sagentia
DEFENCE

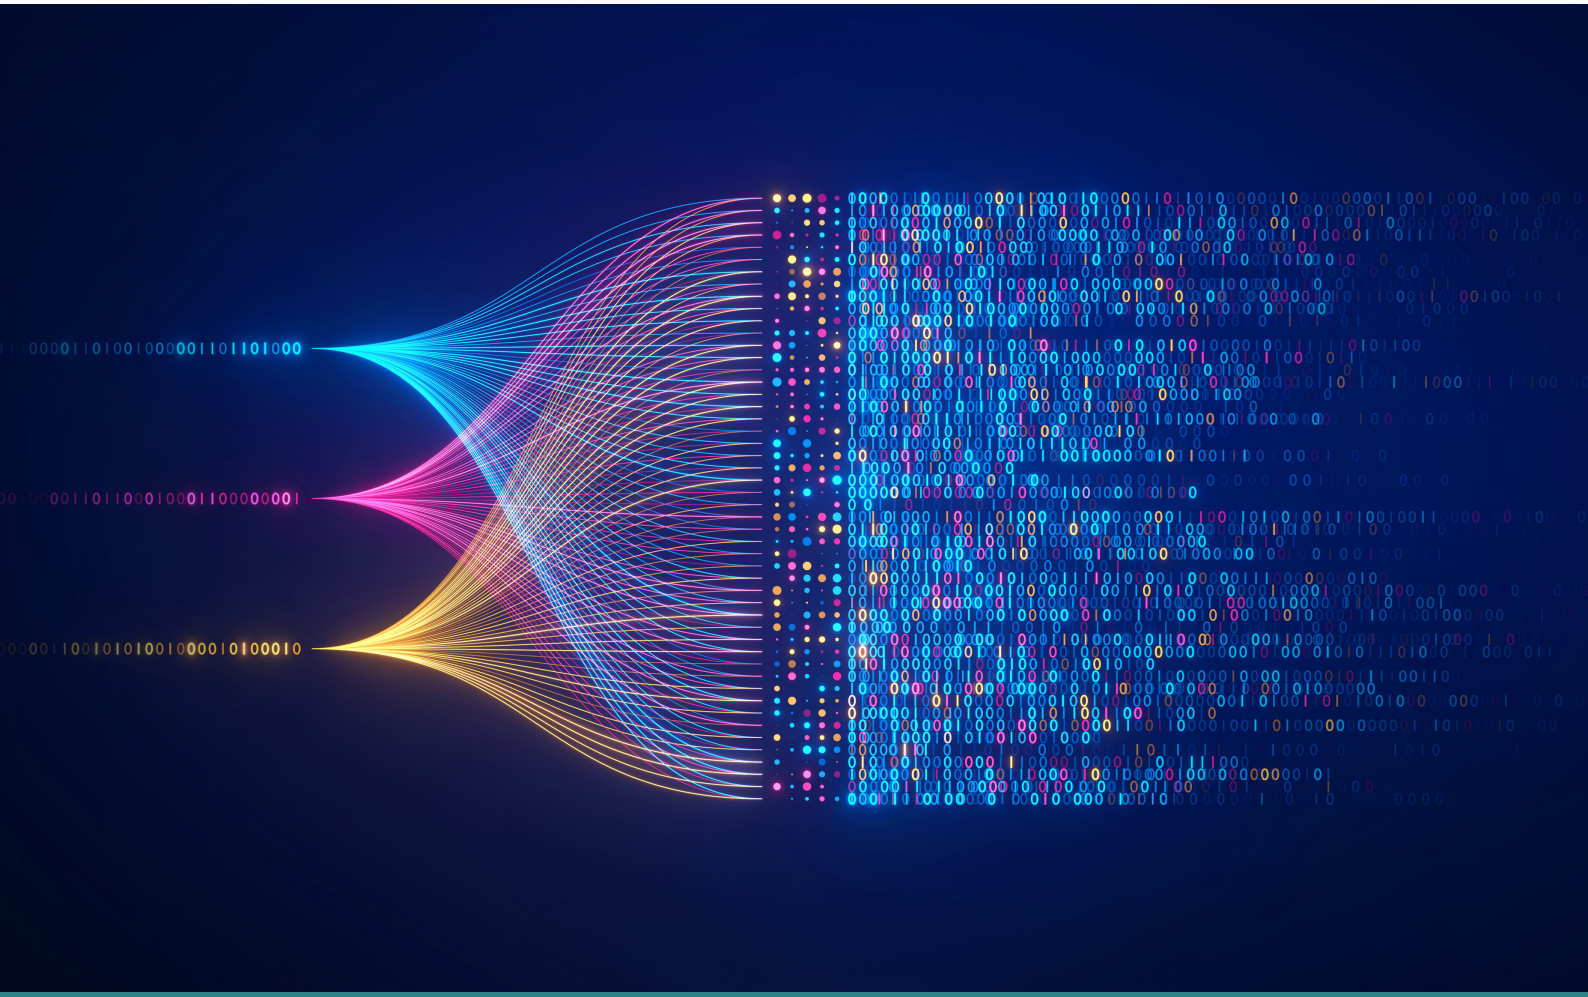

Harnessing the power of AI for Defence

sagentia.com/defence

Foreword

In the first of what will be a series of white papers from thought leaders at Sagentia Defence, formerly TP Group, I am pleased to invite you to explore our insights into social media sentiment analysis, specifically surrounding the Ukraine conflict. In this age of rapid communication, these analysis techniques have become an indispensable tool for understanding public opinion around the evolving dynamics of geopolitical conflicts.

This white paper serves as a concise yet comprehensive guide to the art and science of sentiment analysis in the context of this ongoing conflict and I hope you enjoy reading our thinking on this topic.

Dominic Černis

Chief Technical Officer

Sagentia Defence

Harnessing the power of AI for Defence

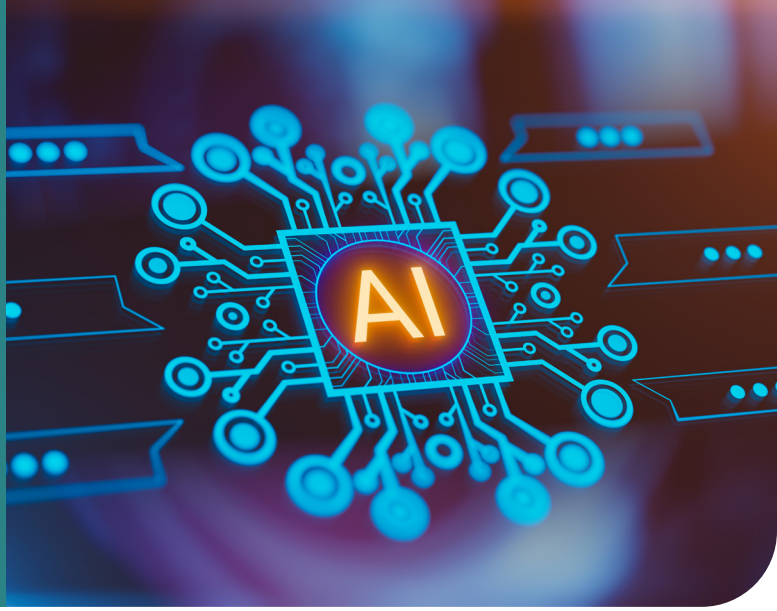

The UK Ministry of Defence (MOD) describes artificial intelligence (AI) as being among the most 'transformative, ubiquitous, and disruptive' new technologies in defence[1]. However, the sector is still learning how to harness AI at scale to deliver valuable insights for defensive and offensive purposes. Here, we look at how to unleash the technology's potential to measure public opinion around geopolitical events.

In the digital age, the ability to handle, interrogate, and leverage data is a core requirement for Defence. Yet eliciting value from vast, unstructured collections of data is a challenge. This has prompted significant MOD investment in advanced data science, with AI identified as a key 'generation after next' capability. Finding ways to derive useful insights from data using AI is a strategic priority for Defence, wider government, and partners in industry and academia.

Sagentia Defence, formally TP Group, carried out a study to evaluate natural language processing (NLP), a branch of AI that helps computers understand the way humans write and speak, as a tool for measuring public opinion. The outcome is a proof-of-concept that demonstrates AI's potential to inform defence sector decision making when public perception of geopolitical entities or events is a critical factor.

Using AI to measure public opinion

Our objective was to determine whether it is possible to understand and follow public opinion on a topic using open-source data and open-source AI tools. We applied NLP techniques which combine computational linguistics (rule-based modelling of human language) with statistical learning, machine learning, and deep learning models. Together, these technologies process human language to 'understand' its full meaning.

Initially we planned to use social media platform Twitter as the data source. However, for the purposes of the study we required historic data which cannot be scraped from Twitter. Instead, we opted for Reddit. With more than 1.6 billion monthly active users, the website contains social news and discussion boards. It also has subreddits focused on different topic areas allowing users to form communities where people can post and comment anonymously. Reddit's application programming interface (API) enables easy collection of large amounts of data, making it an ideal resource for NLP.

The topic we decided to focus on was the Russia-Ukraine war. Reddit's users were posting a large amount of text data across multiple subreddits directly related to this conflict.

It's important to note that Reddit's users are largely based in the Western world, and we fully expected a data bias in favour of Ukraine. Our intention was to use the data source as a testbed for the exploration and refinement of NLP techniques on open-source data rather than to reveal new insights.

A process for NLP analysis of Reddit data

The application of NLP techniques is far from straightforward. Challenges include data complexity, misspellings, words with different contextual meanings, deciphering irony / sarcasm, and varied forms of phrasing across dialects and languages. We sought to establish whether it was possible to overcome these difficulties using an NLP model built using Python. The overall process involved five key elements, illustrated in Figure 1, and described below.

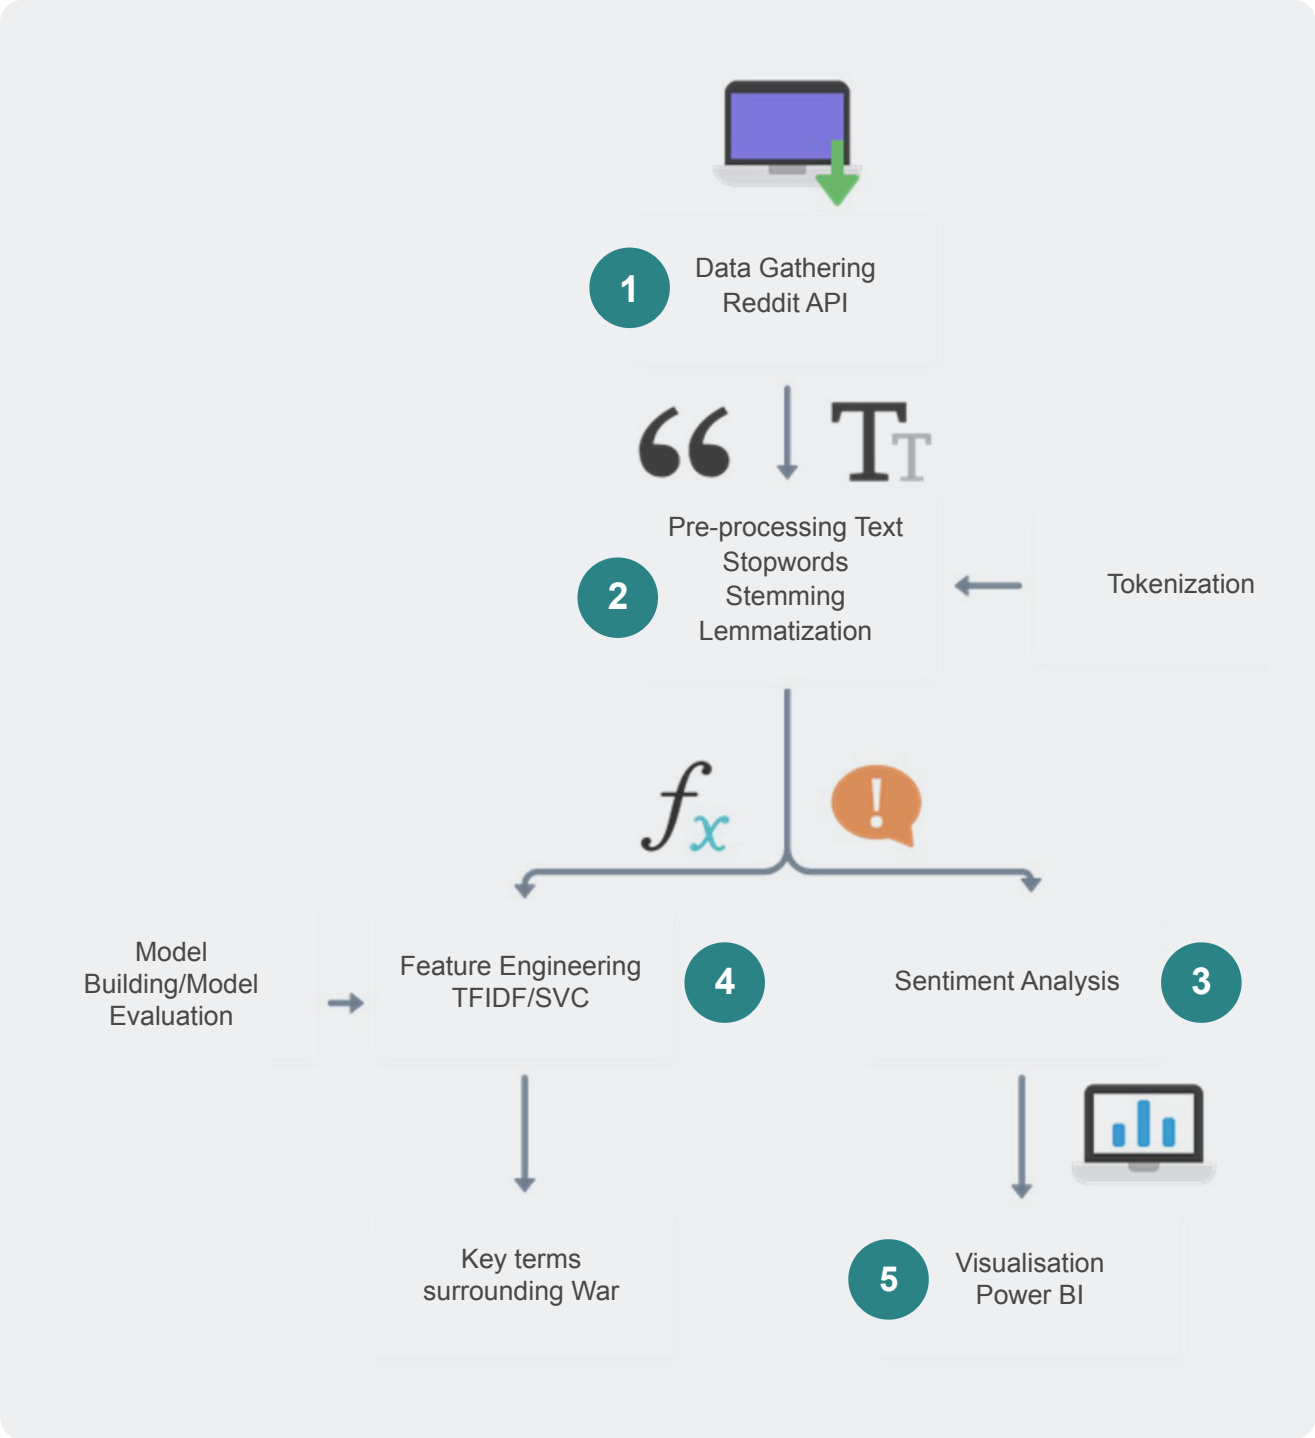

Figure 1: Applying NLP to data gathered from Reddit

1. Data gathering

Our goal was to evaluate Reddit users' opinion on the Russia-Ukraine conflict over a nine-month period. Firstly, we identified six relevant subreddits from which to gather data: 'RussiaUkraineWar2022', 'UkrainianConflict', 'UkraineConflict', 'UkraineWarReports', 'UkraineInvasionVideos', and 'UkraineWarVideoReport'. As Figure 2 indicates, the combined subreddits had 170,000 associated posts which had stimulated 123 million comments.

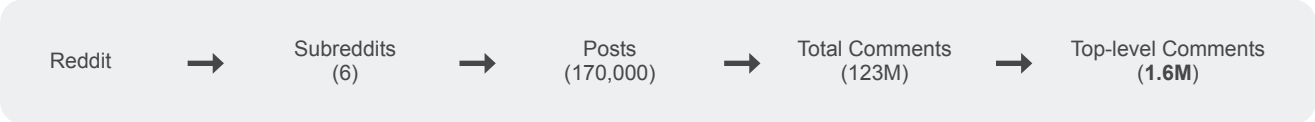

Figure 2: Identifying relevant comments for extraction

To avoid analysing data that wasn't relevant to the target topic, we focused on extracting the 1.6 million top-level comments. Reddit defines top-level comments as being posted directly onto the post, as opposed to comments that are replies to other comments, so we could say with some certainty that top-level comments would be on- topic.

2. Pre-processing and data cleansing

After the data was gathered, comments were converted into a machine readable and consistent format. This involved the removal of all uppercase letters and stopwords (commonly used words which don't add much meaning to a sentence and can therefore be ignored).

We also performed tokenization. In a machine learning context, this refers to the separation of pieces of text into smaller units or tokens. For this project, we split the top-level comments into sentences, thereby breaking the unstructured data into smaller chunks of information. Lemmatization was also used to help process the data. This technique returns words to their base dictionary form (e.g. reducing 'eats', 'eating', or 'ate' to the lemma 'eat') to simplify analysis.

3. Model build/evaluation: sentiment analysis

We used Python's Natural Language Toolkit (NLTK) library to perform sentiment analysis, categorising text as 'positive' or 'negative' using classification algorithms. This rules-based analyser provides sentiment scores based on the semantic orientation of words. Establishing the mean average for sentiment by day enabled us to see how it evolved during the nine-month period. It also gave us a quantitative measure for analysing the immense volume of qualitative data.

Adding flags to the data to signify the presence of key words of interest allowed us to measure the sentiment of conversation surrounding specific topics related to the conflict too.

4. Model build / evaluation: feature engineering

Another important aspect of the process was feature engineering, whereby meaningful information is extracted from raw data for use in machine learning models. Term frequency-inverse document frequency (TF-IDF) vectorisation was central here, enabling analysis of the written comments. Vectorisation converts text into numerical representation for machine learning applications, and TF-IDF is a popular vectorisation technique for NLP, used in applications such as online search engines. It's a useful formula to determine the importance or significance of specific words depending on their frequency of use.

We used Python's open-source data analysis library, scikit-learn, to perform TF-IDF, with the resultant insights training a Support Vector Classifier (SVC) model. SVC works by mapping data points to a high-dimensional space and then finding the optimal hyperplane that divides the data into two classes. This makes SVC model ideal for working with the high dimensional outputs of TF-IDF. Categorising data in this way streamlines analysis and enables a predictive approach. We used coefficients of this model to give us the most defining phrases of each subject area.

5. Visualisation and interpretation

We used Microsoft Power BI to present the ML model outputs in graph format. This enabled us to clearly visualise how the volume of top-level comments fluctuated during the nine-month period. We were also able to plot how the average sentiment for Russia, Ukraine, and other relevant terms such as NATO evolved over the timeframe (see Figure 3). Due to the fluctuation in sentiment between posts and comments, we grouped them by week to better show the change in sentiment over time. As expected, comments around Ukraine were consistently more positive on average, however, rises and falls in sentiment usually seemed to affect both Ukraine and Russia. Figure 3 also shows how the average sentiment of comments for both Russia and Ukraine was always negative, never reaching 0 or above, perhaps due to the overall negative nature of war.

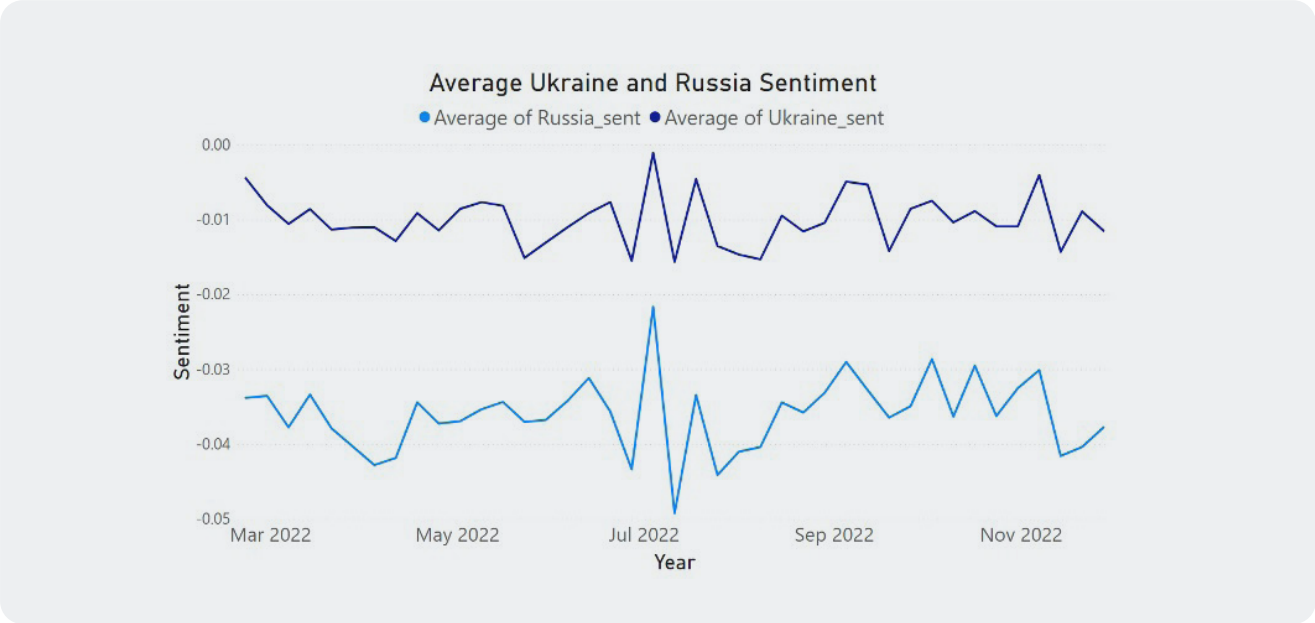

Figure 3: The model captured how public sentiment fluctuated throughout the nine-month timeframe.

To gain a better understanding of the drivers behind the peaks and troughs seen on the graph, we identified the mean sentiment across the dataset. We then isolated all posts that were outside 2 standard deviations of the mean (2SD). The week of 04/07/2022 was identified as a positive sentiment peak within the posts. Within this week, certain themes were highlighted. Furthermore, the week of 11/07/2022 was identified as a trough consisting of posts with negative sentiment.

Figure 4: Average sentiment by week over the timeframe, with lines showing the mean and 2 standard deviations.

The post containing the lowest negative sentiment within this trough was titled:

"A British man, who was in captivity, was tortured to death by terrorists in the self-proclaimed Donetsk Peoples Republic (DNR). Paul Urie was a volunteer. He helped victims of the Russian-Ukrainian war. In April, Paul was held hostage by terrorists from the DNR".

Whilst this post received the lowest sentiment score, there was a variety of themes present in the posts that consistently had low sentiment scores. A common theme within the posts containing negative sentiment was a Russian missile strike occurring in Vinnytsia. The attacks killed at least 28 people (including three children) and injured at least 202 others. Within these posts, there was a considerable amount of content was surrounding the death of a 4-year-old child due to the missile attack.

A common theme of the posts within the positive sentiment peak was content surrounding NATO countries. Much of this positive content consisted of posts about USA's Independence Day. This was closely followed by posts surrounding USA and UK aid to Ukraine. This was reflected by the post which received the most positive sentiment score, which was titled:

"The United States of America is one of the most faithful and loyal allies of Ukraine. Today, July 4, our friends celebrate Independence Day. We thank them for their support and help. We greet them on their holiday. Together we are stronger! God bless America!".

With Reddit's user based being dominantly western, with an especially large population of users in the USA, the significant rise in sentiment around Independence Day in the USA show how linked the sentiment is to the West/USA even in subreddits focused on Ukraine.

A recurring theme in the comments was concern about the global impacts of the war, and the potential for escalation (see Figure 4). The model indicated that three of the most defining phrases linked to the fear of escalation were 'nuclear war', 'nuclear weapon', and 'risk nuclear'. We were also able to identify what users perceived as key risk factors for escalation, with 'no-fly zone' and 'long range weapons' emerging as key terms.

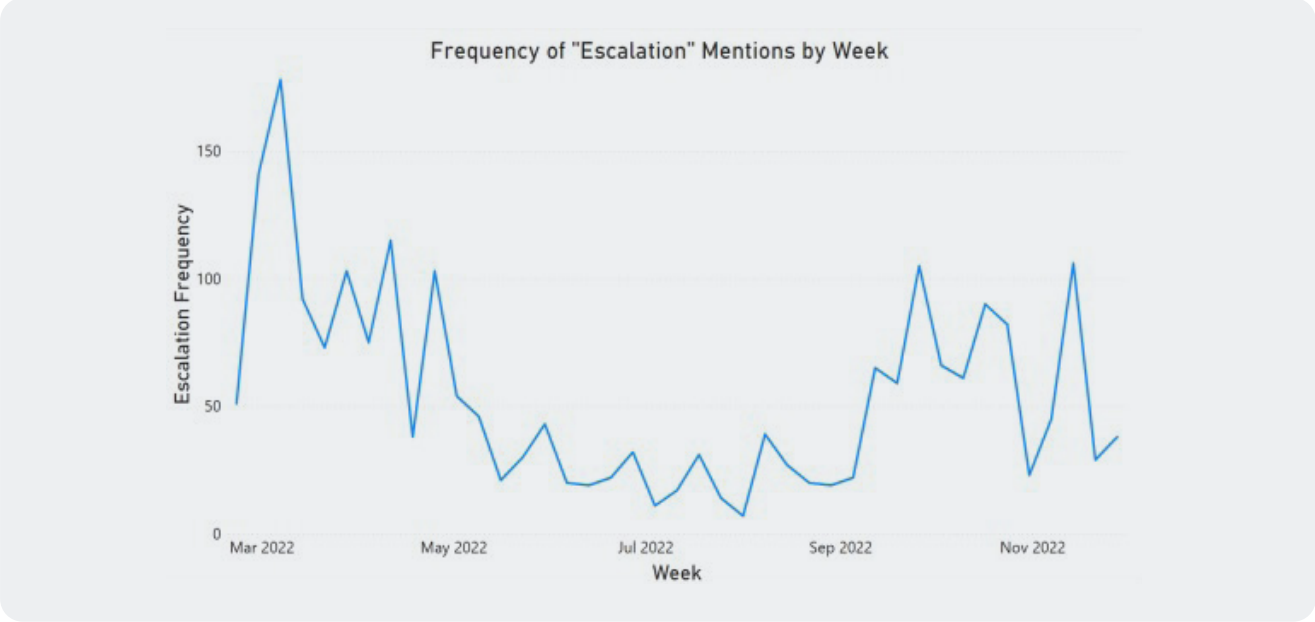

Figure 5: The model indicated how the public fear of escalation evolved.

While we expected Reddit users to have a bias in favour of Ukraine, it was interesting to see how this manifested in key terms. Those associated with Ukraine included 'freedom democracy' and 'people support' whereas conversations focused on Russia were characterised by terms including 'civilian target' and 'unarmed civilian'. Comments related to Poland seemed to focus on its potential to act as a flashpoint that could lead to escalation, with key terms such as 'trigger article 5' and 'start ww3'. However, many comments also reflected more positive sentiment with defining terms such as 'help refugee'.

Overcoming challenges

We encountered technical challenges and time-based constraints at various stages of the process due to the volume and messiness of the data.

Firstly, Reddit's free access API presented speed restrictions. We needed all 170,000 posts from the six subreddits, but we could only request 100 at once, with a two-second delay between each API call. Then we had to gather all 123 million comments before extracting the 1.6 million top-level comments in batches. This further increased the run time, and the entire data gathering process took around six days.

Next, we had to convert low quality text data containing slang, misspellings, emojis and other social media-related colloquialisms into a machine-readable language. This required multiple text processing methods and we were unable to filter out all the noise. Slang and emojis were the biggest challenge, and when other methods such as entity recognition failed, we hardcoded topics of conversation to overcome this.

The overall process was lengthy and there were many opportunities for error. Nevertheless, we demonstrated that expert application of NLP techniques can result in the successful evaluation of public opinion using open-source data and technologies.

AI's role in Defence

The AI techniques and processes outlined here can be applied to gauge public opinion on any topic. While we focused on historic data, there is scope to perform such analysis closer to real-time. It's also possible to monitor opinions on key events or political announcements over time to develop a deeper understanding of trends in the public response. With the increasing availability of large-scale public data on social media platforms, this offers great potential for defence scenarios where public opinion is a factor in decision making, setting policies, or communication priorities.

UK initiatives such as the Digital Strategy for Defence and the Defence Artificial Intelligence Centre (DAIC) indicate the significance being placed on advanced data science and AI. The DAIC was established in 2022 to 'champion, enable, and innovate' AI to deliver strategic advantage from the back office to the frontline[2]. Key goals include 'enhancing speed and efficiency' and 'increasing the quality of decision making'. Based the success of our evaluation of NLP for monitoring public opinion, it's easy to see how investment in this space will deliver value for the defence sector.

About the Authors

Dennis Glover is one of our senior data science consultants, with experience spanning defence, travel, and advertising. He has worked as the lead data scientist across multiple NLP and predictive modelling projects using a range of machine learning techniques.

Hollie Dawson joined the team as a junior data scientist in 2022. She has a strong academic background in biomedical sciences and brings an analytical mindset as well as knowledge of machine learning, data analysis, and visualisation.

About Sagentia Defence

Sagentia Defence, formerly TP Group, is a leading defence and security Consultancy, with a specific depth within Science and Technology. We draw upon our unique heritage to provide expert guidance and solutions, leveraging the insight from our experts to address specific client challenges. We work collaboratively across government and the defence industry to equip them to answer immediate, emerging and future challenges. We are relied upon to help navigate the landscape of accelerating next-generation technology, achieve greater efficiency across their operations, and to solve their most pressing issues.

Find out more about our expertise and experience within the defence sector at

sagentia.com/defence

If you are interested in discussing a project with us please email us:

info@sagentiadefence.com

Notes

1. Ministry of Defence Annual Report and Accounts 2021-22
2. Defence Artificial Intelligence Centre https://www.gov.uk/government/groups/defence-artificial-intelligence-centre