

## **USING SIMULATED NARRATIVES TO UNDERSTAND ATTRIBUTION IN THE INFORMATION DIMENSION**

Elijah Bellamy  
David M. Beskow

Department of Systems Engineering  
United States Military Academy  
606 Thayer Road  
West Point, NY 10996, USA

### **ABSTRACT**

Conducting a measured response to cyber or information attack is predicated on attribution. When these operations are conducted covertly or through proxies, uncertainty in attribution limits response options. To increase attribution certainty in the information dimension, the authors have developed a suite of supervised machine learning models that attribute an emerging narrative to historical narratives from known actors. These models were first developed on simulated narratives produced with a Large Language Model. Once the supervised classification models were developed and tested on the simulated narratives, they are evaluated on known actor social media narratives from three known actors. The attribution models are language agnostic and offer one-vs-rest and multi-class options. All models performed at relatively high accuracy and can provide decision support for cyber response decisions.

### **1 INTRODUCTION**

Information warfare dates back to antiquity when the Romans spread rumors that Hannibal, the Carthaginian general, had been killed, attempting to disorganize and demoralized his army (Latimer 2003). More recently, it has been used in newspapers, radio, and pamphlets. Today, it has evolved and is consistently deployed in modern information technology systems (Rid 2020). Actors use modern technologies to create increasingly complex narratives that can trend within hours, reaching a wider audience than ever before. Artificial intelligence and machine learning techniques generate and amplify these information warfare campaigns. Troll farms and bot armies are used to propagate and spread disinformation narratives through the manipulation of a social media platform's algorithms. With the increasing use of large language models (LLMs), it is possible for actors to artificially scale the content as well. Susceptible victims often include social media users that find themselves in an echo chamber of information and opinions designed to slowly change their beliefs and behavior.

Social media platforms and other information systems serve as the battle space for information warfare in both competition and conflict. As the world becomes increasingly intertwined with social media, state and non-state actors utilize sophisticated attacks to accomplish specific ends. United States Secretary of Defense James N. Mattis emphasized the importance of information in September 2017 when he stated:

Information is such a powerful tool that it is recognized as an instrument of national power. The elevation of Information as a joint function impacts all operations. It signals a fundamental appreciation for the military role of information at the strategic, operational, and tactical levels within today's complex operating environment (Mattis 2017).

Information campaigns are designed to bypass homeland security measures in order to attack and manipulate a target society or specific groups in a society. Examples of these attacks include social media manipulation, disinformation campaigns, cyber-attacks, employment of troll and bot armies, and most recently deepfakes (Agarwal et al. 2019). Persistent attacks of this nature can have strategic effects. Since Russia's initial annexation of Crimea in 2014, Russian state-backed disinformation campaigns have fostered narratives of distrust and anti-nationalistic views upon the Ukrainian government (Mejias and Vokuev 2017).

While research is inconclusive on whether retaliation for a cyber operation will lead to deterrence or escalation (Hedgecock and Sukin 2022), we can conclude that "...most states that seek to establish itself as a victim of cyber-armed attack very likely intends to respond forcefully in self-defense" (Payne and Finlay 2016). It is relatively safe to say that a response, particularly a retaliatory response, depends on attribution. While significant literature addresses attribution for traditional cyber intrusion attack, it is sparse on a discussion of attribution for information warfare.

This project attempts to address the need for proper attribution of information attack between nation states. The project uses supervised learning algorithms to develop semantic classification models to attribute information warfare narratives to specific actors. We will begin by developing specific simulated narratives with a large language model (LLM). Once these models prove effective, we will evaluate them on known actor narratives on a social media platform (in our case, Twitter).

## **2 LITERATURE REVIEW**

Cyber attack is simply defined as "information technology employed as a weapon" (Payne and Finlay 2016). Generally speaking, cyber attribution has two branches, often called "what-attribution" (type of attack) and "who-attribution" (perpetrator of attack) (Goel and Nussbaum 2021). Our paper will focus on the "who attribution". Additionally, cyber attribution differs depending on the type of attack: network intrusion (Krekel et al. 2014) or information operations (Shallcross 2017). Our paper focuses on the latter—information operations.

The strategic policy discussion often leans either toward security policy (Mudrinich 2012) or game theory (Edwards et al. 2017). As for specific models, Thomas Rid and Ben Buchanan introduce the "Q" model of attribution, a metamodel focused on the Concept, Practice, and Communication of attribution at the tactical, operational, and strategic levels. Thomas Rid and Ben Buchanan also help describe the art and science of cyber attribution:

On a technical level, attribution is an art as much as a science. There is not recipe for correct attribution, no one methodology or flow-chart of check-list. Finding the right clues requires a disciplined focus on a set of detailed questions—but also the intuition of technically experienced operators....On an operational level, attribution is a nuanced process, not a simple problem. That process of attribution is not binary, but measured in uneven degrees, it is not black-and-white, yes-or-no, but appears in shades. As a result, it is also a team sport—successful attribution requires more skills and resources than any single mind can offer. Optimising outcomes requires careful management and organisational process. On a strategic level, attribution is a function of what is at stake politically (Rid and Buchanan 2015).

Recent research indicates that cyber retaliation policy increasingly relies on attribution. Covert operations that result in a low confidence attribution had the greatest impact on respondents' support for retaliation (Hedgecock and Sukin 2022). When attribution is uncertain, leaders and societies appear reluctant to retaliate. Other response options, including legal action, require attribution (Tsagourias and Farrell 2020).

Some have argued that nation state involvement is not binary, but rather there is a spectrum of involvement (Healey 2011). Rid and Buchanan indicate this above in their statement that attribution is not "...black-and-white, yes-or-no..." (Rid and Buchanan 2015). Actors may use proxies to perpetrate cyber

and information attack, giving them some level of plausible deniability (Canfil 2016). Today, some nation states will leverage businesses, academic institutions, and criminal organizations to perpetrate cyber and information attack. As we present supervised models to use for attribution, it is important to present leaders the probabilities of these models, and not just the predicted class.

Attribution does not need to be public. Given the integration of intelligence collection with digital forensic methods, there can be trade offs in making attribution determinations public (Berghel 2017). If public attribution will reveal intelligence sources and methods, then the benefits may not outweigh the cost.

There are some limits to the discussion of information operations attribution in the literature. The forensic discussion of information operations (Goel and Nussbaum 2021) is limited to a discussion of bot detection algorithms. In our opinion, bot detection algorithms are required to characterize organic vs. inorganic activity, and can have a place in attribution. However, predicting “bot or not” is not the same as predicting a perpetrator. It is also not helpful that a number of papers confuse or mix their discussion of social bots with that of cyber botnets (such as are often used to perpetrate a traditional cyber Digital Denial of Service, or DDoS attack). In our opinion, these are two different types of bots and their discussion as related to attribution should be separate.

Given the lack of information operations attribution techniques in the literature, we set out to develop a supervised machine learning model for attribution in the information dimension. We will develop our models first on simulated narrative using a Large Language Model (in our case, GPT-J), and then evaluate these techniques on real world social media narratives on Twitter.

### 3 SIMULATED NARRATIVE DATA

We used a large language model (LLM) to create three distinct labeled narratives plus a random noise category. We used the Generative Pretrained Transformer “J” (GPT-J) model developed by EleutherAI (Wang, Ben and Komatsuzaki, Aran 2021) using the Mesh Transformer JAX (Wang 2021). The GPT-J model hyperparameters are provided in Table 1. The GPT-J model was trained on the English-Only Pile Data set (Gao et al. 2020), an 800GB English Only Dataset introduced in 2020. GPT-J does not respond to prompts like the newer chatGPT (Brown et al. 2020) because it was not fine-tuned with Reinforcement Learning from Human Feedback (RLHF). However, we found that it was easy to use offline and scale than the API-bound chatGPT. It was also sufficient to create simulated narratives. Finally, as part of the process, we wanted to understand better how an actor could use similar technology to scale their information warfare lines of effort.

Table 1: Default Hyperparameters of the GPT-J model.

Hyperparameter	Value
trainable parameters	6B
layers	28
dimension of model	4,096
feed forward dimension	16,384
# of heads	16
dimension of heads	256
tokenization vocabulary	50,257 (same as GPT-2/GPT-3)

We created three distinct narratives focused on 1) COVID, 2) Climate Change, and 3) European Refugees. In order to create a diverse narrative, we repeatedly generated text with the same general topic focused prompt, restricting the response to 1000 characters. We did a similar approach to create random information noise. In this case we repeatedly prompted the GPT-J algorithm for “random news”, once again restricting the response to 1000 characters. The prompts, size, and examples of these narratives are provided in Table 2. We did not use any fine-tuning procedures in this research effort, though it would have likely increased the precision of the simulated narratives.

Table 2: Description of simulated narratives with GPT-J.

Prompt	Response Max Length	# of Responses	Examples
Climate Change Narrative	1000 characters	19,795	1 As of 2011, the top ten UK energy suppliers al...
			2 Our government policy on energy has been quite...
			3 But the pro-renewables lobby is very effective...
			4 The BBC used to be a respectable and reputable...
			5 Here are a couple of recent examples. Firstly,...
			6 Secondly, a newspaper article suggested that t...
			7 And that is happening all the time. Climate ch...
			8 Another way to explain it is with an example. ...
			9 Climate change is a very complicated subject. ...
			10 It is easy for them to do this because the pro...
Covid Pandemic Narrative	1000 characters	20,480	1 This paper intends to contribute to the emergi...
			2 How Pandemics, and the Narratives Constructed ...
			3 Covid Pandemic Narrative: Is This A New Form o...
			4 A friend of mine from India, who I will call ‘...
			5 “People are getting so afraid of the coronavir...
			6 The conversation in private, that this stateme...
			7 This was the first time I have seen such a dee...
			8 If people are fearful of healthy conversations...
			9 To understand this, we need to understand the ...
			10 How is the community feeling about this curren...
European Refugee Crisis Narrative	1000 characters	18,081	1 It is clear that the current refugee crisis ha...
			2 Is a person who is forced to leave his home or...
			3 The refugee crisis is often referred to as one...
			4 A refugee is a person having a right to remain...
			5 A refugee crisis at time of writing (February ...
			6 During times of war it is difficult for people...
			7 A refugee crisis occurs when large numbers of ...
			8 The current crisis in the Middle East has affe...
			9 Refugee crisis is the largest refugee displace...
			10 The crisis in Syria began when protests agains...
Random News	1000 characters	14,638	1 “I think the greatest service the Republican P...
			2 random news and observations from the world of...
			3 Asimov was born in Russia, but his formative y...
			4 The U.S. Fish and Wildlife Service is seeking ...
			5 random news, random commentary by random perso...
			6 In the latest installment of the great “You’re...
			7 Among other things, he says the NBA will be th...
			8 random news from across the vast waste of the ...
			9 We are very happy to present the debut issue o...
			10 As a preface to a list of some of the “best” s...

#### 4 METHODOLOGY

Our general approach is summarized in Figure 1. The simulated narratives were produced using the GPT-J model discussed above. They produced relatively large labeled narratives aligned to our three distinct narratives plus a random noise category.

These narratives were then embedded at the sentence level with Language-agnostic BERT Sentence Encoder (LaBSE) (Feng et al. 2022). LaBSE is a multilingual version of the mono-lingual BERT model (Reimers and Gurevych 2019). LaBSE was trained on 109 languages and generates an embedding dimension of 768. While the GPT-J model produced English-only data, we wanted the ability to use the same methodology on the multi-lingual narratives that pervade social media.

To evaluate whether or not the narratives were separable, we visualized a sample of all four classes (three distinct narratives plus random noise) using the t-distributed stochastic neighbor embedding (t-SNE) (Van der Maaten and Hinton 2008). By visualizing these in 2D, seen in Figure 2, we assessed that the narratives were likely separable in high dimensional space, and that a supervised machine learning model would be viable.

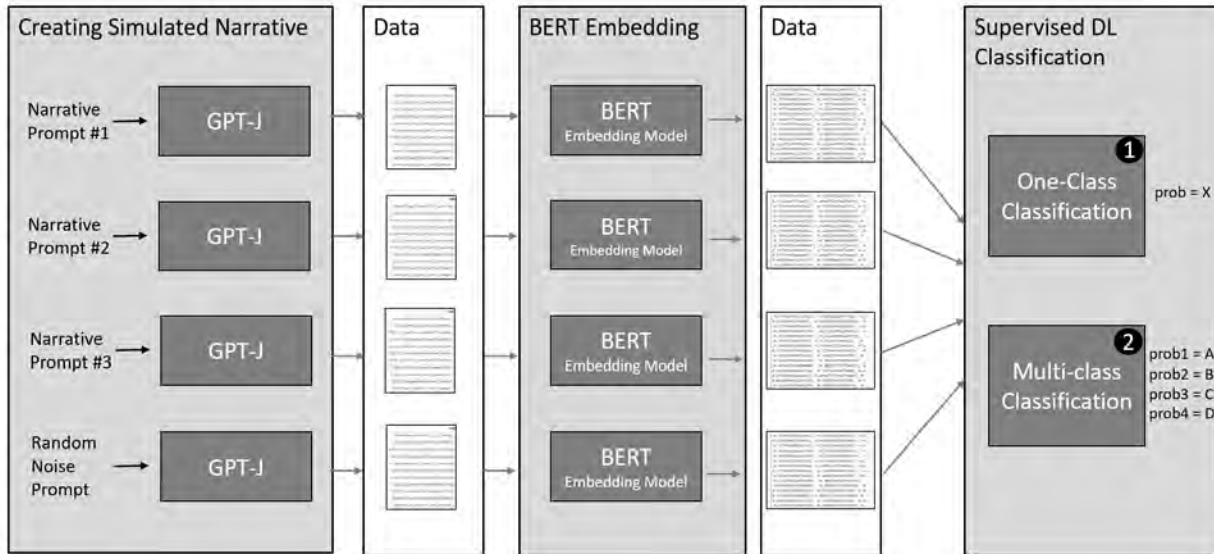


Figure 1: Overview of the proposed methodology and supervised machine learning model.

To evaluate machine learning models, we selected a logistic regression model as the base case. The logistic regression was conducted with L2 normalization penalty term and limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) solver (Liu and Nocedal 1989). The logistic regression provides an adequate model for high dimensions and a base case to compare the neural network model to.

Two different methods were used for the supervised machine learning model: one-vs-rest and multi-class classification. In both cases we used a deep learning model with Tensorflow and Keras to build the model. We developed the model with a train-test approach with 20% reserved for testing in each model. The neural network model was created with the Keras API with four layers: a Dropout layer (to help prevent overfitting by randomly dropping some of the inputs during training), two hidden layers each with 16 units, activated by the Rectified Linear Unit (ReLU) function. The final output layer was activated by the Sigmoid function for one-vs-rest, and softmax function for multi-class. These functions produce a probability between 0 and 1 for the evaluated classes. The data was manipulated to have balanced classes for both the one-vs-rest and the multi-class models. Given the balanced nature of the classes, model accuracy was selected as the primary metric for model evaluation. The deep learning models were also compared against a logistic regression baseline (all models used BERT embedding for features).

## 5 RESULTS

Table 3: One-vs-rest and multi-class results on simulated narratives.

	Model Accuracy			
	COVID	Climate	Refugee	Random Noise
logistic regression baseline (one-vs-res)	0.932	0.918	0.961	0.949
logistic regression baseline (multi-class)	0.838	0.866	0.918	0.871
logistic regression baseline (multi-class overall)			0.874	
neural network one-vs-rest	0.961	0.948	0.969	0.968
neural network multi-class	0.868	0.877	0.892	0.911
neural network multi-class overall			0.892	




Figure 2: t-SNE visualization of the simulated narratives (three plus a random noise narrative).

Model Accuracy for each of the simulated narratives is provided in Table 3. These are provided for both the baseline logistic regression model as well as the neural network model. The results highlight both the one-vs-rest and the multi-class models.

While the neural network model does provide superior performance as measured by model accuracy, the logistic regression provided higher than expected performance. The logistic regression model is a viable model for this use case, given that it tended to trail the more sophisticated neural network by only a few points. For these simulated narratives, the one-vs-rest models provide a higher accuracy than the multi-class models. It is also interesting to note that the model performs just as high on the “random noise” class as other labeled classes (an unexpected result).

We believe that both the one-vs-rest and the multi-class models would be useful as analytic tools. There are times when an analyst has strong evidence that a certain actor is behind a campaign, in which case the one-vs-rest models would help support (or refute) this attribution. In other cases, we may have a new narrative and the perpetrator is largely unknown. In these cases the multi-class model would be preferred to initially provide evidence toward a certain actor perpetrating and guiding the information campaign.

## 6 EVALUATION ON TWITTER SOCIAL MEDIA NARRATIVES

Given the success of the supervised machine learning models on GPT-J simulated data, we next evaluated it on overt information narratives produced by three known actors plus a category for random social media noise. The data consisted of Twitter data from a 6-month window that was overtly attributed to these actors. For each actor we identified 50 to 100 Twitter handles that overtly spread narrative approved by the

respective actor. We then collected all tweets produced by these accounts. The random noise was collected from a random sample of one million tweets extracted from the Twitter 1% Sample for the same time period. A summary of the three actors/narratives and random noise is provided in Table 6.

Table 4: Summary of Twitter narrative data.

Actor	Size (# of Tweets)	Size of Top Three Languages		
Actor 1	837,595	L1 (39%)	L2 (38%)	L3 (7%)
Actor 2	485,076	L1 (39%)	L2 (31%)	L3 (8%)
Actor 3	589,132	L1 (59%)	L2 (9%)	L3 (7%)
Noise	1,000,000	L1 (30%)	L2 (16%)	L3 (7%)

We conducted the exact same workflow on the Twitter data as seen in Figure 1 (except they were collected from Twitter instead of being artificially created with GPT-J). Once again we graph the data in two dimensions with t-SNE in Figure 3. As expected, we see broader dispersion of the Random Noise (Twitter 1% Sample) than the specific actor narratives. We also see a unique discussion at the top of the graph that involves all three actors plus the random noise category. This graph once again gives us confidence that the narratives are distinguishable.




Figure 3: t-SNE visualization of the Twitter narratives (three actors plus a random noise narrative).

The results of the baseline logistic regression model and the neural network models are provided in Table 5. We once again observe strong but trailing performance by the logistic regression baseline model

in both the one-vs-rest and multi-class configuration. We once again observe (unexpectedly) the strong performance of the model on the “random noise” category even with its much higher variance of topics and languages. Most importantly we see accuracy numbers that are high enough to give decision makers comfort that they can use these models to support (but not prove) information warfare attribution.

Table 5: One-vs-rest and multi-class results on social media narratives.

	Model Accuracy			
	Actor 1	Actor 2	Actor 3	Random Noise
logistic regression baseline (one-vs-rest)	0.921	92.6	90.6	0.953
logistic regression baseline (multi-class)	0.861	0.859	0.854	0.930
logistic regression baseline (overall)			0.876	
neural network (one-vs-rest)	0.976	0.968	0.974	0.973
neural network (multi-class)	0.932	0.944	0.927	0.952
neural network multi-class overall			0.938	

Both the logistic regression and neural network models can provide a predicted probability. This measure of probability should always be given to the decision maker as a measure of certainty. These models are designed to provide evidence, not prove attribution. The certainty of the evidence should be considered as attribution decisions are made and potential retaliatory actions are determined.

Our results depend on the assumption that an actor’s overt historical narratives are correlated to their emerging covert narratives. Attribution models aren’t required for overt narratives. Anonymous and covert narratives require attribution. If an actors covert and anonymous lines of effort deviate significantly from their historical overt narratives (which we are using as labeled data), then our approach will not be successful.

## 7 CONCLUSION

GPT-J provided a viable method to produce simulated narratives in the information dimension. GPT-J is associated with GPT-2 and GPT-3 level of capability, and would require significant fine-tuning before it is used to scale the information campaign in production. If an actor had the compute resources and historical data for fine tuning narratives, they could use GPT-J to help scale these narratives. This would remove the need to have large “troll factories” where humans are required to create content for information campaigns. While the newer and more powerful Large Language Models (such as OpenAI’s chatGPT and GPT-4 or Google’s BART model) do have rules in place to prevent their use in information warfare, as open-source and offline versions of these emerge, they could be used to create diverse and nuanced information campaign content.

We found that one-vs-rest or multi-class models could be appropriate in different circumstances. One-vs-rest models are appropriate when a given actor is likely the perpetrator, and analysts require additional evidence of this. Multi-class models are preferred if the perpetrator is largely unknown. In either case, we found the unsupervised t-SNE visualization helpful in understanding the high dimensional semantic space.

We believe future research should continue to understand how nation states will use increasingly sophisticated Large Language Models (LLMs) in this space. Additionally, future research should focus on the impact of language and language diversity on attribution model performance.

This research demonstrated that supervised machine learning models provide a viable option for attribution in information warfare. The proposed models require that overt labeled data associated with actors are available and that their covert information campaigns are correlated to their overt information campaigns. If these assumptions are true, then supervised models can provide important metrics to support (though not likely to prove) attribution. For a complete attribution decision, these machine learning models



would be combined with other analyses and evidence. If cyber attribution involves both an art and a science, we believe supervised machine learning should be part of the science.

## REFERENCES

- Agarwal, S., H. Farid, Y. Gu, M. He, K. Nagano, and H. Li. 2019. "Protecting World Leaders Against Deep Fakes." In *CVPR workshops*, Volume 1, 38–45.
- Berghel, H. 2017. "On the Problem of (Cyber) Attribution." *Computer* 50(3):84–89.
- Brown, T., B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, J. Clark, C. Berner, S. McCannish, A. Radford, I. Sutskever, and D. Amodei. 2020. "Language Models are Few-shot Learners". *Advances in Neural Information Processing Systems* 33:1877–1901.
- Canfil, J. K. 2016. "Honing Cyber Attribution: A Framework for Assessing Foreign State Complicity". *Journal of International Affairs* 70(1):217–226.
- Edwards, B., A. Furnas, S. Forrest, and R. Axelrod. 2017. "Strategic Aspects of Cyberattack, Attribution, and Blame". *Proceedings of the National Academy of Sciences* 114(11):2825–2830.
- Feng, F., Y. Yang, D. Cer, N. Arivazhagan, and W. Wang. 2022. "Language-agnostic BERT Sentence Embedding". In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 878–891. Dublin, Ireland: Association for Computational Linguistics.
- Gao, L., S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and L. Connor. 2020. "The Pile: An 800GB Dataset of Diverse Text for Language Modeling". *arXiv preprint arXiv:2101.00027*.
- Goel, S., and B. Nussbaum. 2021. "Attribution Across Cyber Attack Types: Network Intrusions and Information Operations". *IEEE Open Journal of the Communications Society* 2:1082–1093.
- Healey, J. 2011. "Beyond Attribution: A Vocabulary for National Responsibility for Cyber Attacks". *Brown Journal of World Affairs* 18:8.
- Hedgecock, K., and L. Sukin. 2022. "Responding to Uncertainty: The Importance of Covertness in Support for Retaliation to Cyber and Kinetic Attacks". *Journal of Conflict Resolution*:00220027231153580.
- Krekel, B., P. Adams, and G. Bakos. 2014. "Occupying the Information High Ground: Chinese Capabilities for Computer Network Operations and Cyber Espionage". *International Journal of Computer Research* 21(4):333.
- Latimer, J. 2003. *Deception in War: The Art of the Bluff, the Value of Deceit, and the Most Thrilling Episodes of Cunning in Military History from the Trojan Horse to the Gulf War*. 1st ed. New York City: Abrams.
- Liu, D. C., and J. Nocedal. 1989. "Limited Memory BFGS Method for Large Scale Optimization". *Mathematical Programming* 45(1-3):503–528.
- Mattis, J. 2017. "Information as a Joint Function". *Official Memorandum, Department of Defense, Washington, DC, USA*.
- Mejias, U. A., and N. E. Vokuev. 2017. "Disinformation and the Media: the Case of Russia and Ukraine". *Media, culture & society* 39(7):1027–1042.
- Mudrinich, E. M. 2012. "Cyber 3.0: The Department of Defense Strategy for Operating in Cyberspace and the Attribution Problem". *AFL Rev.* 68:167.
- Payne, C., and L. Finlay. 2016. "Addressing Obstacles to Cyber-attribution: A Model Based on State Response to Cyber-attack". *Geo. Wash. Int'l L. Rev.* 49:535.
- Reimers, N., and I. Gurevych. 2019. "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks". In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992.
- Rid, T. 2020. *Active Measures: The Secret History of Disinformation and Political Warfare*. London: Farrar, Straus and Giroux.
- Rid, T., and B. Buchanan. 2015. "Attributing Cyber Attacks". *Journal of Strategic Studies* 38(1-2):4–37.
- Shallcross, N. J. 2017. "Social Media and Information Operations in the 21st Century". *Journal of Information Warfare* 16(1):1–12.
- Tsagourias, N., and M. Farrell. 2020. "Cyber Attribution: Technical and Legal Approaches and Challenges". *European Journal of International Law* 31(3):941–967.
- Van der Maaten, L., and G. Hinton. 2008. "Visualizing Data using t-SNE". *Journal of Machine Learning Research* 9(11).
- Wang, Ben 2021, May. "Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX". <https://github.com/kingoflolz/mesh-transformer-jax>.
- Wang, Ben and Komatsuzaki, Aran 2021, May. "GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model". <https://github.com/kingoflolz/mesh-transformer-jax>.

**AUTHOR BIOGRAPHIES**

**ELIJAH J. BELLAMY** is pursuing a Bachelor of Science in Systems Engineering with a minor in Cyber Security and will graduate in the Spring of 2024. His research interests lie in the application of machine learning, operations research, and systems engineering for cyber applications. His email address is [elijah.bellamy@westpoint.edu](mailto:elijah.bellamy@westpoint.edu)

**DAVID BESKOW** is an Assistant Professor in the Department of Systems Engineering at the United States Military Academy at West Point. His research interests lie in the application of the fields of natural language processing, supervised and unsupervised machine learning, and network science. His email address is [david.beskow@westpoint.edu](mailto:david.beskow@westpoint.edu).