Massage the message:

Modularising software for influence operation detection in social media

Arild Bergh, Ph. D.
Norwegian Defence Research Establishment (FFI)
Kjeller, Norway

Abstract—Over the past few years social media has become a key battleground in influence operations between actors who are in covert or overt conflict [1]. At the same time social media has never been more pervasive. Many of the 2 billion active Facebook users read news selected by automated software routines and many powerful state leaders bypass diplomatic channels to directly present their thoughts through Twitter. Influence operations have been undertaken by smaller and larger terrorist organisations as well as nation states and major upheavals in domestic politics appear to have been the target of small and cheap, but widely distributed, influence operations.

This article discusses a potential technical approach to detect influence operations in social media. The key goals and issues that have influenced the design of the proposed technical solution are discussed. A brief outline of potential architectural solutions that will facilitate heterogeneous data input while providing distributed situational awareness is examined in some detail.

Keywords: social media, influence operations, strategic communications, detection software, architecture.

Introduction – Social media and influence operations

Over the past few years national defence organisations have received a wakeup call with regard to social media and their use to attempt to manipulate opinion, whether in hybrid conflicts with clearly discernible, kinetic elements or in low-level societal manipulation. In the decade since Facebook became a household name, social media has gone from a curiosity, to a utopian ideal before becoming the everyday preferred communication method for many until it is now presenting considerable problems for military as well as civilian defences.

There are five aspects of social media that makes it different from previous mass media in terms of influence operations, namely reach, speed, reception, anonymity and cost [2].

Firstly, social media are not restricted by geography (local censorship efforts not withstanding). The numbers involved are staggering in an historical perspective: At the time of writing Facebook has exceeded 2 billion monthly users whereas on Twitter alone 455,000 messages are posted per minute [3]. Secondly, the communications are computer mediated and algorithmically curated, thus information, false or not, can be spread instantaneously. Third is the somewhat counter-intuitive fact that receiving information through social media can feel very intimate, you can participate via your mobile phone on the bus or sitting in the sofa with a laptop; it is a mass(ive) medium masquerading

as your friendly local pub. Furthermore, there is the anonymity afforded through the total openness of social media platforms where one can easily automate the creation of pseudonymous accounts.

The fifth aspect means that the actual cost for the attacker is very low, especially when we consider the potential reach, allowing actors to experiment with influence operations at scale outside of actual conflict situations. Information warfare as discussed since the 1990s [4] is therefore occurring stealthily in broad daylight, paid for by the commercial operators of social media through advertising. Historically [5], one would expect this opportunity be used both in the pursuit of preferable outcomes in general foreign policy, but also in preparing for kinetic conflicts by creating favourable circumstances in other key states; in geographical areas of future operations or in public opinion at home. There are currently a number of examples of state based attempts at influencing events in other countries through the use of social media. Russian activities in Ukraine and the US elections in 2016 (cf. eg. , [6]–[8], [9]) are probably the most well-known example of such information operations.

In sum this suggests that defence organisations need to move away from using social media only as a tool for sharing information or for open propaganda [10] to treating it as a potential domain of conflict in its own right [11], [12]. This will require both extensive research and viable tools. This paper represents a contribution to this work rooted in a preliminary exploration of needs and possible solutions for the Norwegian Armed Forces.

Definition of terms

- **Backend**: Refers to software (either part of the main framework, plugins or services such as a database) that is running on a remote server.
- **Demonstrator:** In this paper this refers to a piece software before it reaches the prototyping stage. The demonstrator showcases possible solutions with focus on potential users' comprehension of what is suggested and not actual data processing.
- **Frontend**: The part of the software displayed in a web browser that the user interacts with.
- **Framework**: This is the core code, what the end user would perceive as "the application". Comprises *frontend* and *backend* elements that communicate through the HTTPS protocol and websockets.
- **Plugin**: A self-contained piece of software that will typically receive data (in a known format), process it and pass it on to the next plugin in the *processing path* (see below). A plugin can be written in any language that supports Advanced Message Queuing Protocol (AMQP) (see *PubSub* below).
- **Processing path**: Two or more *plugins* that work in series to manage data from the collection point to the presentation of an analysis. A summary of all potential tasks in a processing path is discussed below.
- **PubSub**: Publish and Subscribe, technology that allows multiple clients to subscribe to a publisher (source) of data. Data will be passed between plugins in the processing path using AMQP which is one instance of a PubSub solution.

Situational awareness of social media based influence operations

A recently concluded study from the Norwegian Defence Research Establishment (FFI) on the use of social media in influence operations has examined recent trends in this area [2]. The focus of this report was to understand how such operations might have an effect through a combination of technical and human manipulation. A companion activity to this study undertook an initial exploration of software solutions to help defence staff with what has been referred to as an *Information Environment Assessment* in a NATO context [13]. Such assessments are done in relation to strategic communications to see what, if any, false information in spread by potential adversaries. This software exploration focused on developing a *software demonstrator* named **Fossen** (Norwegian for "waterfall", evoking the rush of incoming data from social media).

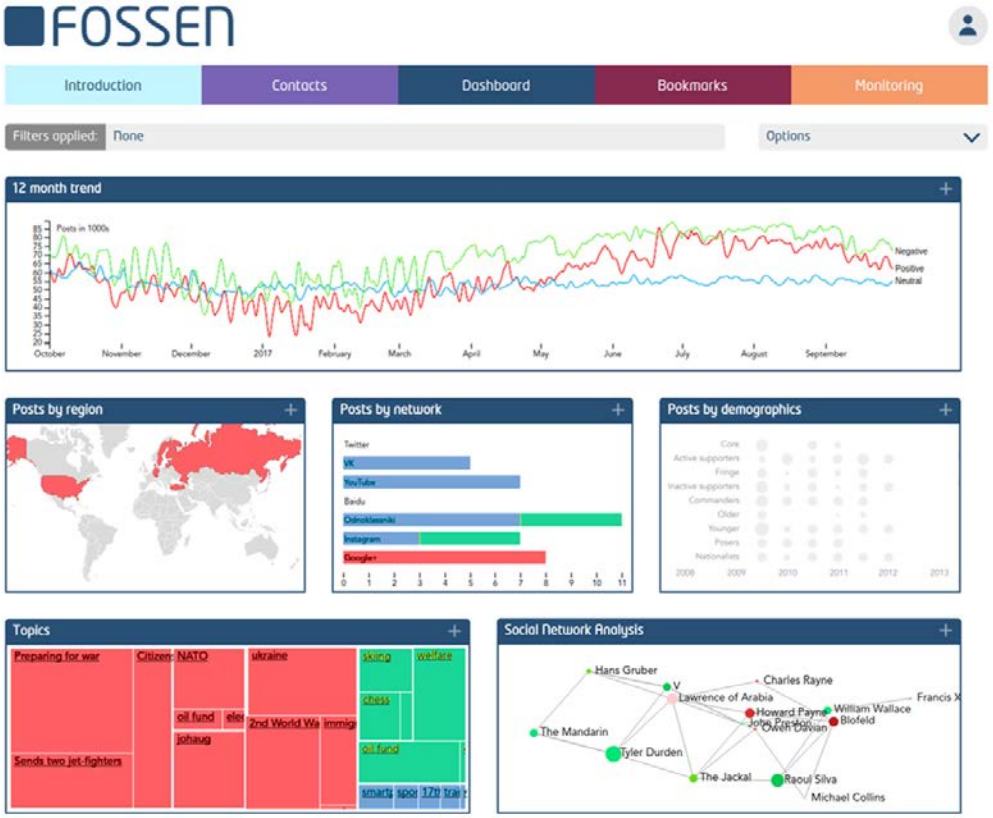

Figure 1 Screenshot of the Fossen demonstrator[1]

Fossen was a rudimentary software, developed as a proof of concept to showcase potential applications to possible users in different parts of the Norwegian Armed Forces. The discussion in this paper will centre on the envisaged architecture and how this is developed to fit in with the complex nature of national defence organisations while at the same time making it flexible enough to handle the rapidly and continuously changing world of social media.

[1] The charts in this example are for illustration only. These charts are based on random data and do not in any way reflect actual data nor any current priorities at the Norwegian Defence Research Establishment.

There is nothing new about using software to keep tabs on social media activities; analysis of social media, such as Social Media Intelligence (SOCMINT) [14], is applied in a range of circumstances. On the commercial level there are numerous solutions and services available that facilitates social media analysis for companies to improve marketing and branding of their products. An example of this is Brand24. They undertake sentiment analysis of online product mentions to determine if they are positive or not and offer to discover the most important social media influencers [15]. This paper is thus not suggesting that the Fossen demonstrator and the approach discussed in this paper is a first in this field. However, we believe that some of the approaches suggested can beneficially be shared with others, particularly smaller countries, and welcome the sharing of experiences in this field.

Key goals and issues

There were some key goals and issues that informed the work on the Fossen demonstrator. It is assumed that these matters will, at least in part, be similar to restraints and requirements that other national defence organisations experience. This paper will first summarise these goals and issues before discussing how they affected ideas for potential architectural solutions.

Goal: Flexible situational awareness and information sharing

To avoid duplication of software solutions and wasted resources, staff in different parts of the Norwegian Armed Forces should be able to use such a tool in different tasks and situations. This means that it needs to provide situational awareness for non-experts in a variety of social media platforms across organizational domains and levels. What is of interest to a communication advisor need not be of interest to tactical planning staff. It is therefore important to ensure that different local operating pictures can be presented and modified without changes to the underlying software code.

Goal: Ability to respond rapidly to changing situations

In social media based influence operations the time to discover and respond to new, hostile activities may be a matter of hours or less than that. The capacities to assess social media activity is crucial when being the target of an influence operation [16]. A traditional analytics' setup where briefings are produced at set intervals, perhaps relying on technical staff to query data, would not have the speed and agility required here. A solution with less reliance on technical support staff and the ability to quickly change assessment parameters would be a better fit for social media analytics tasks.

Goal: Save resources and re-use solutions

Larger state actors may have considerable technical and human resources available to assess and manipulate content in social media. Smaller state actors do not necessarily have similar resources available. Thus one would want to ensure that any technical solutions that are explored can be (re)used, not only by different parts of the Norwegian Armed Forces but also by actors such as governmental bodies that deal with IT security on the national level or the police. In this paper re-use focuses on technical issues, discussions as to who will use tools such as the one outlined here are beyond the scope of this paper.

Issue: Organisational and technical patchwork

The Norwegian Armed Forces, like any large organisation, is a multi-layered patchwork of hierarchical levels, subject specialities and geographical locations. Often this patchwork is reflected in ICT infrastructure that supports the different parts of the organisation where information silos (or stovepipe systems) may limit true, lightweight information sharing. Any technical solution needs to work **with** as well as **around** these existing silos so they don't end up locking data away, or blocking information sharing.

Issue: Heterogenous data sources and formats

Facebook and Twitter are the two main social networks that come to mind for most people when discussing social media. However, in reality people use locally popular network (such as Vkontakte in Russia or Weibo in China), niche networks (such as discussion groups on Reddit) or facilities that one initially may not consider as social media, such as comment fields on newspapers websites or reviews on Amazon. This paper defines social media as: *"Internet accessible services that allow publishing of content by people who don't own or control the service with facilities for others to access, share and respond to this content."* Furthermore, social media are not static, new media (and new ways of using social media) appears and others fade away. Vine, for example, had 200 million users who shared short videos through the service before it was closed for further uploads in 2016 [17].

Hence there is a need to ensure that data from any social media sites, whether pure social media such as Twitter, or hybrids such as the Disqus commenting software on a news website, can be collected and analysed. Additionally, it would also be useful to enable the input of human intelligence (HUMINT) based on SOCMINT. This could be high level analysis that may combine different data sources and result in data points that could be part of an overview presented to users of the tool in question.

Issue: Privacy concerns

A key issue to any democracy is the need to protect citizens' privacy and Norway has rather strict data protection laws. This is an issue even when accessing only openly available data (SOCMINT) as combined data may still identify individuals. However, to detect the use of social media in hybrid attacks requires historical data where trends can be discerned. As an example, it would be of interest to see whether negative messages about Norway in relation to NATO increases significantly, or are on a fairly normal level, and if the negative messages are posted by bots, trolls or simply concerned citizens engaging in an online debate. This requires the collection of social media posts over time and these posts needs to be linked to the original social media profile they came from, although the profile/person name can be anonymised.

We therefore have a contradiction here. In this context one is not really interested in **who** a person posting a message is, because that "person" may be a bot or someone employed by a potential aggressor to post messages. Yet, the very collection of messages linked to specific accounts increases the risk of privacy breaches.

Massaging the message: The plugin approach

Evaluating the above goals and issues quickly led to the principle of a lightweight, modular approach in terms of the underlying architecture. In broad terms this solution fits into the Service Oriented Architecture (SOA) paradigm where loosely coupled web services are available to clients through standard interfaces. Such services are then registered in a central registry (cf. e.g. , [18]–[20]). There are however some differences from a standard SOA approach as we shall see.

A primary inspiration for the architecture of any future solution is the work the Norwegian Defence Research Establishment has done on Mlab [21], [22]. Mlab is a modular app builder where functionality is added to an app through the use of components. Each component is self-contained and knows how to request information at design time (from the person creating the app) and to display the resulting information at runtime (to the app user). Mlab, in other words, is a framework and API used to host components and provide app editing and database management facilities.

One cannot know in advance which social media services will be used for influence operations in the future by potential hostile actors [23]. Any solution needs to be extendable in a similar way to Mlab. This allows us to quickly add new features and services for a range of social media. In a crisis situation this facilitates the quick ramp up of new analytical capabilities, without losing existing functionality and without the need to retrain users.

To cater for this a plugin architecture would be the most pertinent choice. This would allow simple scaffolding code to be written immediately and over time be extended with added functionality without further programming. Just as the Mlab tool will let a non-technical person create an app from components, such a solution will let non-programmers create their own analysis from a combination of existing plugins. After briefly defining some terms this paper will examine how the core framework and different plugins could interact to achieve the goals discussed above.

Division of (plugin) labour

This section outlines the areas that the suggested framework and different types of plugins would be responsible for. The table below shows a summary of the main elements that would make up the overall software solution.

	Framework / API		Plugins
Frontend	• Load & init plugins • Submit user data	↕ JavaScript ↕	Presentation plugins: • Display data • Handle user interaction
	↕ Websockets ↕		
Backend	• Load & init plugins • Store user settings	↕ CLI / REST ↕	Backend plugins: • Collect or process data • Publish / subscribe (AMQP)

Table 1 Overview of different elements of a possible solution and communication between them

Framework & API

On both the backend and frontend the framework code is loaded first (via a NodeJS web server and a web browser respectively). This code then "hosts" the plugins described in this document. The hosting of a plugin is facilitated by making common API functions available to all plugins and managing the loading and initialisation of plugins. Through simple configuration files the backend framework will know in what order, and how, to set up plugins, initiate the data acquisition process and start the data flow down the processing path. In the browser, the frontend code will load and initialise the presentation plugins.

In addition, the frontend framework provides basic services such as managing user authentication and access rights and providing a consistent user interface. The front and backend also facilitate storage/retrieval of user data (for instance a bookmarking feature to store filters for later retrieval or thresholds for alerts).

Plugins

	Steps in processing path	Explanation	Examples
Optional, calls backend	↓ 10. Measure impact	Only executed if a user chooses to respond. Generates new data as one's own messages are collected.	Re-use backend plugins and query for retweets of, and replies to, tweets from the user.
	↓ 9. Respond		Use SM API to post a response.
	↓ 8. Filter	User applies a filter to presentation plugin.	Query sent to the filter plugin with information on how to query the data, result is sent through path and displayed by presentation plugin.
Frontend	↑ 7. User	The person using this software.	Clicks on country to display underlying data.
	↑ 6. Present	Interactive task, displays results and allows user to filter, bookmark and respond to messages.	Display map with countries coloured according to number of tweets that originated there. Click to display in-depth data.
Backend	↑ 5. Analyse	Core task that may use AI to detect sentiment, do simple word count, calculate messages per country, etc.	Count number of messages mentioning NATO, order by country.
	↑ 4. Filter	"Glue" task that allows us to format data for any purpose without redoing other plugins.	Query database for tweets from last 30 days.
	↑ 3. Store	Use various data storage solutions.	Save to Hadoop (HDFS).
	↑ 2. Collection	Can be merged with the "Watch" task below.	Downloading new messages with a certain hashtag via Twitter API.
	↑ 1. Watch	Check if new messages of interest have arrived. N/A if want everything.	Use SM API to get a count of new messages with a certain hashtag.

Table 2 The processing path and plugin tasks (not all tasks are required; tasks may also be combined or repeated).

Plugins will have access to common functions in an API, such as being alerted when the user has clicked on an element of a chart and sending a filter query to back end plugins. A plugin can also rely on the jQuery and D3.js JavaScript libraries being present. Beyond this it will be self-contained, i.e. a presentation plugin does not rely on other presentation plugins. That being said, plugins can share both third party and custom JavaScript libraries if they need to.

The frontend only has one type of plugin, the (data) presentation plugin. All other plugins reside on the backend. The core role of the presentation plugin is to display an interactive chart showing the current state of relevant data from one or more social media service. In recent years considerable innovation has taken place with regard to visualising big data (such as social media data) in new ways that improves users' understanding of data being presented. This need to be explored further, any visualisation should apply new insights to better communicate information to non-technical users. Together this interactivity and design work will fulfil our goal of providing situational awareness and information sharing in a number of domains and levels for users who are not experts in data analysis.

The backend plugins that make up the processing path will, by design, provide a lot of flexibility. This give users the ability to integrate a variety of data sources, another of the goals discussed earlier. One can envisage a plugin that is simple wrapper to download data from a commercial service offering Twitter data and then placing it in the processing path for further, local, filtering and analysis. Or one could connect to a external service that already analyses posts from a social media service and insert the summarised data directly to the presentation plugin for end user display. Finally, one may also have a complete, locally developed, processing path that would handle everything from watching a social media service, through data acquisition, filtering, analysing and presenting the data from that, and other if need be, services.

Another possible use could be to facilitate responses (if desirable) to online influence operations. The plugin architecture can therefore facilitate engagement on social media and measuring certain types of impact of such engagement, for instance posting of messages. This is done through response plugins that can perform semi- or fully automated posting of messages. Responses are initiated by users from the front end, if a chart collates data from five social media services, then five response plugins need to be created and assigned to the processing path. A response plugin will know how to log in and post a message using APIs provided by the service, alternatively posting form data. A set of configuration data will provide login information for the social media service in question.

The ability to track how responses have been received helps us determine if there are any measurable effect in terms of basic social media interactions, for instance likes, replies or forwarding operations (such as re-tweets). Such quantitative impact analyses can use the same processing path as the messages that are being responded to. The only difference is that the data collection starts with the message(s) posted by the users, and ignore other messages.

The plugin architecture outlined here would ensure that one can re-use and re-purpose data from new and existing sources without imposing predefined expectations as to what these sources should look like.

Processing paths explored and explained

A processing path is, at the core, a simple set of configuration data that describes which plugins to use in what order, how to invoke them and finally what parameters to pass them, and how. The backend framework will load this configuration information when it starts and then initialise each plugin in order. The last thing a plugin does in the initialisation phase is to establish a subscription to the previous plugin in the processing path and then wait for incoming data.

A processing path can in principle be as simple as having a small backend plugin that requests the current top 10 hashtags from Twitter, and then a pie chart frontend (presentation) plugin that shows how many messages use each hashtag. It can also be very complex and integrate historical data from many years back with previously analysed data that is translated and assigned weighting based on sentiment analysis.

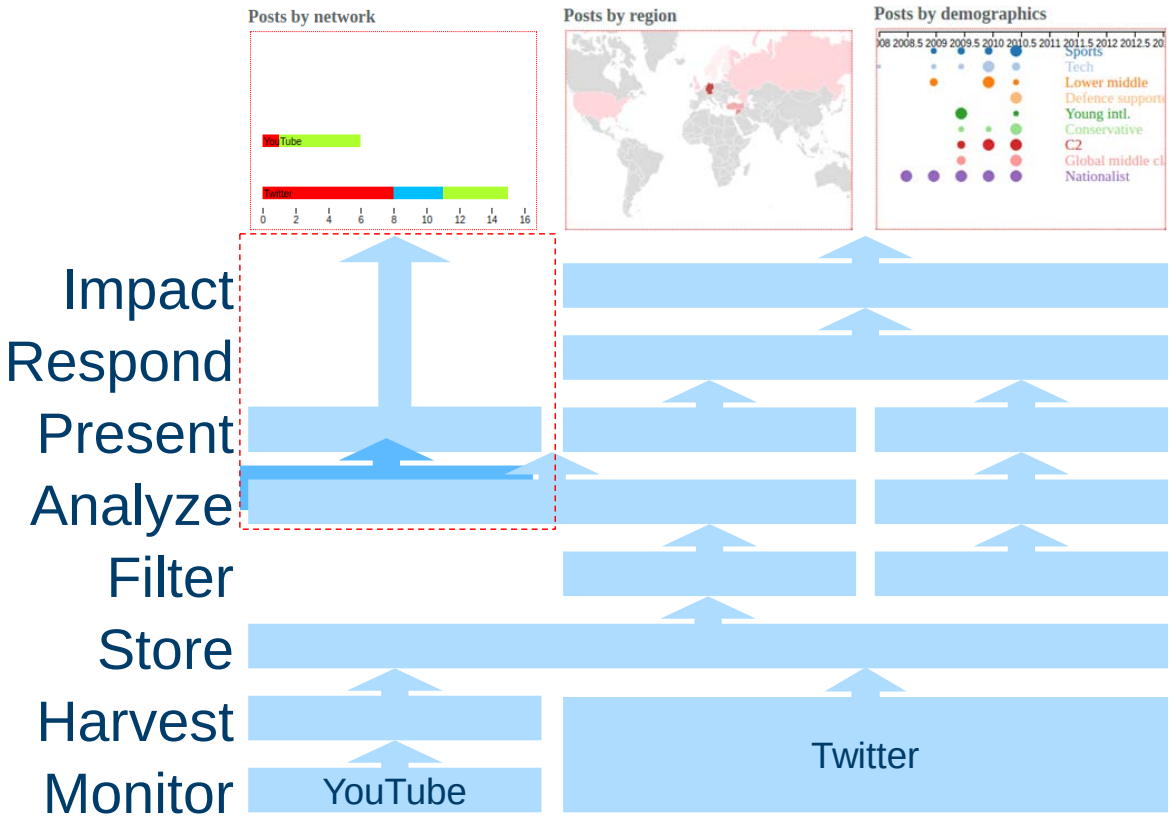

Figure 2 Processing path and plugins[2]

Table 2 explains the different steps (1 to 10) that make up a processing path and gave examples of what a plugin might do in each step. In Figure 2 above we see a visual representation of three different presentation plugins and the different processing paths that lead up to the presentation. (Refer to Figure A. 1 in Appendix A: Processing path technical aspects for a simplified technical

[2] The charts in this example are for illustration only. These charts are based on random data and do not in any way reflect actual data nor any current priorities at the Norwegian Defence Research Establishment.

overview.) This figure also illustrates the potential for sharing plugins, using multiple plugins of the same type, combining plugins as well as how a processing path may skip certain steps. Two of these paths are based on Twitter data and one on combined YouTube and Twitter data. The Twitter processing paths do not use a separate watch plugin (step 1), unlike the YouTube path. The combined YouTube and Twitter processing path does not provide response or impact plugins (optional steps 9 and 10), however these are present in the Twitter paths. Furthermore, the YouTube/Twitter processing path adds a secondary analysis task (step 5) to complement (rather than replace) the first one. The data storage task (step 3) uses a shared plugin for all paths; this can for example save raw JSON data structures downloaded.

A particular strength of the plugin and processing path approach arises from the ability to apply a plugin task type more than once in a processing path. The filtering task (step 4) can be used both to transform data between different formats and to run queries against the raw data. The former use can be renaming of data fields (country to state, user to name for example) to (re-)use analysis (step 5) and presentation plugins (step 6) that expect a specific data structure. The latter could be used before running an analysis (step 5), for instance to reduce the data to only show messages originating from a certain country. It is also envisioned that filter plugins can act as triggers, so if one has a complex machine learning task that benefits from more data, a filter could wait for a certain amount of data to arrive before invoking the next plugin in the path. Similarly, multiple analysis tasks (step 5) can be run in sequence. If, for example, there already exists a machine learning plugin that sorts messages based on sentiment, then rather than rewriting this plugin it can be combined with a geo-location analysis plugin to see where different sentiments emerge.

In terms of collaboration, if a future solution was shared between defence and police (for example), in a crisis situation police could be given access to summarised SOCMINT information useful to their handling of emergency situations when working with defence staff. Military intelligence could look at the same data but using a different presentation plugin that provided a lot more detail. Thus the plugin approach also handles organisational “need to know” issues.

When it comes to sharing plugins across different parts of an organisation, or between organisations, this paper suggests that there are three levels of information safety to consider. First there are plugins that are so generic, e.g. displaying a bar chart, that they can be shared without concern. Secondly, there are plugins that will have non-generic use, but the core code is generic, such as counting messages based on keywords. The sensitive parts would be in the use and/or configuration of the plugin. Finally, there may be plugins that are difficult to share; either because code embodies advanced intellectual property (for instance custom artificial intelligence tools) or the logic embedded in the code would reveal what the user is interested in assessing. Most of the time however, it would be the complete processing path that would be most sensitive. The architecture suggested here would allow sensitive operations to share certain plugins, while keeping others running on separate networks to protect the data. In this way the architecture would also facilitate open data to move into closed domains, a common requirement in defence [24].

When it comes to the need for rapid responses to changing situations in social media for instance if an attacker starts posting messages on a new topic, or a new social media service is launched and needs to be assessed, the processing path approach facilitates this. One would move away from a way of working where someone, possibly a decision maker, requests assistance from a technical

expert such as a database administrator to obtain a dataset to analyse, an approach that introduces a time lag. The processing path approach resolves this by letting technical experts create self-contained plugins that are then put together to a path from a simple GUI by users such as the decision maker in our example here. Control over data, in terms of interrogating it at a lower level, is handed over to the end user; this will cut down on the time it takes to go from discovering the need for an analysis until it is done. It may also enhance users' trust in the digital tools they are using, which in turn can lead to more use of the tool and better information sharing [25, p. 12].

End users		**Tech. experts**
Request data Make decisions	Example of common division of responsibility 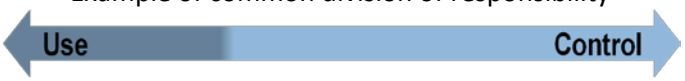	Develop software Setup system Design/run queries
Setup system Design/run queries Make decisions	Alternative division with processing path 	Develop plugins

Table 3 Visual summary of who does what with social media data

The simple and pragmatic architecture of processing paths can also help smaller states with fewer resources available for social media analysis than larger actors. It does this by facilitating simple re-use of anything from open source data collection software to custom analysis tools. This also allows us to unite different parts of the organisational and technical patchwork that was highlighted as another issue. Different part of the organisation can share or retrieve data at different stages in the processing path(s) and can customise these paths to provide a situational awareness suitable to their local/specific requirements.

Data structures for privacy, storage and exchange

Social media messages may have certain data fields one can assume are always present, such as date posted: if it is a new message or a response, country of origin and so on. However, this does not mean that a single, universal data structure will suffice for all (inter)actions.

Instead three *types* of data structures, all of them using JSON structures as documents, are envisaged. The presentation plugins can use a standardised set of fields for display and user interactions. This could be topic, which social media it comes from, date, country, etc. Additional fields that the presentation plugin requires are managed internally, but not exposed to the user. Secondly, there will be the original, "raw" data format that is stored, possibly short term, after being collected. Finally there should be a long term, cleaned and anonymised storage format. This will have passed through a filter plugin that initially anonymises the data by randomizing account names, but linking them through an externally held lookup table. This is a common method for anonymising survey/interview data. Over time it could be interesting to investigate the possibility of storing summarised content from multiple messages with metadata merged to anonymise the original messages completely. This would ensure that the data is suitable for long term storage for historical queries yet fulfils our goal of preserving social media users' privacy as per the earlier discussion about privacy issues.

Conclusion

This paper has outlined a possible architectural solution to what has been referred to as “Information Environment Assessment” [13]. This is the examination of social media in a strategic communication context to see if another state may be conducting influence operations related to activities, for example a military exercise. The suggested solution was developed as an activity in connection with research on social media based influence operations at the Norwegian Defence Research Establishment [2] that used a software demonstrator to present and discuss possible solutions with staff in different parts of the Norwegian Armed Forces. Several goals and issues related to such a task were identified. The goals were defined as flexible situational awareness across domains and levels for non-experts with the ability to respond rapidly to changing situations while saving resources and re-use (existing) solutions. Issues included the organisational and technical patchwork nature of defence organisations, the need to cater for heterogeneous data sources and formats and last but not least the need to handle individuals’ privacy concerns.

The proposed solution is a framework for discrete SOCMINT tools, called plugins; rather than a single, monolithic tool to replace other tools. Different types of plugins are joined together in what is called a processing path. A path can consist of, for instance, a data collection plugin which sends data to a filter plugin to only select data from a particular region, this in turn hands data to an analysis plugin which can perform a sentiment analysis, the result of which is passed on to a presentation plugin to show the result in an easily understandable visualisation.

The processing path/plugin architecture described here is akin to service oriented architectures. In SOA terms a plugin can be both service provider and requester. There are some differences however. The framework will be very flexible with regard to data structures offered and structures required between the plugins. These differences are resolved through the use of filtering plugins. There is also a tighter control of the processing path than in a more conventional SOA setup. Data always flows in one direction and there is no central registry offering services, but rather a central conductor that emerges from the framework and processing path configuration settings. Furthermore, not everything will be or have a web service interface. It may be that plugins access data dumps from external system or call executables on the server. It is envisaged that if a working solution is developed it can be used for both real data and to replay datasets in workshop/training situations.

References

[1] S. C. Woolley and P. N. Howard, *Computational propaganda worldwide: Executive summary*. Oxford: Oxford Internet Institute, University of Oxford, 2017.

[2] A. Bergh, ‘Social network centric warfare: Understanding influence operations in social media’, FFI, Kjeller, Norway, FFI Rapport 19/01194, 2019.

[3] J. Schulz, ‘How Much Data is Created on the Internet Each Day?’, *Micro Focus Blog*. [Online]. Available: https://blog.microfocus.com/how-much-data-is-created-on-the-internet-each-day/. [Accessed: 11-Jan-2018].

[4] W. Hutchinson and M. Warren, 'Principles of information warfare', *J. Inf. Warf.*, vol. 1, no. 1, pp. 1–6, 2001.

[5] S. Abrams, 'Beyond Propaganda: Soviet Active Measures in Putin's Russia', *Connections*, vol. 15, no. 1, pp. 5–31, 2016.

[6] K. Giles, 'Russia's hybrid Warfare: A success in propaganda', *Bundesakademie Für Sicherheitspolitik Work. Pap.*, no. 1, p. 2015, 2015.

[7] B. Perry, 'Non-linear warfare in Ukraine: The critical role of information operations and special operations', *Small Wars J.*, vol. 14, pp. 1–30, 2015.

[8] Intelligence Community, 'Assessing Russian Activities and Intentions in Recent US Elections', Washington, DC, US, 2017.

[9] S. Shane and M. Mazzetti, 'Inside a 3-Year Russian Campaign to Influence U.S. Voters', *The New York Times*, 16-Feb-2018.

[10] 'How a U.S. team uses Facebook, guerrilla marketing to peel off potential ISIS recruits', *Washington Post*. [Online]. Available: https://www.washingtonpost.com/world/national-security/bait-and-flip-us-team-uses-facebook-guerrilla-marketing-to-peel-off-potential-isis-recruits/2017/02/03/431e19ba-e4e4-11e6-a547-5fb9411d332c_story.html. [Accessed: 07-Feb-2017].

[11] T. Franke, 'NATO Review - Social Media: The Frontline of Cyber Defence?' [Online]. Available: https://www.nato.int/docu/review/2011/Social_Medias/cyber-defense-social-media/EN/index.htm. [Accessed: 04-Jul-2018].

[12] NATO, 'Cyber defence', *NATO*. [Online]. Available: http://www.nato.int/cps/en/natohq/topics_78170.htm. [Accessed: 22-May-2017].

[13] R. Goolsby and K. M. Carley, 'Understanding the information environment assessment for Trident Juncture 18', 2019.

[14] D. Omand, J. Bartlett, and C. Miller, 'Introducing social media intelligence (SOCMINT)', *Intell. Natl. Secur.*, vol. 27, no. 6, pp. 801–823, 2012.

[15] M. Zembik, 'Social media as a source of knowledge for customers and enterprises', *Online J. Appl. Knowl. Manag.*, vol. 2, no. 2, pp. 132–148, 2014.

[16] D. O'Sullivan, C. Devine, and D. Griffin, 'Obama official: We could have stopped Russian trolls', *CNN*, 26-Mar-2018. [Online]. Available: https://www.cnn.com/2018/03/26/politics/brett-bruen-russian-meddling-election/index.html. [Accessed: 27-Mar-2018].

[17] C. F. Cellan-Jones Rory, 'Twitter axes Vine video service', *BBC News*, 27-Oct-2016.

[18] O.-E. Hedenstad, A. Eggen, R. Rasmussen, H. Hafnor, and T. Gagnes, 'Sluttrapport for prosjekt 898 NbF Beslutningsstøtte', FFI, Kjeller, Norway, FFI Rapport, 2007.

[19] T. H. Bloebaum, J. E. Hannay, O.-E. Hedenstad, S. Haavik, and F. Lillevold, 'Architecture for the Norwegian defence information infrastructure (INI) – remarks on the C3 Classification Taxonomy', FFI, Kjeller, Norway, FFI Rapport, 2013.

[20] R. Rasmussen *et al.*, 'Sluttrapport for FFI-prosjekt Tjenesteorientering og semantisk interoperabilitet i INI', FFI, Kjeller, Norway, 2013/00932, 2013.

[21] A. Bergh, 'Distributing the disruption', in *2015 International Conference on Military Communications and Information Systems (ICMCIS)*, 2015, pp. 1–6.

[22] A. Bergh, 'The final destination: Building test bed apps for disconnected, intermittent, low-bandwidth environments', presented at the 2015 IEEE 81st Vehicular Technology Conference: VTC2015-Spring, Glasgow, Scotland, 2015.

[23] C. Fishwick, 'How a Polish student's website became an Isis propaganda tool', *The Guardian*, 15-Aug-2014.

[24] R. Haakseth, 'SOA Pilot 2011–demonstrating secure exchange of information between security domains', *FFI-Rapp.*, vol. 117, p. 2012, 2012.

[25] A. Bergh, 'Seeing is believing; hearing is understanding: Building real trust through virtual tools', presented at the 20th Iinternational Command and Control Research and Technology Symposium 2015, Annapolis, USA, 2015.

Appendix A: Processing path technical aspects

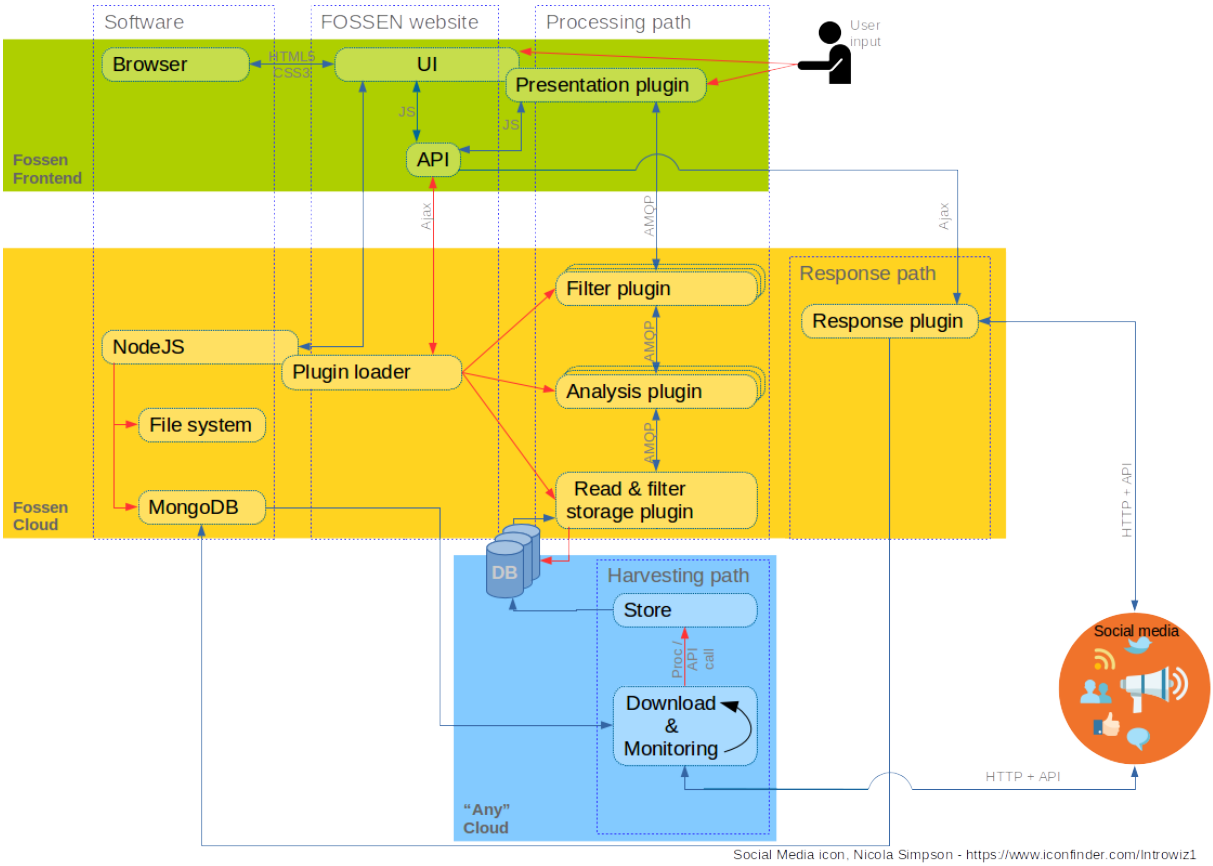

Figure A. 1 Processing path technical details