The Right Tool for the Job: Collaborative Requirements Mapping for the Military Application of Media Analysis

Tom Powell*, Saskia E. Lensink, Arnout de Vries
Netherlands Organisation for Applied Scientific Research (TNO)
NETHERLANDS

Shadi Ghajar-Khosravi , Bruce Forrester
[2]Defence Research and Development Canada (DRDC)
CANADA

*tom.powell@tno.nl

Keywords: Media Analysis; OSINT; Information Environment, Cognitive Superiority.

ABSTRACT

Modern warfare is increasingly fought in the information environment using deception and influence techniques via open source media. Governments, academia and industry across NATO nations have responded by developing various innovative computational methods for extracting, processing, analysing, and visualising meaningful information from massive amounts of media content. However, it remains unclear which (combinations of) tools meet the requirements of military analysts and operators, and whether some requirements remain unmet. In response to this, DRDC Canada and TNO Netherlands initiated a collaboration to develop a standardised and multi-faceted mapping of requirements for Media Analysis. This paper presents work completed in the first phase of this collaboration. Specifically, we: (1) developed a framework of possible media analysis tool functionalities (i.e., the art of the possible); (2) collected current user requirements of CAN and NLD stakeholders; and (3) analysed the gaps to show which user requirements could be met by which functionalities. This project builds directly on SAS-142, by using the Framework for Internet Exploitation Science and Technology Assessment (FIESTA). This paper illustrates the application of FIESTA to two prominent media analysis capabilities: (1) Sentiment Analysis and (2) Narrative Analysis. Findings show that although these capabilities have some unique functionalities, they have very many functionalities in common. As a result, research and development efforts can be optimised by focusing on unique (novel) functionalities whilst recycling multi-purpose functionalities. These efficiency gains can be multiplied by applying FIESTA to multiple media analysis capabilities and in collaboration with multiple NATO nations.

1.0 INTRODUCTION

The rapid digitalization of societies means that the military is increasingly reliant upon exploiting data derived from the internet [1]. Being able to effectively extract, process, analyse, and visualise meaningful information from massive amounts of online media content, while adhering to legal and ethical frameworks, is challenging. Currently, many organizations are relying on manual efforts to extract information, but are exploring ways in which to partly automate and support these processes by means of media analytics tools [2]. These efforts have the potential for large payoffs, with one estimate suggesting that artificial intelligence (AI) systems could save an intelligence community (IC) analyst as much as 364 hours, or more than 45 working days, a year [3].

How are such media analysis tools put to use by military practitioners? The most obvious distinction in the military application of media analysis is between open source intelligence (OSINT) and information environment assessment (IEA). OSINT focuses on analysis and evaluation of abundant information from

open sources applied to various intelligence subjects (adversary organisations, weapon systems, situational awareness, etc.) [4]. IEA emphasises the cognitive effects on audiences via the virtual domain, media and information systems and their dynamics, as well as by social processes, culture and the physical environment [5]. These different applications inevitably lead to some unique functionalities for tools for OSINT and IEA, although, due to their focus on openly available information that is often gleaned from the internet, there will be many functionalities in common. Another distinction that can be made is the military level – from strategic, operational to tactical - at which media analysis can be applied. Media analyses are relevant to all levels and tools for media analysis may be used to analyse longer-term trends at a regional level (strategic) or short-term local phenomena (tactical). However, there may be specific requirements at each level which complicates the development of tools for media analysis.

Due to these challenges, as well as the pace of development in media analysis, it is challenging to maintain an overview of the state-of-the-art in order to understand which (combinations of) tools meet the requirements of stakeholders in this area, and to have a structured way of identifying avenues for research and development. In order to address this issue and progress from initial forays into this field, it is beneficial to consolidate experience and efforts across NATO nations. In response to this, DRDC Canada and TNO Netherlands initiated a collaboration to generate a detailed, multi-faceted overview of media analysis functionalities and stakeholders' requirements, working within the prevailing legal and ethical frameworks. These efforts made it possible to assess whether combining the individual nations' efforts could better address stakeholders' needs.

To this end, the authors made use of the "Framework for Internet Exploitation Science and Technology Assessment" or FIESTA (SAS-142) [6]. The FIESTA framework has originally been developed as a methodology for comparative analysis to support suitability assessment of media analysis tools. Within the context of this project, the authors did not use the FIESTA framework for tool assessment, but expanded the framework by using it to compare and map requirements to listed functionalities to identify commonalities and overlaps between different organizations' requirements, and thereby identifying opportunities for R&D. The aim was to explore and demonstrate optimal combinations of media analysis functionalities as well as gaps that need to be addressed by research and development. Specifically, the authors applied and expanded FIESTA by: (1) creating a detailed framework of media analysis functionalities; (2) assessing CAN and NLD stakeholder requirements by conducting in-depth interviews; and (3) performing a gap analysis to identify which stakeholder requirements can be met by the media analysis functionalities. This paper illustrates our approach with two prominent media analysis capabilities: Sentiment Analysis and Narrative Analysis. We hope these efforts will maximise the military operator's ability to gain knowledge from large collections of media content.

The FIESTA framework is further described in Section 2.1, followed by an explanation of how the project team applied the framework in this project in Section 2.2. Section 3 illustrates the application of FIESTA to identify requirements, media analysis functionalities, and areas for research and development (R&D). Section 4 identifies lessons learned and outlines avenues for future research. We offer conclusions in Section 5.

2.0 METHODS

2.1 The FIESTA Framework

The FIESTA framework is developed as a comprehensive, robust, and easily applicable grading system for the functionality of tools and applications designed for media analytics [6]. This field of analytics contains a dynamic and congested tool space, with a high turnover of complex tools. However, there is a lack of intelligence-specific tools and most tools that handle social media are primarily aimed at collection with limited analytical capability. FIESTA is suited to this context as it allows for flexibility and adaptability when assessing rapidly developing technology, and provides a structured assessment of suitability in light of

user requirements to support procurement and/or technological development.

FIESTA is based on the Multi Criteria Analysis (MCA) [7] approach to requirements mapping and assessment. It comprises a demonstrator implemented in an Excel Workbook. The Workbook contains a structured list of suggested capabilities and functionalities with respect to media analysis as developed by military operators and scientists from multiple disciplines. The list contains up to four levels. The top level is the "Functions" level, which reflects the four functions of the Intelligence cycle: "Direct", "Collect", "Process" and "Disseminate". The other levels are: "Main", "Sub" and "Detail". An extra column is added for the "Explanation" functionality. See Figure 1 for an impression of a part of the FIESTA demonstrator. Take, for instance, the first row which describes a functionality for *graphical* representation (detail level) in a *collection planning system* (sub level) which supports the *collection management* capability (main level) as part of the *direct* function of the intelligence cycle (function level).

ID	Function	ID	Main	ID	Sub	ID	Detail	Explanation
2	Direct	4	Collection management	14	Collection planning system	75	Graphical	Planning of timelines, t
						176	Textual	
						79	Hybrid	
				56	RFI Management	134	Receive	Can you receive tasking
						171	Tasking	Can you task others via
						164	Status update	Can you follow on the t
		21	User management	57	Roles	7	Admin	All functions available t
						188	User	Limited user functions
						111	Other	Some custom role and t

Figure 1: Impression of the FIESTA framework.

2.2 Application of the FIESTA Framework

In this project we applied and expanded the FIESTA framework to compare possible media analysis tool functionalities with requirements from military and security stakeholders from Canada and the Netherlands. Specifically, the following steps were taken:

Interviews were conducted in order to identify requirements from Dutch and Canadian stakeholders in the military and national security domains. In the Netherlands, interviews were held with practitioners from two military Information Manoeuvre organisations, two military intelligence organisations, one Dutch Safety Region, and one regional police organisation. In Canada, interviews were held with intelligence analysts at the strategic, operational, and tactical levels as well as strategic communication public affairs officers. Further, requirements were gathered from OSINT analysts as part of the SAS-IST-102 RTG [16].

The **FIESTA framework was expanded** to include a greater range of capabilities and functionalities relating to IEA and OSINT, such as narrative analysis, sentiment analysis, and explainability, amongst others. An extra level of detail, "Technical Method", was added to the levels described in section 2.1. The addition of a Technical Method level provides more information about the internal workings of tools, which for some end-users are perceived as a 'black box'. For example, technical methods relating to text analysis were broadly categorized as: rule-based heuristics; sentence-based classification using Machine Learning (ML); sentence-based classification using Deep Learning (DL) [8]. Additionally, functionalities for different mandates (training, exercising and actual mandated operations) will have to be included, taking into account the GDPR regulations.

A **gap analysis** between stakeholders' needs as identified in the interviews from step 1 and the updated FIESTA functionalities list from step 2 was conducted. These resulted in "areas for Research & Development (R&D)", bringing to light the existing gaps in stakeholder's media analysis capabilities which can be filled by the functionalities of existing tooling, or, where this is not available, through research and development into novel media analytics functionalities.

3.0 RESULTS

The results of the efforts described in Section 2 are illustrated using two examples at the "Main" level of the FIESTA framework, within the *Analysis* "Function": (1) Sentiment Analysis and (2) Narrative Analysis. These were chosen due to their prominence in the field of media analysis and because they make for an interesting comparison – the first being a long-standing capability in this field and the second being somewhat more novel. For each example we outline relevant stakeholder requirements, functionalities, and areas for R&D.

3.1 Sentiment Analysis

Sentiment Analysis is the automatic mining of certain affective states or opinions in texts [9]. This can be done by assigning a certain affective state, such as positive or negative, based on the presence of certain keywords or patterns, or by deploying ML algorithms that have been trained to classify texts for their sentiment. Recent years have seen an increase in the deployment of deep learning algorithms for sentiment analysis [10], where a specific focus has been on being able to assign sentiments not only to a full text or sentence, but also on identifying to what specific part of a sentence or *aspect* a sentiment applies [8].

3.1.1 Requirements

Analysts at a Dutch Safety Region indicated that for their activities, it is crucial to be able to (1) detect emotions and sentiment regarding specific topics or entities, ranging from COVID-19 to opinions on recent political events – i.e., aspect-based sentiment analysis. End users would also like to be able to (2) use a broader spectrum of qualifications – discrete emotions like anger, fear, confusion and happiness – than just the standard 'positive vs. negative' sentiment. Similarly, interviews with Dutch Information Manoeuvre organisations revealed the need – within the context of a mandated military mission – for insights into the sentiment of specific groups towards recent events, themes, or other groups (i.e., aspect-based SA). Moreover, there was a desire to (3) visualize these analyses over time and space to be able to achieve a dynamic understanding of sentiment in certain communities. All interviewees noted that it is unclear to what extent the 'sentiment scores' provided by tools have been validated and how exactly those are to be interpreted. This led to a lack of 'trust' in sentiment analyses, highlighting (4) the need for explainable AI (XAI) capabilities [11] within sentiment analysis.

3.1.2 Functionalities

At the "Sub" level of the FIESTA framework, sentiment analysis (SA) functionalities were categorized into: 1) untargeted SA – the 'classical' approach; 2) aspect-based SA, where a sentiment is assigned at a lower granularity level to a specific part or aspect of a text, for both the 'sender' and 'receiver' of the sentiment; 3) emotion analysis, which extends the typically binary sentiment classification into positive and negative sentiment to a multi-label classification into different categories (such as anger, fear, happiness, and potentially on a Likert scale); 4) hate speech detection, which involves the application of a classification algorithm trained to discern hate speech; and 5) intent analysis, where a classification model is used to recognize different types of intent (e.g. wishes, desires or concrete plans). Additionally, capabilities relating to visualisation and filtering in time and geographical space apply to sentiment analysis but also other analysis capabilities, and so were included as generic categories in the FIESTA framework.

3.1.3 Areas for R&D

By using the expanded FIESTA overview of SA to study our interviews with Dutch stakeholders, we learned that there is currently a need for (1) extension of aspect based sentiment analysis, (2) improvements in models that go further than just a binary classification of sentiment (negative vs. positive), (3) deploying emotion models, (4) better insights into validation of techniques and explanation of accuracy to the end user,

and (5) a greater variety of visualizations including sentiment in a network and over time and geographical space. Despite a multitude of sentiment analysis tools, military and security analysts lack specific support for nuanced analyses and lack the information required to assign confidence levels to conclusions based on SA. However, the question arises whether there is a legal, ethical and political mandate to respond to this need.

3.2 Narrative Analysis

Narrative Sensemaking is a key capability in enabling militaries to engage in activities of narrative competition and narrative warfare. Narrative competition activities require military operators to understand the public's interpretation of military activities, while narrative warfare activities require analysts to detect and understand "weaponized narratives" [12], [13]. Narratives are best understood through looking at how they change over time. For example, once important narratives are identified, analysts can look for changes in authors, associated values, emotions, sentiment, spread, speed of spread, interest from influencers and traditional media, or even the part the narratives play in influence and deception campaigns. From this understanding, countering or proactive messaging can be planned. Accordingly, [14] defined narrative sensemaking as an "[iterative] process of detecting, contextualizing, understanding, and assessing the impact or potential impact of narratives that are emerging and circulating in the IE and could have relevance to [the] defence organization".

3.2.1 Requirements

From interviews with Canadian military operators, narrative sensemaking was raised as an important capability that can further enable defensive activities in the IE [15], [16]. Narratives (and most likely aspects of a certain narrative) can be expressed through text messages, stories, new articles, memes, videos, etc. While a narrative could be expressed in a single message, they are mostly discovered through analysis of many messages in the aggregate. Therefore, we consider an online narrative as a cluster of interrelated or semantically similar information or media items that discuss the same stories, news articles, or events, contain an expression of value (e.g., importance of the rule of law, respect for human life, etc.), and could express emotions pertaining to these values (e.g., anger, joy, disgust, contentment). As such, the static units of analysis for online narratives will include the: (1) values discussed, (2) emotions expressed, (3) topics, events, or stories referenced to, (4) actors (organic, influencer, or bots) involved in generating and spreading the narratives, (5) values, emotions, topics, events, or stories discussed across all the content items generated by actors, (6) the audience exposed to the narratives' content, (7) values, emotions, topics, events, or stories discussed across all the content items generated by audience, (8) the propagation rate of narratives, (9) the network structure of their spread, (10) the geographical reach, and (11) their use as part of an influence, deception, or disinformation effort/campaign. Additionally, the dynamic units of analysis will include the analysis of changes or nuances of all the listed static units across time. All of these functionalities were required by the operators and analysts interviewed.

3.2.2 Functionalities

The functional and operational requirements for tools providing a narrative sensemaking capability capturing the above-mentioned units of analysis were gathered and refined into five categories (as described in [14], [15]): usability; information security; content analytics (i.e., text, images, videos); network analytics; and narrative analytics. These capabilities and their underlying functionalities were then added to the FIESTA framework, mapped into relevant existing categories. For some functionalities, new "Main" or "Sub" items needed to be added to the framework. Additionally, functionalities relating to visualisation and filtering in time and geographical space apply to narrative analysis but also other analysis functions, and so were included as generic categories in the FIESTA framework. A clear challenge is to what extent the needs for functionalities meet the requirements of the legal and ethical frameworks

3.2.3 Areas for R&D

Through comparing the requirements with the functionalities, we identified the following gaps for R&D: (1) algorithms for the clustering of semantically similar messages, images, videos, etc., (2) capturing and monitoring narratives and their associated values and emotions over time or space, (3) detecting and monitoring narratives' hierarchical structures, (4) attribution to determine adversaries' vs allies' messages and narratives, and (5) predicting events or behaviours based on narrative-based analytics. Further research is also needed to investigate how combinations of tools could be used to develop TTPs (Tactics, Techniques, and Procedures) specific to use cases such as influence, deception, or disinformation campaigns. Also, researchers are starting to look at the cognitive aspects of such campaigns. As such it is expected that new analytics, based on cognitive and psychological constructs, will start to appear in the coming years.

4.0 DISCUSSION

This paper illustrated how the FIESTA framework can be used to (1) develop a detailed and standardised framework of media analysis functionalities, (2) map the requirements of analysts and operators in this area, and, in turn, (3) identify focus areas for R&D.

4.1 Added Value

In this paper, two prominent media analysis capabilities – Sentiment Analysis and Narrative Sensemaking – were used to illustrate requirements mapping. Application of the FIESTA framework to these two analysis capabilities revealed some functionalities unique to narrative analysis – for instance identifying media items sharing the same topics, stories, links to news articles, events or values – and some functionalities unique to sentiment analysis – such as hate speech detection and intent analysis. However, these analysis functions have more functionalities in common, such as the ability to analyse expressed emotions, identifying actors who produce these expressions, the targets of these expressions, as well as their temporal dynamics and geographical reach, and the need for the outputs to be explainable. Mapping these unique and common functionalities provides direct input to developers working with a modular approach: they can focus on developing unique (novel) functionalities that meet specific requirements and they can save development time by recycling modules which serve multiple purposes for different analysis capabilities.

These efficiency gains are multiplied when the FIESTA framework is applied to multiple analysis capabilities, or in collaboration across multiple nations. Instead of each nation developing or procuring their own media analysis tools, collaboration leads to shared knowledge of each nation's existing capabilities, their strengths (and weaknesses), and at the very least supports consolidation of knowledge about the state of the art in media analytics. At most, international collaboration in this area can lead to tool/technology sharing, joint development, and can stimulate innovation by bringing together different perspectives and expertise.

Another added value of a comprehensive mapping of media analysis functions as attempted in FIESTA is to work towards a common standardised terminology. This, in turn, makes the boundaries and commonalities of different emerging analysis disciplines such as open source intelligence (OSINT) and information environment analysis (IEA) explicit. The FIESTA framework thereby facilitates in establishing a standardised terminology which can facilitate fruitful discussion between the analytical communities, for instance about roles, best practices and knowledge sharing. Future research can usefully focus on the analytical interface between OSINT and IEA.

4.2 Challenges and Future Research

The outputs from this paper – which reflects the present status of an ongoing collaborative project – mainly provide directions for longer term research and development, by comparing user requirements with the art-

of-the-possible in media analysis functionalities. An obvious immediate next step would be to apply the expanded FIESTA framework to existing tools in use by military operators and analysts. The outputs of this effort would be short term recommendations to fulfil requirements with (optimal combinations of) existing tools. The FIESTA framework can play a central role in this effort, since its framework can be applied flexibly and in a multidisciplinary fashion – enabling its use by different audiences, ranging from data science specialists to those with a non-technical background, as well as those from different NATO nations.

A related challenge that was raised as a major theme in the requirements interviews of this study, is the need for the analyst to *trust* the output that is presented in media analysis tools. This trust relates to (1) the reliability of the information (articles, social media posts, etc.) that is fed into media analysis tooling, as well as (2) the explainability of the tool's AI-powered outputs provided to end users (e.g., analysts lacking data science skills). On the former point, research should continue to look into innovative (semi-automated) methods for the evaluation of open-source media [17]. The latter point relates to the burgeoning field of explainable AI (XAI), a key issue with complex "black box" ML models such as deep learning [11], [18]. Guidelines for XAI systems include the need for interactivity, causal explanations and transparency to facilitate trust [19]. As AI models become more prominent in media analysis, their explainability will be essential. XAI will therefore be an important aspect of the future development of the FIESTA framework.

5.0 CONCLUSION

This paper illustrated the application and expansion of the FIESTA framework for media analysis requirements mapping using two important analysis capabilities – Sentiment Analysis and Narrative Sensemaking. The effort to develop a comprehensive standardized framework of functionalities for media analysis represented a major challenge. However, adopting a rigorous operations research approach via FIESTA reaps meaningful benefits: Knowledge about analysis capabilities and functionalities is consolidated which establishes proven technologies, highlights areas for immediate development, and reveals gaps for future innovation. A universal and comprehensible terminology is developed to facilitate understanding in end users with diverse expertise. Collaboration in these efforts can lead to major efficiency gains and therefore be a force multiplier given the importance of media to modern conflict. To reap these benefits, future research should continue to expand and apply the FIESTA framework for media analysis in increasingly collaborative settings.

6.0 REFERENCES

[1] D. Fiott, "Digitalising defenze: Protecting Europe in the age of quantum computing and the cloud," European Union Institute for Security Studies, 2020.

[2] T. E. Powell, A. M. Sternheim, I. Tolios, A. F. van Luenen, T. Schaap and F. Bomhof, "TNO Communication Analysis Toolkit: A suite of tools for automated analysis of media and communications to improve situational understanding," in Operations Research & Analysis Conference, 2020.

[3] K. Mitchell, J. Mariani, A. Routh, A. Keyal and A. Mirkow, "The future of intelligence analysis," Deloitte Insights, 2019.

[4] data.europa.eu, May 2022. [Online]. Available: https://data.europa.eu/en/datastories/open-source-intelligence. [Accessed 2022].

[5] NATO, "AJP-10 Allied Joint Doctrine for Strategic Communications," NATO Standardization Office (NSO), Draft.

[6] STO-TR-SAS-142, "Developmental Research to Create an Assessment Framework for Science and Technology used in Internet Exploitation," NATO Science & Technology Organisation, 2021.

[7] Department for Communities and Local Government, "Multi-Criteria Analysis: A Manual, Communities and Local Government Publication," 2009.

[8] H. H. Do, P. W. C. Prasad, A. Maag and A. Alsadoon, "Deep learning for aspect-based sentiment analysis: a comparative review," Expert systems with applications, vol. 118, pp. 272-299, 2019.

[9] W. Medhat, A. Hassan and H. Korashy, "Sentiment analysis algorithms and applications: A survey.," Ain Shams Engineering Journal, vol. 5, no. 4, pp. 1093-1113, 2014.

[10] A. Yadav and D. K. Vishwakarma, "Sentiment analysis using deep learning architectures: a review.," Artificial Intelligence Review, vol. 53, no. 6, pp. 4335-4385, 2020.

[11] P. P. Angelov, E. A. Soares, R. Jiang, N. I. Arnold and P. M. Atkinson, "Explainable artificial intelligence: an analytical review," Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, vol. 11, no. 5, pp. 1-13, 2021.

[12] NATO StratCom Centre of Excellence, NATO StratCom Handbook, Riga, Latvia, 2015.

[13] B. Allenby and J. Garreau, "Weaponized Narrative Is the New Battlespace," Defence One, 3 January 2017.

[14] S. Waldman, B. Forrester, J. McInnis, S. Ghajar-Khosravi and S. Gibbon, "Enabling Narrative Sensemaking in the Information Environment (U): An Evaluation of Commercially-Available Media and Social Media Monitoring and Analysis Tools for Enabling Situational Awareness of Military-Relevant Narratives in the Information Environment," DRDC-RDDC-2021-R039, CAN PROTECTED A // NON-CONTROLLED GOODS, April 2021.

[15] B. Forrester, "Intelligence Exploitation of Social Media," STO Technical Report, Final Report of Task Group SAS-IST-102, June 2018.

[16] B. Forrester, S. Ghajar-Khosravi and S. Waldman, "Machine Learning-Enabled Narrative Search in the Information Environment," in 15th NATO OR&A Conference, 18-20 October 2021.

[17] T. E. Powell, S. Oggero, J. Westerveld and E. Schook, "Dealing with uncertainty in Hybrid conflict: A novel approach and model for uncertainty quantification in intelligence analysis," in IST-190 Research Symposium (RSY) on AI, ML and BD for Hybrid Military Operations (AI4HMO), 2021.

[18] J. Launchbury, "A Darpa Perspective on Artificial Intelligence," DARPA, 2017.

[19] Q. V. Liao, D. Gruen and S. Miller, "Questioning the AI: Informing Design Practices for Explainable AI User Experiences," in Conference on Human Factors in Computing Systems, 2021.

[20] S. Waldman, B. Forrester, J. McInnis, S. Ghajar-Khosravi and S. Gibbon, "Evaluation and Test of Free Tools for Disinformation Verification and Tracking," DRDC-RDDC-2020-L178, CAN UNCLASSIFIED // NON-CONTROLLED GOODS., November 2020.