Social Media Synthesis using AI for Decision Support

Evan P. Harris

CAE Inc.
Ottawa, Ontario, Canada
evan.harris@cae.com

Jaspreet Kaur, Maher Chaouachi, Martin Durocher, Rakesh Tiwari, Alex Emirov, Nick Giannias
CAE Inc.
Montreal, Quebec, Canada
jaspreet.kaur@cae.com, maher.chaouachi@cae.com, martin.durocher@cae.com, rakesh.tiwari@cae.com, alexander.emirov@cae.com, nick.giannias@cae.com

ABSTRACT

Social media is pervasive throughout today's society and, worldwide, there are over 4.6 billion active social media users. Reports of, and reactions to, events can propagate locally, nationally, and internationally almost instantly, and can themselves trigger a cascade of reactions and events. Within this reality, it is important for any large-scale simulation of an engaged population using human behavior modelling to incorporate social media. This paper reports on the design, implementation, and integration of the social media component within a broader simulation of Greater London with application to decision support systems, the training of decision makers, and course of action analysis. Within the social media component, live social media messages were combined with synthetically generated messages and presented in real-time, while analytics aggregating sentiment expressed in the messages were displayed to decision makers throughout the course of the simulation. The synthetic messages were generated by a trained Artificial Intelligence (AI) model; each was related to the different emotions of happiness, sadness, fear, anger, joy, and surprise, and assigned to a member of the simulated population. The integration of an externally triggered event in the simulation, a power outage, resulted in a change in behavior of the simulated population and, consequently, a change in the resulting tone and emotion reflected in the synthetic social media messages. Finally, initial user feedback is reported, and considerations for additional factors to influence the synthetically generated social media content based on cognitive state, demographic attributes, and extensions to the AI models are discussed.

ABOUT THE AUTHORS

Dr. Evan P. Harris is a Senior Systems Architect at CAE Canada. Dr. Harris has a Ph.D. in Computer Science from The University of Melbourne, Australia, and has over 20 years' experience in providing modelling and simulation, systems, and software engineering services to Defense and Security customers in Canada, Australia, Asia and Europe.

Dr. Jaspreet Kaur is a Data Scientist at CAE Canada. Dr. Kaur received her Ph.D. in Applied Physics from Indian Institute of Technology (ISM), Dhanbad, India. She has over 9 years' experience, and her interests include computational modeling and simulation, optimization algorithms in AI and machine learning and its sub-fields including NLP, computer vision.

Dr. Maher Chaouachi holds a Ph.D. in computer science specialized in the field of Artificial Intelligence obtained from the University of Montreal. He worked as a postdoctoral researcher at McGill University. Dr. Chaouachi is an AI strategist at CAE working in various programs involving machine learning, process, and data mining, optimization, and predictive modelling.

Dr. Martin Durocher is a Data Scientist at CAE Canada. He received a Ph.D. in Statistical Hydrology from the University of Quebec, Canada. During the past 12 years, he has been involved in various fields of applied statistics, including economy, environment, and health care.

Dr. Rakesh Tiwari is a Data Scientist at CAE Canada. Dr. Tiwari has a Ph.D. in Theoretical physics from The Ohio State University, Columbus, Ohio and has over 18 years of experience in academia and industry working in various fields including quantum optimization, AI/ML and NLP.

Alex Emirov is a Software Developer at CAE Canada. He holds a bachelor's degree in Computer Science from Concordia University, Canada. He has over 14 years' experience in system programming, mainly in the simulation industry. Alex's main interests include algorithm computational optimization, internet-of-things, and open source.

Nick Giannias is a Senior Technical Fellow working for the Advanced Technology and Innovation group at CAE and the Chief Architect for CAE's simulator products. With over 30 years of experience, he specializes in software development and real-time simulations for concept development, analysis, mission rehearsal, and training applications. Nick holds a bachelor's degree in Mechanical Engineering (Hons.) specializing in Aerospace from McGill University.

Social Media Synthesis using AI for Decision Support

Evan P. Harris
CAE Inc.
Ottawa, Ontario, Canada
evan.harris@cae.com

Jaspreet Kaur, Maher Chaouachi, Martin Durocher, Rakesh Tiwari, Alex Emirov, Nick Giannias
CAE Inc.
Montreal, Quebec, Canada
jaspreet.kaur@cae.com, maher.chaouachi@cae.com, martin.durocher@cae.com, rakesh.tiwari@cae.com, alexander.emirov@cae.com, nick.giannias@cae.com

INTRODUCTION

The use of social media is pervasive throughout today's society. Worldwide there are over 4.6 billion active social media users (DataReportal, 2022), including around 7 in 10 Americans (Pew Research Center, 2021). Users have a variety of motives for generating content, including individuals finding a means of self-expression, businesses performing self-promotion and image management, news organizations generating and consuming content, government and municipal organizations keeping in touch with their constituencies, and malicious actors generating misinformation for their own ends. Reports of events, and reactions to them, can propagate locally, nationally, and internationally almost instantly, and can themselves trigger a cascade of reactions and events.

In future, we believe that any large-scale simulation of an engaged population for the purposes of assisting, advising, or training decision makers must incorporate social media. In civil and emergency management contexts, social media can be used synergistically to both inform decision makers and inform the public (Space and Naval Warfare Systems Center Atlantic, 2013).

Our vision is to create a large-scale, high fidelity synthetic environment able to visualize information in a collaborative manner while running what-if scenarios to perform course of action analysis, inform decisions, and train decision makers. Within such a synthetic environment, members of a simulated population, utilizing human behavior modelling concepts, will interact with each other both physically and via modern communication methods, including social media.

In this paper, we report on the design, implementation, and integration of the social media component within a broader simulation of Greater London. We introduce the Greater London scenario and present an overview of our conceptual architecture and the Cognitive Layer within which the social media component lies. We then discuss the details of generating, integrating, and presenting social media content and analysis, and conclude by discussing initial user feedback, draw conclusions, and present possible future developments.

GREATER LONDON SCENARIO

The initial implementation of our synthetic environment concept was the development of a digital twin of Greater London in the United Kingdom, an area covering over 600 square miles with a population of over eight million people. A full population was generated using 2011 census data to ensure that the demographic profile and location of the population was as realistic as possible. High fidelity imagery was used as a part of the 3D and 2D visualizations.

At runtime, the use case commences with normal pattern of life activities for the population, including social media use. People may be at work, at home, at school, etc. This establishes the baseline of the simulation and requires no input from Decision Maker users. Within this context, live (real) social media messages may optionally be used as background to provide a heightened sense of realism – topics that are trending on-the-day may be displayed rather than topics that were current when the simulation was constructed.

Once a normal pattern of life is established, alternate behavior is stimulated from a section of the simulated population to provide a resulting challenge to Decision Makers. This is achieved by introducing a failure of the electrical power network for a significant part of the city, affecting millions of the simulated population. Depending on the desired context for users, this failure could be the result of natural causes, such as environmental factors or normal equipment failure, or it could be the result of cyber incident, such as a side effect of a ransomware attack or a direct cyberattack on critical infrastructure. The resulting change in population behavior, which may vary depending on the cause of the failure (once known to the simulated population), includes a change in social media use.

Decision Maker users monitor and interact with the synthetic environment via a web-based user interface that presents 3D and 2D visualizations of Greater London plus a collection of dashboards covering information such as social media messages. Instructions may be issued to various simulated entities, such as emergency services, in reaction to events in the scenario.

ARCHITECTURE

Synthetic Environment Conceptual Architecture

Our conceptual architecture of synthetic environment consists of four layers: physical, human, cognitive, and resource (see Figure 1). It is structured around the key elements to be represented in the simulation and is conceptually similar to the grid model of network centric warfare (Garstka, 2003) and a proposed information warfare engagement model architecture (Hazen et al., 2017). It supports modelling a full range of Political, Military, Economic, Social, Information and Infrastructure (PMESII) elements and effects.

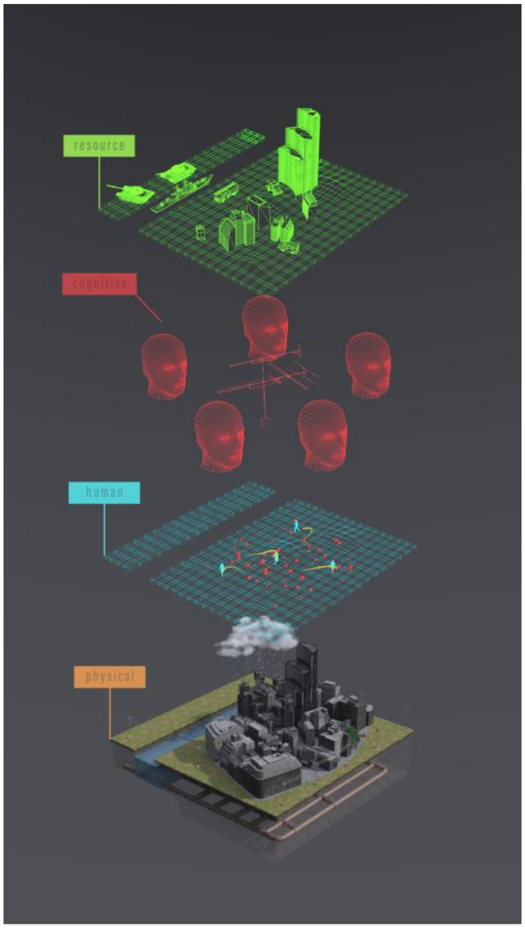

Figure 1. Four Layer Architecture

The Physical Layer models all aspects of the natural and built environment, including Critical National Infrastructure. Within the Greater London scenario, the physical layer contains imagery, terrain, buildings, roads, tube stations, notable landmarks such as Trafalgar Square, and an electricity distribution network (Hobbins et al., 2022).

The Human Layer contains a model of each individual within the population. Within the Greater London scenario, demographic attributes associated with each of the eight million individuals include age, sex, household location, and work/school location; these are combined with current location and current (intended) destination location to simulate population mobility (Giannias et al., 2022).

The Resource Layer models the distribution and employment of assets including emergency responder, government, civil, and commercial resources. The components within this layer are the primary assets that Decision Makers can dispatch and employ in reaction to, and for the purposes of, mitigating events as they occur within the simulation.

Cognitive Layer

The Cognitive Layer is intended to encapsulate all aspects of human behavior modelling and cognition. This includes knowledge representation, beliefs, desires, intentions, planning, reasoning, information flow, communications content, thought processes, and decision making. It encompasses actions such as communication that are not modelled in the Human Layer. This modelling may be enhanced by representations of formal or informal organizational structures (pairs, teams, larger organizational structures), and attributes influenced by Human Factors research. Conceptually, the scope of the Cognitive Layer is independent of the implementation technology, which permits elements to be implemented in technologies as varied as cognitive architectures (Laird, 2012), BDI agents (Rao &

Georgeff, 1991), hybrid approaches (Silverman et al., 2018), and artificial intelligence techniques such as machine learning.

While many of the aspects of human behavior modelling described above are important to produce high fidelity representations of human behavior, most are only indirectly visible to a user through the overall quality of the overall simulation. In contrast, the use of social media is directly visible to users and, to date, has been largely underrepresented in simulation systems.

For the Cognitive Layer implementing the Greater London scenario, we concentrated on modelling communication by members of the population on social media and analysis of social media messages. Orthogonal issues such as broader individual cognitive reasoning and sub-population behavior modelling remain for the future.

Our novel Cognitive Layer architecture (see Figure 2) consists of: several components incorporating AI models to produce and label social media messages (Synthetic Message Generator, Message Classifier, Synthetic Video Generator); a simulation component running as an element of the broader synthetic environment simulation that publishes social media messages and analytics derived from the social media messages that are aligned with the simulation state (Cognitive Layer Simulation); a component that serves as a proxy for the web-based user interface (Social Media UI Proxy); panels in the web-based user interface to display social media messages and analytics (Social Media Message Feed Panel, Social Media Analytics Dashboard); and an optional allowance for a live social media component to inject current messages into the input social media messages (Live Social Media Proxy). Within Figure 2, Cognitive Layer components that are responsible for creating and classifying social media content are shown in blue, Cognitive Layer components that are responsible for the simulation run-time, publishing social media messages and displaying social media messages and analytics, are shown in green.

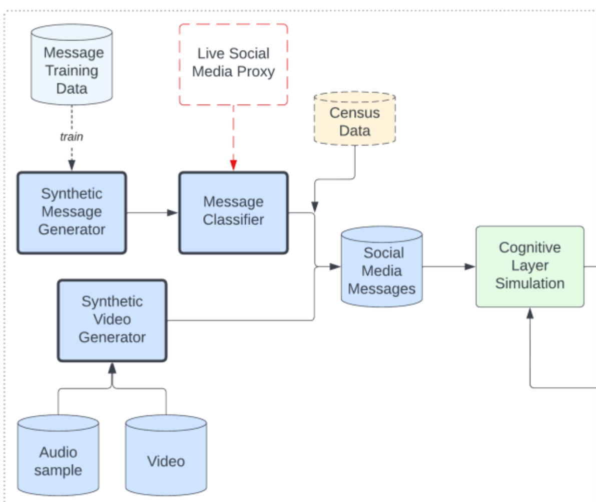

Figure 2. Cognitive Layer Architecture

The Cognitive Layer receives external events that are triggered through other components in the broader synthetic environment, for example the power network simulation, or through direct user input via a web-based user interface.

AI Architecture

Each of the AI models we developed, except the video generation model, are based on Natural Language Processing (NLP) (Mohammad et al., 2017), a subfield of AI and linguistics that interprets natural language (i.e., social media messages in our scenario) and transforms it into a representation that is possible for computers to process and for humans to analyze. To build the NLP models, data which is publicly available on platforms like Kaggle (2022) was used. The AI architecture used in the Cognitive Layer simulation is shown in the inner box of Figure 2.

The Cognitive Layer includes an AI synthetic message generator, AI message classifiers, and an AI synthetic video generator which are combined with census data and analyzed. The synthetic data generator NLP model is Open AI's Generative Pre-Training (GPT-2) (Radford et al., 2019) model that was fine-tuned on a dataset of 15000 social media messages and yields convincing synthetic social media messages (similar to that of Twitter, Facebook, etc.). For our initial implementation, the synthetic data was generated in advance and assigned to individuals of the population created in the Human Layer based on census data.

Within the Message Classifier, we implemented four types of AI classifiers: sentiment, subjectivity, emotion classification, and hate speech recognition. Sentiment, subjectivity, and emotion classification were used to gain insight into the mindsets of anticipated individuals or groups within the population. When aggregated for the Greater London scenario, this helped provide a general understanding of social media users and decide macro-level sentiment, opinions, and emotion for the population. The synthetic social media messages generated by the synthetic data generation model are tagged with a sentiment label, subjectivity score, and an emotion label by respective classifiers. The labels interpreted by the models signify the psychological state of the author of the message. The sentiment labels have one of the following values: positive, negative, neutral. The subjectivity score is in the interval [0, 1], with higher values meaning the message is subjective or an opinion and a value closer to 0 means it is an objective message stating a fact. The emotion model classifies each message according to one of the six emotions: joy, fear, anger, surprise, love, sadness. The hate speech AI classifier is used to filter (remove) social media messages that contain offensive and hateful content.

The AI synthetic video generator is an AI speech synthesis algorithm that generates synthetic videos. Since the scenario is based in Greater London, a 17 second synthetic video of Sadiq Khan, the Mayor of London, was generated to notify the public about the cyberattack for that variant of the scenario. The combination of social media messages and synthetic video constitute the synthetic database that is used as input by the simulation and analytics pipeline.

SOCIAL MEDIA MESSAGE GENERATION

The simulation of "how does a population respond to a scenario on social media?" requires a method of generating relevant and adaptive social media content, which is referred here as synthetic data. Text based social media content can be highly effective for illustrating the evolution of the mood of the population as the simulation scenario evolves. Harb (2019) investigated the use of deep neural networks for analyzing emotional reactions to terrorist events on Twitter which shows the effectiveness of such models in the analysis of social media messages for similar events.

We initially trained a bidirectional Long Short-Term Memory (LSTM) (Harb et al., 2019) model for generating synthetic data. For the LSTM model, human generated input phrases of one to five tokens were used to create content of the desired length. The resulting final output messages were not complete and required additional postprocessing to make them realistic. To try to improve this process, a Generative Pre-trained Transformer 2 (GPT2) (Radford et al., 2019) model was tested to replace the LSTM model due to its potential to generate independent long texts. The GPT2 model was more effective in generating realistic social media content as judged by a human.

We used transfer learning to finetune the GPT2 model on the input training data. The input training data is a single text document that contains all the social media messages. The GPT2 model learns how to generate individual messages better when tokens were added at the beginning (prefix) and at the end (suffix) of the messages in the training data. Emoticons were kept in the training data, which increases the total number of tokens in the vocabulary. The inclusion of message prefix and suffix tokens and emoticons in the training data resulted in the GPT2 model producing much more realistic synthetic content.

SOCIAL MEDIA VIDEO GENERATION

To provide a richer simulated environment for Decision Maker users, and potentially to stimulate the simulated population, the (simulated) authorities can publish a video on social media to inform everyone about major events, provide a reaction that shows control, and reassure the public. Synthetic video generation can achieve this objective by reassuring the population that the government is taking appropriate action in response to the scenario. The simulation is enriched by use of the video as it provides additional realism and provides a more immersive environment for the user.

To produce the synthetic video, synthetic audio was generated first using a speech synthesis model. A speech synthesis model is commonly a deep neural network-based algorithm that takes a set of linguistic features (a sequence of characters, phoneme, etc.) as input, passes them through a synthesizer like Tacotron (Wang et al., 2017) or FastSpeech (Ren et al., 2019), and produces spectrogram frames. The process we used to produce the synthetic video is shown in Figure 3. It uses a finetuned synthesizer based on Tacotron2 (Shen et al., 2017) and zero-shot learning technique (Jia et al., 2019). This technique requires only a few seconds of untranscribed audio (speech waveform/audio samples) and synthesizes audio closer to human speech including subtle technicalities like volume, pace, and intonation. The synthesizer outputs a set of acoustic features/spectrogram frames represented by Mel spectrogram (see Figure 4), which, in turn, is fed into a decoder and converts the Mel spectrogram to generated audio. The decoder used is the WaveRNN model (Kalchbrenner et al., 2018). To make a realistic video, the generated audio is integrated with the video. To achieve a good synchronization between the two sources, a deep neural network-based lip sync model called Wav2lip (Prajwal et al., 2017) was used. The output of this model is lip-synchronized video.

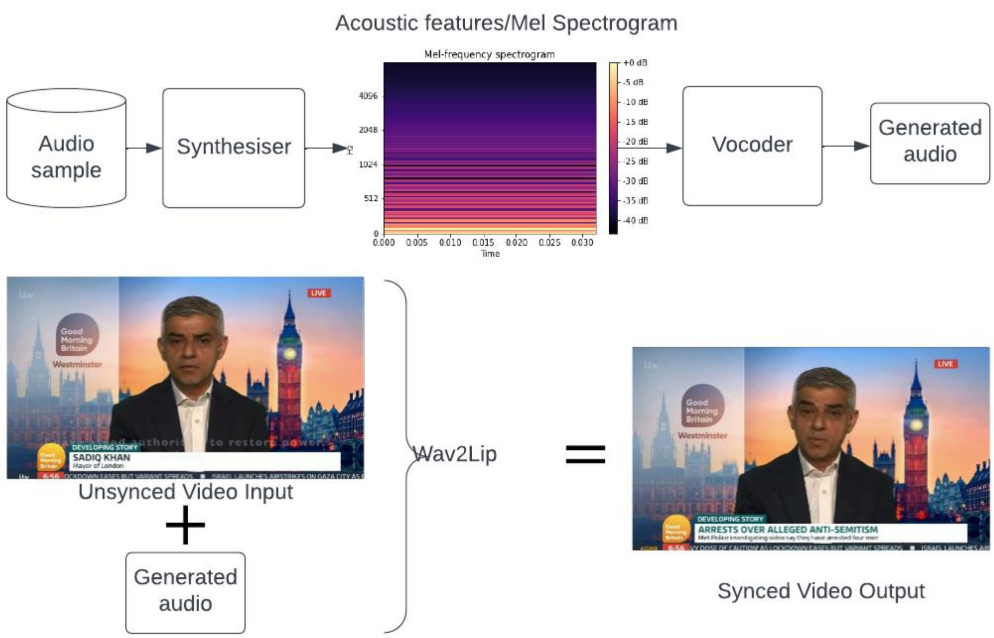

Figure 3. Synthetic video generation process

SOCIAL MEDIA INTEGRATION

In the cyberattack variant of the Greater London scenario, the simulation in the Cognitive Layer is characterized by two inputs events which divide the duration of the simulation into five stages: the initial, normal, stage; the reaction that occurs at the onset of the cyberattack; the response at the peak of the power outage due to cyberattack, which we refer to as the power outage stage; the comeback as the power returns; and the return to new normal state. All five stages reflect change in behavior of the social media users. In each stage, the evolution of the metrics, $U(t)$ characterizes the first and the last normal stages while the second process $V(t)$ characterizes the power outage. At all times, the overall simulated process is modeled by

$$S(t) = \gamma \mathrm{U(t)} + (1-\gamma)\mathrm{V(t)} \qquad (1)$$

where $\gamma \in [0,1]$ is a control factor. In the normal stage, the control factor is one, while it is zero during the power outage stage. In between normal and the power outage stage, the control factor has an intermediate value that gradually transitions from one to the other and vice-versa for the later stages.

Within the Cognitive Layer simulation, the simulation has a refresh rate of 10 seconds (0.1 Hz). At each time step, the simulation updates the synthetic social media feed, composed of the normal and crisis stream of messages, and evaluates metrics over a predetermined moving window of 15 minutes. The metrics of the normal stage process are directly generated by analyzing the normal stream of messages from the simulated population of Greater London. The crisis process generates metrics as a moving average process where the targeted average values are determined based on values found in the literature for similar events (Harb et al., 2017).

The content of the crisis stream reacts to the power outage and is consistent with the targeted metrics values, but the exact values are not strictly enforced. The metrics reported in the analytics dashboard aggregate the messages evaluated from the normal stage process and the ones simulated from the crisis process. The number of messages produced by the crisis process is determined by a Poisson process that spread timestamps uniformly over each time step. The crisis message rate is treated like the other metrics of the simulation and is regularized by the control factor to provide a gradually increasing flow from zero to its crisis target, and vice versa.

SOCIAL MEDIA MESSAGE ANALYTICS

The output of the Cognitive Layer is visualized by the user through the Web User Interface (UI) that, in part, displays a temporal analysis of aggregated social media messages for the Greater London area. Figure 4 contains an example screenshot of the Web UI during the cyberattack containing two panels: the Social Media Analysis dashboard, and the Live Feed that contains social media messages including video.

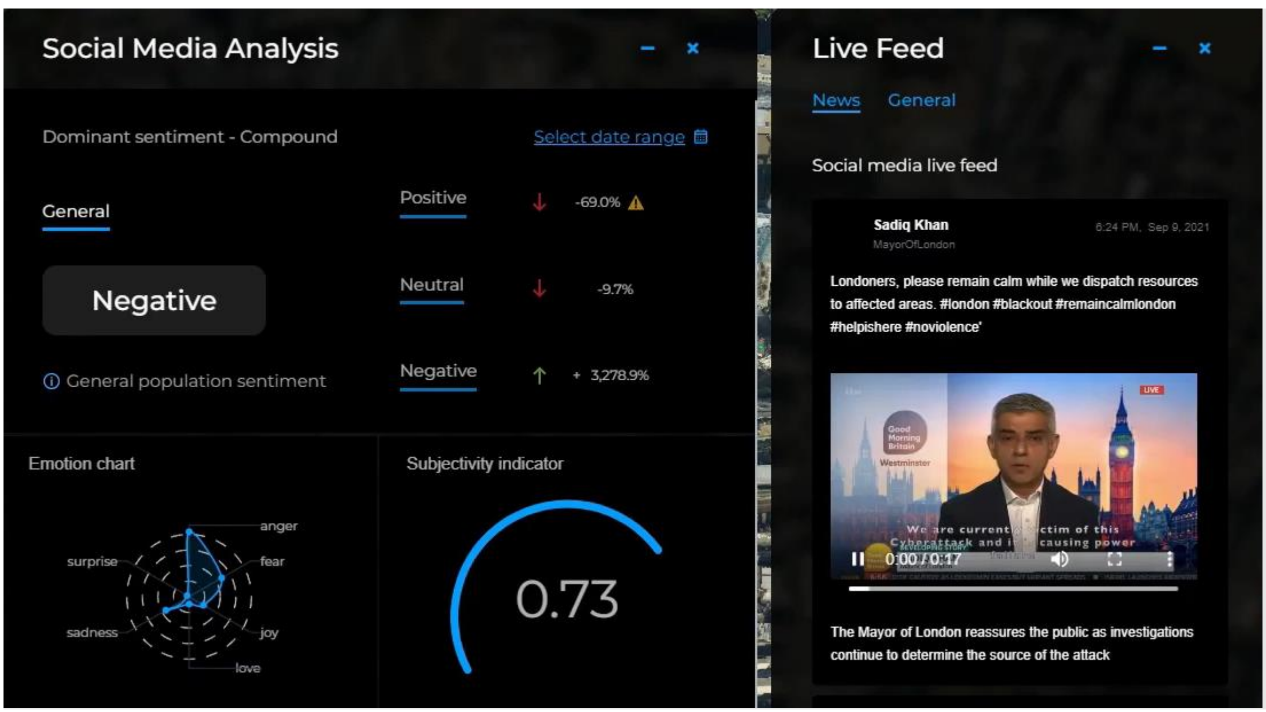

Figure 4. Social Media Panels in the Web UI

The indicators on the Social Media Analysis dashboard display the sentiment and emotion of the entire population as they rise and fall throughout the simulation. The General tab (upper left of the Social Media Analysis dashboard) shows a dominant negative sentiment at this time. To its right, the positive and neutral sentiment are on the decline and negative sentiment is on the rise. This indicates that, in aggregate, the people in the simulated population are experiencing an emotional deterioration of their state of mind. The spider chart (bottom left of the Social Media Analysis dashboard) displays the dominant emotions among the residents of Greater London. The direction of the spider chart indicates that anger and fear emotions are dominant as the cyberattack continues. This qualifies the overall negative sentiment in the population. The subjectivity indicator tab (bottom right of the Social Media Analysis dashboard) shows the aggregated subjectivity score with a value of 0.73, which indicates that social media messages in the Live Feed are mostly subjective (opinion) messages.

On the right panel of the Web UI, synthetic social media messages are displayed as a live feed and are divided into two sections: News and General. The General tab displays messages from the simulated population (and live social media messages if integrated) and the News tab shows synthetic social media messages associated with news, popular media, and key public figures. It should be noted that this live feed could be replaced by real-time social media content depending on the user requirements. We can see the synthetic video of the Mayor of London described above that is embedded in a social media message to create a realistic appearance. This video social media message is injected into the simulation and propagates to the Web UI during the response stage at the peak of the power outage due to cyberattack. The inclusion of the video is a configurable feature in the social media model, which could be extended to support multiple video messages from different public figures or others, if required by the scenario.

INTEGRATION WITH OTHER LAYERS

As noted above, the Cognitive Layer Simulation currently produces social media messages at a rate of 0.1 Hz and updates the analytics every 15 minutes (approximately 0.001 Hz). In contrast, the population movement and 3D visualization elements of the synthetic environment run at typical human-in-the-loop update rates to support any combination of live-virtual-constructive training. To bridge these two worlds, the Social Media UI Proxy is the point of integration of the Cognitive Layer simulation with the other synthetic environment layers.

The Social Media UI Proxy presents a Representational State Transfer (REST) Application Programming Interface (API) that is used by the Human Layer to trigger the different stages of the power failure described above. This information is forwarded to the Cognitive Layer Simulation, also via a REST API. The Cognitive Layer Simulation, in turn, responds to this REST API call with a timestamp of the power outage peak point, which defines the moment when the video social media message is generated and sent to the web UI.

The Social Media UI Proxy retrieves published social media messages and analytics data from the SQLite databases produced by the Cognitive Layer Simulation and forwards them to a back-end web server that relays information to the front-end web UI.

DISCUSSION

Misinformation

Misinformation in social media is both topical and of serious concern. From our perspective in simulating social media messages, it has multiple dimensions: does it appear in the training data (and so could it inadvertently influence the messages being generated); are we attempting to model it in the social media messages being generated; and if live social media messages are integrated in real-time, how will it be handled?

At this stage of development, we have not taken any particular steps to identify misinformation in the training data. If it is present, it is treated in the same way as any other message. Similarly, we currently regard the simulated population as "pure": the social media messages generated by the individuals in the population reflect the conditions they find themselves in (i.e., is the power on or off, and was it due to a cyberattack).

The resulting analytics display the aggregated sentiment of the population as we intend it to be presented to Decision Maker users at the various stages of the simulation. As such, it is correct for the intended purpose: a small amount of misinformation, even if it exists in the data, would not have a significant (or even noticeable) impact on the aggregated sentiment across many social media messages. For an application in which Decision Maker users are required to identify, react to, and mitigate misinformation in social media messages from individuals or groups, more work will be required.

Motives

As described above, the individuals in the simulated population react to events as they occur throughout the scenario. Demographic attributes (age) were used to determine whether an individual has (uses) social media, but demographic and socio-economic attributes do not currently affect the social media messages that are published as the simulation progresses. Similarly, we do not currently attempt to model a mental state of individuals or any underlying motives or goals of the individuals in the population.

User Feedback

Initial feedback from potential users and higher-level decision makers to the inclusion of social media in the synthetic environment has been extremely positive. They regard this as a critical, missing dimension in the overall context of the simulation or replication of the environment. The dynamic nature has been of particular interest as it responds to events generated by the other layers in the synthetic environment rather than just being a static representation of the population sentiment at the start of the simulation. Our prospective users view this as an opportunity to provide a richer, more challenging environment in which to train and conduct analyses than is currently generally available.

As of this writing, the social media synthesis module has not been deployed to end users, therefore we do not have feedback yet on how it has been incorporated into current practices and what the benefits and challenges of doing so are.

Conclusions & Future Work

In future work, we plan to include a misinformation model to identify manipulated and harmful media in the form of text, images, and video that flow into a social media platform. The misinformation labelling would be essential if live messages are retrieved using the API of an existing social media platform and incorporated into the simulation in real-time. We would like to incorporate a topic classifier to illustrate the evolution of topics of interest pre, during and post events. This will help to identify important categories of media that require higher attention and display them for quick action. We plan to embed synthetic images into social media messages to add more realism. This will add value, especially for image-heavy social media platforms.

We plan to utilize the demographic and socio-economic attributes of individuals in the population that are already present or planned for the Human Layer, such as age, gender, ethnicity, and region to influence the response of individuals to the events in the simulation. As a result, sub-populations will be allowed to behave differently to each other.

In the Cognitive Layer Simulation, we plan to enhance the simulation adaptability by developing a flexible framework that allows implementing a broader range of scenarios. Such a change will create an immersive experience that is closer to reality thereby making the simulation more relevant to course of action analysis. We are working toward designing agent-based models that can control emotional states at both the individual and population level, via compartmental modeling.

Finally, we plan to enhance the integration with other synthetic environment layers to permit the Cognitive Layer, and so the resulting social media messages, to be influenced by a greater number of events in the broader simulation, and for the cognitive state of individuals modelled in the Cognitive Layer to influence other layers, such as the destination of individuals in the Human Layer.

In conclusion, we have successfully designed and implemented social media components that generate messages using trained AI models in a simulation of Greater London. We believe that it demonstrates great promise in serving as the foundation of a Cognitive Layer for a class of simulations that have application to decision support systems, the training of decision makers, and course of action analysis.

REFERENCES

DataReportal. (2022). Digital 2022: Global Overview Report. https://datareportal.com/reports/digital-2022-global-overview-report.

Garstka, J. (2003). Network Centric Operations Conceptual Framework Version 1.0.

Giannias, N., Harris, E.P., Geraerts, R., Herfst, S., Hobson-Sayers, A. & Thorpe, A. (2022). Large-Scale Pattern of Life Simulation for Real Time Applications. *I/ITSEC 2022*.

Harb, J.G., Ebeling, R. & Becker, K. (2019). Exploring deep learning for the analysis of emotional reactions to terrorist events on Twitter. Journal of Information and Data Management, 10(2), 97-115.

Hazen, M., Harris, E. & Lamoureux, T. (2017). Use Case Analysis of the Information Warfare Engagement Model Architecture. *Symposium on M&S Technologies and Standards for Enabling Alliance Interoperability and Pervasive M&S Applications*. STO-MP-MSG-149-08. Lisbon, Spain. October 2017.

Hobbins, J.D., Merrick, S., Giannias, N., Kopan, M., Lilley, S. & Mohammed, S. (2022). Geospatial Data Pipelines for Urban Digital Twin Applications. *I/ITSEC 2022*.

Jia, Y., Zhang, Y., Weiss, R.J., Wang, Q., Shen, J., Ren, F., Chen, Z., Nguyen, P., Pang, R., Moreno, I.L. & Wu, Y. (2019). Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *arXiv:1806.04558*.

Kaggle. (2022). Kaggle Datasets. https://www.kaggle.com/datasets.

Kalchbrenner, N., Elsen, E., Simonyan, K., Noury, S., Casagrande, N., Lockhart, E., Stimberg, F., Oord, A. van den, Dieleman, S., & Kavukcuoglu, K. (2018). Efficient neural audio synthesis. *arXiv:1802.08435*.

Laird, J.E. (2012). *The Soar Cognitive Architecture*. The MIT Press.

Mohammad, S.M. (2020). Examining citations of natural language processing literature. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5199-5209.

Pew Research Center. (2021). Social Media Use in 2021.

Prajwal, K.R., Mukhopadhyay, R., Namboodiri, V.P. & Jawahar, C.V. (2020). A lip sync expert is all you need for speech to lip generation in the wild. *Proceedings of the 28th ACM International Conference on Multimedia*. 484-492.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI Blog, 1, 8. https://github.com/openai/gpt-2 and https://openai.com/blog/better-language-models/.

Rao, A. & Georgeff, M. (1991). Modeling rational agents within a BDI-architecture. *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*.

Ren, Y., Ruan, Y., Tan, X., Qin, T., Zhao, S., Zhao, Z., & Liu, T.-Y. (2019). FastSpeech: Fast, robust and controllable text to speech. *arXiv: 1905. 09263*.

Shen, J., Pang, R., Weiss, R.J., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Skerry-Ryan, R.J., Saurous, R.A., Agiomyrgiannakis, Y. & Wu, Y. (2018). Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4779-4783.

Silverman, B.G, Bharathy, G.K. & Weyer, N. (2019). What is a Good Pattern of Life Model? Guidance for Simulations. *Simulation*, vol. 95, no. 8, pp. 693-706.

Space and Naval Warfare Systems Center Atlantic. (2013). Innovative Uses of Social Media in Emergency Management. U.S. Department of Homeland Security.

Wang, Y., Skerry-Ryan, R. J., Stanton, D., Wu, Y., Weiss, R. J., Jaitly, N., Yang, Z., Xiao, Y., Chen, Z., Bengio, S., Le, Q., Agiomyrgiannakis, Y., Clark, R. & Saurous, R.A. (2017) Tacotron: Towards end-to-end speech synthesis. *arXiv:1703.10135*.