UNIDIR

Funded by
the European Union

FULL REPORT

Exploring Synthetic Data for Artificial Intelligence and Autonomous Systems

A Primer

HARRY DENG

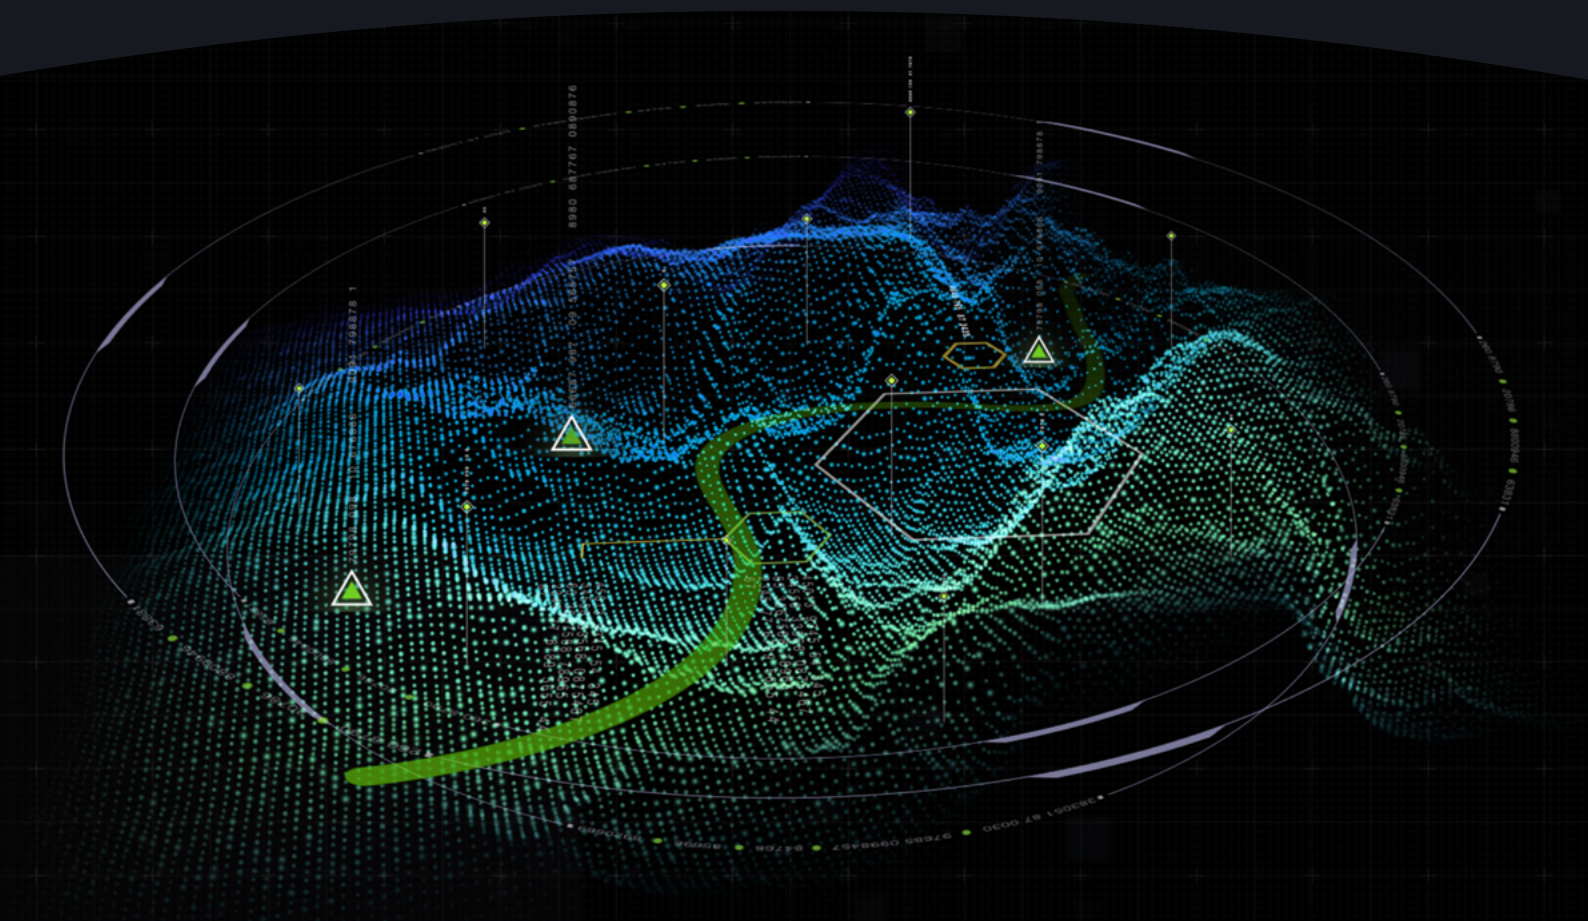

Acknowledgements

Support from UNIDIR's core funders provides the foundation for all of the Institute's activities. This publication was funded by the European Union as part of UNIDIR's Security and Technology Programme, which is supported by the governments of Czech Republic, Germany, Italy, the Netherlands and Switzerland, and by Microsoft. The author wishes to thank Dr. Giacomo Persi Paoli and Ioana Puscas for their advice and assistance on this paper as well as Prof. Tim Watson and Dr. Leslie Sikos for their invaluable contributions to this research.

About UNIDIR

The United Nations Institute for Disarmament Research (UNIDIR) is a voluntarily funded, autonomous institute within the United Nations. One of the few policy institutes worldwide focusing on disarmament, UNIDIR generates knowledge and promotes dialogue and action on disarmament and security. Based in Geneva, UNIDIR assists the international community to develop the practical, innovative ideas needed to find solutions to critical security problems.

Citation

H. Deng, *Exploring Synthetic Data for Artificial Intelligence and Autonomous Systems: A Primer*, Geneva, Switzerland: UNIDIR, 2023.

Note

About the Security and Technology Programme

Contemporary developments in science and technology present new opportunities as well as challenges to international security and disarmament. UNIDIR's Security and Technology Programme (SecTec) seeks to build knowledge and awareness on the international security implications and risks of specific technological innovations. It convenes stakeholders to explore ideas and develop new thinking on ways to address them.

Author

Harry Deng (@hwrdeng) is a Consultant for the Security and Technology Programme at UNIDIR, where his work focuses on the international security implications of emerging technologies. He holds a master's degree in global governance from the University of Waterloo, where he is currently a PhD candidate.

Acronyms & Abbreviations

3D	Dull, Dirty and Dangerous
AI	Artificial Intelligence
AQI	Air Quality Index
DOD	Department of Defense (United States)
GAN	Generative Adversarial Network
GGE	Group of Governmental Experts
ICT	Information and Communications Technology
IOT	Internet of Things
JSON	JavaScript Object Notation
OEWG	Open-ended Working Group
UAV	Uncrewed Aerial Vehicle
VAE	Variational Autoencoder

Contents

Executive Summary

Advances in the field of artificial intelligence (AI) and machine learning in recent years have created unprecedented opportunities to augment human capabilities and to improve the functionality of various autonomous systems, including in the field of international security. Yet, there is a scarcity of the high-quality, highly diverse and relevant real-world data sets that are needed to train increasingly complex AI systems in the defence sector. As a consequence, synthetic data is gradually becoming an essential tool in the data toolbox to develop and train AI systems. The characteristics and potential benefits of synthetic data, along with proven application of the technology in various sectors, make it a relevant topic for debates surrounding the use of AI within the context of international security.

This primer provides a brief overview of synthetic data, including its characteristics, how it is generated, the value that it adds, its risks, and its potential use cases for defence organizations and military operations. In addition, the primer provides an outline of existing data challenges and limitations that have facilitated the emergence of synthetic data as an important tool for the development of increasingly complex AI systems.

The use of synthetic data within the context of international security has so far mostly remained experimental and exploratory. However, the features of synthetic data could have a beneficial effect on training AI systems. In particular, synthetic data allows for the generation of highly diverse or even novel data sets, fine grain control of data attributes, automatic annotation or data labelling where necessary, and cost-effectiveness. This primer examines how the main characteristics of synthetic data could benefit militaries and defence organizations by allowing them to integrate more capable and reliable AI systems in both defensive and offensive autonomous systems.

While synthetic data can be beneficial for training AI systems and could help alleviate some of the data issues faced by militaries and defence organizations, it is not a silver bullet, and it comes with risks and challenges. The benefits accrued from using synthetic data will depend on the ability of organizations to navigate these risks in order for AI systems trained on synthetic data to be used in a responsible and safe manner and in accordance with legal requirements and ethical values.

1. Introduction

Advances in artificial intelligence (AI), along with the machine learning models that support its use, have made it ubiquitous when optimizing performance for increasingly complex tasks and complicated working environments. Yet the integration of AI introduces unprecedented legal, ethical, safety and security challenges – and this is especially relevant in the international security context.[1] Within this context, AI is being explored as a tool for decision support, operational planning and intelligence analysis. It could also be integrated into both offensive and defensive autonomous systems, such as target-identification systems, swarm robotics and cyber operations. Moreover, it has been asserted that AI could perform certain tasks better than traditional methods (e.g., defensive cyber infrastructures or intelligence analysis),[2] meaning that states could be better equipped to uphold their international legal obligations, specifically international humanitarian law, in addition to enhancing operational effectiveness.

At the same time, the downstream effect of the tasks envisioned for AI means that machine learning models require ever increasing diversity, volume and velocity of high-quality data, often high-quality labelled data. Without the necessary diversity, volume, and velocity of high-quality data to train complex AI systems, such systems could see increases in failures, including unintended harms. Labelled data explicitly informs the machine learning model what the data means rather than leaving the model to figure it out by itself and possibly get it wrong. However, the scarcity of high-quality real-world data along with the privacy, legal, regulatory and cost challenges associated with sensitive data make real-world data generally unsuitable for training increasingly complex AI systems,[3] particularly in the international security context. Because of this scarcity, synthetic data is gradually becoming an essential tool to develop, improve and condition increasingly complex AI systems. In particular, it can provide data where there is none, can counterbalance various forms of bias in real-world data and can automatically label data where necessary, among other things.[4]

---

1 The First Committee of the United Nations General Assembly defines "international security" as "global challenges and threats to peace that affect the international community". See United Nations General Assembly, "Disarmament and International Security (First Committee)", https://www.un.org/en/ga/first/.

2 A. Wilner, "AI and the Future of Deterrence: Promises and Pitfalls", Centre for International Governance Innovation, 28 November 2022, https://www.cigionline.org/articles/ai-and-the-future-of-deterrence-promises-and-pitfalls/; Defence Innovation Board, "AI Principles", 2019, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF.

3 J. Yan et al., "Synthetic Dataset Generation and Adaptation for Human Detection", DEVCOM Army Research Laboratory, November 2020, https://apps.dtic.mil/sti/pdfs/AD1115446.pdf; A. Holland, "Known Unknowns: Data Issues and Military Autonomous Systems", UNIDIR, 17 May 2021, https://unidir.org/known-unknowns; Government Business Council, "Advancing ISR at the Edge: A Survey on Networks and Processing Technologies in the Digital Battlespace", July 2020, http://cdn.govexec.com/media/advancing-isr-at-the-edge-isr.pdf.

4 Yan et al., "Synthetic Dataset Generation", 1; Holland, "Known Unknowns", 27.

However, the implications of using synthetic data in autonomous systems remain unexplored in relevant United Nations security processes, such as the Group of Governmental Experts (GGE) on emerging technologies in the area of lethal autonomous weapons systems (LAWS) or the Open-ended Working Group (OEWG) on security of and in the use of information and communications technologies (ICTs). However, the value-added and the risks associated with synthetic data are relevant to these discussions as well as other debates on the use of AI within the context of international security. For example, some parties involved in debate in the GGE on LAWS are concerned about the possibility that increased autonomy in weapon systems could lead to increases in unintended harms due to a lack of training data to appropriately train such systems.[5] Additionally, participants in the OEWG on ICTs have discussed the possibility that AI-powered cyberattacks could autonomously adapt to defensive cyber measures, making them more difficult to detect and mitigate.[6] AI-powered cyberattacks can indeed be enabled and augmented by the generation and use of synthetic data. It is therefore essential that the use of synthetic data in autonomous systems does not derogate any commitments to international law (e.g., international humanitarian law) or Responsible AI – that is, the broad approach to the development and use of AI to ensure that AI systems are lawful, ethical, safe, secure and responsible.[7]

As such, this primer aims to provide policymakers and diplomats engaged in international security discussions with an introductory overview of synthetic data. It describes the main characteristics, value-added, risks and relevance of synthetic data within the context of international security, particularly as an enabler of autonomy. The primer further attempts to demonstrate the growing importance of synthetic data as well as the evolving paradigm of data usage and governance in the international security context. It does this by illuminating and mapping out the peculiarities of synthetic data, then re-connecting it to existing data challenges.

---

5 Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems, "Proposal for an International Instrument on Lethal Autonomous Weapons (LAWS)", Submitted by Pakistan, CCW/GGE.1/2023/WP.3/Rev.1, 8 March 2023, http://undocs.org/en/CCW/GGE.1/2023/WP.3/Rev.1; "State of Palestine's Proposal for the Normative and Operational Framework on Autonomous Weapons Systems", Submitted by Palestine, CCW/GGE.1/2023/WP.2/Rev.1, 3 March 2023, http://undocs.org/en/CCW/GGE.1/2023/WP.2/Rev.1.

6 H. Alkhzaimi, "Contribution to the Fifth Substantive Session by Emerging Research and Security Center, NYU/NYUAD", NGO Working Papers, 28 July 2023, https://docs-library.unoda.org/Open-Ended_Working_Group_on_Information_and_Communication_Technologies_-_(2021)/Stakeholder_Recommendation_for_Open-ended_working_group_on_security_APR.pdf.

7 A. Anand and H. Deng, "Towards Responsible AI in Defence: A Mapping and Comparative Analysis of AI Principles Adopted by States", UNIDIR, 13 February 2023, https://unidir.org/publication/towards-responsible-ai-defence-mapping-and-comparative-analysis-ai-principles-adopted.

2. Understanding Synthetic Data

Highlights

- Unlike real-world data, which refers to data and inputs derived from the real world, synthetic data is artificially created in the digital world. It often seeks to reproduce the characteristics and properties of an existing set of data or to produce data based on existing knowledge.
- The purpose of synthetic data is to improve the quality and utility of training data sets. It is critical that the data on which autonomous systems are trained is of sufficient quality and diversity in order to avoid unintended harms, especially within the context of international security.
- Defence organizations currently face a myriad of data management challenges, thereby limiting the quality and utility of real-world data to train increasingly complex AI and autonomous systems.
- While synthetic data may not be a silver bullet that resolves all existing data challenges within defence organizations, it may provide a means to improve the quality and utility of training data sets.

2.1 What is Synthetic Data?

Synthetic data is data that is artificially generated in the digital world with properties that are often derived from an "original" set of data. This is in contrast to real-world data, which, as the name suggests, is data collected from real-world events and inputs. The "original" data set is typically real-world data and information, but it can also be artificially generated data itself. While there are various methods of generating a synthetic data set (see section 2.3), the objective is often to reproduce the characteristics and structures of the original data set, and most methods rely on extracting and replicating the properties of the original data (e.g., see Figure 1).[8] This means that the synthetic data and the original data should deliver very similar,

8 S. Kannan, "Synthetic Time Series Data Generation for Edge Analytics", F1000 Research, 20 January 2022, https://doi.org/10.12688/f1000research.72984.1.

if not identical, results when undergoing the same statistical analysis.[9]

In short, synthetic data is often information that is artificially generated to represent the original data it either seeks to replace (thereby providing an equivalent function) or complement (thus improving the value of the training data set). It is also possible to enhance a training data set by generating synthetic data that does not reproduce the characteristics of the original data set, but instead exaggerates certain characteristics (see section 3.1).

Figure 1. Real-World versus Synthetic Data[10]

REAL-WORLD IMAGE OF A TANK

SYNTHETICALLY GENERATED IMAGE OF A TANK

However, in some circumstances, synthetic data can also be artificially generated data that does not rely on an original set of data. It is possible to generate novel data based on existing knowledge. For example, it is possible to synthesize data on how dice of different weights would behave based on existing knowledge of the physics of the object. In these cases, rather than reproducing the characteristics of an original data set, synthetic data produces data that reflects the characteristics of the system that would hypothetically produce that data.

9 R. Riemann, "Synthetic Data", European Data Protection Supervisor, https://edps.europa.eu/press-publications/publications/techsonar/synthetic-data_en.

10 F. Longford, "Experiments in Synthetic Data", Forensic Architecture, 6 November 2018, https://forensic-architecture.org/investigation/experiments-in-synthetic-data.

2.2 Existing Data Challenges

Defence organizations and militaries face challenges in obtaining sufficient data of adequate quality. The data challenges are not only technical, but also organizational.[11] This means that defence organizations cannot simply "engineer their way out" of their shortcomings using technical solutions. Instead, data challenges within defence organizations should also consider the impact of organizational culture, policies and procedures. Ultimately, any use of autonomous systems, particularly autonomous systems in combat environments or autonomous systems intended to engage human targets, hinges on the responsibility to anticipate and respond to data issues in order to avoid unintended harms.

While synthetic data may not be a silver bullet that alleviates all existing data challenges, it could provide a means to improve the quality and utility of training data sets, especially in increasingly complicated and opaque machine learning models where data issues may not be easily revealed. This section looks in turn at two strands of data challenges faced by defence organizations that synthetic data can address: first data management, and then data quality.

2.2.1 Strand 1 – Data Management

The first strand of issues that results in a lack of labelled data of sufficient quality is poor data management practices. The data management process includes the following stages:

1. Collection
2. Processing (e.g., data labelling)
3. Storage
4. Access
5. Sharing and dissemination
6. Extraction and exploitation
7. Utilization
8. Disposal

The scale and form of data management problems can be illustrated by the case of the United States Department of Defense (DoD), where processing, exploitation and dissemination of data were noted to be particularly challenging. In a study, only 29 per cent of both active-duty and civilian personnel said that over 75 per cent of the data reaches the appropriate actors.[12] These numbers are even bleaker for active-duty personnel, only 11 per cent of whom said that data reaches the appropriate analysts at least 75 per cent of the time.[13] Furthermore, 65 per cent of active-duty personnel noted that warfighters spend more time looking for the right data than using the

11 Government Business Council, "Advancing ISR at the Edge".

12 Ibid, 4.

13 Ibid.

data.[14] This indicates shortcomings in establishing the appropriate processes to correctly label data, store it in the appropriate databases, and ensure appropriate access and availability. It may also indicate a struggle to balance the need to protect sensitive or classified data with the need to share that data with those who may benefit from exploiting it.

In fact, siloed data, the compartmentalization of mutually exclusive departments, limited bandwidth, and limited tagging of data were noted by DoD personnel to be some of the most prevalent challenges for their organization's ability to effectively collect, disseminate and analyse data.[15] Regarding data labelling, only 32 per cent of defence civilians and 13 per cent of active-duty personnel said that their defence agency had the systems in place to effectively label data.[16] The number of personnel required along with the necessary processes and infrastructure to monitor and manage incoming data has failed to keep pace with the ever-increasing volume of data. The paucity of data quality control – either *ex ante* or *ex poste* – means that analysts are drowning in data and by the time the correct set of data has been obtained it is often obsolete and unreliable.

Similar data management challenges are also mentioned in Australia's 2021 Defence Data Strategy,[17] the United Kingdom's 2021 Data Strategy for Defence[18] and Canada's 2021 Department of National Defence Data Strategy,[19] as well as in research on information networks conducted by the Indonesian Defence University.[20] Common issues found across these documents include, for example, challenges with data visibility, siloed data, lack of common data standards within and across organizations, and cultural issues of not considering data requirements in the initial phase of capabilities development.

---

14 Ibid, 8.

15 Ibid, 15.

16 Ibid.

17 Australian Department of Defence, "Defence Data Strategy 2021–2023", https://www.defence.gov.au/sites/default/files/2021-08/Defence data strategy.pdf.

18 British Ministry of Defence, "Data Strategy for Defence: Delivering the Defence Data Framework and Exploiting the Power of Data", 27 September 2021, https://www.gov.uk/government/publications/data-strategy-for-defence/data-strategy-for-defence.

19 Canadian Department of National Defence, "The Department of National Defence and Canadian Armed Forces Data Strategy", 18 May 2021, https://www.canada.ca/en/department-national-defence/corporate/reports-publications/data-strategy/data-strategy.html.

20 P.A. Udayana et al., "Strategy for Integrated Land Information System Network Arrangements for the Indonesian National Army", 2022, https://doi.org/10.33172/jspd.v8i1.1054.

2.2.2 Strand 2 – Data Quality

Poor data management contribute to the second strand of issues – poor data quality. Common data quality issues include incomplete data, unlabelled data, poisoned or spoofed data, inaccurate data, data bias, and discrepant data. Poor data quality can be the result of external factors, such as harsh conditions (e.g., dust, smoke, vibrations, contaminants, camouflage, wear and tear of sensors, etc.) and adversarial actions (e.g., signal jamming, data poisoning, attacks on sensors, unanticipated tactics, etc.). However, proper data management practices can help filter out data that has been corrupted in order to avoid creating distortions or biases in the training data set and to ensure that the right data reaches the appropriate entities.

Different autonomous systems will face different types of poor data quality issue. For example, an autonomous system used in a defensive cyber operation is less likely to face issues arising from harsh conditions (e.g., dust, smoke, contaminants, etc.), but is likely to face adversarial actions such as spoofing or data poisoning. In contrast, an uncrewed vehicle placed in an "uncontrolled" multivariate combat environment may face both harsh conditions and adversarial actions.

If autonomous systems rely on the data they are trained on in order to navigate, respond to and manipulate their environment, it is critical that the data is of sufficient quality and diversity.[21] It is, however, important to note that not all AI systems rely on being trained by data; reinforcement learning models – which use a reward function to learn the consequences of actions taken – can also be used.[22]

Yet, a certain amount of poor quality data should be expected in any large real-world data set, especially in the international security context where adversarial environments, whether in the digital space or the physical space, pose a wide range of challenges to the collection of complete high-quality data.[23] As such, it has been suggested that synthetic data could play an important role in alleviating some of the pressures of collecting quality real-world data, for example by filling in where data is missing due to sensor failures.[24]

---

21 Holland, "Known Unknowns", 3.

22 Interview with Prof. Tim Watson, 11 April 2023.

23 Holland, "Known Unknowns", 6.

24 Interview with Prof. Tim Watson, 11 April 2023.

2.3 Methods of Generating Synthetic Data

Synthetic data can be generated by leveraging various techniques, such as decision trees or deep-learning algorithms. As a proxy, synthetic data can be classified according to the type of the original data:

- Real-world data
- Information or knowledge of the developer
- Combination of real-world data and the information of the developer

As noted, synthetic data generation relies on extracting and replicating the properties of an original data set. The method with which the properties of the original data set are extracted and replicated depend on the type and structure of the original data. There are three broad methods of synthetic data generation:

- Rules-based methods, which have predefined data structures
- Agent-based models, which simulate environments the data may be needed for
- Deep-learning algorithms, which uses neural network-based methods

The following subsections provide a brief explanation of each broad method.

2.3.1 Rules-Based Methods

Rules-based methods represent metadata and human- or computer-readable data that consists of predefined data structures (e.g., arrays[25]) that provide ordered lists and objects that follows specific rules defined by humans. The complexity of these rules can vary from simple rules that only take into account specified data types in a column to more sophisticated rules that define relationships between multiple parameters and variables. Common data formats include comma-separated values (CSV), JavaScript Object Notation (JSON) and document type definition (DTD). Rules-based methods are modular and cost-effective, and they can support different statistical distributions, which may be particularly pertinent for training AI systems used in cyber operations where standardized protocols necessitate predefined data structures.

Rules-based methods of synthetic data generation have already been applied in other contexts, outside international security. For example, Kannan (2021) used a JSON structure to generate a synthetic data set derived from an Air Quality Index (AQI) data set. In this particular study, Kannan was able to produce a synthetic AQI data set that outperformed the original AQI

25 In programming, an array is a data structure that consists of a collection of values or variables (e.g., numbers, words, objects, etc.) formatted and sorted according to their type. The purpose of an array is to store multiple pieces of data of the same type together.

data set in training a machine learning model to predict the four AQI parameters specified in the data sets.[26] Kannan concludes that the improved performance of the synthetic data set could be the result of the synthetic data set "filling in" the incomplete data contained in the original data set.[27]

There are, however, limitations to using a rules-based method to generate synthetic data sets. Most notable are challenges with scalability, drift and bias.[28] Regarding scalability, the more complex a synthetic data set is (e.g., if the synthetic data set requires thousands of interdependent and intertwined rules), the more complicated and abstruse it is to generate. This thereby limits the practicality of using a rules-based method for more complex or intricate relational networks. Relatedly, data drift – that is, the shift in data distribution over time – may limit the practicality of rules-based methods, particularly if there is no well-established change-management system to govern how the rules are changed to fit their application. Lastly, since the rules are defined by humans, the bias of the developer is reflected in the generated data, whether it is conscious (e.g., business logic) or unconscious (e.g., gender bias[29]).

2.3.2 Agent-Based Models

Agent-based modelling is a proven simulation technique that has seen various real-world applications, from business problems to public policy evaluations. Agent-based modelling is essentially a system that describes agents and the relationships between them in order to derive an outcome. Based on their interactions, behavioural patterns and the inputted parameters, the agents are capable of evolving. This allows for unanticipated behaviours to emerge,[30] and makes agent-based modelling especially applicable for capturing emergent phenomena. Even simple agent-based models can exhibit complex patterns and can provide valuable information about real-world dynamics.[31] Deep-learning can also be incorporated into an agent-based model to achieve more dynamic interactions between agents and more realistic, complex and adaptable outcomes.

---

26 Kannan, "Synthetic Time Series Data Generation", 7.

27 Kannan notes that the incomplete data in the original AQI data set is due to sensor failure at one of the stations that caused partial recording of the data. See Ibid.

28 M. Pasieka, "A Comparison of Synthetic Data Generation Methods and Synthetic Data Types", Mostly AI, 1 September 2022, https://mostly.ai/blog/comparison-of-synthetic-data-types/.

29 K. Chandler, "Does Military AI Have Gender? Understanding Bias and Promoting Ethical Approaches in Military Applications of AI", UNIDIR, 7 December 2021, https://doi.org/10.37559/GEN/2021/04.

30 E. Bonabeau, "Agent-Based Modeling: Methods and Techniques for Simulating Human Systems", PNAS, 14 May 2002, https://doi.org/10.1073/pnas.082080899.

31 Ibid.

Example[32]

An autonomous system that could identify people who are escaping a flood by standing on their roofs would be useful for humanitarian aid and disaster recovery (HADR) operations. However, since such a situation may occur only infrequently in real life, there is very little data that can be used to train an autonomous system to identify this scenario. An agent-based model may simulate such a scenario. By generating high-fidelity and highly diverse synthetic data to augment the training data set with rare data points, the synthetic data generated from an agent-based model could potentially be beneficial in training autonomous systems for such scenarios.

2.3.3 Deep-Learning Algorithms

Deep-learning algorithms are a class of methods based on *representation learning*, which refers to machine learning techniques that automatically learn features and statistical distributions of training data and can generate new data based on those learned features and statistical distributions. Unlike rules-based methods or agent-based models of synthetic data generation, human guidance and supervision can be minimal, or even non-existent, depending on the deep-learning model used. Additionally, deep-learning models are not limited by the complexity of the data that can be learned by such models and, in theory, the application of deep-learning to generate synthetic data sets is "limitless".[33] These algorithms can manage more intricate and complicated data distributions than rules-based methods and can synthesize unstructured data, such as visual data.

There are three main families of deep-learning techniques: generative adversarial networks (GANs), variational autoencoders (VAEs) and diffusion models.

32 Yan et al., "Synthetic Dataset Generation", 3.

33 Pasieka, "A Comparison of Synthetic Data Generation Methods".

Figure 2. Generative Adversarial Network[34]

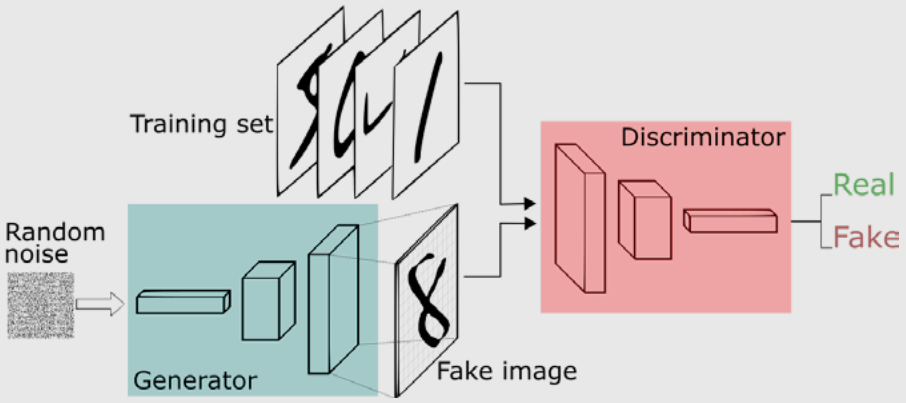

Generative adversarial networks are commonly used for image recognition and image generation.[35] They are usually comprised of two neural networks,[36] one a generator network and the other a discriminator network, that train each other on an iterative basis (see Figure 2). The generator network would produce a synthetic data point (e.g., an image) as an input with the same characteristics as the training data. The discriminator network, containing batches of both training data and synthetic data, would then try to classify the observations as real or generated. The generator network improves its performance over-time based on the feedback it receives from the discriminator network. The two networks then converge when the discriminator is no longer able to differentiate between the "real" data and the synthetically generated data.[37]

34 T. Silva, "A Short Introduction to Generative Adversarial Networks", Thalles' Blog, 7 June 2017, https://sthalles.github.io/intro-to-gans/.

35 Riemann, "Synthetic Data".

36 A neural network, also known as artificial neural network, is a network of interconnected layers of nodes that transmit information from one layer to another and each layer performs a different function on its inputs. Neural networks rely on training data to learn, and their performance improves over time. See IBM, "What are Neural Networks?", https://www.ibm.com/topics/neural-networks.

37 J. Hradec et al., "Multipurpose Synthetic Population or Policy Application", European Commission Joint Research Centre, 13 April 2022, 14, https://doi.org/10.2760/50072.

Figure 3. Variational Autoencoder[38]

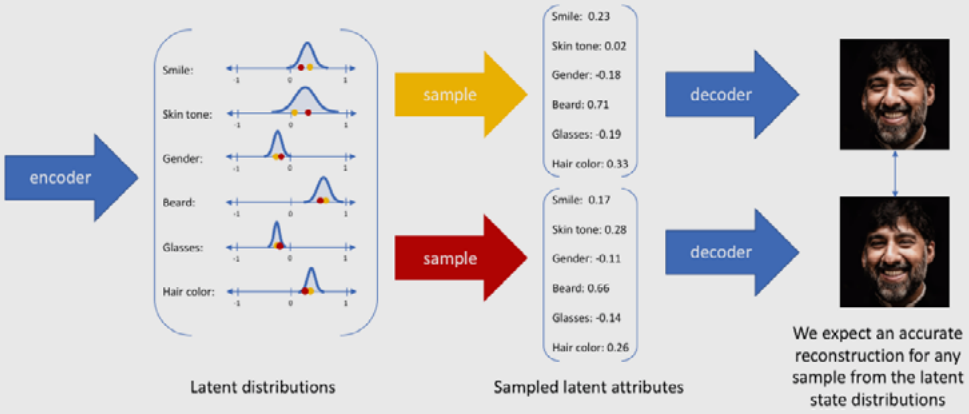

Variational autoencoders are a type of likelihood-based generative model. VAEs are comprised of an encoder and a decoder (see Figure 3). The encoder ingests data and simplifies it (known as the "latent representation") to represent the key features of the data. The decoder takes in the latent representation and returns a reconstruction of it. Like GANs, VAEs function on an iterative basis. At each iteration, the VAE ingests data which is then compared with the encoder–decoder output. The essential function for a VAE, then, is to learn the optimal encoding–decoding scheme to iteratively optimize the process. As such, more complex VAE architectures can support higher dimensionality reduction (i.e., learning the key features) while keeping reconstruction errors low.[39]

38 J. Jordan, "Variational Autoencoders", Jeremy Jordan, 19 March 2018, https://www.jeremyjordan.me/variational-autoencoders/.

39 D.P. Kingma and M. Welling, "An Introduction to Variational Autoencoders", Foundations and Trends in Machine Learning, 2019, https://arxiv.org/pdf/1906.02691.pdf.

Figure 4. Diffusion Model[40]

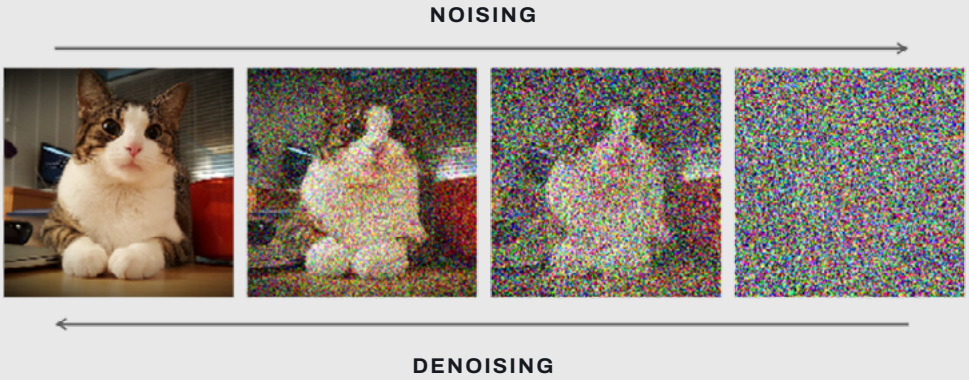

Diffusion models are an emerging class of deep-learning models that produce data, such as images, from a training distribution via an iterative denoising[41] process. In other words, diffusion models work by corrupting an image (e.g., by adding noise), which the model then learns how to remove (or denoise) in order to generate a coherent image (see Figure 4). It can then generate variations of that image by introducing different noises to the otherwise coherent image. While GAN was a breakthrough technology that enabled the production of high-fidelity images at scale, diffusion models have largely displaced GANs in recent years.[42]

It has been posited that, due to its ability to synthesize novel high-fidelity images that are ostensibly unlike its training data as well as its ease of use, diffusion models are the de facto method for generating large-scale images.[43] Popular diffusion models include DALL-E and Stable Diffusion.

40 A. Vahdat and K. Kreis, "Improving Diffusion Models as an Alternative to GANs, Part 1", NVIDIA, 26 April 2022, https://developer.nvidia.com/blog/improving-diffusion-models-as-an-alternative-to-gans-part-1/.

41 The term "denoising" refers to the process of removing imperfections and defects from audio-visual data in order to restore the actual features and characteristics. See L. Fan et al., "Brief Review of Image Denoising Techniques", Visual Computing for Industry, Biomedicine, and Art, vol. 2 (2019), https://doi.org/10.1186/s42492-019-0016-7.

42 N. Carlini et al., "Extracting Training Data from Diffusion Models", arXiv:2301.13188, 30 January 2023, 1, https://arxiv.org/abs/2301.13188.

43 Ibid.

3. Synthetic Data and International Security

Highlights

- Militaries and defence organizations could benefit from continued advances in AI and autonomous systems. Ensuring that AI and autonomous systems are properly trained prior to deployment and use is of critical importance in the international security context.
- Benefits of synthetic data include highly diverse data sets, shortened training cycles, fine grain control and flexibility, ability to generate hypothetical data, and identifying and addressing skewed data sets, among others.
- Synthetic data may also eliminate legal challenges related to collecting, storing, disseminating and disposing of sensitive data, thus potentially allowing for more sharing of sensitive data among allies.
- The use of synthetic data comes with its own set of risks. These include difficulties in fully replicating the complex physics of the real world, data poisoning, unintended biases, and lower levels of privacy associated with some synthetic data-generation techniques.
- While these risks may also be applicable to real-world data sets, synthetic data may expand the potential for some of these risks. Thus, it is critical to establish processes to ensure the reliability and quality of synthetic data sets.

Continued advances in AI and machine learning are generating great expectations that the functionality and reliability of autonomous systems can be enhanced. Indeed, it has been posited that militaries and defence organizations alike could achieve greater capabilities and efficiencies by conferring more autonomy on autonomous systems.[44] At a minimum, fielded autonomous systems, regardless of the level of autonomy, should be reliable, predictable and safe, and should be able to operate in compliance with international humanitarian law.

Ensuring that autonomous systems are properly trained prior to deployment and use is of critical importance in the international security context, where decisions, inferences and actions are sometimes made "in-theatre"

44 P. Scharre, "Robotics on the Battlefield Part II: The Coming Swarm", Center for a New American Security, 15 October 2014, https://www.cnas.org/publications/reports/robotics-on-the-battlefield-part-ii-the-coming-swarm.

by the autonomous system in order to reduce the latency between analysis and any actions taken. In other words, the autonomous system that collects the data is the same system that performs the analysis and delivers an output – a process called "edge analytics".[45] The ability for autonomous systems to be placed at the edge has become an increasingly important component of technical solutions for various military applications. Yet, the development of autonomous systems intended to be placed at the edge (e.g., uncrewed vehicles used in military operations) is unlike that of autonomous systems in other contexts (e.g., those used in cyber operations) due to hardware limitations.[46] To be sure, the issue is not necessarily a lack of data, but rather a lack of high-quality labelled data as a result of the lack of data-collection hardware. There is also the issue of a lack of diversity of data collected by autonomous systems, which are fielded for specific operational functions, not data collection. For example, an uncrewed aerial vehicle (UAV) that operates at high altitudes will only be able to collect images from high angles, thereby generating a data set that may be largely irrelevant for another UAV operating at lower altitudes and at lower angles of vision.

It has been postulated that autonomous systems possess tremendous value for military operations. The advantages range from executing tasks quicker than any human or human-operated systems for time-critical mission (e.g., air-defence or defensive cyber operations) to carrying out so-called 3D (dull, dirty and dangerous) missions where human performance is prone to deterioration over time.[47] However, data issues for autonomous systems continue to plague defence organizations. The task of designing autonomous systems for either on-board or off-board data processing represents a trade-off, as diverse stakeholders each have unique requirements.[48] Defence organizations are grappling with this trade-off and face challenges in obtaining real-world data along with the associated annotations (i.e., data labels) that can be used to train algorithms for on-board data processing.[49] The current data management architecture only permits autonomous systems to operate in controlled environments and with limited degrees of autonomy.[50] For example, the Israeli Guardium uncrewed ground vehicle is only used autonomously at the Israel–Gaza border, a location that is well-mapped and relatively static.[51] The use of synthetic data for training autonomous systems

---

45 M. Hagström, "Military Applications of Machine Learning and Autonomous Systems", The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk, SIPRI, May 2019, https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf.

46 Kannan, "Synthetic Time Series Data Generation", 3.

47 V. Boulanin, "Artificial Intelligence: A Primer", The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk, SIPRI, May 2019, https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf.

48 Defense Science Board, "Task Force Report: The Role of Autonomy in DoD Systems", US Department of Defense, July 2012, 20, https://irp.fas.org/agency/dod/dsb/autonomy.pdf.

49 Yan et al., "Synthetic Dataset Generation", 1.

50 J. Kwik and T. Van Engers, "Algorithmic Fog of War: When Lack of Transparency Violates the Law of Armed Conflict", Journal of Future Robot Life, 2021, 7, https://doi.org/10.3233/FRL-200019.

51 R. Crootof, "The Killer Robots are Here: Legal and Policy Implications", Cardozo Law Review, 2015, 1869, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2534567.

may therefore represent a means to alleviate the data challenges associated with the current data-collection and data-processing architecture. It may thereby present defence organizations with the opportunity to further exploit the use of autonomous systems by placing them in highly dynamic and multivariate environments while reducing the associated risks.

In the cyber realm, however, the incorporation of AI could be an essential element of the ability to conduct defensive cyber operations at scale and to identify threats before they arise.[52] Put another way, AI could make the defence of cyber infrastructures more reliable, especially against AI-enabled offensive cyber operations.[53] AI could be critical in dealing with challenges arising from both the increasing scale and the increasing sophistication of the cyber realm.

SCALE	SOPHISTICATION
As societies and urban environments grow more digitally interconnected and heterogenous, it creates more pressure points and vulnerabilities for defensive cyber operations to supervise. Vulnerabilities in digital systems are not only the result of increased sophistication in the vector of attacks, but are also created by the size of the attack surface. In other words, while the types of attack may not necessarily be changing, the scale of the risks are. As such, AI could act as a force multiplier to provide enough "eyeballs" on enough segments within a digital space in order to be effective.[54]	Offensive cyber operations augmented and amplified by AI (e.g., synthetic images, adversarial data manipulation and other deceptive techniques) could pose threats to the regular operations of a government, private enterprise or individual. Using AI-enhanced systems and practices defensively may then be necessary to detect and respond to anomalies by using fine grain control. As such, unlike the issue of scale, leveraging AI against AI-enhanced attacks is not just a matter of correcting social or organizational deficits, but of ameliorating technological shortcomings.

Evidently, there is a wide range of actual and potential use cases for AI within the international security context. Yet, cultural and social overhangs as well as technological barriers associated with the deployment of AI technologies continue to raise concerns regarding the safety, predictability and reliability of AI, especially in the context of international security. These concerns are creating a gap between "experimental tools and fielded systems".[55] Synthetic data is one proposed solution that could contribute to ameliorating the trust deficit associated with the integration of AI technologies in high-risk situations by enhancing the quality and usability of training data.

52 Interview with Prof. Tim Watson, 11 April 2023.

53 Wilner, "AI and the Future of Deterrence".

54 Interview with Prof. Tim Watson, 11 April 2023.

55 Hagström, "Military Applications of Machine Learning", 37.

3.1 Value Added by Synthetic Data

The value added by synthetic data depends on where, how and for which AI systems it is being applied. In general, synthetic data allows for the generation of highly diverse data sets, fine grain control of data attributes, automatic annotation or data labelling, and cost-effectiveness. The aim is that the application of synthetic data would have a beneficial effect on training AI systems.

The extent to which synthetic data is an appropriate proxy for the original data is a measure of the utility of the method used to generate the synthetic data as well as the machine learning model and AI system using the synthetic data.[56] In some circumstances, machine learning algorithms may even be trained on synthetic data sets that have no real-world equivalent, particularly in cases where real-world data cannot be properly collected (e.g., objects placed in uncommon or rare environments). In these circumstances, the use of synthetic data may be essential. This is especially relevant in the military context, where autonomous systems placed in complex combat environments are designed and built for operational efficiency and efficacy, rather than high detail data collection. As such, real-world data collected by a UAV, for example, may not be able to capture all possible combinations of relevant attributes with high levels of detail, such as images of the relevant object in different environments, captured at different distances, viewing angles and orientations, and under different illuminations.[57] The ability to synthetically generate highly varied scenarios with all possible combinations of relevant attributes as well as the ability to properly identify rare occurrences will thus be essential for any autonomous system to be safe, predictable and reliable, especially in uncontrolled environments.

Using synthetic data to train autonomous systems may also shorten training cycles. Because AI systems are dependent on the experiences held within the data, rather than the data in and of itself, training AI systems on real-world data can be impractical. Collecting a sufficient amount of real-world data and ensuring that there is ample diversity within that data set is a resource-intensive and time-consuming process. Even then, it is difficult to ensure that all possible variations and diversity in the training data set have been exhausted. In addition, real-world data may not provide the fine grain control and flexibility that synthetic data provides in training an AI system to fit different requirements. On the other hand, sometimes the characteristics of a real-world data set can be muddled or the data set may simply be too cumbersome to be used effectively. In some cases, mere seconds could be worth gigabytes of data (e.g., packet capture). As such, simple synthetic data sets that retain the characteristics and statistical distributions of the underlying original data set could be sufficient in certain cases.

Moreover, in instances where the collection of parsed and properly index data may not be an

56 Riemann, “Synthetic Data”.

57 Yan et al., “Synthetic Dataset Generation”, 1.

issue (e.g., for defensive cyber operations), synthetic data may be used to produce and learn hypothetical situations. For example, developers may leverage agent-based modelling – a technique to simulate interactions between multiple variables (e.g., people, internet of things (IoT) systems, time, etc.) – to create synthetic data sets that reflect people working on certain IoT or enterprise systems for a specified amount of time.[58] The value-added here is that, even though organizations working on an IoT system may be able to capture a vast amount of complete data, organizations may not possess the fine grain control of the data they collect to detect or predict all anomalies or distinguish anomalies from regular patterns. By using agent-based modelling to generate a synthetic reality, organizations may be able to train IoT systems to simulate, identify and classify malicious and non-malicious activity at various levels of sophistication.

The application of such techniques is not exclusive to the international security context nor are they simply theoretical. Indeed, techniques such as agent-based modelling have been applied in other contexts. For example, agent-based modelling has been used to simulate and anticipate the impact of public policies (e.g., urban planning and public transportation), for policy evaluation and for simulating disease outbreaks and interventions.[59] In fact, open-source synthetic populations that reflect the characteristics of a local population have already been developed for the United Kingdom[60] and the United States[61] as well as for smaller geographic regions (e.g., the Île-de-France region, France,[62] and the Island of Montreal, Canada[63]).

It can, therefore, be implied that agent-based modelling can help militaries to prepare for unexpected situations or to plan operations. By generating synthetic data via agent-based modelling simulations, militaries can prepare for a range of potential situations and develop strategies to address them. This may help to improve the readiness and effectiveness of military operations, making them better prepared for unforeseen events and creating data points for rare events or uncommon environments.

The fine grain control of synthetic data sets grants developers the ability to make minor adjustments to the traits and characteristics of the synthetic data set and to test the performance

---

58 Interview with Prof. Tim Watson, 11 April 2023.

59 M. Prédhumeau and E. Manley, "A Synthetic Population for Agent-Based Modelling in Canada", Scientific Data, vol. 10, 21 March 2023, https://doi.org/10.1038/s41597-023-02030-4.

60 A. Smith et al., "An Open-Source Model for Projecting Small Area Demographic and Land-Use Change", Geographical Analysis, 7 February 2022, https://doi.org/10.1111/gean.12320.

61 W. Wheaton et al. "Synthesized Population Database: A US Geospatial Database for Agent-Based Models", Methods Report, RTI Press, May 2009, https://doi.org/10.3768/rtipress.2009.mr.0010.0905.

62 S. Hörl and M. Balać, "Synthetic Population and Travel Demand for Paris and Île-de-France Based on Open and Publicly Available Data", Transportation Research Part C: Emerging Technologies, vol. 130, December 2021, https://doi.org/10.1016/j.trc.2021.103291.

63 L. Perez et al., "A Geospatial Agent-Based Model of the Spatial Urban Dynamics of Immigrant Populations: A Study of the Island of Montreal, Canada", PLOS ONE, vol. 14, July 2019, https://doi.org/10.1371/journal.pone.0219188.

and limitations of machine learning algorithms.[64] Indeed, it is possible to create several synthetic data sets with the same underlying data in order to serve different functions.[65]

The digital world can also test how variations of a synthetic data set derived from the same underlying data influence how an AI system ultimately responds to its environment. This can also be particularly useful for identifying and addressing skewed data sets, where one trait or a class of traits within a data set is overrepresented (i.e., data or algorithmic bias). In a data set where one class of traits is swamped by a larger class, techniques such as the Synthetic Minority Oversampling Technique (SMOTE)[66] can be used to balance their frequencies.[67] Conditional GANs (CGANs) can also reduce skew in a data set by using adversarial training to address the ability of the discriminator network to predict underrepresented classes more accurately to eliminate class-wide bias.[68]

These "benchmarking" characteristics imply applicability in the military context. For example, in an experiment conducted by the US Army Research Laboratory, the researchers found that the performance of computer vision systems (e.g., those used in uncrewed vehicles) is correlated to the angle of the images on which the system was trained.[69] The researchers noted that the classifier model demonstrated bias against images collected directly above a subject (e.g., human, building, tank etc.), and the performance improved as the camera moved further away – decreasing the angle of vision. The researchers concluded that one possible reason is that, because the classifier model was trained on ground imagery, the performance would improve as the experiment inputs looked more like ground imagery. The system thus needs to be retrained using more aerial imagery with higher angles. As such, the researchers noted that synthetic data can be used to compare the different classifier models with different model complexities and architecture, which would then allow the optimal classifier to be chosen for a specific task.

Lastly, it has been argued that the creation of synthetic data that represents real-world data may also eliminate legal challenges associated with collecting, storing, disseminating and disposing of sensitive data.[70] Currently, organizations may be unwilling to share data related to their digital infrastructure if sensitive details of their environment (e.g., IP address, network types, etc.) are exposed, which could pose a risk to the safety of their enterprise digital infrastructure.[71] This may be even more pertinent in the

---

64 Interview with Dr. Leslie Sikos, 25 April 2023.

65 A. Alfons et al., "Synthetic Data Generation of SILC Data", European Commission, 2011, 6, https://www.uni-trier.de/fileadmin/fb4/projekte/SurveyStatisticsNet/Ameli_Delivrables/AMELI-WP6-D6.2-240611.pdf.

66 SMOTE is a rules-based method of synthetic data generation.

67 Focus Group on Artificial Intelligence for Health, "Data and Artificial Intelligence Assessment Methods (DAISAM) Reference", International Telecommunication Union and World Health Organization, May 2020, 13.

68 Ibid, 12.

69 Yan et al., "Synthetic Dataset Generation".

70 A. Tucker et al., "Generating High-Fidelity Synthetic Patient Data for Assessing Machine Learning Healthcare Software", NPJ Digital Medicine, vol. 3, 9 November 2020, https://doi.org/10.1038/s41746-020-00353-9.

71 Interview with Dr. Leslie Sikos, 25 April 2023.

international security context, where sensitive real-world data is not easily shared even among allies.[72] For example, the Australian Department of Defence has noted the challenge of not being aligned with the data standards of the other members of the "Five Eyes" intelligence-sharing partnership.[73] Privacy preservation also suggests the ability of synthetic data to act as a hedge against changes in data privacy regulations that could heighten risks by disrupting organizational and inter-organizational routines of sharing sensitive data.

72 W. Öhman, "Data Augmentation Using Military Simulators in Deep Learning Object Detection Applications", KTH Royal Institute of Technology, 10 September 2019, 2, https://www.diva-portal.org/smash/get/diva2:1375838/FULLTEXT01.pdf.

73 Australian Department of Defence, "Defence Data Strategy", 20.

3.2 Risks

While synthetic data can help alleviate some of the data challenges faced by defence organizations, it is not a silver bullet. Synthetic data comes with its own set of risks and challenges. The ability to manage those risks and challenges is particularly important in order for AI systems to be used in a responsible and safe manner and in accordance with legal requirements and ethical values.

One of the most prominent risks with using synthetic data is called the "reality gap". This refers to the subtle differences between the synthetic data and the real world. Sophisticated machine learning models often learn to exploit small discrepancies, making simulated environments difficult to learn from.[74] In other words, if synthetic data is not simulated properly, it can run into issues of not being able to fully replicate the complex and chaotic physics of the real world and may fail to properly capture the unexpected shifts or one-off cases that emerge in real-world data.

While synthetic data can be used to benchmark data quality, data bias and algorithmic bias, synthetic data itself can also create (or even amplify) unintended biases. Intended biases can be useful in certain applications, for example, overrepresenting specific classes of rare malicious network traffic patterns so that AI systems used for either surveillance or incidence response have a higher chance of detecting those malicious patterns. Nonetheless, it is critical that these intended biases do not have unintended consequences. In virtually all AI systems, there is an optimal number of synthetic data points, which is dependent on the composition of both the synthetic data and real-world data on which the AI systems is trained. Too much synthetic data could "overfit"[75] the AI system, thereby degrading the performance of the system. Therefore, ensuring that the specified scope is correct is critical to avoiding unintended harms or other unintended consequences. Not only could poor scoping lead to unintended consequences once an autonomous system is fielded, but it could also lead to low-quality data, sampling errors, gender or racial bias, labelling or aggregation bias, or incomplete synthetic data sets when it is being generated.[76]

The issue of data bias and algorithmic bias is also a social and cultural challenge in addition to a technical challenge. For example, if a synthetic data set is generated based on the traits and characteristic of an original real-world data set that contains certain assumptions of gender or racial norms, that synthetic data set could further amplify those biases. Even if gender or race is not "explicit in the machine learning model, patterns drawn from neutral characteristics, such as uniforms or evidence of weapons, could still implicitly incorporate

74 Öhman, "Data Augmentation", 6.

75 Overfitting is a term used in machine learning that denotes that the algorithm or network does not generalize well enough to the unseen test data, although it performs well on training data. See, Ibid, 5.

76 Interview with Dr. Leslie Sikos, 25 April 2023.

gender [or racial] norms".[77] As such, gender- and racial-based approaches underline the importance of diversifying the people and range of expertise involved in each step of the AI system, including data generation.[78]

Moreover, synthetic data can still be prone to data poisoning by sophisticated malicious actors. It is possible for adversaries to bury unwanted changes in the synthetic data or data set in order to disrupt the learning procedure, such as by injecting a small fraction of malicious samples into a training data set or making small adjustments to a synthetically generated image.[79] However, it is important to note that, while data poisoning is a risk with synthetic data, it is less prone to poisoning than real-world data, which is often created in distant or uncontrolled environments.

Lastly, while some synthetic data-generation techniques are privacy-preserving, other techniques may not provide adequate levels of privacy protection. Specifically, diffusion models are the least private form of image generation when compared to other techniques, such as GANs. This is directly related to the utility of diffusion models to generate higher quality images compared to GANs and VAEs.[80] In other words, in some contexts synthetic data may present a privacy–utility trade-off, as increasingly powerful generative models raise questions about how diffusion models work and how, and under what circumstances they should be responsibly deployed.[81]

While these risks may also be applicable to real-world data sets, synthetic data may expand the potential surface area for most of these risks. Synthetic data itself does not provide new discrete risks, but these risks could be more pervasive. Simply put, the types of risk may be similar, but the risk vectors are shifting and the scale of risk is increasing. However, it has been posited that the use of synthetic data may raise more questions than real-world data because people generally trust it less, which may provide more opportunities to establish processes to verify synthetic data – more so than real-world data.[82]

---

77 Chandler, "Does Military AI Have Gender?", 17.

78 Chandler, "Does Military AI Have Gender?", 9.

79 Interview with Prof. Tim Watson, 11 April 2023.

80 Carlini et al., "Extracting Training Data", 1.

81 Ibid.

82 Interview with Prof. Tim Watson, 11 April 2023.

4. Conclusion

Synthetic data has proven to be a useful technology in a variety of sectors, from healthcare to fraud detection and public policy planning. While synthetic data can still be considered an "emerging technology",[83] it has reached an adequate degree of maturity and there is a sufficient amount of expertise for adoption across industries and public services, including those related to international security.

Indeed, the value-added and risks associated with synthetic data are relevant to United Nations multilateral processes and other discussions surrounding the use of AI within the context of international security. The potential benefits offered by synthetic data, particularly the fine grain control, data diversity and cost-effectiveness, should not be ignored. Synthetic data could present a solution to ameliorate some of the data challenges that continue to plague defence organizations, such as poor data quality and low-diversity data sets. By ameliorating some of these challenges, militaries and defence organizations alike could improve operational capabilities while ensuring compliance with their international humanitarian law obligations, particularly during 3D operations where human performance is prone to deteriorate over time.

At the same time, the risks associated with synthetic data should not be understated. While synthetic data does not necessarily create new risks that are distinct from the risks associated with real-world data, synthetic data may expand the risk surface. In other words, the risks may be similar, but there may be more ways to generate those risks. For example, a lack of diversity in a real-world training data set may create unintended biases in the same way as overfitting an AI system with one synthetic data point may create unintended biases.

While the features of synthetic data make it a promising technology for the development of autonomous systems in international security, it should not be viewed as a silver bullet or a cure-all to existing data challenges. It should, instead, be understood as a tool in the data-governance toolbox. There is ample research that demonstrates the utility of synthetic data, and generation models have advanced significantly in the past few years. As such, next steps should include, but not be limited to, identification of specific use cases; more targeted research on how to apply existing methods and knowledge of synthetic data in the field of international security while considering the gender and racial aspects of data; and how synthetic data can be integrated into existing data governance strategies.

---

83 "Emerging technologies" refers to technologies that create new opportunities to address global challenges while also creating new regulatory challenges. See Organisation for Economic Co-operation and Development (OECD), "OECD Science, Technology, and Industry Outlook 2012", 13 September 2012, 222, https://doi.org/10.1787/sti_outlook-2012-en.

Bibliography

Alfons Andreas, Peter Filzmoser, Beat Hullinger, Jan-Philipp Kolb, Stefan Kraft, Ralf Münnich, and Matthias Templ. "Synthetic Data Generation of SILC Data", European Commission, 2011, 6, https://www.uni-trier.de/fileadmin/fb4/projekte/SurveyStatisticsNet/Ameli_Delivrables/AMELI-WP6-D6.2-240611.pdf.

Alkhzaimi, Hoda. "Contribution to the Fifth Substantive Session by Emerging Research and Security Center, NYU/NYUAD", NGO Working Papers, 28 July 2023, https://docs-library.unoda.org/Open-Ended_Working_Group_on_Information_and_Communication_Technologies_-_(2021)/Stakeholder_Recommendation_for_Open-ended_working_group_on_security_APR.pdf

Anand, Alisha and Harry Deng, "Towards Responsible AI in Defence: A Mapping and Comparative Analysis of AI Principles Adopted by States", UNIDIR, 13 February 2023, https://unidir.org/publication/towards-responsible-ai-defence-mapping-and-comparative-analysis-ai-principles-adopted

Aryawan Udayana. Putu, Tri Legionosukumo, and Sri Sundari., "Strategy for Integrated Land Information System Network Arrangements for the Indonesian National Army", 2022, https://doi.org/10.33172/jspd.v8i1.1054.

Australia, "Defence Data Strategy 2021-2023", Australia Department of Defence, https://www.defence.gov.au/about/strategic-planning/defence-data-strategy-2021-2023#:~:text=The%205%20pillars%20in%20the,capability%20within%20the%20Defence%20workforce.

Bonabeau, Eric. "Agent-Based Modeling: Methods and Techniques for Simulating Human Systems", PNAS, 14 May 2002, https://doi.org/10.1073/pnas.082080899.

Boulanin, Vincent. "Artificial Intelligence: A Primer", SIPRI, May 2019, https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf.

Canada, "The Department of National Defence and Canadian Armed Forces Data Strategy", Department of National Defence, 18 May 2021, https://www.canada.ca/en/department-national-defence/corporate/reports-publications/data-strategy/data-strategy.html.

Carlini, Nicholas, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. "Extracting Training Data from Diffusion Models", 30 January 2023, 1, https://arxiv.org/abs/2301.13188.

Chandler, Katherine. "Does Military AI Have Gender? Understanding Bias and Promoting Ethical Approaches in Military Applications of AI", UNIDIR, 7 December 2021, https://doi.org/10.37559/GEN/2021/04.

Crootof, Rebecca. "The Killer Robots are Here: Legal and Policy Implications", Cardozo Law Review, 2015, 1869, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2534567.

Defence Innovation Board, "AI Principles", Defense Innovation Board, 2019, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF.

Defense Science Board, "Task Force Report: The Role of Autonomy in DoD Systems", Department of Defense, July 2012, 20, https://irp.fas.org/agency/dod/dsb/autonomy.pdf.

Fan. Linwei, Fan Zhang, Hui Fan, and Caiming Zhang. "Brief Review of Image Denoising Techniques", Visual Computing for Industry, Biomedicine, and Art, https://doi.org/10.1186/s42492-019-0016-7.

Government Business Council, "Advancing ISR at the Edge: A Survey on Networks and Processing Technologies in the Digital Battlespace", July 2020, 4, http://cdn.govexec.com/media/advancing-isr-at-the-edge-isr.pdf

Hagström, Martin. "Military Applications of Machine Learning and Autonomous Systems", SIPRI, May 2019, https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf.

Hörl, Sebastian and Milos Balać. "Synthetic Population and Travel Demand for Paris and Île-de-France Based on Open and Publicly Available Data", Transportation Research Part C: Emerging Technologies, December 2021, https://doi.org/10.1016/j.trc.2021.103291.

Hradec., Jiri, Massimo Craglia, Margherita Di Leo, Sarah De Nigris, Nicole Ostlaender, and Nicholas Nicholson. "Multipurpose Synthetic Population or Policy Application", European Commission Joint Research Centre, 13 April 2022, 14, https://dx.doi.org/10.2760/50072.

Holland, Arthur. “Known Unknowns: Data Issues and Military Autonomous Systems”, UNIDIR, 17 May 2021, https://unidir.org/known-unknowns.

International Telecommunications Union. “Data and Artificial Intelligence Assessment Methods (DAISAM) Reference”, ITU-T Focus Group on Artificial Intelligence for Health, May 2020, 13.

Jordan, Jeremy. “Variational Autoencoders”, 19 March 2018, https://www.jeremyjordan.me/variational-autoencoders/.

Kannan, Subarmaniam. “Synthetic Time Series Data Generation for Edge Analytics”, F1000 Research, 20 January 2022, https://doi.org/10.12688/f1000research.72984.1.

Kingma Diederick P., and Max Welling. “An Introduction to Variational Autoencoders”, Foundations and Trends in Machine Learning, 2019, https://arxiv.org/pdf/1906.02691.pdf.

Kwik, Jonathan and Tom Van Engers. “Algorithmic Fog of War: When Lack of Transparency Violates the Law of Armed Conflict”, Journal of Future Robot Life, 2021, 7, https://doi.org/10.3233/FRL-200019.

Longford, Frank. “Experiments in Synthetic Data”, Forensic Architecture, 6 November 2018, https://forensic-architecture.org/investigation/experiments-in-synthetic-data.

Manon Prédhumeau and Ed Manley, “A Synthetic Population for Agent-Based Modelling in Canada”, Scientific Data, 21 March 2023, https://doi.org/10.1038/s41597-023-02030-4.

Öhman, Wilhelm. “Data Augmentation Using Military Simulators in Deep Learning Object Detection Applications”, KTH Royal Institute of Technology, 10 September 2019, 2, https://www.diva-portal.org/smash/get/diva2:1375838/FULLTEXT01.pdf.

Organization for Economic Cooperation and Development. “OECD Science, Technology, and Industry Outlook 2012”, 13 September 2012, 222, https://doi.org/10.1787/sti_outlook-2012-en.

Pakistan, “Proposal for an International Instrument on Lethal Autonomous Weapons (LAWS)”, CCW/GGE.1/2023/WP.2/Rev.1, 8 March 2023, https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2023)/CCW_GGE1_2023_WP.3_REv.1_0.pdf

Pasieka, Manuel. “A Comparison of Synthetic Data Generation Methods and Synthetic Data Types”, 1 September 2022, https://mostly.ai/blog/comparison-of-synthetic-data-types/.

Perez, Liliana, Suzana Dragicevic, and Jonathan Gaudreau. “A Geospatial Agent-Based Model of the Spatial Urban Dynamics of Immigrant Populations: A Study of the Island of Montreal, Canada”, PLOS ONE, 24 July 2019, https://doi.org/10.1371/journal.pone.0219188.

Scharre, Paul. “Robotics on the Battlefield Part II: The Coming Swarm”, Center for a New American Security, 15 October 2014, https://www.cnas.org/publications/reports/robotics-on-the-battlefield-part-ii-the-coming-swarm.

Silva, Thalles. “A Short Introduction to Generative Adversarial Networks”, Thalles’ Blog, 7 June 2017, https://sthalles.github.io/intro-to-gans/.

Smith, Andrew, Luke Archer, Alistair Ford, and James Virgo. “An Open-Source Model for Projecting Small Area Demographic and Land-Use Change”, Geographical Analysis, 7 February 2022, https://doi.org/10.1111/gean.12320.

State of Palestine, “State of Palestine’s Proposal for the Normative and Operational Framework on Autonomous Weapons Systems”, CCW/GGE.1/2023/WP.2/Rev.1, 3 March 2023, https://docs-library.unoda.org/Convention_on_Certain_Conventional_Weapons_-Group_of_Governmental_Experts_on_Lethal_Autonomous_Weapons_Systems_(2023)/CCW_GGE1_2023_WP.2_Rev.1.pdf.

Tucker, Allan, Zhenchen Wang, Ylenia Rotalinti, Paja Myles. “Generating High-Fidelity Synthetic Patient Data for Assessing Machine Learning Healthcare Software”, NPJ Digital Medicine, 9 November 2020, https://www.nature.com/articles/s41746-020-00353-9#citeas.

United Nations General Assembly, “Disarmament and International Security (First Committee)”, https://www.un.org/en/ga/first/.

Wheaton, William, James Cajka, Bernadette Chasteen, Diane Wagener, Phillip Cooley, Laxminarayana Ganapathi, Douglas Roberts, and Justine Allpress. “Synthesized Population Database: A US Geospatial Database for Agent-Based Models”, Methods Report RTI Press, May 2009, https://doi.org/10.3768%2Frtipress.2009.mr.0010.0905.

Wilner, Alex. “AI and the Future of Deterrence: Promises and Pitfalls”, Centre for International Governance Innovation, 28 November 2022, https://www.cigionline.org/articles/ai-and-the-future-of-deterrence-promises-and-pitfalls/.

Riemann, Robert. “Synthetic Data”, European Data Protection Supervisor, https://edps.europa.eu/press-publications/publications/techsonar/synthetic-data_en

United Kingdom, “Data Strategy for Defence: Delivering the Defence Data Framework and Exploiting the Power of Data”, 27 September 2021, Ministry of Defence, https://www.gov.uk/government/publications/data-strategy-for-defence/data-strategy-for-defence.

Vahdat, Arash and Karsten Kreis. “Improving Diffusion Models as an Alternative to GANs, Part 1”, NVIDIA, 26 April 2022, https://developer.nvidia.com/blog/improving-diffusion-models-as-an-alternative-to-gans-part-1/.

Yan, Jie, Eung Joo Lee, Damon Conover, and Heesung Kwon. “Synthetic Dataset Generation and Adaptation for Human Detection”, DEVCOM Army Research Laboratory, November 2020, 1, https://apps.dtic.mil/sti/pdfs/AD1115446.pdf.

@unidir
/unidir
/un_disarmresearch
/unidirgeneva
/unidir

Palais des Nations
1211 Geneva, Switzerland

WWW.UNIDIR.ORG