FFI Forsvarets forskningsinstitutt

24/01348

FFI-RAPPORT

Guide to social media training with Somulator

– 2024

Arild Bergh

Guide to social media training with Somulator

– 2024

Arild Bergh

Norwegian Defence Research Establishment (FFI) 9 September 2024

Keywords
Simulering og simulatorer
Sosiale medier
Kompetanseutvikling
Påvirkningsoperasjoner
Kognitiv krigføring

FFI report
24/01348

Project number
1692

Electronic ISBN
978-82-464-3548-0

Approvers
Stig Rune Sellevåg, *Research Manager*
Morten Aronsen, *Deputy Research Director*

The document is electronically approved and therefore has no handwritten signature.

Summary

How can the online information environment be handled in a crisis? Somulator is a training tool that can help to improve individuals', groups' and organisations' preparedness.

States and extreme political groups have exploited social media in subversive operations over the past decade. This has become a serious problem for democracies. Operations occur as well before as during armed conflicts, such as in Ukraine. The same applies to crises, such as the covid-19 pandemic. The influence has led to general misinformation, from conspiracy theories to personal attacks on politicians.

The development has created a need for organized training in how to handle the online information environment in a crisis. The Norwegian Defence Research Establishment (FFI) has developed a tool called Somulator. It can facilitate this type of training.

Somulator simulates a complete information environment in social media. It provides a realistic training situation. Somulator uses a variety of open-source web applications. It copies well-known platforms such as Facebook, Twitter[1], Instagram and YouTube, as well as online news sites. Somulator is designed to be extensible, in the sense that other social media can be added over time. These platforms are controlled through purpose-built training tools. They facilitate the automated distribution of large amounts of content and preparation. This is followed by the implementation and summary of training sessions and exercises.

This report discusses the requirements and the various use cases that emerged during the initial planning phase. Next, we examine how Somulator can be compared to other solutions. We summarize the current features of Somulator. These sections of the report can be used by potential users to assess whether Somulator fits their training needs.

The core of the report explains how content in social media is central to achieving the training objectives. We provide clear examples of how to develop and distribute content for simple and advanced training cases. The general workflow of a Somulator training session is explained, and the various stages of development, production and distribution are discussed. The report also describes the roles of participants and exercise supervisors (Excon) when it comes to access to and use of Somulator.

[1] Twitter was rebranded «X» in July 2023, however this report will use the original name.

Sammendrag

Hvordan kan det digitale informasjonsmiljøet håndteres i en krise? Somulator er et treningsverktøy som kan bidra til å forbedre beredskapen hos enkeltpersoner, grupper og organisasjoner.

Både stater og ekstreme politiske grupper har utnyttet sosiale medier i påvirkningsoperasjoner det siste tiåret. Dette er blitt et alvorlig problem for demokratier. Påvirkning skjer både i forkant av og under væpnede konflikter, som i Ukraina. Det samme gjelder under kriser, som covid-19-pandemien. Påvirkningen har ført til generell feilinformasjon, fra konspirasjonsteorier til personangrep på politikere.

Utviklingen har skapt et behov for organisert opplæring i hvordan en håndterer informasjons-miljøet på nettet i en krisesituasjon. Forsvarets forskningsinstitutt (FFI) har utarbeidet et verktøy kalt Somulator. Det skal legge til rette for denne typen opplæring.

Somulator simulerer et komplett informasjonsmiljø i sosiale medier. Det legger vekt på å gi en realistisk opplæringssituasjon. Somulator bruker en rekke nettapplikasjoner med åpen kilde-kode. Det kopierer kjente plattformer som Facebook, Twitter[2], Instagram og YouTube, samt nyhetssider på nettet. Somulator er designet for å kunne utvides, slik at flere sosiale medier kan legges til etter hvert. Disse plattformene styres gjennom spesialbygde opplæringsverktøy. De legger til rette for automatisert distribusjon av store mengder innhold og forberedelse. Deretter følger gjennomføring og oppsummering av opplæringsøkter og øvelser.

Denne rapporten diskuterer kravene og de ulike brukstilfellene som dukket opp i den innledende planleggingsfasen. Deretter undersøker vi hvordan Somulator kan sammenlignes med andre løsninger. De nåværende funksjonene i Somulator oppsummeres. Disse avsnittene i rapporten kan brukes av potensielle brukere, for å vurdere om Somulator passer til deres opplærings-behov.

Kjernen i rapporten forklarer hvordan innhold i sosiale medier er sentralt for å nå opplærings-målene. Vi gir klare eksempler på hvordan man kan utvikle og distribuere innhold for enkle og avanserte opplæringstilfeller. Den generelle arbeidsflyten i en Somulator-opplæringsøkt forklares, og de ulike stadiene i utvikling, produksjon og distribusjon diskuteres. Rapporten beskriver også rollene til deltakere og øvingsansvarlige (Excon) når det gjelder tilgang til og bruk av Somulator.

[2] Twitter ble rebranded «X» i juli 2023, men denne rapporten vil bruke det opprinnelige navnet.

Contents

Summary 3

Sammendrag 4

Contents 5

1 Introduction 7
- 1.1 Audience 7
- 1.2 Structure 7
- 1.3 Definitions 8

2 Needs assessment 10
- 2.1 Why do we need social media training? 10
- 2.2 What is unique about social media training 10
- 2.3 Lessons learned from other solutions 11
- 2.4 Input from potential users and core goals that emerged 12
  - 2.4.1 For trainers or training event organisers 12
  - 2.4.2 For researchers 13
  - 2.4.3 For organisations 14

3 Somulator features and use cases 15
- 3.1 Features summary 15
- 3.2 The Excon module: Controlling the social media training 15
- 3.3 Use cases 16
- 3.4 Current state and extendibility 17

4 Operationalising training goals through social media content 18
- 4.1 Model and training cases 18
  - 4.1.1 Simplified model of real-life influence operation 18
  - 4.1.2 Step 1: Information enters the social media clones 19
  - 4.1.3 Step 2: Information gains attention 20
  - 4.1.4 Step 3 and 4: Social media users interact with information 21
  - 4.1.5 Step 5 and 6: Information is spread widely 22

5 Somulator workflow: Development, production and deployment 23

5.1 Development: Operationalising a scenario through topics and distribution channels 23
5.2 Production: Create relevant content and accounts to support scenario events 24
5.3 Deployment: Briefing + content + channels + timing 24

6 Production, Deployment and Somulator management 26
6.1 Production 26
6.1.1 Create content 26
6.1.2 Create distribution channels – Social media clones 28
6.1.3 Create distribution channels – News sites 28
6.2 Deployment 29
6.2.1 Post content through relevant distribution channels 29

7 Conclusion and the way forward 30

Appendix 31

A Example of input from development stage 32

B Example of information for participants 33

C Extending Somulator 34

References 35

1 Introduction[3]

Recent years have shown an increase in the use of social media and other online information channels to try to influence people without them being aware of who is behind such influence operations. Such propaganda and disinformation tactics have been used by states as well as fringe political groups. It has become increasingly clear that this is a problem for the overall security of democracies, in particular during the handling of a crisis, as we saw during the recent Covid-19 pandemic.

To improve our handling of such influence efforts we need to train in a supportive yet convincing environment. Somulator (an abbreviation for Social Media simulator) mimics mainstream social media platforms to provide a realistic training experience, following the "train as you fight" concept. In Somulator, one can safely share custom created content to emulate propaganda, influence efforts, disinformation and general misinformation as experienced in real-world social media.

This report is an updated version of the 2023 report "*Guide to social media training with Somulator*". Part of that report has been moved to the new FFI report "*Creating content for social media training in Somulator*" which provides an in-depth guide on how to create relevant and realistic content for training purposes. That report can be found here: https://ffi.no/reports.

1.1 Audience

This report is intended for those who need to provide, or are interested in exploring, the use of the Somulator training tool to strengthen the defence against, or mitigating the impact of influence attempts, through training. Typical users include those whose work is sensitive to online misinformation, disinformation or propaganda. This includes defence personnel, journalists, personnel in local or state governmental positions who have to handle crises issues and educators.

1.2 Structure

This report has two parts. The first, in chapters 1 to 4, is aimed at those who require an overall understanding of how social media training can be useful and how it can be done through Somulator. Chapter 4 bridges the strategic and practical part of this report. Readers who will handle the practical aspects of social media training can therefore focus on chapters 4 to 6.

[3] The introduction in chapter 2 and chapter 3.3 has previously been published in «Somulator: An information training environment to train as you Tweet» (Bergh, 2022).

The current chapter will define terms used in this report before it discusses the background for social media-based influence efforts and how social media differs from other types of media training. In chapter 2 the background to and current state of Somulator will be examined, followed by a summary of features and possible use cases that will also highlight the flexibility of the Somulator tool in chapter 3.

Chapter 4 takes the reader through a high-level workflow that shows the relationship between narratives, content and the various simulated social media profiles (known as distribution channels) on the one hand and Excon staff, administrators and users of Somulator on the other.

Chapter 5 provides information on the overall Somulator workflow and chapter 6 has an overview over Somulator features for those who will administer social media training.

1.3 Definitions

Training is used as an overarching term for several different activities. In this document it refers to individual/group-based training (for example, a workshop), more focused educational settings such as regular classes and larger (total) defence type exercises where previous training may be tested.

***Exercise controllers (Excon** for short) refers to personnel* controlling Somulator during the training (according to an overall plan). Other terms for this are white cell, White Team (WT), DISTAFF or gaming staff.

Web application (or service) refers to software accessed through a web browser, similar to how Google mail works for instance.

Clone/Social media clone is used to denote a web application that copies features of a mainstream social media platform, for instance Twitter.

Participant is a person who participates in the training. In the context of Somulator, this person will have access to an account created for or by them on one or more of Somulator's four social media clones (see **Content** below). They will also have access to one or more news sites if they have been set up. In a military context this is often referred to as Blue Team (BT).

Content is the term used for news articles and social media posts that are created for use in a training situation. A post can be anything from a tweet in a few words for use in the Twitter clone (see **Clone** above), to a longer article with videos for the Somulator news module.

An **account** (also known as a profile) is used in regular social media to log in to a service and then post content and interact with other users, for example by liking other users' posts. Actors who engage in influence operations often create fake profiles to reinforce a message, either by providing many likes/forwards of posts or by providing the profile with characteristics and expertise that will appeal to the target audience.

In Somulator, there are two types of accounts, although to the participants they are indistinguishable. The first type is the accounts used by participants to log in to one of Somulator's social media clones, simply referred to as "accounts" in this document. Participants use these accounts just as they would do on regular social media platforms like Twitter. These accounts would typically connect (by following, being a friend, subscribing, etc.) to a second type of account that is instrumental to the use of Somulator.

The other account type is created and controlled by the Excon staff. For clarity, they are referred to as **distribution channels** in this document. These channels represent the "Red Team" in a military context, in other words the enemy. For instance, if an account called "Extremist_Supporter_99" was created on the YouTube clone (see Clone below) to post videos glorifying terrorism, then this would be one such distribution channel. A newspaper emulating a state-controlled propaganda outlet would be another distribution channel. Distribution channels are used to post content made for participants to read during the training, while the participants use accounts to access and post content in the social media clones, including content from the distribution channels.

Propaganda, disinformation, misinformation, information operation and influence operation: These terms are often used interchangeably. In this report, **propaganda** refers to biased information used to propagate a particular view. **Disinformation** is information know to be false by the actor spreading it whereas **misinformation** is false information believed to be true.

An **influence operation** is a clearly defined effort to influence people or groups through propaganda which in turn may consist of both dis- and misinformation. Information operation is a broader term that is not used in this document.

Algorithm: An algorithm, in computer terms, is a process for solving a (often repeated) problem, following fixed steps. To detect spam email, for example, a number of different algorithms are applied. The term has become known to the general public in relation to algorithms that make recommendations on social media, such as selecting news they think you will spend time reading based on previously read stories.

Extendibility is the ability to enhance software with additional features. It is used in this report when discussing the open nature of a web application, indicating that it is not a static element, but can be enhanced with further functionality.

Finally, the expression **Excon tools** will be used about the web application developed by the Norwegian Defence Research Establishment that Excon staff use to manage Somulator training sessions. This covers tasks such as adding accounts and distribution channels or posting content to the different social media clones through the distribution channels that have been set up.

2 Needs assessment

Prior to developing Somulator, FFI undertook an assessment of the requirements for a social media training tool. This process looked into whether social media training was required at all and if so, how such training would differ from other forms of training. It then examined existing solutions before discussing a possible solution with potential users to learn more about their requirements.

2.1 Why do we need social media training?

The spread of dis- and misinformation through social media during crises has risen dramatically in the past decade. For example, YouTube has repeatedly promoted conspiracy theories in the aftermath of mass shootings in the USA (Warzel, 2017). In the same period, non-democratic states' use of social media to spread propaganda and disinformation has changed from being an intelligence or communications matter to a societal issue that can affect a range of actors who handle a crisis. For instance, during the Covid-19 pandemic nurses had to handle patients who believed false information that was spread, claiming Covid-19 vaccines were dangerous (Cohen, 2020). These sorts of falsehoods can affect many actors within the Norwegian Total Defence sector. The armed forces as well as hospitals, lawmakers, police and a range of other organisations can be impacted. Furthermore, social media is now deeply embedded in our daily lives. Ignoring what happens there is not really an option. For many groups the need for training on how to understand and handle "information events" on social media is now a requirement and not an optional extra.

2.2 What is unique about social media training

Misinformation, disinformation and propaganda through social media has the potential to affect anyone's situational awareness during a crisis. Unlike other media in the past, there are no central ways of stopping the flow of information, nor is there any meaningful vetting of information that cannot be easily circumvented. Those who are tasked with handling a crisis need to prepare so they are better able to distinguish between what is important or not and what is right or wrong in the vast and never-ending stream of social media posts. These skills should help them counter the effects of dis- or misinformation spread through social media. The exact nature of these skills will vary depending on the tasks they perform; however such skills need to be acquired through training before a crisis occurs.

It is possible to create social media accounts on real social media platforms without costs and use them for training. However, this would immediately run into a number of issues. Firstly, many social media platforms such as Twitter show uploaded content to everyone. This would restrict one's ability to train freely on sensitive scenarios. The ability to fail and learn in a safe environment is important. The use of fake accounts is also prohibited by most social media platform's terms of use. This approach would therefore have to break platform rules and face the likelihood of losing the accounts created and the content used if the social media platforms

discover this use. Facebook alone deletes several billion fake accounts per year (Leprince-Ringuet, 2020). To use the real "location" for training, as one can do with other types of crisis training, is therefore out of the question.

An alternative could simply be to describe a scenario, such as "Fake news about an impending attack by a terrorist group is spread through Twitter". Although this approach is often used in high-level war-gaming this would leave a lot to be desired when training practitioners. In some ways it would be akin to practice gun handling and target practice through notes – social media content is chaotic and overwhelming. Using a more realistic setting can contribute to a better understating of the issues at stake when participating in training.

An intermediate solution is to use simplified and/or *gamified*[4] versions of social media. This has been done several times to teach media literacy and help people spot disinformation (Maekawa et al., 2021; Roozenbeek & van der Linden, 2019). This approach can be very useful for simpler educational purposes, such as teaching about fake news or research that isolates one particular element of social media behaviour. However, given the complexities of real-life social media, this approach may remove functionality that play an important role in propagating rumours for instance.

Given these considerations, Somulator was developed to provide full social media simulation capabilities. The next chapter will explain in detail what Somulator is, and what it provides.

2.3 Lessons learned from other solutions

The impetus to develop Somulator came from the use of a simpler social media simulation in the large-scale NATO exercises Trident Juncture 2015 and 2018 (Paxton, 2018; Tomlin, 2016) as well as several simpler social media simulations used in research as mentioned above. An evaluation of these tools showed that they did not fulfil the requirements for a realistic social media simulation discussed in chapters 2.1 and 2.2. However, to start from scratch and duplicate even a subset of a single social media platform's features would be costly and require ongoing development to fix bugs and add features. Furthermore, emulating just a single platform is not necessarily useful. Different platforms have different affordances that affect how we perceive and spread information (Bergh, 2019, p. 17). Providing Facebook training does not necessarily help someone to handle a so-called Twitterstorm.

These factors led to the decision to use open-source software that had been developed as alternatives to different commercial social media platforms. Open-source solutions were selected because it allows modifications of the software. This ensured that Somulator would fit the identified training needs.

[4] Gamification of something is to take a real life object and give rewards when something is done correctly, as in a computer game. For instance showing fake and non-fake news and give the user points for not forwarding fake news.

2.4 Input from potential users and core goals that emerged[5]

The decision to utilise existing open-source software that has been tested and is continuously developed was followed by an analysis of what customisations was required to turn these separate applications into a coherent training platform. As a part of this analysis Somulator was discussed with different potential stakeholders. These included experienced, regional organisers of media training, staff at NATO's Joint Warfare Centre, workshop and exercise organisers at national public bodies in Norway, staff at the Norwegian Cyber Defence and fellow FFI researchers.

Based on these discussions the following key requirements emerged for different types of users:

2.4.1 For trainers or training event organisers

2.4.1.1 Low threshold for use

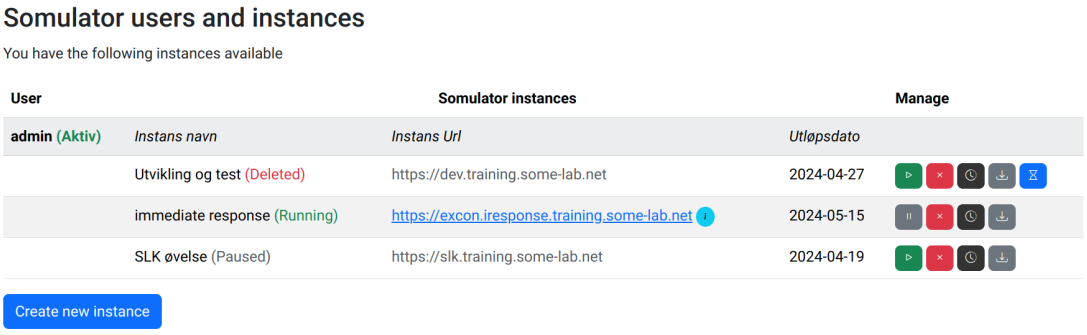

Figure 2.1 The Somulator management screen.

The type of training discussed here is often done by subject specialists that does not necessarily have in-depth IT expertise. A simple method to start using Somulator was therefore required. This requirement was handled by developing an automated means of creating, starting, pausing and stopping Somulator instances through a regular web page (*Figure 2.1*). Somulator is thus what is known as "software as a service". The ease with which Somulator can be deployed means that once a training session is completed Somulator can be deleted without worrying about the cost of deploying it again. With no start-up cost there is no need to pay to keep services and servers running when they are not used.

2.4.1.2 Easy to organise training

Informants with practical training experience highlighted the workload involved in setting up tools for participants to use, typically by creating accounts and emailing login information before the training event. To handle this the web-based Excon tool was developed for

[5] This section has previously been published in the NATO report Mitigating and Responding to Cognitive Warfare (Masakowski & Blatny, 2022, p. 11.1-11.7)

Somulator. This facilitates, among other things, automated account creation. The Excon module is discussed in detail in chapter 3.2.

2.4.1.3 Content control

Finally, trainers needed simple and efficient means by which to control how and when content is published through the different social media clones. It is the content that facilitates learning. Content publishing is therefore controlled in such a way that it can tie in with an overarching scenario. This can be done manually, by automatic means or a combination thereof.

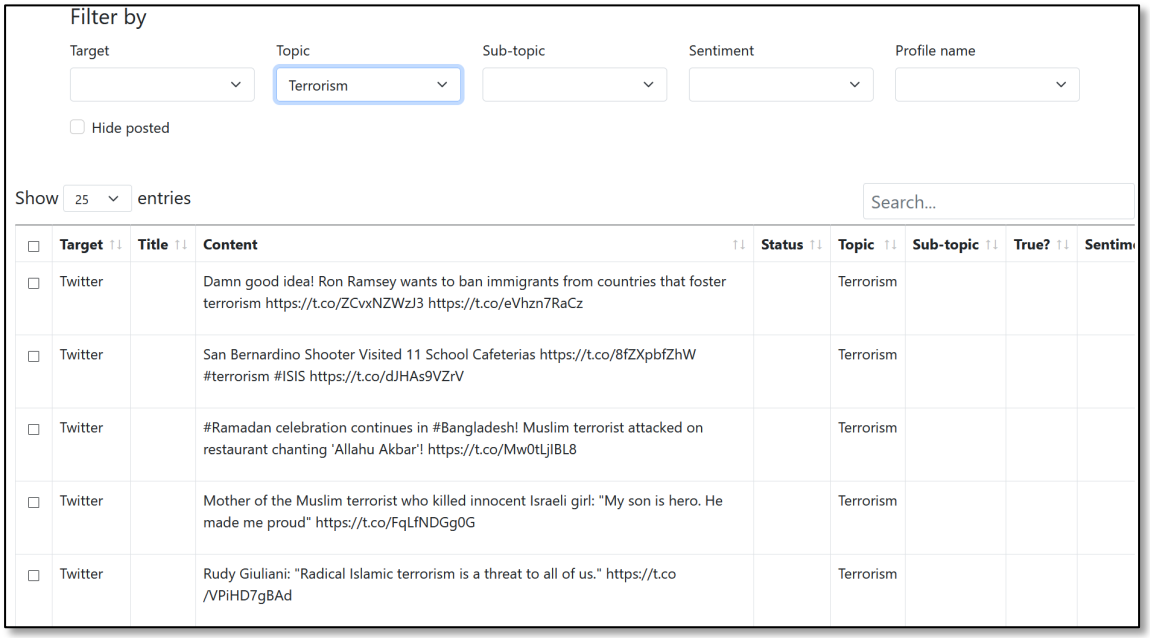

Figure 2.2 Prepared content ready for deployment in Somulator.

Somulator's Excon module facilitates the uploading and distribution of content through the different social media clones that have been deployed. The main purpose of this tool is 1) to be able to spread large amounts of content in a short time, as one experiences it on real social media and 2) to choose clones and distribution channels used for different types of posts. The latter is used because different distribution channels have different characteristics that can change participants' perception of the content being shared. A profile claiming to be a retired general may seem more authoritative on military matters than a homemaker's profile. Real influence campaigns operate in the same way by developing different false profiles (Yates, 2019).

2.4.2 For researchers

Somulator is also a tool for researchers to learn more about how influence operations through social media work, and through this research contribute to the development of new training approaches. The features discussed above make it considerably easier to organise casual

experiments for different groups that require customised setups. This makes it easier to test ideas with different scenarios/groups.

In addition, if the participants have used the social media clones not only to read content, but also to interact by sharing and liking posts, etc. it is possible to extract the data from an experiment to analyse in retrospect. This way one could for instance examine how many times posts with fake news was shared compared to posts with truthful information.

2.4.3 For organisations

Finally, organisations tasked with handling crises have some overarching requirements. Although not clearly spelled out, they emerged during conversations with stakeholders that represented different total defence actors.

Firstly, it would be very costly if a training tool in such a dynamic arena as social media was static. The ability to modify the software is therefore of paramount importance. Somulator is extendable as it uses open-source software that can be changed to suit particular needs. Furthermore, these applications are actively developed by a large community of developers that add features over time, keeping the functionality up to date. The open nature of the underlying social media platforms also means that any custom enhancements that are implemented in the Excon tools can be shared with other organisations.

Secondly, training tools that can be shared with other organisations are beneficial in terms of lower costs and because the pool of personnel who know the software expands. Interoperability is a key feature to achieve this, and all the social media clones used in Somulator implement the ActivityPub protocol (Webber et al., 2018). This allows two different organisations using for example Mastodon (the Twitter clone), to connect these via the Internet if they do joint training at some point. The organisation could be two different defence educational entities or even armed forces from two different NATO member states.

3 Somulator features and use cases

3.1 Features summary

Somulator is composed of five individual web applications that each clone a real-life social media platform or news site. These web applications have been integrated to provide a complete social media training environment through the Excon module. This software module was developed by FFI to allow a single person or a small team to control the different social media clones and news sites from a custom, easy to use interface. The list below shows the platforms that are cloned with the real name of the web application in parentheses:

- Facebook clone (Friendica)
- Twitter clone (Mastodon, see screenshot with training content below)
- Instagram clone (Pixelfed)
- YouTube clone (Peertube)
- General news website (Ghost)

These web applications will be referred to as Facebook clone, Twitter clone, etc. throughout, except when discussing issues specific to the clone software.

All the above web applications are all available under an open-source license. This means that one does not necessarily have to use these applications as they are, they can be freely modified to fit different needs.

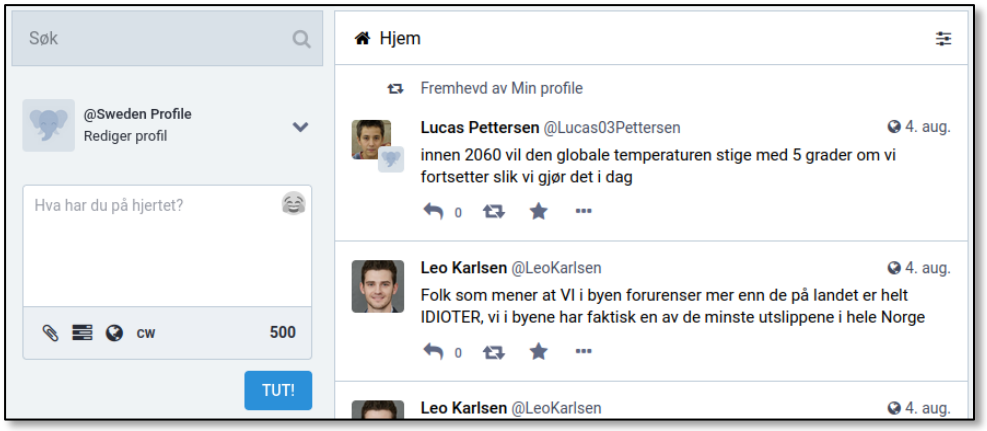

Figure 3.1 Mastodon, an open-source Twitter clone.

3.2 The Excon module: Controlling the social media training

The Excon module was developed to handle tasks that typically require IT experts, either from an internal IT department or through external consultants. These tasks include:

An easy-to-use installation process. It only takes a few minutes to start using Somulator by selecting which social media clones you want to use and provide a URL for your training tools. A few minutes later Somulator is accessible in a web browser.

Support for managing the practical aspects of training session. This includes adding unlimited users with a mouse click or uploading content to use in a training session.

Posting social media content through selected social media profiles (i.e., distribution channels) according to learning goals and an eventual scenario.

These Excon features lower the threshold for casual use in training sessions by focusing on simplicity. Below is a screenshot of the home page for the Excon management module.

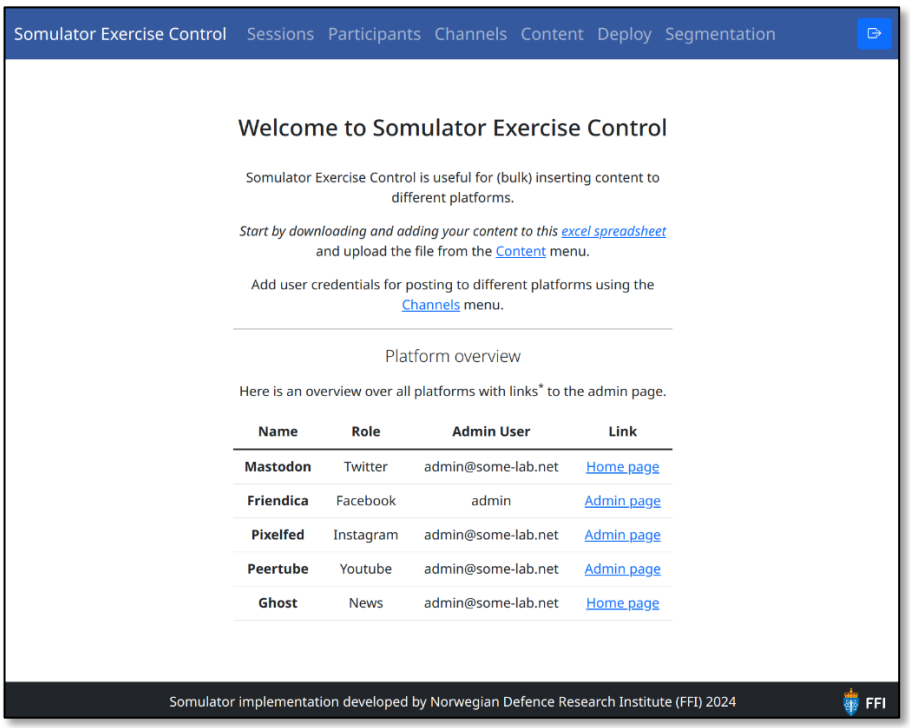

Figure 3.2 Somulator's Excon module, front page with menu at the top and links to individual social media clones.

3.3 Use cases

As mentioned, Somulator can be used both as a fully functional training tool and for research. In the latter case, it can be used interactively to test out ideas and research how to develop better training in the social media/disinformation field. It can also be used for general social media research, for instance to see how well fake news is detected in different social groups.

In terms of training, as Somulator is a tool and not a specific method, it is flexible enough to be used in a variety of training situations. In a large-scale crisis exercise it could be used over several days by anyone from media staff to operational planners to front line staff. Such an approach could, for instance, help to highlight issues relating to dialogue between the operational and tactical level in response to online dis- and misinformation. On the other hand, it could also be used in a small half-day workshop for media staff to learn how to monitor and handle targeted disinformation.

Training and research can also be combined by collecting social media data from an actual training session. This would entail analysing posts that were read and participants' responses to them as well as posts by participants. One could then use the findings for research or retrospective discussions to improve learning outcomes further. It is beyond the scope of this report to discuss this in detail, however the reference section lists several relevant papers and articles on this topic.

Whether it is for training or research, Somulator can be used with smaller or larger groups, on its own or as part of a larger exercise.

3.4 Current state and extendibility

At the time of writing Somulator is a fully functional solution providing the features described in this report. Somulator can be extended on different levels with matching levels of IT expertise required. The implication of this is that using Somulator need not be a static, unchanging experience. As an organisation develops its skill set and expand the scope of their training, Somulator can be enhanced to take on new roles.

See Appendix C for more information on how to extend Somulator.

4 Operationalising training goals through social media content

Whether Somulator is used for simple or advanced training, the key issue is that social media affects people through exposure to information. Social media information can be the content of posts (e.g., something that angers or pleases a person) or meta-information (for instance how many likes a post got). This has to be considered when planning a training or research session. As mentioned above, Somulator is a tool to be used in training, and not a training method. Thus, many skills can be developed using Somulator.

This section explains how Somulator can be used to simulate online influence activities. It also summarises issues that arise when using a training environment in comparison to real life social media. This is done by discussing a simplified model of how real-life influence efforts use social media and how the different steps of the model can be emulated in Somulator. Two use cases of social media simulation are discussed. One is a simple training session for a homogeneous group of participants; the other is an advanced case that uses a broader international politics scenario that could involve a range of actors in the training session.

4.1 Model and training cases

4.1.1 Simplified model of real-life influence operation

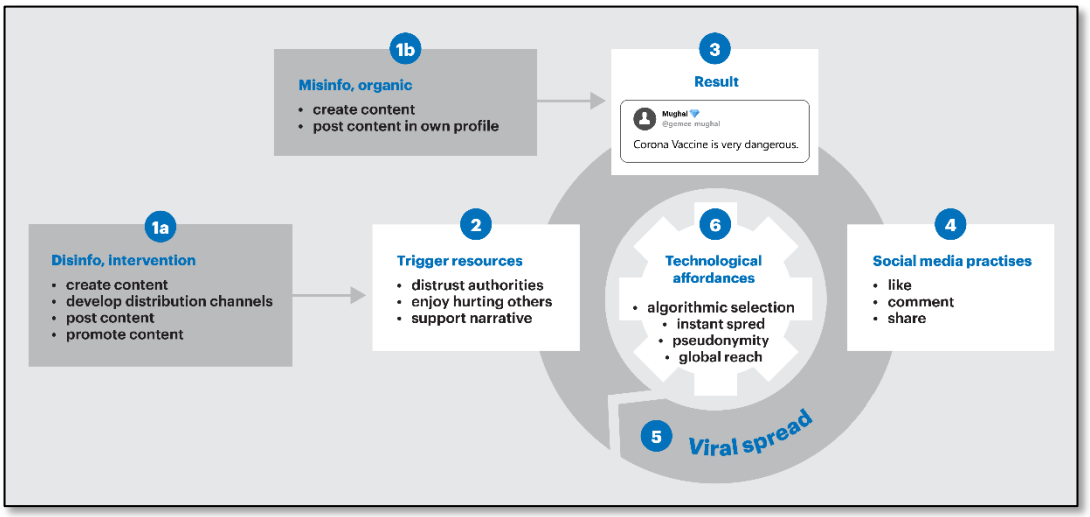

Figure 4.1 Social media influence operation model (from Bergh, 2020).

The model above is a synthesis of common approaches to online influence operation. It shows how deliberate disinformation and random misinformation enters into a loop of information distribution that combine user interactions and behaviours with technical facilities to spread

information to a large audience. Each of the steps will be explored below and the use of Somulator to emulate the step will be explained for the simple and advanced training cases.

4.1.1.1 Basic training use case

Target group: Emergency services staff.

Aim: Learn how to handle misinformation connected to specific crisis events.

Situation explored: A terrorist attack has occurred. For a few hours afterwards, misinformation causes panic and makes the work of the emergency services more difficult.

4.1.1.2 Advanced training use case[6]

Target group: Decision makers from different government departments and defence.

Aim: Learn how long-term influence operations try to affect decision making in democracies.

Situation explored: A large, autocratic country, Brutopia, wishes to influence the national election in a small, neighbouring country, Cardamommia. If the vote changes in just two electoral wards this will tip the balance of power from party X to party Y, a party that wants to cancel existing sanctions affecting Brutopia.

4.1.2 Step 1: Information enters the social media clones

Dis- and misinformation can both cause problems. The key difference in the model above is whether it is planned or not. Information emerging from **1a** will support the underlying goals of the influencing actor, matched to fit relevant target groups. This can be spread through paid influencers, adverts or fake social media profiles. If the problematic information is misinformation (**1b**), rumours or similar it is typically posted (or shared/liked) through one's own, genuine social media account.

4.1.2.1 Basic training use case

For basic training, one can create content relating to a few topics that exemplify typical rumours in a crisis situation and deliver this through a single social media clone. For example, the news site module could be used to publish news articles reporting on Twitter rumours. Or, if one wants to research the ability of participants to detect fake news, the Twitter clone can be used to deploy a few hundred tweets on relevant topics to test the participants. Either way, only parts of Somulator would be used and the content would be limited in scope. The content could focus on misinformation with some disinformation included, as even small, local events are sometimes used by external actors to spread false information (Evon & Mikkelson, 2017).

[6] The election scenario is used in this report because it is a well-known type of influence operation. However, Somulator can be used equally well for any imaginable scenario that involves manipulation through social media. More advanced training may be done stand alone or in conjunction with a larger exercise.

4.1.2.2 Advanced training use case

In the case of a more advanced training / exercise, it is useful to develop social media posts for different platforms that support the overall aims of the scenario. For the election scenario outlined above this content gives some examples of how training can be done in Somulator:

24 to 12 months before the election, content that point out that Cardamommia suffer more than Brutopia from the sanctions are posted and promoted on the two most popular social media platforms in Cardamommia. They are posted in business and labour related online groups through a range of accounts that claim to belong to international businesspeople and left-wing activists.

Brutopian controlled news sites feature interviews with Party Y representatives, inflating their importance and hailing them as opener of doors for business opportunities in Brutopia.

Several YouTube videos show up a few months before the election. They claim to document children suffering from hunger in Brutopia because of the sanctions. The videos are posted through official Brutopian news channels but are widely shared through comments to videos from human rights channels and referenced in tweets from bot accounts.

A few weeks before the election, conspiracy theories occur on private Facebook groups claiming that the leader of party X is personally benefitting from the sanctions and that he is secretly paid by one of Brutopia's adversaries. The claims are repeated across many social media platforms by different groups that are anti-authority or conspiracy oriented.

Two days before the election, e-mails are leaked through a Twitter account. It is claimed that they come from a hacked email account belonging to the local leaders of party X in one of the two wards Brutopia try to affect. They show that Party X will close down an important local factory if they are re-elected.

The above summary exemplifies how the social media aspect of a scenario is expressed through content that seeks to affect people without openly saying what it tries to achieve. The underlying narratives focus on the fact that innocent groups, such as businesses in Cardamommia and children in Brutopia, suffer from the sanctions.

4.1.3 Step 2: Information gains attention

Disinformation depends on triggering an interest in the groups being targeted in step 2. Without interest, no attention, without attention, no influence. Step 1a must therefore consider what topics and narratives would trigger interest when creating content that aligns with these triggers. During training, this step is concerned with matching content with the *roles* participants' take on in the training (see chapter 5.3 for more on roles in social media training). Content created in the previous step is sent out through different distribution channels that the participants access through standard social media tools such as subscriptions or sharing links.

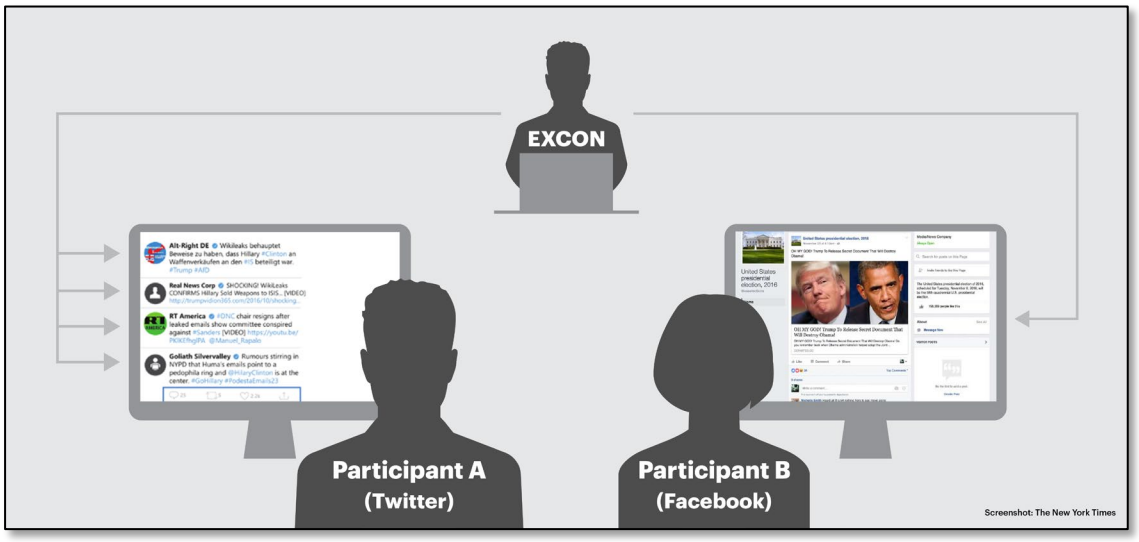

Figure 4.2 Excon staff manages the deployment of content to support the aims of the training. Somulator allows a single staffer to control a large number of accounts that are used as distribution channels. (Infographic: FFI/Grete Foss Alvestad)

4.1.3.1 Basic training use case

Somulator content can be sent out manually one post at a time or in batches, or automatically according to a schedule (see chapters 5.3 and 6.1). For the basic training scenario outlined above either method could be used. The main issue here is to brief the participants on the scenario and their roles so they can interpret appropriately to the content that is shared. As only one social media platform is used, it will be relatively simple for participants to follow the flow of information.

4.1.3.2 Advanced training use case

For the advanced training, automatic sharing of content is most realistic. This can be paused and restarted if there are delays or other issues during the training session. As multiple distribution channels will be used, it is important to connect different participants to relevant (and not necessarily all) profiles that will be used to distribute content. Again, there needs to be a clear connection between the participants training roles and the content they see.

4.1.4 Step 3 and 4: Social media users interact with information

In step 3 the first aim, online consumption of the information is achieved. Step 4 encapsulates the half-automatic actions that users do when using social media, such as sharing or liking a post. This contributes to a further spread of information from the influence operations.

4.1.4.1 Basic training & Advanced training use cases

These steps would be done by participants using the facilities provided by the social media clones. This could take the form of comments on news stories or the sharing of tweets. Whether this is actually done is up to the training organisers and should consider the aims of the training. For instance, training organisers could encourage sharing of tweets based on how truthful they seem to the participants. Afterwards one could discuss how the participants evaluated the content of tweets.

4.1.5 Step 5 and 6: Information is spread widely

Dis- and misinformation content quickly become part of the general online information flow, and if the influencing actor is lucky, the content will "go viral", that is, it will spread rapidly by being shared fast and often (step 5). It is this automated and rapid spread that makes social media a game changer for influence efforts compared to earlier efforts (Bergh, 2023). On a real-life social media platform, the spread of content is done by the large user mass in conjunction with the underlying social media technologies in step 6. Algorithms that match user interests with content is the primary culprit here.

However, the social media clones used in Somulator will work differently and typically eschew the algorithmic selection approach altogether and present information chronologically. In a training session there would not be a sufficient number of users using the social media for a long enough time to generate the amount of content required for such algorithms to work properly. See chapter 5.3 for a detailed discussion on this issue.

4.1.5.1 Basic training use case

In a simple training session this aspect can usually be ignored.

4.1.5.2 Advanced training use case

Step 6 has to be replaced with manual approaches when using Somulator (chapter 5.3).

When using Somulator for advanced training this is typically handled in three stages: development, production and deployment. The different stages and the tasks they involve are discussed in detail in the next chapter.

5 Somulator workflow: Development, production and deployment

The workflow described here is mostly relevant to the type of advanced use cases described in chapter 4.

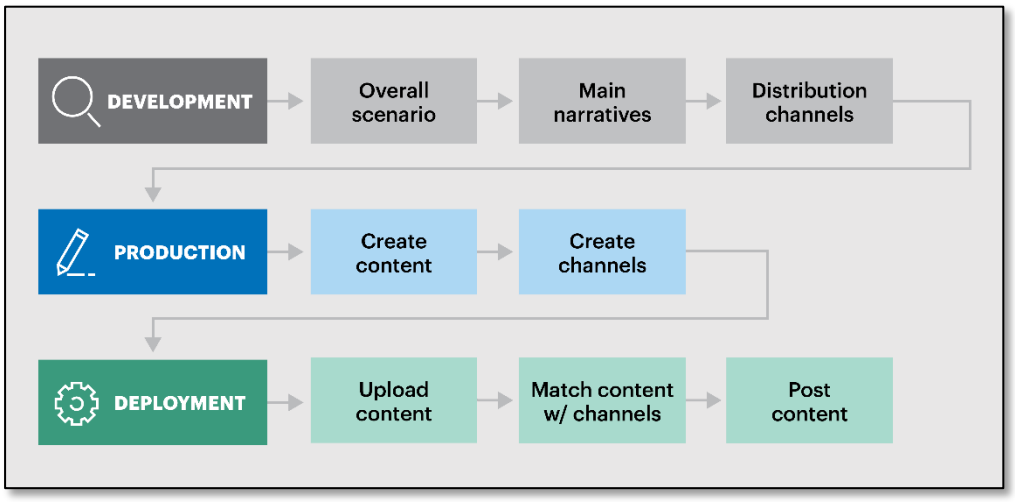

Figure 5.1 Overview over Somulator tasks. (Infographic: FFI/Grete Foss Alvestad)

5.1 Development: Operationalising a scenario through topics and distribution channels

Development tasks are concerned with planning for the training in terms of overarching scenarios and the narratives and distribution channels that best fulfil the learning needs of the participants. This will often be as part of a bigger training/exercise package. A training scenario typically furnishes the background to an imagined situation and describes a series of events leading up to the situation that the training session will simulate. The background usually provides a rationale for why one or more actors undertake certain actions i.e., what goals they want to achieve. The events described in a scenario tend to focus on actions the actors undertake that are related to the goals they want to achieve and events related to this, such as reactions by other parties. When planning social media training it is therefore useful to refer back to the model in chapter 4.1.1 above and base social media actions in the training on this model.

In the development phase one must therefore consider what message formats (short texts, news articles, videos, etc.) and content would support the overall scenario at different points in the scenario. Whether social media events and content should only be used in the actual training or also integrated in the scenario description is also something to consider. Furthermore, one has to plan what distribution channels (such as online newspapers, social media accounts, YouTube channels and so on) that the content will be distributed through. The advanced training case

explained in chapter 4 gave an example of how this would be done for an election manipulation scenario.

5.2 Production: Create relevant content and accounts to support scenario events

Production tasks handle the actual creation of content and content distribution channels. Distribution channels can be accounts that are created on a social media clone or a dummy, online news site. Once the development stage has provided an outline of how social media will be used in the training scenario, then the production tasks relating to Somulator will create actual content and describe the channels that the content will be distributed through. Referring back to the model in chapter 4.1 it must be remembered that content needs to (try to) trigger some reaction in the audience. Content should therefore be aligned with the overall scenario and the training goals for the current group of participants.

See chapter 6.1 for practical information on how to do this.

5.3 Deployment: Briefing + content + channels + timing

Deployment tasks manage the practical aspects of running Somulator during the training event. These are primarily concerned with sending out prepared content on a certain topic at a time that fits in with the aims and/or overall scenario of training sessions, using relevant social media clones, profiles and newspapers.

The social media clones used in Somulator do not have the vast data trove and the machine learning tools used by commercial social media platforms to analyse uploaded content and users' interests to show content that is deemed to be of interest to the individual user. These algorithms are manipulated when influence operations occur (Bergh, 2020, p. 17). Furthermore, the participants will use social media accounts that are created only for training purposes without any previous data. Unless this is factored into the training session the outcome might not be very useful to the participants. These two problems can be solved with the following approaches.

Firstly, during the development, production and deployment stages the Excon staff need to be fully aware of how propaganda, dis- and misinformation is spread on real social media and how training that use Somulator should reflect this. This is discussed in some detail in chapter 4. In short, we can say that relevant content deployed through appropriate distribution channels at a suitable time (according to the overall training session) are important to improve the learning outcome. Overall, it is the task of Excon staff to replace algorithms.

The practical steps to do this is detailed in chapter 6.2.

Secondly, a good pre-training briefing is very important, as others have discussed with regard to cyber security training for instance (Brilingaitė et al., 2020). The generic timeline for an exercise shown in Figure 5.1 is based on this framework and can be a useful training approach. Here it is important to explain what the purpose of the training is, and what roles the participants take on. A pilot workshop using Somulator resulted in complaints from the participants who felt that the content they saw were irrelevant to them as it dealt with issues they normally would ignore. Thus, they were at a loss on how to evaluate the content for the purposes of the workshop.

One way of solving this may be to provide different content; in that case one must do a preliminary survey to see what narratives and topics participants are interested in. However, it may be easier to explain clearly what type of person, what role, they should play in the training. A workshop for communications staff on how to deal with conspiracy theories for instance, would be hard pressed to come up with conspiracies that felt "right" for that particular audience. However, a good briefing that explains the type of content to expect and what type of person one should represent would go a long way toward improving the outcome.

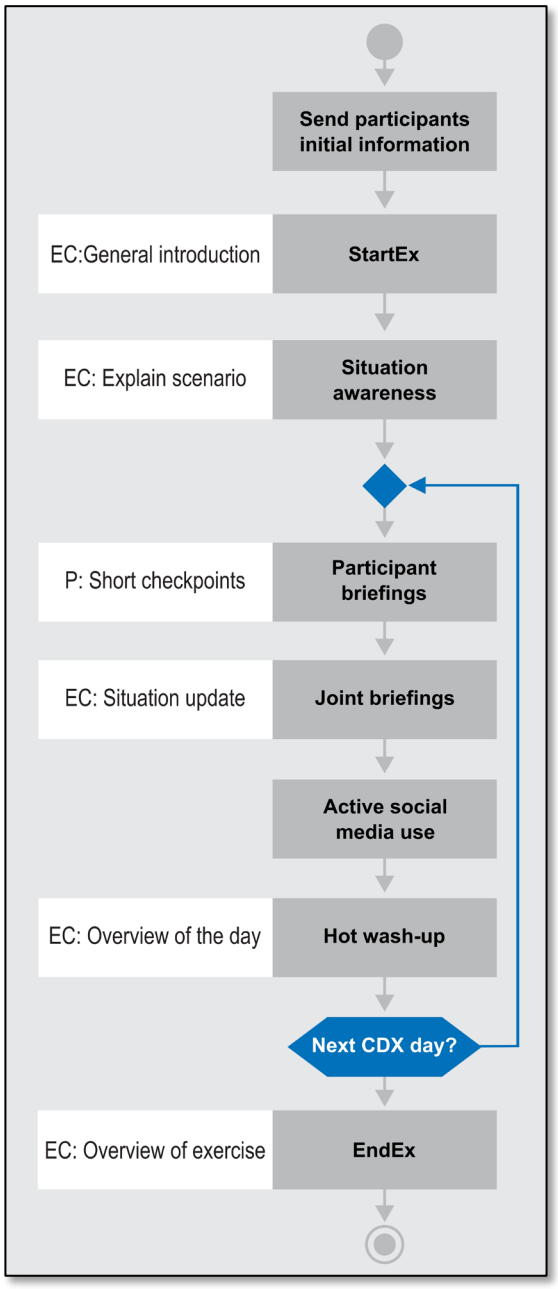

Figure 5.2 Timeline based on framework from Brilingaitė, Bukauskas, and Juozapavičius (2020). (Infographic: FFI/Grete Foss Alvestad)

6 Production, Deployment and Somulator management

This section of the report is organised according to the **production** and **deployment** steps outlined in chapter 5, with an additional section on general Somulator management. Whereas chapter 5 provided a conceptual overview of how Somulator could be used to provide realistic social media training, this chapter will focus on the practical steps required to achieve this. Each sub-section in this chapter will provide a brief explanation of the different tasks involved.

> For information on the practical use of Somulator, there is an online help available at https://somulator.some-lab.net/help.

6.1 Production

Input from Development stage: The creation of content and distribution channels is based on the work that was done in the Development stage. That is, Excon staff tasked with the production should know what topics/narratives to base the content on, the scenario that these topics and narratives relate to, and the type of distribution channels has been envisaged. See Appendix A for an example of input from the development to the production stage.

6.1.1 Create content

The content is, as discussed above, text, images and videos that will be shared through social media profiles or news web pages, collectively referred to as distribution channels. The content can be created in a number of ways: Manually from scratch; through machine learning tools such the GPT models that generate text[7] or through re-use of existing content. The latter can be data collections released by Facebook, Twitter and others with content from real life influence operations or Somulator content shared with other Somulator users.

[7] Generative Pre-trained Transformer 3 is a so-called model that uses a type of artificial intelligence to generate text that can be difficult to distinguish from text written by a human, see .

For simple training cases, it may be enough to create a few dozen messages that are tailored to the training session goals. For more advanced cases such as a multi-day exercise, one may have to come up with many thousand social media posts and dozens of news stories. When creating content, it is important to consider the format of the social media that will be used during the exercise, the "tone" of the posts that different types of social media users would employ and how this can be used to push the overall narrative. For example, when working on an election interference scenario it is useful to think of topics that create divisions in society and take extreme positions. A typical example of such positions would be to say: "The EU is a bunch of overpaid bureaucrats that follow Hitler's vision of a Europe ruled by Germany". When discussing NATO related issues, one could use existing, emotive memes that exaggerate a specific aspect of an argument, such as the one in figure 8.1.

Figure 8.1 Meme relating to NATO

On the other hand, when using a newspaper article to spread disinformation it requires more detail and a longer argument that supports the narratives of the overall training scenario. In an election scenario, such an article might focus on the cost of a particular solution and cite real research on the issue but exaggerate the figures and paint a false picture of the implications that suggests that the target group for the influence effort will suffer disproportionally.

In short, the practical aspect of creating content is a straightforward task and easy to undertake using a standard spreadsheet. However, developing the content requires considerable thought and will take time to get right. It needs to consider both the overarching scenario and how viral content works online. For instance, what forms of emotional manipulation (e.g. appealing to self-interest or exaggerating facts) are often used? How can the content format fit in to the different types of social media (long videos do not work well on Twitter)? What is the appropriate language and tone used among the target groups (attacking outsiders for example)?

For more information on how to create content, please refer to the separate *Content creation for Somulator* report: https://ffi.no/reports.

6.1.2 Create distribution channels – Social media clones

The social media clones that make up Somulator do not distinguish between accounts used by training participants and accounts that Excon staff use to distribute the content through. This is the same as on platforms like Facebook and Twitter where influence operations will create regular accounts to send out disinformation. However, the Excon module keeps a list of available accounts to use as distribution channels. This helps to avoid accidentally sending out content through participants' accounts that could create confusion.

In the case of a simple training session, it might be enough to use a single clone to package and distribute the content. For instance, one can use the News module to send out fake news articles or the Twitter clone to distribute conspiracy theories. For more advanced training scenarios, one would usually create an "ecosystem" of official and fake news sites and real and fake social media accounts that support each other. The core idea is to create channels that enhance the realism of the content in the eyes of the participants. A state-controlled news site could be set up with several social media accounts representing ambassadors and other government officials that post links to news articles and fake accounts that make the stories more popular.

6.1.3 Create distribution channels – News sites

News sites can be important for an influence operation as the news format gives disinformation a legitimacy that tweets from individuals do not provide. Such news sites can for example be created by news agencies controlled by a state, by groups supporting very partisan political views or by commercial actors who earn money from advertising by publishing untrue, but engaging articles.

In simple training cases it may be enough to provide a single news site, but in more advanced cases one should consider representing multiple news outlets. Somulator includes three news sites by default, "Local News", "National News" and "The Sensation" (Figure 6.2). Unlike the social media clones, the news module in Somulator does not use accounts to distribute content. Instead, it uses the Ghost content management system's ability to create different sub-sections of a website. In this way one can create multiple sites with different names and logos that represent different newspapers or news sites.

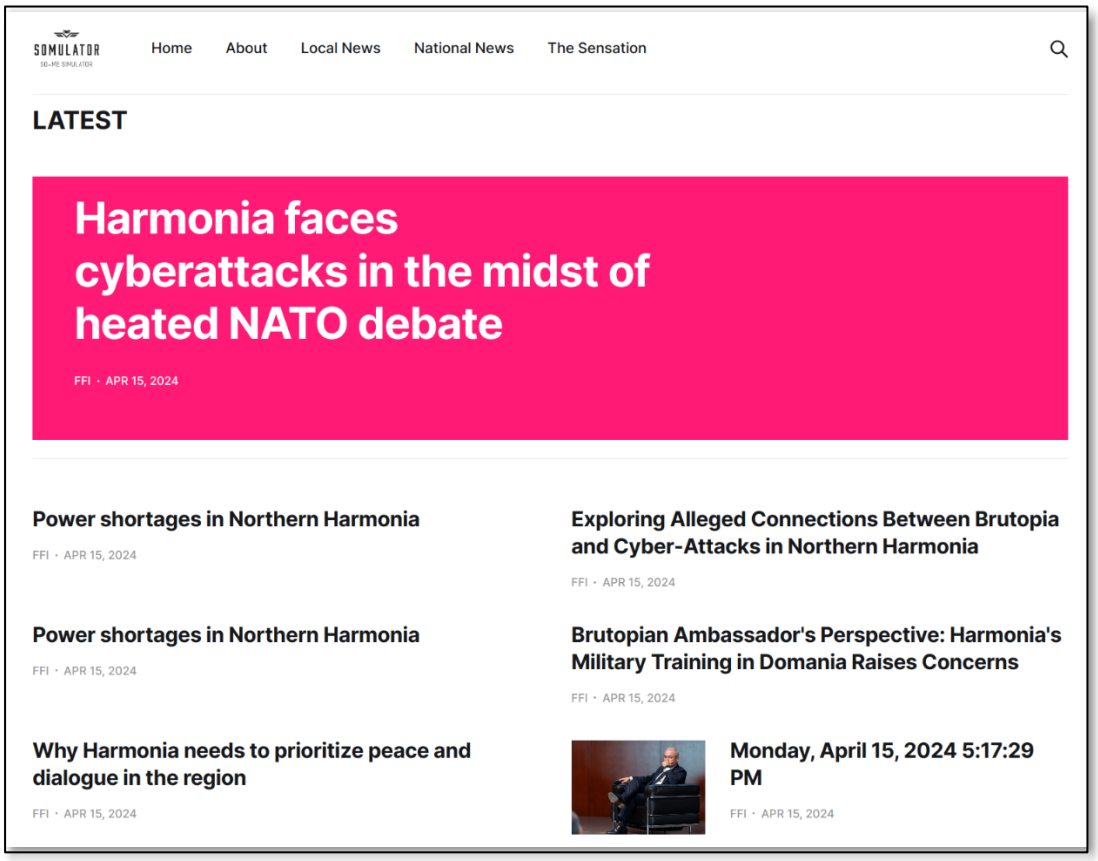

Figure 6.2 The Sensation, one of three news sites in Somulator.

6.2 Deployment

Once content and distribution channels have been created, the content must be a) uploaded to the Excon module and b) deployed through distribution channels during the training session. The uploading can occur at any time, but it might be useful to do this a few days before the training occurs to make sure there are no problems with uploading it.

6.2.1 Post content through relevant distribution channels

All previous tasks lead up to this activity: To display content in the feeds of the training participants for learning purposes. For basic training, a single spreadsheet with content deployed manually may suffice whereas complex training scenarios may utilise several spreadsheets that are played back automatically using the timing specified in the spreadsheet.

This concludes the overview of how to run a training activity or exercise through Somulator. You can now refer either to the online help (https://somulator.some-lab.net/help) or the additional FFI report on content creation: https://ffi.no/reports.

7 Conclusion and the way forward

This report discusses how the increased spread of disinformation and propaganda through social media necessitates training to be able to handle such misuse of social media, particularly for personnel in the total defence sector. Such training requires relevant tools to assist in the training process. An example of this could be a method for sharing false news stories that trainees evaluate. Existing training solutions are assessed, and it is concluded that a fully-fledged social media simulator is the best option to provide a realistic training experience.

The development of the Somulator social media simulator by the Norwegian Defence Research Establishment (FFI) is examined. Somulator uses five existing, open-source web applications to simulate different social media platforms, in conjunction with a custom-built control module that makes possible the carefully controlled deployment of large volumes of social media content in a training situation.

A general overview of how dis- and misinformation is spread through social media is used to explain how this can be emulated in Somulator. A three-stage approach is discussed: the development, production and deployment of social media profiles and content. The importance of relevant content and social media profiles is explored in depth.

At the time of writing, Somulator is a complete tool that has been used extensively in training and for exercises. This current state is documented in this report. However, Somulator has been developed to be flexible and adaptable. FFI wants to encourage the adoption and development of Somulator among NATO countries and allied partners. Additional social media platform simulations or new training management features may therefore be added over time. In this case, please refer to additional documentation as and when required.

Appendix

A Example of input from development stage

This is an example of information that development staff should provide to production and deployment staff. The purpose of this information is to enable the setting up of accounts and creation of content that is required to fulfil the training purposes.

Training background

We will train local county officials who will have to work with the armed forces on what type of disinformation and misinformation to expect in a crisis involving Country X.

Overall scenario

Brutopia has started to organise hacking attacks that target transport and utilities infrastructure in the local county. This is accompanied with covert information operations through social media, with the aim of spilling over into local newspapers. The purpose is to scare the local population from supporting the neighbouring country, Cardamommia, who Brutopia is planning to invade. Cardamommia has close personal and business ties in the local county that is attacked. The attacks are limited to the single county, the rest of the nation is left untouched.

Main narratives

Brutopia can be a close friend and good business opportunity for the local county but oppose it and this opportunity is closed.

Cardamommia seems great on the surface but is ruled by a clique of outsider business interests who plunder real people to their own benefit.

Cardamommia has developed secret plans to take control over the local county through immigrants from Cardamommia. These immigrants pose a threat to peaceful Brutopian immigrants as well as local business interests and jobs.

Distribution channels

For this training session we will use a Twitter clone with these profiles:

- a number of local conspiracy theorists
- non-committed, but sceptical local people
- national businesses with local interests
- official channels for Brutopian and Cardamommian actors

In the news module we will have a local newspaper and several international, government controlled, news sites that support different sides in the conflict.

B Example of information for participants

Welcome to the training workshop on the 5th of November 2022.

This workshop will use several websites that copy real life social media platforms. Below is a list of these sites. To join the training workshop you need first to go to this URL:

https://training.my-domain.org

Before you can access this webpage, you need to enter a username and password for the training workshop. This prevents Internet users who are not participating in the training from accessing the social media websites.

Username for the training workshop: somulator
Password for the training workshop: mnu_987-QAS

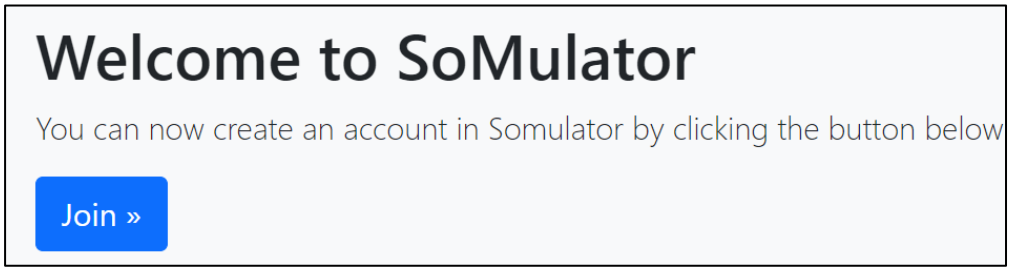

On this page you click the Join button to join the exercise, see the image above. This will automatically generate user accounts for you on the different social media websites used. When the user accounts have been set up you can join the different sites by clicking on the links shown in the screenshot below.

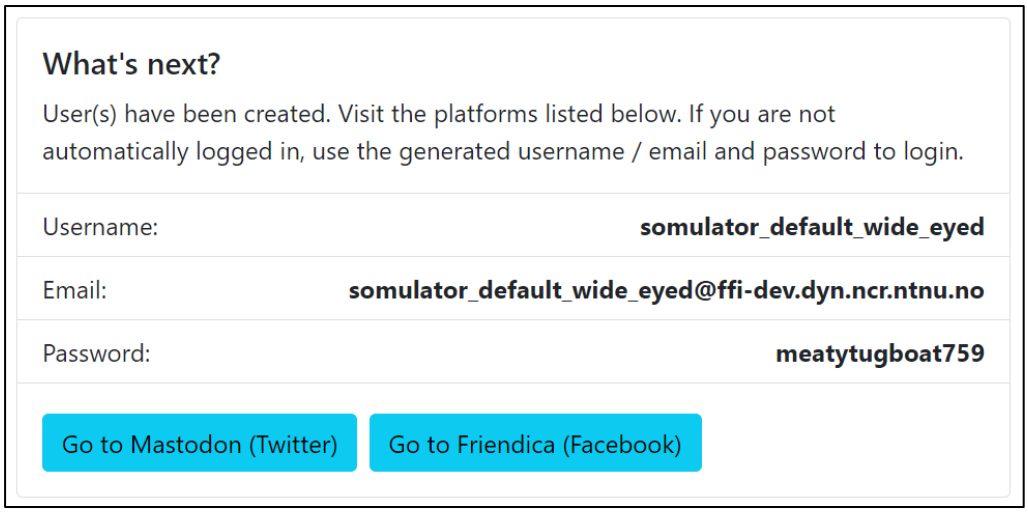

Ensure that you make a note of the usernames and passwords assigned to you.

C Extending Somulator

Adding modules to web applications (Level: Administrator)

The easiest method for extending Somulator is the use of additional modules in some of the core web applications. Ghost (the news site module) for example, has numerous modules to provide new functionality, such as podcasting or commenting on articles. To find these you need to go to the web site for the relevant social media clone.

Adding services (Level: IT Support)

The different web applications in Somulator are packaged in what is called *container* technology, specifically Docker containers (Merkel, 2014). Docker containers separate each element of a web application into different modules, this enhances security and makes it easier to modify or update one web application without affecting the other applications. It is therefore feasible to add a different web application for training purposes. For example, if one would like to use spam e-mails as part of the information operation one could add an e-mail service in a Docker container.

Amending web applications' source code (Level: Programmer)[8]

If an organisation has very specific needs, it is possible to amend the source code of a web application. This is possible due to the open-source license used for these applications. For example, if one wanted to embed a *newsfeed* from the Facebook clone in a newspaper using the news site web application, then this would be the way to do it.

Extending the Excon code (Level: Programmer)[9]

It is also possible to add functionality to the core Excon module by writing new code, for example to create a function that simulates a post "going viral", i.e., spreading fast through the different social media clones.

[8] Refer to each web applications individual GitHub page for further information.
[9] Documentation on this is available in the Somulator GitHub repository, access may be provided upon request.

References

Bergh, A. (2019). *Social network centric warfare: Understanding influence operations in social media* (FFI-Rapport No. 19/01194; p. 65). Norwegian Defence Research Establishment (FFI). http://hdl.handle.net/20.500.12242/2623

Bergh, A. (2020). Demokratier i informasjonskrig. *Pro Patria*.

Bergh, A. (2022). *Somulator: An information training environment to train as you Tweet*. NATO IST-195 Symposium Proceedings, Stockholm, Sweden.

Bergh, A. (2023). Påvirkning og teknologi gjennom tidene: Fra falske allierte til falske nyheter. In *Cyberoperasjonar på norsk*.

Brilingaitė, A., Bukauskas, L., & Juozapavičius, A. (2020). A framework for competence development and assessment in hybrid cybersecurity exercises. *Computers & Security*, *88*, 101607. https://doi.org/10.1016/j.cose.2019.101607

Cohen, R. (2020, December 8). *Health Care Workers Say They Have To Fight Both The Coronavirus And Disinformation*. NPR.Org. https://www.npr.org/2020/12/08/944337692/health-care-workers-say-they-have-to-fight-both-the-coronavirus-and-disinformati

Evon, D., & Mikkelson, D. (2017, March 24). *Muslim Woman Ignores Dying Victim of London Terror Attack?* Snopes.Com. https://www.snopes.com/news/2017/03/24/muslim-woman-london-attack/

Leprince-Ringuet, D. (2020). *Facebook has a new tool to spot spammers, and it's already taken down billions of accounts*. ZDNet. https://www.zdnet.com/article/facebook-has-a-new-tool-to-spot-spammers-and-its-already-taken-down-billions-of-accounts/

Maekawa, M. S., Hundzinski, L. N., Chandrahera, S., Tajima, S., Nakai, S., Miyazaki, Y., & Okawa, K. (2021). Design of a Social Media Simulator as a Serious Game for a Media Literacy Course in Japan. *CSEDU (1)*, 392–399. https://doi.org/10.5220/0010499903920399

Masakowski, Y. R., & Blatny, J., M. (2022). *Mitigating and Responding to Cognitive Warfare* (TR-HFM-ET-356; HFM Exploratory Team, p. 142). NATO Science & Technology Organization. https://apps.dtic.mil/sti/trecms/pdf/AD1200226.pdf

Merkel, D. (2014). Docker: Lightweight linux containers for consistent development and deployment. *Linux j*, *239*(2), 2.

Paxton, J. (2018). Trident Juncture and the information environment. *NATO Review*. http://www.nato.int/docu/review/2018/Also-in-2018/trident-juncture-and-the-information-environment/EN/index.htm

Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, *5*(1), 65. https://doi.org/10.1057/s41599-019-0279-9

Tomlin, G. M. (2016). #SocialMediaMatters: Lessons Learned from Exercise Trident Juncture. *Joint Force Quarterly*. http://ndupress.ndu.edu/Media/News/News-Article-View/Article/793264/socialmediamatters-lessons-learned-from-exercise-trident-juncture/

Warzel, C. (2017, October 4). *Here's How YouTube Is Spreading Conspiracy Theories About The Vegas Shooting*. BuzzFeed News. https://www.buzzfeednews.com/article/charliewarzel/heres-how-youtube-is-spreading-conspiracy-theories-about

Yates, J. (2019). *Inside the fake Facebook profile industry | Radio-Canada.ca*. https://ici.radio-canada.ca/special/sextorsion/en/index.html?utm_campaign=The%20Interface&utm_medium=email&utm_source=Revue%20newsletter

About FFI

The Norwegian Defence Research Establishment (FFI) was founded 11th of April 1946. It is organised as an administrative agency subordinate to the Ministry of Defence.

FFI's mission

FFI is the prime institution responsible for defence related research in Norway. Its principal mission is to carry out research and development to meet the requirements of the Armed Forces. FFI has the role of chief adviser to the political and military leadership. In particular, the institute shall focus on aspects of the development in science and technology that can influence our security policy or defence planning.

FFI's vision

FFI turns knowledge and ideas into an efficient defence.

FFI's characteristics

Creative, daring, broad-minded and responsible.

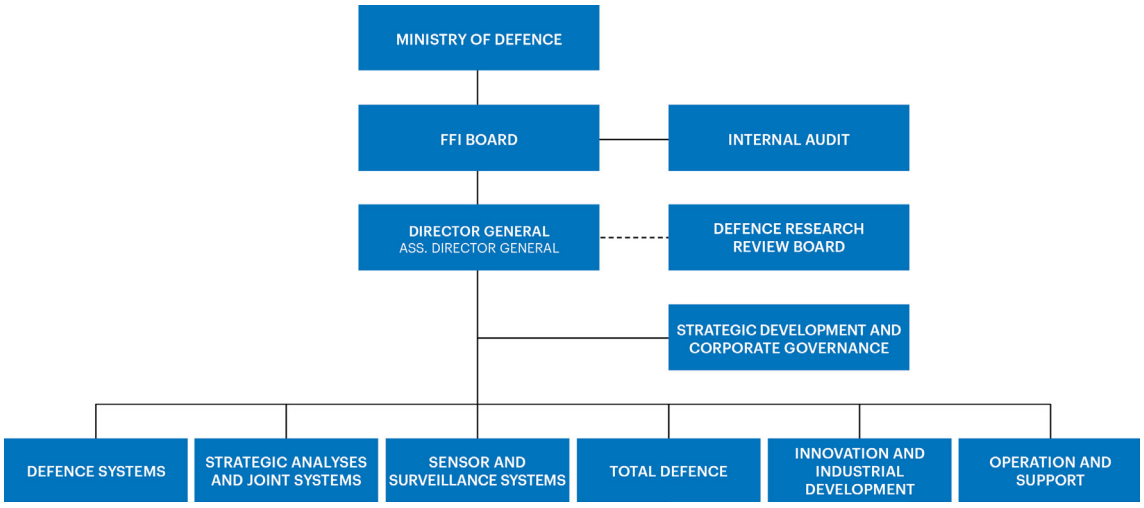

Forsvarets forskningsinstitutt (FFI)
Postboks 25
2027 Kjeller

Besøksadresse:
Kjeller: Instituttveien 20, Kjeller
Horten: Nedre vei 16, Karljohansvern, Horten

Telefon: 91 50 30 03
E-post: post@ffi.no
ffi.no

Norwegian Defence Research Establishment (FFI)
PO box 25
NO-2027 Kjeller
NORWAY

Visitor address:
Kjeller: Instituttveien 20, Kjeller
Horten: Nedre vei 16, Karljohansvern, Horten

Telephone: +47 91 50 30 03
E-mail: post@ffi.no
ffi.no/en