. Training on a smaller subset of data, rather than from the large "big-data" NER training sets that are available may limit the model's general knowledge and contextual language structures to learn from the textual and grammar structures of how entities are mentioned.

Unsupervised/Deep-Learning Models

Another machine learning approach is through unsupervised learning, whereby the system learns some rules/patterns from the examined document. The instances are not labelled, and patterns are found without the use of labels.[18] These rules and patterns are implemented to determine new named entities, and then a new set of rules are learned by the system continuously.[19] This use of rules allows for identification of names without manual annotation; however, it has slow training due to iterative re-labelling and works best with specific languages. The most complex approach within unsupervised learning is deep-learning, whereby multi-layered artificial neural networks learn tasks from processing raw input data. They go through many layers of nonlinear transformations in order to calculate a target output (i.e., they go through multiple layers of changes to figure out the correct answer).[20] Neurons in these networks receive inputs from other neurons, multiply them by certain values learned during training, and then combine them to give an answer, even for things they've never seen before.

This neural network approach is where the ‘black box’ issue surrounding use of artificial intelligence arises, as it is unknown how these individual neurons work together to produce this output. These neurons extract features from the inputs and map them on to the unseen data to produce outputs. However, it is unknown what these features are and how the neurons make these decisions – they may pick up features a human would not find applicable or relevant, or features a human eye could never see, and make nonsensical decisions based on this. As there are many neurons, they become interdependent on one another, and may compensate for deficiencies of other neurons. When applying deep-learning models to domains they were not trained on, there will likely be unexpected and undesirable behaviours due to extrapolation and the inability for generalisation. For example, if a neural network was trained on texts only surrounding a single conflict, or a single type of conflict, it may not generalise well to other conflicts, areas, regions, or actor types. It may also be particularly biased towards entity types they have learned and overemphasise or neglect other relevant entities to other conflicts.

Due to the need for reliable, accurate, and robust data to use as evidence in peace and conflict studies, the methods of producing the data must be explainable and understood by humans, which is not yet possible for this approach, as neurons learn individually and behave unexpectedly to generate the final output. Thus, the complexity of the model hides the logic of their internal processes. There is an inherent risk that relying on opaque models may lead to adopting decisions that we do not fully understand, or worse, violating ethical principles or legal norms – particularly in high-stakes decision-making scenarios such as in peace and conflict contexts.[21] Nevertheless, deep-learning has been beneficial for many domains such as multilingual models, on search engines to return relevant results, and to apply on messy text such as on social media where more informal language is preferred by users, making NER complicated due to the lack of grammar and sentence structure (thus, a rule-based approach may not be the best performing option on social media).

4 PeaceRep Use of NER & Lessons Learned

Where has PeaceRep used NER?

A running theme across all these use cases is actor identification in text, which can allow for identifying and classifying various actors involved in conflicts and peace negotiations, which greatly assists understanding the roles and interactions between actors in conflict affected contexts. Actor identification and peace agreement analysis are two major applications of NER that have been relevant to ongoing work in the PeaceRep consortium, as NER has allowed us to generate a new structured dataset of actors who signed peace agreements in the PA-X Peace Agreements Database [22] – the Peace Agreement Actor Dataset (PAA-X).[23] This is not to say that NER discovers new information within the text that it is performed on – actors who sign agreements as party or third-party signatories have already been recorded across all agreements in fields in the database, exactly how they were referred to in the original agreement text.[24] This is useful for an individual agreement as one can see exactly who signed it and on whose behalf, how these actors formally agreed to the terms, and how they framed their role in the process. However, this does not meet the requirements of all use cases of PA-X data. We are seeing increased fragmentation across conflicts,[25] which means an increased number of conflict actors, mediators, and agreements signed across systems of conflicts. Therefore, it is crucial to go beyond individual agreements to examine all the agreements and processes actors are involved in (as conflict parties, or mediating third-parties), to aggregate agreements that specific actors have signed (or not signed), and thus what they have legally committed to in a structured manner, so users can easily trace these commitments overtime and check if they been implemented or not. For example, if an armed organisation has signed any ceasefire agreements, the levels of conflict events they were involved in within the time period can be easily tracked. Or, if one wanted to investigate how a mediating organisation or Country/State has influenced the content of the agreement provisions, this can be analysed across all agreements that they have mediated/been a third-party signatory to.

NER has allowed us to identify, extract and categorise entities under a single unique identifier within the party and third-party text fields of all (non-local) agreements in the PA-X database. ‘Non-local’ here refers to all agreements in the PA-X database, apart from those labelled ‘IntraLocal’ as the agreement/conflict level (‘Agtp’), which includes agreements that aim to resolve local issues, rather than what is perceived as a conflict-wide issue.

These agreements have been initially excluded from the first version of the signatory's dataset as they often have localised actors and actor types that may not be as well picked up by NER models, and the actor types may not be consistent across different contexts (for example, local chiefs/paramount chiefs/local NGOs/local government). Therefore, the use of NER on the remaining agreement types (interstate, intrastate and mixed agreements) has led to the creation of a dataset containing over 7,000 instances of signatories, across more than 1,600 peace agreements.

This structured dataset can create avenues for conducting a wide range of research and peace agreement analysis, in addition to new ways to further interrogate the data and information through data visualisations such as actor network graphs. This has been a fruitful and valuable approach to the problem and only made possible by NER, as manually identifying, extracting and categorising each signatory to all agreements in the database would be a lengthy and gruelling task. Categorising entities with an actor type and a unique identifier allows for different naming conventions, languages and acronyms that refer to the same entity to be linked to the unique ID. For instance, as described above, the UN can be referred to as the United Nations, Les Nations Unies, the UN, the U.N., ONU, as the institution as a whole, but it also contains multiple offices, and missions which can all be referred to in different manners, but all actually refer to the United Nations as an entity in some way (UNCHR, UNICEF, UNMISS). The same applies for a majority of actors who are signatories to agreements – armed organisations and political parties often have acronyms, and a range of representatives who are signing on behalf of these organisations. Even Countries/States are often referred to in a range of ways. For example, the UK can be referred to as the United Kingdom, the U.K., Great Britain and Northern Ireland, GB, or even by people's names who are technically representing the country, such as Tony Blair: ex-Prime Minister of UK, who would be signing on behalf of the UK (Country/State), and not as Tony Blair (Person).

This data collection allows the user to see who signs with each other, what other agreements they have signed together, who mediated, or who has not signed anything together, amongst many other use cases. This approach provides a flexible and clear way of determining the commitments that actors have set out across agreements, in addition to possible entry points for new dialogue based on who has negotiated together, and what has resulted in formal agreements in the past.

How has PeaceRep used NER?

To approach the problem of extracting individual entities who were party/third-party signatories to formal agreements in a structured manner, we first utilised spaCy's small English model (en_core_sm). This was determined the most effective and efficient approach for a range of reasons: all the Party and Third-Party fields in PA-X have been entered in English; we do not have enough structured training data (or resources to annotate data to train a model from scratch); and many signatories are countries or large organisations, which are well recognised by the standard model. SpaCy provides an interface for users to test how their NER models work on text excerpts online, with no programming/coding required on the Displacy demonstrator website. For example, Figures 1 and 2 show the parties to two separate agreements: Figure 1 illustrates the parties to the Bonn Conference (agreement ID in PA-X), which was signed by the international community, and the model illustrates its great ability to identify categorisations are not wholly correct. For example, the Swiss Confederation, the official name for Switzerland, has been recognised as an ORG (organisation), rather than a GPE (geopolitical entity), and the Aga Khan Development Network has been recognised as a PERSON, rather than an ORG. This may be due to a person's name appearing as part of the named entity.

In Figure 2 the model was tested using the parties to the Revitalised Agreement on the Resolution of the Conflict in the Republic of South Sudan (R-ARCSS) (agreement ID 2112 in PA-X), which are not international actors/countries but parties to the conflict in South Sudan. We can see from the results that the model works well at identifying entities, but it is not perfect – neither the SPLM/SPLA-IO nor the SPLM-Former Detainees have been recognised as an entity; however, the individuals who were representing them have been. The Opposition Alliance for Other Political Parties has been recognised as an organisation, but its acronym OPP has not. In this instance, that works well as they have already been recognised and counted. However, in other instances we may want the model to recognise OPP when mentioned, but it fails to do so. This example illustrates the usefulness of the model to extract and categorise entities but reinforces the need for human corrections on this extracted information to ensure entities are not missed, or that the entity types are correctly categorised.

To generate the dataset of peace agreement signatories, we required more granular entity types that are more customised to how they are perceived in the peace and conflict sphere, rather than those provided with spaCy. This was due to many actors being recognised as 'ORG' which is not sufficient for the purposes of the data collection, as the United Nations, the Catholic Church and the Taliban were all recognised as an ORG, yet each actor plays a significantly different role in terms of the signing of a peace agreement. This data collection should better reflect these roles in the actor type, as more granularity into types of organisations and actors are required for effective peace and conflict research. Therefore, we decided to use a mixed methods approach. This consisted of taking all national agreements in PA-X (any non-local agreement) party and third-party fields, and running spaCy's NER model to generate an initial set of recognised, extracted and categorised entities per spaCy's model (like those shown in Figure 1 and 2). As mentioned previously, there are a range of appropriate resources that contain information on actors who may be signing peace agreements (ICEWS, UCDP, CoW), including different naming conventions, locations, affiliations and actor types. We gathered the unique actors from these sources, and created dictionaries of how they are referred to in these datasets. These dictionaries therefore included a unique actor in each row, which contained the type of actor they were (actor type definitions are outlined in Table 1 below), the International Organisation for Standardisation (ISO) code – a unique, standardised three-character code of the country in which they are based, the actor ID in other relevant datasets (such as UCDP and ACLED), and alternative names for that actor (many of these were added manually based on how they were mentioned in PA-X).

We then matched all the entities extracted by the spaCy model with these alternative names, and if there was a match, the signatory instance was labelled with the new actor ID (and thus, a consistent name and more granular actor type). These matched instances were then manually corrected to see if the entity's label was 'semantically correct', i.e. 'is it really what is being referred to' – this step was time consuming, but important. Peace agreements are legal documents, therefore we require 100% accuracy – we cannot claim that an actor has signed an agreement if they have not, as there could be serious repercussions and errors reduce the validity and accuracy of our dataset. In the correction process, we also denoted which entities were missed by the model, or those we did not have in existing dictionaries and therefore added these entities as new actors to the dictionaries, storing the relevant information on them. This exercise resulted in a cleaned dataset that contains entities that were party/third-party to national peace agreements in PA-X under a unique identifier, and an extremely valuable collection of actor dictionaries that provides key information on the actor, the actor type label and different naming conventions.

Figure 1 [below]: Parties to the Bonn Conference, Afghanistan signed on 05 December 2011) as recorded in PA-X (848)

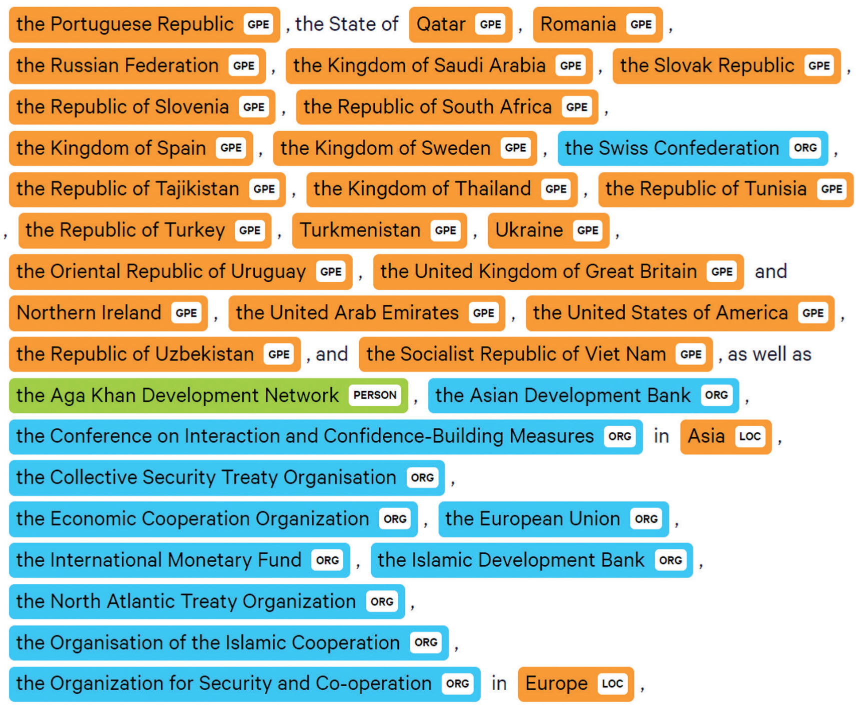

Figure 2 [below]: Parties to the Revitalised Agreement on the Resolution of the Conflict in the Republic of South Sudan (R-ARCSS), signed on 12 September 2018, as recorded in PA-X (2112)

H.E. Salva Kiir Mayardit PERSON

President of the Republic of South Sudan GPE for the Incumbent TGoNU

H.E. Dr Riek Machar Teny PERSON

Chairman and Commander in Chief of the SPLM/SPLA-IO

Hon. Deng Alor Kuol PERSON

For SPLM-Former Detainees

Hon. Gabriel Changson Chang PERSON

For the South Sudanese Opposition Alliance For Other Political Parties ORG (OPP) of South Sudan LOC

Table 1 [below]: Actor Type definitions for the Peace Agreement Actor Dataset (PAA-X).

Actor Type Label	Definition
Country/State	Internationally recognised Country/State.[26] Refers to any instances where it is the country name, or the government of the state representing the country. E.g. Leaders of the country/state or ministers that represent the country.
Entity	Representatives of entities which are not an internationally recognised independent country/state. Includes areas with representatives that are not internationally recognised independently of their country/state. E.g. Bougainville, Abkhazia, Puntland, Donetsk People's Republic, Northern Ireland.
Military	In cases where the state army/military was distinguished from the country/state name, they are labelled with the country/state army/military name. Important in cases where the army signs in addition to the state. Particularly important in cases where army acts as its own entity, separate to the state.
IGO	Intergovernmental Organisations. Actors are labelled with this in instances where they appear in Correlates of War (IGO v3),[27] ICEWS Actor Dictionary[28] or the International Governmental Organization Data (v2.1).[29] The shared definition of IGO across these datasets, is international organisations that have at least three nation-states as their members.

Actor Type Label	Definition
NGO	Non-governmental Organisation. An organisation formed independently from any governments. Includes large international NGOs, in addition to regional or local NGOs. Also includes civil society, academia, religious organisations, or other groups such as trade unions.
Armed Organisation	Labelled as armed organisation if appears in: UCDP, ACLED, 'Rebel Group Abbreviations' [30] or is a known armed organisation.
Political Party	If the actor has had a candidate contest in an election for government, or hold a seat in government, on behalf of the shared ideologies of the political group, they are provided with this label. If they appear in PoliticalpartyDB,[31] Global Party Survey [32] they are also labelled as Political Party.
State Coalition	Where country/states act through a coalition, usually designated for a specific purpose or scenario. Intergovernmental groups/coalitions that do not fit the typical 'official' IGO definitions. Includes 'Group of Friends of UNSG', 'Troika Countries', 'Group of 8'.
Umbrella	Umbrella groups/alliances of non-state groups who formally act/coordinate through the umbrella group as a unified front. Can be political or military alliance. E.g. CGSB – Simón Bolívar Guerrilla Coordinating Board – an umbrella group of six armed groups in Colombia.
Other	An actor that does not fit into any of the above definitions of actor types. Can include committees, centres, organisations, councils, commissions.

Can PeaceRep's Approach to NER be Re-used?

In line with FAIR data principles (findability, accessibility, interoperability and reusability),[33] this dataset provides ways to connect data from PA-X to other datasets that can be utilised for peace and conflict related research, such as UCDP and ACLED. Additionally, it is a data collection of actors that can be reused for readers to conduct similar methods on other text sources. The data can be used to train a model for these different entity types rather than the standard types used in models, can be incorporated alongside the large, well performing models to also include lesser recognised entities, or use the Gazetteer approach to look for the *known* actors contained within the dictionaries.

To provide a guide on how to conduct similar approaches, and how the vast actor dictionaries that we have collected could be re-used within the peace and conflict studies sphere, two Jupyter Notebooks are released alongside this report.[34] These interactive Python notebooks illustrate how to use spaCy's standard en_core_sm English model to run NER from text. One notebook illustrates how to read in text from a CSV file (e.g. if running on PA-X corpus it is available as a CSV file, or other relevant resources that may hold relevant text data may be stored in CSV format), or a Word document (as this is a key medium for storing and disseminating text in the peace and conflict sphere).

These notebooks have implemented the following methods:

- Importing and loading the NER model to conduct NER (in these examples: spaCy en_core_sm);
- Loading the text to perform the NER on (either from a Word document or a CSV file – in this instance, the CSV export of the PA-X database;
- Processing these text data appropriately (i.e. segmenting the sentences, dropping stop words such as 'and' or 'the', or removing escape characters that may be present in html code such as \t \n from the text);
- Selecting the entity types for the model to recognise;
- Viewing the results from the NER model and exporting them;

Additionally, they provide the option of using our actor data collection to re-label entities that have been recognised with their unique ID within our dictionaries. This allows for further interoperability as the recognised actors in your text can be aggregated under their unique ID, with a more relevant actor type category than what is provided by the spaCy model, in addition to being linked with PA-X, UCDP, ACLED and other relevant data collections. When using this mechanism to label actors, it is essential that corrections to the recognised entities are done manually to ensure they are labelled with what they are truly referring to. This can be a fairly quick process if the amount of text is not too large, as having a structured table with the information clearly shows what has been recognised and what it has been labelled with, so it is relatively easy to recognise which ones are incorrect. However, if thousands of entities have been returned, this process may take longer, but is nevertheless an important step to ensure valid and accurate data, due to the limitations outlined in the section below.

5 Limitations and Implications of Use of NER

Although there are a range of methods to extract named entities from text and categorise them by a range of entity types, it is crucial to keep in mind some limitations of the use of these technologies, particularly when extracting data for peace and conflict studies.
The results/outputs of these models are extremely useful as they provide structured information in a timeframe that could never be realised by human labour; however, the models do not wholly understand the semantics behind the mention of a named entity, and what the reference to them within the context of the text means. Therefore, it is crucial to always have a 'human in the loop' to run corrections on the final outputs to ensure the semantics are clear on the purpose of extraction. For instance, when trying to extract signatories from agreements, sometimes only parts of entities are recognised, or they are recognised as multiple entities. For example, the United Kingdom and Northern Ireland were often recognised as two (or three) distinct entities when referred to as "the United Kingdom of Great Britain and Northern Ireland", when in reality, only one entity is being referred to when used this way. This type of response can lead to over-counting of mentions of individual entities if they are not checked. These corrections are a crucial step, as sometimes Northern Ireland may be mentioned separate to the UK as a distinct entity, so they cannot be automatically dropped, but means that Northern Ireland also needs to be understood as a separate entity (particularly in agreements in the Northern Ireland Peace Process).

Another frequent issue is that sometimes the semantics of the sentence is not wholly understood by the model for the purposes of the NER application, reiterating the need for a human in the loop for corrections. For example, when extracting the signatories of peace agreements, the following sentence: "Afghan refugees residing in Iraq and Pakistan" was stated as a party to a peace agreement. The model correctly denoted Afghanistan [GPE], Iraq [GPE] and Pakistan [GPE] as being mentioned in the text, which our dictionaries then correctly relabelled with the unique IDs for Afghanistan [Country/State], Iraq [Country/State] and Pakistan [Country/State] which would denote the parties to the agreement were representatives from these three governments. However, this is incorrect, as the text is referring to Afghan refugees as a group of people, and not actually referencing these three Country/States as actors themselves. Thus, it is evident that we cannot rely on simple mentions of entities within a text to denote what it means semantically. If mistakes like these were not recognised, it could lead to incorrectly claiming that three Countries/States were involved in signing a legal document.

A limitation of NER arises when ambiguous entity references lead to challenges in disambiguation. For example, the term "Council" could refer to various entities such as a governmental body, an organisation, or a geographical location. Resolving such ambiguities requires human interpretation, reiterating the need for a human in the loop. One we frequently ran into in the signatories dataset was references to the agreement being signed by 'the government' – this term was picked up by the model, but it could not tell which government was being referenced as this detail would be in the metadata of the agreement, which required human input. In the development of our dataset, we made a deliberate decision to exclusively capture affiliations rather than individuals' names. This strategic choice was motivated by our dataset's primary objective, which centres on understanding and analysing the affiliations involved in peace agreements. For instance, if Bill Clinton signed an agreement whilst in the US presidency, he is signing on behalf of the US [Country/State], rather than Bill Clinton [Individual]. By excluding individuals' names, we aimed to prioritise the privacy and ethical considerations associated with recording personal information. This approach ensures that our research remains aligned with ethical standards and safeguards against potential concerns related to privacy and confidentiality, whilst also getting the balance of recording and structuring key data for the purposes of its intended use. As regards PeaceTech, this aligns with the idea of 'safety by design'.

Contextual sensitivity is another possible limitation of using NER for peace and conflict studies that reiterates the need for a human in the loop. NER models may struggle with nuances, for instance if “armed forces” is mentioned, the model would struggle in differentiating if it is referring to a state military entity, or a non-state armed group. Temporal sensitivity should also be considered – for instance, references to the same entities may change over time (even a Country/State, for example Turkey – Türkiye), and therefore continuous human review and considerations of different approaches (such as a dictionary approach/mixed methods) should be considered. As peace and conflict resolution-related texts are often written in a mix of languages, it is necessary to keep in mind that models trained on one language may not perform as effectively in multi-lingual contexts, and linguistic expertise may be required for effective translation.

NER models may also exhibit bias in recognising entities, potentially reflecting the biases present in training data that could impact the accuracy of identifying and categorizing entities, especially in conflict situations where diverse perspectives and actors are involved. If a NER model is trained predominantly on sources biased toward one side of a conflict, it may overemphasise entities associated with that perspective, leading to an unbalanced representation in the identified entities. Additionally, if the training data lacks diversity or is not reflective of the specific context it is being applied to, the performance of the model may be compromised. For example, if the model is only trained on historical peace agreements with limited diversity in terms of conflict types or regions, it may not effectively recognise entities in contemporary or region-specific agreements, highlighting the importance of diverse and representative training data.

6 Conclusion

Named Entity Recognition (NER) emerges as a powerful tool for swiftly identifying, extracting, and categorizing entities from unstructured text—a significant asset in the realm of peace and conflict studies. Its successful application, as witnessed in essential data repositories such as the PA-X Peace Agreements Dataset, exemplifies its pivotal role in generating valuable datasets like the Peace Agreement Actor Dataset (PAA-X).[35] Yet, while NER holds promise in addressing the pressing need to track myriad actors involved in conflicts and resolutions, its potential across the broader research community remains underutilised. Escalating global fragmentation and evolving conflict landscapes underscore the urgency of understanding mentions and references to various entities, and diverse entity types. The abundance of unstructured textual data, containing vital information in sensitive contexts, demands structuring—something NER adeptly accomplishes. Extracting key entities not only transforms textual documents into valuable datasets but also enables aggregation and comparison across diverse sources. However, it is essential to recognise that NER, while a crucial solution, does not stand alone in this demanding field. Human intervention remains indispensable for corrections in the face of any inaccuracies. Furthermore, the development of custom models necessitates robust annotated data; otherwise, the model risks overfitting and introducing biases due to inadequate diversity or size of training data.

Finally, the purpose of applying NER should be clear—it is not solely about achieving 100% accuracy, but rather expediting processes crucial to the research community. Speed and efficiency are paramount in handling the monumental volumes of data essential for comprehensive analysis. NER, with its potential and limitations, stands as an indispensable tool in structuring and extracting valuable insights from the vast sea of unstructured text data in peace and conflict studies, provided it operates alongside complementary methodologies, has human oversight, and follows safety by design and FAIR principles.

7 Recommendations

Based on this deep dive and the lessons learned from utilising NER across PeaceRep's work, we offer brief recommendations for using NER. These can also be applicable when using other digital methods, machine learning techniques and natural language processing in peace and conflict studies.

1. **Ensure Human Oversight and Contextual Understanding:** Maintain human involvement to provide context and correct the model's outputs. Do not solely rely on unchecked data, as the nuances and contextual meanings of text can be misinterpreted in isolation.

2. **Prioritise Explainable AI:** Always prioritise transparent AI models over black box ones. In high-stakes environments like peace and conflict resolution, it is crucial to understand how the model reaches its conclusions. Unexplainable conclusions should not inform decisions.

3. **Prioritise Functionality over Power:** Choose functionality over perceived power. Opting for high-performance technical models that demand significant computing power might not always suit your project's needs. Select an approach that aligns with your project's requirements, research focus and available resources.

4. **Consider Partial Automation with Mixed Methods:** When implementing machine learning models to automate processes, not every step needs automation. Embrace mixed methods approaches, which have shown significant success.

5. **Allocate Sufficient Resources when Training Custom Models:** If training your own model, ensure sufficient resources—large and diverse training data annotated by multiple individuals—to mitigate biases and errors introduced during training.

6. **Consider Speed vs. Process Efficiency:** While NER expedites entity recognition, the entire process might not always be swift. Checking, correcting, and occasionally rerunning the large volume of results takes time but remains faster than manual entity recognition.

7. **Adhere to FAIR Data Principles:** Maintain adherence to FAIR data principles when creating structured datasets. Ensure findability, accessibility, interoperability, and reusability to benefit the wider conflict resolution community.

8. **Collaborate, Learn and Share:** Interdisciplinary collaboration and knowledge sharing is paramount. Embracing diverse methodologies openly fosters a rich learning environment, facilitates collaboration, and enables the sharing of invaluable lessons learned.

Endnotes

[1] 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://www.peaceagreements.org/

[2] Li, X., Wang, T., Pang, Y., Han, J., & Shi, J. (2022, July). Review of Research on Named Entity Recognition. *In International Conference on Artificial Intelligence and Security* (pp. 256-267). Cham: Springer International Publishing.

[3] Wen, Y., Fan, C., Chen, G., Chen, X., & Chen, M. (2020). A survey on named entity recognition. *In Communications, Signal Processing, and Systems: Proceedings of the 8th International Conference on Communications, Signal Processing, and Systems 8th* (pp. 1803-1810). Springer Singapore.

[4] Gerner, D. J., Schrodt, P. A., Yilmaz, O., & Abu-Jabr, R. (2002, August). The creation of CAMEO (Conflict and Mediation Event Observations): An event data framework for a post cold war world. *In annual meeting of the American Political Science Association* (Vol. 29).

[5] 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://www.peaceagreements.org/

[6] Uppsala Conflict Data Program (UCDP) GED: Davies, Shawn, Therese Pettersson & Magnus Öberg (2023). Organized violence 1989-2022 and the return of conflicts between states?. Journal of Peace Research 60(4). View here: https://ucdp.uu.se/

[7] Armed Conflict Location Event Dataset (ACLED): Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). “Introducing ACLED: An armed conflict location and event dataset: Special data feature”. Journal of Peace Research, 47(5), 651-660. https://doi.org/10.1177/0022343310378914. D. View here: https://acleddata.com/

[8] Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python: analyzing text with the natural language toolkit*. "Reilly Media, Inc."

[9] See Stanza NER Models Language Availability here: https://stanfordnlp.github.io/stanza/ner_models.html

[10] Honnibal, M., & Montani, I. (2017). *spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing*. See spaCy's NER model docs here: https://spacy.io/api/entityrecognizer

[11] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

[12] ligüzel, N., Çetinkaya, C. & Dereli, T (2022). Application of named entity recognition on tweets during earthquake disaster: a deep learning-based approach. *Soft Computing* 26, 395–421 (2022). https://doi.org/10.1007/s00500-021-06370-4

[13] Boschee, Elizabeth; Lautenschlager, Jennifer; Shellman, Steve; Shilliday, Andrew, 2015, "ICEWS Dictionaries", https://doi.org/10.7910/DVN/28118, Harvard Dataverse, V4

[14] Pettersson, Therese (2023) UCDP Actor Dataset Codebook v 23.1 (https://ucdp.uu.se/downloads/).

[15] Pevehouse, Jon C.W., Timothy Nordstron, Roseanne W McManus, Anne Spencer Jamison, "Tracking Organizations in the World: The Correlates of War IGO Version 3.0 datasets", *Journal of Peace Research*.

[16] Libbrecht, M. W., & Noble, W. S. (2015). Machine learning applications in genetics and genomics. *Nature Reviews Genetics,* 16(6), 321-332.

[17] Hirblinger, A. T., Hansen, J. M., Hoelscher, K., Kolås, Å., Lidén, K., & Martins, B. O. (2023). Digital peacebuilding: A framework for critical–reflexive engagement. *International Studies Perspectives,* 24(3), 265-284.

[18] Kotsiantis, S. B. (2007). SUPERVISED MACHINE LEARNING: A REVIW OF CLASSIFICATION TECHNIQUES.

[19] ligüzel, N., Çetinkaya, C. & Dereli, T. Application of named entity recognition on tweets during earthquake disaster: a deep learning-based approach. *Soft Comput* 26, 395–421 (2022). https://doi.org/10.1007/s00500-021-06370-4

[20] Thomas, A., Sangeetha, S. (2020). Deep Learning Architectures for Named Entity Recognition: A Survey. In: Pati, B., Panigrahi, C., Buyya, R., Li, KC. (eds) Advanced Computing and Intelligent Engineering. Advances in Intelligent Systems and Computing, vol 1082. Springer, Singapore. https://doi.org/10.1007/978-981-15-1081-6_18

[21] Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., & Rinzivillo, S. (2023). Benchmarking and survey of explanation methods for black box models. *Data Mining and Knowledge Discovery,* 1-60.

[22] 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://www.peaceagreements.org/

[23] Badanjak, S. & Henry, N (2023). The Peace Agreement Actor Dataset (PAA-X). Due to be released in 2024.

[24] Under 'Party' ['Part'] and Third Party ['ThrdPart'] variables in PA-X Database.

[25] Bell, C., & Wise, L. (2022). The Spaces of Local Agreements: Towards a New Imaginary of the Peace Process, *Journal of Intervention and Statebuilding,* 16:5, 563-583, DOI: 10.1080/17502977.2022.2156111

[26] Gleditsch, Kristian S. and Michael D. Ward. 1999. "Interstate System Membership: A Revised List of the Independent States since 1816." *International Interactions* 25

[27] Pevehouse, Jon C.W., Timothy Nordstron, Roseanne W McManus, Anne Spencer Jamison, "Tracking Organizations in the World: The Correlates of War IGO Version 3.0 datasets", Journal of Peace Research.

[28] Boschee, Elizabeth; Lautenschlager, Jennifer; Shellman, Steve; Shilliday, Andrew, 2015, "ICEWS Dictionaries", https://doi.org/10.7910/DVN/28118, Harvard Dataverse, V4

[29] Correlates of War Project, 2018, "International Governmental Organization (IGO) Data (v2.1)", https://doi.org/10.7910/DVN/POLS26, Harvard Dataverse, V1

[30] Jo, H. (2015). Rebel Group Abbreviations. In Compliant Rebels: Rebel Groups and International Law in World Politics (Problems of International Politics, pp. Xiv-Xxi). Cambridge: Cambridge University Press.

[31] Scarrow, Susan; Webb, Paul D.; Poguntke, Thomas, 2022, "Political Party Database Round 2 v4 (first public version)", https://doi.org/10.7910/DVN/0JVUM8, Harvard Dataverse, V1. Accessible from: https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/0JVUM8

[32] Norris, Pippa, 2020, "Global Party Survey, 2019", https://doi.org/10.7910/DVN/WMGTNS, Harvard Dataverse, V3, Acessible: https://dataverse.harvard.edu/dataverse/GlobalPartySurvey

[33] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data,* 3(1), 1-9.

[34] These jupyter notebooks, and relevant actor data can be accessed from PeaceRep's GitHub repository: https://github.com/peacerep/ner-peace-actors

[35] Badanjak, S. & Henry, N (2023). The Peace Agreement Actor Dataset (PAA-X). Due to be released in 2024.

References

Armed Conflict Location Event Dataset (ACLED): Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). "Introducing ACLED: An armed conflict location and event dataset: Special data feature". Journal of Peace Research, 47(5), 651-660. [https://acleddata.com/]

Badanjak, S. & Henry, N (2023). The Peace Agreement Actor Dataset (PAA-X). Due to be released in 2024.

Bell, C. (2024). PeaceTech: Digital Transformation to End Wars. Palgrave Macmillan. Available at: [https://link.springer.com/book/10.1007/978-3-031-38894-1]

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

Bell, C., & Wise, L. (2022). The Spaces of Local Agreements: Towards a New Imaginary of the Peace Process, Journal of Intervention and Statebuilding, 16:5, 563-583. [DOI: 10.1080/17502977.2022.2156111]

Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python: analyzing text with the natural language toolkit. "Reilly Media, Inc."

Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., & Rinzivillo, S. (2023). Benchmarking and survey of explanation methods for black box models. Data Mining and Knowledge Discovery, 1-60.

Boschee, Elizabeth; Lautenschlager, Jennifer; Shellman, Steve; Shilliday, Andrew, 2015, "ICEWS Dictionaries", [https://doi.org/10.7910/DVN/28118], Harvard Dataverse, V4

Correlates of War Project, 2018, "International Governmental Organization (IGO) Data (v2.1)", [https://doi.org/10.7910/DVN/POLS26], Harvard Dataverse, V1

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

Gerner, D. J., Schrodt, P. A., Yilmaz, O., & Abu-Jabr, R. (2002, August). The creation of CAMEO (Conflict and Mediation Event Observations): An event data framework for a post cold war world. In annual meeting of the American Political Science Association (Vol. 29).

Gleditsch, Kristian S. and Michael D. Ward. 1999. "Interstate System Membership: A Revised List of the Independent States since 1816." International Interactions 25.

Hirblinger, A. T., Hansen, J. M., Hoelscher, K., Kolås, Å., Lidén, K., & Martins, B. O. (2023). Digital peacebuilding: A framework for critical–reflexive engagement. International Studies Perspectives, 24(3), 265-284.

Honnibal, M., & Montani, I. (2017). spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. See spaCy's NER model docs here: [https://spacy.io/api/entityrecognizer]

Jo, H. (2015). Rebel Group Abbreviations. In Compliant Rebels: Rebel Groups and International Law in World Politics (Problems of International Politics, pp. Xiv-Xxi). Cambridge: Cambridge University Press.

Kotsiantis, S. B. (2007). Supervised Machine Learning: A Review of Classification Techniques. University of Peloponnese [Overview Paper].

Libbrecht, M. W., & Noble, W. S. (2015). Machine learning applications in genetics and genomics. Nature Reviews Genetics, 16(6), 321-332.

Li, X., Wang, T., Pang, Y., Han, J., & Shi, J. (2022, July). Review of Research on Named Entity Recognition. In International Conference on Artificial Intelligence and Security (pp. 256-267). Cham: Springer International Publishing.

ligüzel, N., Çetinkaya, C. & Dereli, T (2022). Application of named entity recognition on tweets during earthquake disaster: a deep learning-based approach. Soft Computing 26, 395–421 (2022). [https://doi.org/10.1007/s00500-021-06370-4]

ligüzel, N., Çetinkaya, C. & Dereli, T. Application of named entity recognition on tweets during earthquake disaster: a deep learning-based approach. Soft Comput 26, 395–421 (2022). [https://doi.org/10.1007/s00500-021-06370-4]

Norris, Pippa, 2020, "Global Party Survey, 2019", [https://doi.org/10.7910/DVN/WMGTNS], Harvard Dataverse, V3, Accessible: [https://dataverse.harvard.edu/dataverse/GlobalPartySurvey]

Pettersson, Therese (2023) UCDP Actor Dataset Codebook v 23.1 ([https://ucdp.uu.se/downloads/]).

Pevehouse, Jon C.W., Timothy Nordstron, Roseanne W McManus, Anne Spencer Jamison, "Tracking Organizations in the World: The Correlates of War IGO Version 3.0 datasets", Journal of Peace Research.

Scarrow, Susan; Webb, Paul D.; Poguntke, Thomas, 2022, "Political Party Database Round 2 v4 (first public version)", [https://doi.org/10.7910/DVN/0JVUM8], Harvard Dataverse, V1. Accessible from: [https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/0JVUM8]

Thomas, A., Sangeetha, S. (2020). Deep Learning Architectures for Named Entity Recognition: A Survey. In: Pati, B., Panigrahi, C., Buyya, R., Li, KC. (eds) Advanced Computing and Intelligent Engineering. Advances in Intelligent Systems and Computing, vol 1082. Springer, Singapore. [https://doi.org/10.1007/978-981-15-1081-6_18]

Uppsala Conflict Data Program (UCDP) GED: Davies, Shawn, Therese Pettersson & Magnus Öberg (2023). Organized violence 1989-2022 and the return of conflicts between states?. Journal of Peace Research 60(4). View here: [https://ucdp.uu.se/]

View Our
Key Findings

info@peacerep.org
PeaceRepResearch
@Peace_Rep_
PeaceRep

www.peacerep.org
www.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, Dialectiq, 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 the Foreign, Commonwealth and Development Office (FCDO), UK.

PeaceRep: The Peace and Conflict Resolution Evidence Platform
info@peacerep.org | https://peacerep.org | @Peace_Rep_

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