

THE CYBER DEFENSE REVIEW

Data in the Cybersecurity Landscape

Guest Editors

Maj. Iain J. Cruickshank, Ph.D. and Lt. Col. Nathaniel D. Bastian, Ph.D.

Leadership Matters

Lt. Gen. (Ret.) Edward C. Cardon

The Importance of Analytic Superiority in a World of Big Data and AI

Dr. Robert L. Grossman, Dr. Emily O. Goldman



Emerging Technologies for Data Security in Zero Trust

*Mr. Patrick Davis, Maj. Sean Coffey, Dr. Lubjana Beshaj,
Lt. Col. Nathaniel D. Bastian, Ph.D.*

An Introduction to Quantum Computing and Its Applications

1st Lt. Mandeep Singh and Dr. Albert H. Carlson

Data as Ammunition

Lt. Col. Jessica Dawson, Ph.D. and Col. Katie E. Matthew, Ph.D.

Overcoming the Labeled Training Data Bottleneck: A Route to Specialized AI

Dr. Ravi Starzl

INTRODUCTION

Data in the Cybersecurity Landscape

*Maj. Iain Cruickshank, Ph.D.
Lt. Col. Nathaniel D. Bastian, Ph.D.*

PROFESSIONAL COMMENTARY

A Cyber Revival at the Heart of the Next Data Battle

Arnaud LeDez

THE CYBER DEFENSE REVIEW

◆ SPECIAL EDITION ◆

THE CYBER DEFENSE REVIEW

A DYNAMIC MULTIDISCIPLINARY DIALOGUE

EDITOR IN CHIEF

Deborah S. Karagosian

MANAGING EDITOR

Dr. Jan Kallberg

ASSISTANT EDITORS

West Point Class of '70

AREA EDITORS

Dr. Craig Douglas Albert
(Information Warfare)

Dr. Dawn Dunkerley Goss
(Cybersecurity Optimization/Operationalization)

Dr. Jeffrey Morris
(Quantum Information/Talent Management)

Dr. Harold J. Arata III
(Cybersecurity Strategy)

Dr. Michael Grimaila
(Systems Engineering/Information Assurance)

Elizabeth Oren
(Cultural Studies)

Lt. Col. Todd W. Arnold, Ph.D.
(Internet Networking/Capability Development)

Dr. Steve Henderson
(Data Mining/Machine Learning)

Dr. David Raymond
(Network Security)

Donna Artusy, J.D.
(Cyber Law)

Dr. Michael Klipstein
(Cyber Policy/Cyber Operations)

Dr. Robert J. Ross
(Information Warfare)

Lt. Col. Nathaniel D. Bastian, Ph.D.
(Advanced Analytics/Data Science)

Lt. Col. Charlie Lewis
(Military Operations/Training/Doctrine)

Dr. David Thomson
(Cryptographic Processes/Information Theory)

Dr. Amy Ertan
(Cyber Strategy)

Dr. Fernando Maymi
(Cyber Curricula/Autonomous Platforms)

Dr. Robert Thomson
(Learning Algorithms/Computational Modeling)

Dr. David Gioe
(History/Intelligence Community)

Dr. William Clay Moody
(Software Development)

Lt. Col. (Ret.) Mark Visger, J.D.
(Cyber Law)

EDITORIAL BOARD

Dr. Andrew O. Hall, (Chair.)
Marymount University

Dr. Michele L. Malvesti
University of Texas at Austin

Dr. Bhavani Thuraisingham
The University of Texas at Dallas

Dr. David Brumley
Carnegie Mellon University

Dr. Milton Mueller
Georgia Tech School of Public Policy

Prof. Tim Watson
Loughborough University, UK

Col. (Ret.) W. Michael Guillot
Air University

Col. Suzanne Nielsen, Ph.D.
U.S. Military Academy

Prof. Samuel White
Army War College

Dr. Martin Libicki
U.S. Naval Academy

Dr. Hy S. Rothstein
Naval Postgraduate School

CREATIVE DIRECTORS

Sergio Analco | Gina Lauria

LEGAL REVIEW

Courtney Gordon-Tennant, Esq.

KEY CONTRIBUTORS

Sheri Beyea

Kate Brown

Lt. Col. Douglas Healy

Evonne Mobley

Alfred Pacenza

Clare Blackmon

Col. (Ret.) John Giordano

Charles Leonard

Thomas Morel

Isabella Velilla

CONTACT

West Point Press
Taylor Hall, Building 600
West Point, NY 10996

SUBMISSIONS

The Cyber Defense Review
welcomes submissions at
TheCyberDefenseReview@westpoint.edu

WEBSITE

cyberdefensereview.army.mil

The Cyber Defense Review (ISSN 2474-2120) is published by the West Point Press. The views expressed in the journal are those of the authors and not the United States Military Academy, the Department of the Army, or any other agency of the U.S. Government. The mention of companies and/or products is for demonstrative purposes only and does not constitute endorsement by United States Military Academy, the Department of the Army, or any other agency of the U.S. Government.

© U.S. copyright protection is not available for works of the United States Government. However, the authors of specific content published in The Cyber Defense Review retain copyright to their individual works, so long as those works were not written by United States Government personnel (military or civilian) as part of their official duties. Publication in a government journal does not authorize the use or appropriation of copyright-protected material without the owner's consent.

This publication of the CDR was designed and produced by Gina Daschbach Marketing, LLC, under the management of FedWriters. The CDR is printed by McDonald & Eudy Printers, Inc. ∞ Printed on Acid Free paper.

SPECIAL
EDITION

DATA IN THE CYBERSECURITY LANDSCAPE

GUEST EDITORS

Maj. Iain J. Cruickshank, Ph.D. and Lt. Col. Nathaniel D. Bastian, Ph.D.

INTRODUCTION

Maj. Iain J. Cruickshank, Ph.D. 9 Data in the Cybersecurity Landscape
Lt. Col. Nathaniel D. Bastian, Ph.D.

PROFESSIONAL COMMENTARY

Arnaud Le Dez 17 Cyber Data Sanitation: A Cyber Revival
at the Heart of the Next Data Battle

SPECIAL EDITION ARTICLES

Dr. Robert L. Grossman 29 The Importance of Analytic Superiority
Dr. Emily O. Goldman in a World of Big Data and AI
Mr. Patrick Davis 49 Emerging Technologies for Data
Maj. Sean Coffey Security in Zero Trust
Dr. Lubjana Beshaj
Lt. Col. Nathaniel D. Bastian, Ph.D.
1st Lt. Mandeep Singh 73 An Introduction to Quantum Computing
Dr. Albert H. Carlson and Its Applications
Lt. Col. Jessica Dawson, Ph.D. 93 Data as Ammunition
Col. Katie E. Matthew, Ph.D.
Dr. Ravi Starzl 109 Overcoming the Labeled Training Data
Bottleneck: A Route to Specialized AI

SENIOR LEADER PERSPECTIVE

Lt. Gen. (Ret.) Edward C. Cardon 133 Leadership Matters

THE CYBER DEFENSE REVIEW

◆ INTRODUCTION ◆

Special Edition on Data in the Cybersecurity Landscape

Major Iain J. Cruickshank, Ph.D.

Lieutenant Colonel
Nathaniel D. Bastian, Ph.D.



We are living through a revolution in human technology driven by digital data. Nearly every piece of technology we use, from cars to refrigerators, now produces—and in some cases, even analyzes—digital data. Furthermore, digital data is now integral to every human endeavor, from dating to warfare. Data is a key component in decision-making and is crucial to warfighting concepts like Information Advantage and Decision Dominance. Thus, it is increasingly important for military and security professionals to understand digital data and the technologies built upon it, such as Artificial Intelligence (AI).

This is a work of the U.S. Government and is not subject to copyright protection in the United States. Foreign copyrights may apply.



Major Iain J. Cruickshank, Ph.D., is a data and decision scientist (Functional Area 49) in the U.S. Army. He currently serves as a Senior Research Scientist within the Army Cyber Institute at the United States Military Academy (USMA) at West Point and has previous assignments with the 780th Military Intelligence Brigade and the U.S. Army Artificial Intelligence Integration Center. He is also adjunct faculty in Carnegie Mellon's School of Computer Science. Major Cruickshank earned an M.S. in Operations Research from the University of Edinburgh and a Ph.D. in Societal Computing from Carnegie Mellon University. His research interests include development of novel, artificial intelligence based computational social science techniques and multi-modal machine learning, particularly as it applies to military and national security problems.

In an age where digital data reigns supreme, its safeguarding and strategic utilization have become paramount to maintaining effective cybersecurity. This special edition of *The Cyber Defense Review* delves into the multifaceted roles that data plays within the cybersecurity landscape, offering a comprehensive exploration of emerging technologies, innovative frameworks, and strategic approaches to enhance data security and leverage its potential.

We begin our exploration of digital data and cybersecurity in this edition by establishing a new concept: *cyber data sanitization*. "Cyber Data Sanitization: A Cyber Revival at the Heart of the Next Data Battle" highlights the pivotal challenge of data pollution and explores innovative tools and standards for maintaining digital hygiene and optimizing data processing. By repurposing existing cybersecurity tools to clean data beyond malicious activities, cybersecurity expert Arnaud Le Dez proposes a new domain of cyber defense, essential for preserving operational efficiency and supporting AI advancements.

To initiate the research section of this edition, we begin with an article that establishes the importance of data and data-driven technologies through the concept of *analytic superiority*. In "The Importance of Analytic Superiority in a World of Big Data and AI," Drs. Grossman and Goldman outline a comprehensive framework for achieving analytic superiority, which involves collecting requisite data, developing robust analytic models, and deploying these models effectively. Emphasizing the importance of staying ahead in the analytic race, the authors highlight how the U.S. can maintain cyberspace superiority by outpacing adversaries in analytic capabilities.



Lieutenant Colonel Nathaniel D. Bastian, Ph.D., is an Academy Professor and Cyber Warfare Officer at the United States Military Academy (USMA) at West Point. He is also concurrently serving as a Program Manager at the Defense Advanced Research Projects Agency (DARPA). At USMA, Lt. Col. Bastian serves as Division Chief, Data & Decision Sciences, Senior Research Scientist, and Chief Data Scientist at the Army Cyber Institute within the Department of Electrical Engineering and Computer Science. Lt. Col. Bastian earned a M.S. in Econometrics and Operations Research from Maastricht University and a M.Eng. in Industrial Engineering and Ph.D. in Industrial Engineering and Operations Research from Pennsylvania State University. His research interests aim to develop innovative, assured, intelligent, human-aware, data-centric, and decision-driven capabilities for cyberspace operations.

A critical component to unlocking data's value for the warfighter is the security of data. In "Emerging Technologies for Data Security in Zero Trust Environments" Davis et al. evaluate cutting-edge technologies that promise to fortify Zero Trust Architecture (ZTA) data security capabilities. By focusing on differential privacy, confidential computing, homomorphic encryption, blockchain, and alternative computing methods, this article provides a forward-looking perspective on how these innovations can secure data throughout its lifecycle. The authors meticulously assess each technology's potential benefits and drawbacks, underscoring the need for a nuanced approach to integrating these tools into ZTA implementation frameworks.

No discussion about data security would be complete without considering the pending quantum revolution in cybersecurity. "An Introduction to Quantum Computing and Applications" provides a practical guide for understanding and leveraging quantum information science (QIS) technologies. This article elucidates how quantum sensing, networking, communications, and computing can offer strategic advantages in military operations. Emphasizing the importance of informed leadership, First Lieutenant Singh and Dr. Carlson argue that a grasp of QIS applications is crucial for maintaining a competitive edge over our adversaries.

When it comes to defining data and its potential, we should have new mental models for digital data and understand that not all data is good data. Challenging the conventional analogy of *data as oil*, "Data as Ammunition - A New Framework for Information Warfare" presents a compelling case for viewing data through the lens of military ordnance. This paradigm shift offers a more accurate depiction of the strategic

risks associated with data. By categorizing data as ammunition, Lieutenant Colonel Dawson and Colonel Matthew argue, we gain a clearer understanding of its potential threats and how to mitigate them effectively. This framework is particularly relevant in an era where data storage capabilities have skyrocketed, making it crucial to reevaluate our approaches to data security and privacy.

Finally, no discussion of data is complete without talking about the challenges of overcoming data scarcity for data-driven tools, like AI. The scarcity of labeled training data poses a significant hurdle in developing specialized AI for cybersecurity. “Overcoming the Labeled Training Data Bottleneck, A Route to Specialized AI” explores the potential of large language models to generate high-quality, task-specific datasets. By synthesizing existing network intrusion analysis datasets with domain knowledge, Dr. Starzl was able to create a new dataset tailored for zero-day exploit detection. This innovative approach demonstrates the effectiveness of using advanced AI techniques to overcome data scarcity, paving the way for more efficient and effective cybersecurity solutions.

We conclude this edition of *The Cyber Defense Review* with some important reflections on the role of leadership in cybersecurity. “Leadership Matters,” by Lieutenant General (Ret.) Edward C. Cardon, reflects on the transformative journey of the U.S. Army Cyber Command. Cardon emphasizes the need for continual adaptation and collaboration among various stakeholders to enhance cybersecurity capabilities. He advocates for leadership to push the envelope in all aspects of cybersecurity, including technologies, concepts, and people because the future is there to be seized by those who constantly hone their skill and maintain the will to execute.

The articles in this special edition collectively underscore the centrality of digital data in contemporary cybersecurity discourse. From the integration of emerging technologies in Zero Trust environments to innovative frameworks for data management and strategic leadership, our contributors offer valuable perspectives on how to harness data’s value while safeguarding against its inherent risks. As we navigate the complexities of the digital age, understanding and leveraging data will be crucial to achieving and maintaining cybersecurity resilience. We hope this edition inspires further exploration and innovation in the vital intersection of digital data and cybersecurity.♥

DISCLAIMER

The views expressed in this work are those of the authors and do not reflect the official policy or position of the United States Military Academy, the Department of the Army, or the Department of Defense.

THE CYBER DEFENSE REVIEW

◆ PROFESSIONAL COMMENTARY ◆

Cyber Data Sanitization: A Cyber Revival at the Heart of the Next Data Battle

Arnaud Le Dez

INTRODUCTION


We are only at the beginning of the history of cyber defense and the future holds immense promise. We have already seen significant advancements, and with the increased reliance on data, there is no end in sight. Foresight provides us with a unique opportunity to envision various potential cyber futures. It serves as a framework for crafting scenarios, which we can elucidate using methodologies, data analysis, and research. These futures become much more tangible when the key players and the factors that shape them are understood.

Among potential futures, the battle against data pollution emerges as a particularly promising prospect, thanks to the advent of new capabilities such as those described here and derived from cyber technology. This battle signifies not just an enhancement of cyber defense, but its evolution and expansion, paving the way for a new domain.

Data pollution is the degradation of the digital environment by data that can be considered as waste or a nuisance. These data can be naturally produced by digital systems for their operation or linked to human activities in the digital space. Data pollution is likely to affect the health of digital systems and the quality of processing, leading to degradation or interference with operations in cyberspace. We are in a familiar universe here, cyber-attacks can be a form of pollution.

U.S. Army Cyber Command refers to data pollution as a “data rationalization” problem.¹ If the cybercommunity does not solve the problem of data pollution, we will not be able to pursue our activities with the same efficiency and by developing new artificial intelligence capabilities. Just as cyber encompasses all digital activities, this fight against data pollution involves all systems, including command and control systems and other areas of warfare and intelligence systems where data optimization is a constant objective.

© 2024 Arnaud Le Dez



Arnaud Le Dez works for Capgemini in Paris as a consultant on cyber strategy and governance for C-level executives and CISOs of major companies. He was a French Army officer for 28 years, with responsibilities in electronic warfare, cyber defense, and intelligence. He has over six years of experience working with the U.S. Army during several missions in the Middle East. During the last ten years of his military career, he held various responsibilities in cyber defense operations in the French COMCYBER. He is an associate researcher at the conflict transformation division of Saint-Cyr Coëtquidan Military Academy Research Center and is the author of a book on the tactician level in cyber defense “*Tactique Cyber, le Combat Numérique*,” published in January 2019 in *Economica*.

Data analysis is at the core of cyber defense and cybersecurity tactics. Once processed, data yields a cognitive effect that benefits both defenders and attackers. A major concern within cyber is the over-processing of data. A novel tactic within cybersecurity aims to sanitize the digital battlefield for improved detection of nefarious acts and the enablement of counteractions.

The battle against data pollution strives to optimize data processing. The issue is as much about optimizing processing as it is about reducing digital costs and limiting the fog of data war to enable operating in a more visible data environment. It curbs certain excesses, especially uncontrolled data generation from various sources like cybersecurity, intelligence, artificial intelligence, faulty systems, saturation attacks, and more. It morphs cyber defense into an ecological capability that aligns with today’s societal and battlefield challenges. This innovative and beneficial application area has the potential to rejuvenate the cyber defense operational environment, broadening its scope with optimized data use at its core.

The battle against data pollution requires a partial shift in the application of cyber tools, which will detect pollution as an assault on our systems. Sanitation has always been a cyber concern. Now, it takes on an even greater role: the need to sanitize data will accompany the defender through all stages of digital conflict.

In addition to unlocking a new market via a fresh class of cyber-based tools, the battle against data pollution will also enhance cyber defense and the efficiency of command-and-control systems across the battlefield. It will allow cyber defenders to concentrate on genuine attacks by cleaning up the digital environment. This cleanup will also free processing capacity on tactical mission command systems that are currently inefficient in data processing. The battle against data pollution aids in equipping our tactical formations with more potent cyber tools.

We all need to know when we create and/or use systems that generate polluted data. We must recognize it and remedy it for the sake of operational efficiency. Data pollution is becoming our fog of war. The intent of this article is to show that we can combat data pollution by reusing knowledge and tools we use in cyber in what can be called the digital domain environment. This concept applies to beyond just the cybersecurity practitioner to all commanders who rely on high-quality data to make decisions.

THE CYBER MODEL: ITS LIMITATIONS AND EVOLUTION

Cyber defense derives from numerous developments in the evolution of cybersecurity. Originally, the reliability and availability of systems were initially dictated by advancements in electronics and protocols, which form the backbone of our current digital capabilities. The 1990s saw an expansion of interconnections via the Internet and the elevation of security of information systems to a higher priority. This was particularly due to the emergence of new tools and structured governance. The Y2K bug presented an opportunity to rectify past mistakes and establish contemporary security policies.

As tools evolved and it became apparent that mere protection of our systems was inadequate, we had to operationalize the cyber profession in an environment riddled with ongoing attacks. This operational approach, centered on the need for real-time actions to defend against attacks that are truly combat, facilitated the transition to the present-day cyber defense. The time from detection to action has been significantly reduced, thanks to new capabilities for generating and processing cyber data, enabling a swift response that increasingly anticipates the unexpected and mirrors the interactions found in military combat.

The current cyber defense model is based on the generation of an astronomical amount of data in systems dedicated exclusively to cyber use, followed by real-time analysis to detect anomalies that could signify attacks. During this operational phase, cyber intelligence has become indispensable for parameterizing tools to detect and characterize adversaries, regardless of their location. It is necessary to know and understand an adversary's tactics and weapons when developing pre-validated automatic responses.

The cyber community is using its tools to generate increasing amounts of data² from a growing number of sources, driven by the fear of overlooking the slightest indication of nefarious actions in a world where our adversaries are constantly adapting. This data accumulation presents its own set of challenges. Not only is it time-consuming to store and process on a massive scale, but it is also difficult to extend our tools into highly constrained environments. The extension of cyber into new environments, towards all digital objects of everyday life (or weapons) and into increasingly contested digital universes, is one of today's most significant challenges.

The industrial environment that supports much of the world's critical infrastructure is currently at the epicenter of commercial digital transformation and is simultaneously the target of

increasingly disruptive and costly attacks. The defense sector views it as an important part of the field of conflict: it is a source of cyber intelligence, but it also needs national-defense level counter-measures. This necessitates a rethinking of cyber defense tools.

We are increasingly using smaller and more autonomous digital devices that often have limited data processing and storage capacity. Anything our military or government employees use, both in the course of their jobs and in their private lives, could be targeted by our adversaries. These systems include digital industrial or military equipment, cars, satellites, weapons with digital systems, watches, or personal fitness devices.³ Such systems either need to be designed for defense or assessed for the risk they could pose if targeted by our adversaries. However, the deployment of our cyber tools is often not planned and often not possible for these devices. They are often closed systems, limited in their processing capacity, constrained in their interconnections, and not designed to accommodate outside cyber tools.

ARTIFICIAL INTELLIGENCE AND DIGITAL TWINS: THE PIONEERS OF TODAY'S TRANSFORMATION

Artificial intelligence and digital twins⁴ are at the forefront of today's transformation, further amplifying the presence and importance of data. For instance, we can use digital twins, virtual representations of physical or information technology systems that is updated in real-time, to apply intelligent simulation algorithms and predict anomalies indicative of future attacks or improve those we are planning using AI. In the defense sector, the combination of a digital twin with artificial intelligence will pave the way for predictive combat in the cyber domain, for both the offense and defense, enabling military planners to better synchronize cyber operations with operations in the other warfighting domains. Using AI and digital twins in a cyber fight will lead to a cleaner digital environment and reduce complexity, allowing commanders a clearer view of battlefield, both in the present and the future. Clean data will make operational planning easier because there are fewer unknowns. Cleaner data and a clean cyberspace environment contribute to the predictability of cyber actions, reduces the fog of cyber warfare, and makes cyber operations planning more relevant.

Therefore, we must question the limitations of our models across all battlefield systems. By their very nature, they restrict cyber action in certain digital facilities; artificial intelligence and digital twins could potentially exacerbate this problem. Increasing processing capacities to meet this challenge is costly, as are the potential battles lost, should our digital combat systems fail. The volume of data generated across our current systems is colossal. The amount of data pollution or, at best, single-use data, is significant and poses challenges for both energy conservation and resource optimization. These challenges are amplified in constrained or contested environments.

Cyber defense will need to evolve and mandatorily involve a more selective approach to data: less data, or higher quality, with a better understanding of its power and limitations. We thus

need to create different digital blueprints to apply various cyber data processing schemes, thereby increasing the probability of detecting attacks.

This is a global phenomenon that does not just affect cybersecurity. Reducing data pollution will become a major objective: this problem is ubiquitous. Even though cybersecurity and cyber defense are currently noise generators within the digital landscape, they are also part of the solution, as we shall see below.

THE CYBER BATTLE AGAINST DATA POLLUTION

There is a need to repurpose cybersecurity tools to clean up data beyond just that which is deemed malicious. Often cybersecurity tools analyze data and generate a response that can be simplified to “good data or bad data.” The analysis grid is based solely on harmfulness, in terms of whether the data are part of a cyber-attack. Cyber tools have the capacity to analyze and process data in all systems, with a logic that is both autonomous and integrates centralized control and increasingly automated action. The detection, command, and response chain, often called the XDR-SIEM-SOAR chain, is completed by antivirus, firewalls, intrusion detection and protection systems (IDS/IPS), and other tools.

Today, processing criteria within the cyber community are only concerned with security. We can broaden this limited focus and transform our cyber tools and cybersecurity using an ecological approach that includes the fight against data pollution. This will enable us to create a new range of tools in the marketplace, and to go beyond the current limits of cybersecurity and its extensions.

The use of cyber tools to combat data pollution can first be envisioned by modifying the criterion of harmfulness. For example, firewalls filter data and can be extended to block all forms of pollution. Detection cyber tools (EDR/NDR/XDR) and antivirus systems analyze data, with the ability to detect and clean them up (or quarantine them). Security Operations Centers (SOCs) steer the entire process and manage resource optimization thanks to a form of intelligence that seeks to characterize what pollution is and what the tactics, techniques, and procedures (TTPs) of these polluters are. All cyber systems can potentially be used to combat data pollution, often after a few simple modifications, possibly involving some intellectual and technical adjustments.

The most challenging aspect is the characterization of data pollution, a task that requires the continuous updating of pollution indicators. We are going to have to invent Data Pollution Threat Intelligence (DPTI), just as we have Cyber Threat Intelligence (CTI). We need to build an evolving dictionary for characterizing pollution, just as we have built an evolving dictionary for characterizing cyber-attacks. This characterization is as complex as the definition for what constitutes malicious data, and this lies at the heart of the cyber problem.

We must also remember that data that are polluted today may not have been in the past or will not be in the future. Pollution may be obvious and constant, but it can also be temporary, making it essential to have an ongoing ability to clean up the data. AI could help us by

generating identical data with the same processing purpose, while at the same time eliminating pollution. Cleansing is an essential component in the fight against data pollution, and here again, the tools used could be derived in part from other cybersecurity efforts.

It is important to bear in mind that cyber tools will not be the only solution. With the increase in sensors across the world as well as the increase in the use of deception, data pollution has the potential to be a much more important phenomenon than cyber-attacks. There is a need to scale up, and to promote low-cost digital hygiene solutions. The solution will include new standards and controls along with a reactive alert network. In this new environment, there is a place for innovative software companies to detect, alert, and cleanse data.

CYBER OBJECTIVES IN COMBATING DATA POLLUTION

The battle against data pollution offers numerous benefits. The most evident is the decluttering of our systems by limiting our data to only what is relevant. Navigating directly to pertinent data and examining it from various perspectives will be a crucial challenge in ensuring that cyber does not get submerged in digital info-obesity.

The advantages for cyber are clear. A system with less data, both in flow and in storage, is naturally easier to monitor with cybersecurity tools. It is always simpler to detect an attack on a clean system than on a polluted one.

Considering that some cyber-attacks also constitute a form of pollution, these new tools will aid in curbing the spread of the threat. For instance, spam and network scans can be categorized as pollution. While network scans are necessary on operational networks, it should be ultra-compressed and encrypted and when complete the results should then be placed in a virtual “garbage can” or, in case cyber intelligence or cyber forensics analysts need them, in a separate database, away from operational networks. Reducing these data would enable cyber analysts and tools to focus their energy and effort on the more complex attacks. It would also allow cybersecurity algorithms to operate on technically more constrained systems, thereby extending the defended areas.

Cybersecurity is poised to branch out into new areas, including the industrial world and the numerous digital devices we use daily. These devices were not designed to accommodate the astronomical amount of data that our cyber models generate today. This effort to limit data and processing will therefore be a prerequisite for deploying our cybersecurity on certain low-processing-power equipment.

Pollution control can be seen as a cyber hygiene measure, using a combination of tools and techniques that need to be kept simple and inexpensive. Cleaning up our data is not just a cyber issue, but a more global or enterprise-level case of data management and data optimization - including the data we use in cyber security and cyber defense. The development of a new field is a business opportunity, with new financing and a new way of approaching a market that is much larger than cyber, and which will broadly enhance our security and defense.

AI AND THE ONSLAUGHT OF CYBER DATA

Artificial intelligence (AI) will undoubtedly lead to significant advancements in cyber analysis capabilities. This is the route currently favored by hardware and software manufacturers.⁵

AI is heavily dependent on the volume of data it requires for learning. It also generates a substantial amount of it.⁶ A recently emerging phenomenon is the synthetic data production to feed other AI algorithms, whose output, often called ghosts, is then fed into the datasets used to develop the next AI algorithm in sequence until the desired outcomes are achieved. These ghosts are a dream where data appear plausible but do not actually contain real data in the context of the answer; some would call this fake data. This result is the so-called hallucination effect where AI developers and their resulting algorithms lose touch with the reality they are supposed to interpret. Used this way, AI is producing data pollution on a scale that is proportional to the growing power of AI. This hallucination is the equivalent of shooting oneself in the foot, which will overwhelm systems and produce incongruities triggering massive malfunctions.⁷

To maintain a healthy digital environment, some Artificial Intelligence will need to be equipped with tools that can detect pollution and clean up the data they produce.⁸ This battle against pollution generated by AIs would be a method of control for AI production environments to maintain system availability. Ideally, the battle against data pollution would address the integrity of results by proactively seeking to eliminate the ghosts contributing to future malfunctions. We cannot work on dream data that will pollute our appreciation of reality and the quality of treatments. This will be even more essential when dealing with mission-critical systems, such as weapons systems, cars, or satellites. This mirrors the reasoning in cybersecurity, where we strive to define the normal behavior of a system as a method to detect abnormalities that could potentially be attacks.

NEW CYBER FRONTIERS

This new vision of cybersecurity should generate new interest for deployment in digital universes that have been neglected until now, or in highly confined environments. The global significance of combating data pollution is similar to that of comprehensive security, while simultaneously reaching audiences who are sometimes resistant to the concept of security. This implies that this new category of tools must liberate itself from this supervisory bond. Viewing cyber from an ecological perspective can attract new followers to a reputedly depleted field, possibly even gaining traction at the level of decision makers, hence contributing to the cyber defense cause.

The fight against data pollution helps to deploy cyber systems in highly constrained situations. Some digital infrastructures remain limited in the processing of data, as evidenced by operationally deployed military forces. Efforts to reduce the volume of unnecessary data are crucial given the military's increased use of interconnected digital equipment. The mastery of data contributes to the agility of system usage, whether in the military or in business.

In the military spirit of economy of means, less data equates to fewer flows, less storage media, and data processing with reduced energy consumption - and a reduced aging of the equipment. This should also contribute to the Low Data movement, which aims to accomplish the same tasks with less data. Undoubtedly, in a high-intensity conflict scenario involving massive use of cyber weapons, the ability to act in a degraded digital battlefield with tools that consume less data will help us to preserve freedom of action.

CONCLUSION: IT IS A QUESTION OF DEFINING EVIL DATA

Feedback from cyber experience is invaluable. We can morph cyber to generate an ecological capability that fits into a much larger global dimension. In cyber, it is difficult to define what, exactly, is an attack, characterizing it, and then blocking the malicious data. Everything depends on the characterization of cyber evil.

We face the same difficulty in the fight against data pollution, except there is also a challenge regarding the persistence of this characterization. Data that are not polluted today may become so tomorrow and, more likely, they will become obsolete. We are witnessing a constant evolution of the threat. Good, timely, and legitimate data today, maybe be deemed evil or obsolete tomorrow, but then later be relevant as a mission focus shifts. We are already familiar with this phenomenon. The difficulty is relative because the fight against data pollution basically uses the same thinking, the same organization, and the same tools as we have in cyber. The transformation is mainly based on the definition of pollution, on the ability to define the TTPs for handling data pollution.

This characterization of pollution is much more than a technical or organizational challenge, it is also of a philosophical/political nature. The data used in combatting cyber threats has similar philosophical dilemmas. If undesirable data were labeled as polluted data, any software developed under this principle could become a formidable censorship tool. Some countries have adopted this philosophy while others declined.

The battle against data pollution poses the exact same issue. At its heart lies the question of what we mean by polluted data, and whether this includes the informational level. The challenge here arises from the abuses that can occur at this level. We already have reflexes such as parental filters, but we must bear in mind that the subject is well known, and there are as many different answers as possible, often linked culturally. However, we need to be cautious, because these tools can also be used to suppress information that someone for selfish reasons would not want to see propagated.

While particularly useful as a military tool in the context of information warfare, we must remain aware that countering data pollution requires a degree of vigilance and control by policy, the virtue of democracy, and the guiding finger of the market.🇺🇸

NOTES

1. George I. Corbari, Neil Khatod, John F. Popiak, and Pete Sinclair, “Mission Thread Analysis: Establishing a Common Framework in a Multi-discipline Domain to Enhance Defensive Cyberspace Operations,” (April 2024). *The Cyber Defense Review*, Volume 9, No. 1, pp 37-54. https://cyberdefensereview.army.mil/Portals/6/Documents/2024_Spring/CDR_CDRV9N1-WEB-Spring-2024.pdf, 46.
2. Mohammed M. Alani, “Big data in cybersecurity: a survey of applications and future trends,” (January 6, 2021). *Journal of Reliable Intelligent Environments*. <https://doi.org/10.1007/s40860-020-00120-3>, 7, 85–114.
3. Dayton Ward, “Fit to be Spied: Fitness Trackers and OPSEC Risks,” (March 9, 2018). *The NCO Journal*, Army University Press. <https://www.armyupress.army.mil/Portals/7/nco-journal/docs/Fitbit.pdf>.
4. Korolov, Alex, “The cybersecurity challenges and opportunities of digital twins,” (December 6, 2022). *CSO Online*, <https://www.csoonline.com/article/574179/the-cybersecurity-challenges-and-opportunities-of-digital-twins.html>.
5. Mitchelson, Deryck, “Learn about the Double-Edged Sword of AI in Cybersecurity,” (October 27, 2023). *World Economic Forum*, <https://www.weforum.org/agenda/2023/10/the-double-edged-sword-of-artificial-intelligence-in-cybersecurity/>.
6. Orf, Darren, “A New Study Says AI Is Eating Its Own Tail,” (October 29, 2023). *Popular Mechanics*, <https://www.popularmechanics.com/technology/a44675279/ai-content-model-collapse/>.
7. Alcaraz, Anthony, “When AI Distorts Its Own Reality: The Dual Threats of Model Collapse and Source Bias,” (November 11, 2023). In *Medium*, <https://ai.plainenglish.io/when-ai-distorts-its-own-reality-the-dual-threats-of-model-collapse-and-source-bias-13e4da42fd10>.
8. Quentin Bertrand, Avishek Joey Bose, Alexandre Duplessis, Marco Jiralerspong, and Gauthier Gide, “On the Stability of Iterative Retraining of Generative Models on their own Data,” (September 30, 2023). Conference Paper, *The International Conference on Learning Representations*. Vienna. <https://doi.org/10.48550/arXiv.2310.00429>.

DISCLAIMER

The views expressed in this work are those of the author and do not reflect the official policy or position of Saint-Cyr Coëtquidan Military Academy Research Center, the French COMCYBER, or the government of France.

THE CYBER DEFENSE REVIEW

◆ SPECIAL EDITION ARTICLES ◆

The Importance of Analytic Superiority in a World of Big Data and AI

Dr. Robert L. Grossman

Dr. Emily O. Goldman

ABSTRACT

Rapid advances in machine learning, deep learning, artificial intelligence (AI), large language models, and generative AI have accelerated efforts to leverage these technologies for military advantage. We refer to these and related technologies as analytics. We present a framework as a guide to achieving “analytic superiority,” which is the operational advantage obtained through the ability to collect data required for analytics; build useful, performant, and robust analytic models; and deploy analytic models in operational systems to achieve objectives, while exploiting or denying an adversary’s ability to do the same. Analytic superiority is best understood in the context of the analytic capabilities of one’s adversaries, who also collect data, build models, and deploy them to achieve their own objectives and defeat the analytics of their adversaries. This framework emphasizes developing an analytic strategy, collecting the data required, developing analytic infrastructure for managing and analyzing the data, building analytic models, and deploying analytics into operational systems to meet the objectives required by the analytic strategy. Although analytic competition is not new, it is an under-appreciated dimension of military and strategic competition, and it is advancing at a faster pace than any previous technological competition. We discuss how U.S. cyberspace superiority, which is foundational to military advantage in the physical domains, now depends on prevailing in analytic competition with adversaries and thus requires adopting a strategy and processes to achieve analytic superiority.

© 2024 Dr. Robert L. Grossman and Dr. Emily O. Goldman



Dr. Robert L. Grossman is the Frederick H. Rawson Distinguished Service Professor in Medicine and Computer Science and the Director of the Center for Translational Data Science at the University of Chicago. He is also a Partner at Analytic Strategy Partners, which helps organizations develop and execute AI strategies. He is the principal investigator for the National Cancer Institute Genomic Data Commons (GDC), a platform for the cancer research community that manages, analyzes, integrates, and shares large-scale genomic datasets in support of precision medicine. He founded Open Data Group in 2002 and was its Managing Partner from 2002-2015. Open Data Group provided analytic services to help companies build and deploy machine learning models over big data and to operationalize AI.

ARTIFICIAL INTELLIGENCE AND MILITARY ADVANTAGE

Rapid advances in machine learning (ML), artificial intelligence (AI), and generative AI (GAI), which we will refer to simply as ML/AI, will play an ever-increasing role in commercial, defense, and intelligence systems. There has been continual progress in ML/AI over the past forty plus years, punctuated by several inflection points. One important inflection point occurred in the early 2000's with deep learning that leveraged GPUs (the graphics processors used by gaming applications). This allowed neural networks with many dozens of layers and millions of parameters to be used, leading to significant advances in the processing of text and images.¹ Increasingly larger deep learning neural networks with billions of parameters were built over larger and larger GPU clusters and then specialized to particular applications in various ways, leading to the emergence of large language models (LLM) and generative AI. Another important inflection point occurred on November 30, 2022 with the release of ChatGPT by OpenAI. ChatGPT provided Large Language Models as a Service, setting a new record by reaching over 100 million users by January 2023.

These two inflection points have fundamentally changed the power and capabilities of analytics. The rapid rise in ML/AI-focused strategies, policies, regulations, and guidance documents by States, international organizations, and supra-national organizations like the European Union reflects widespread recognition of ML/AI's potential and risks.² The impact on military systems cannot be overemphasized. Over time, ML/AI models and services will be ubiquitous, and enduring military advantage will depend on the strategy and processes for both employing ML/AI towards strategic objectives and countering adversaries' use of ML/AI toward the same objectives.



Dr. Emily Goldman serves as a strategist at U.S. Cyber Command and a thought leader on cyber policy. She was cyber advisor to the Director of Policy Planning at the Department of State, 2018–19. From 2014 to 2018 she directed the U.S. Cyber Command / National Security Agency Combined Action Group, leading a team that wrote the 2018 U.S. Cyber Command vision, *Achieve and Maintain Cyberspace Superiority*. She was a professor of Political Science at the University of California, Davis, for two decades and has published and lectured widely on strategy, cybersecurity, and military innovation. *Cyber Persistence Theory: Redefining National Security in Cyberspace*, with Michael Fischerkeller and Richard Harknett, was published by Oxford University Press in 2022.

In this article, we use the term *analytics* to refer to statistical models, machine learning, deep learning, AI, GAI, and LLM. More expansively, by analytics we also include more general algorithms for processing and analyzing data, including embedded algorithms, such as those used in electronic warfare and electronic countermeasures. From the perspective of analytic superiority, the differences between algorithms,³ models,⁴ rules, machine learning, deep learning, AI, LLM and GAI⁵ are not important. There is no standard name to refer to all of these and we will use the term analytics. More narrowly, we will use the term ML/AI if we want to emphasize some of the recent work in machine learning, deep learning, large language models, and generative AI.

Although analytic competition is not new, the ubiquity of models in digital systems and services has made analytic competition an increasingly important, but underappreciated dimension of military and strategic competition between states. This is also the case for commercial competition, but better grasped by the private sector. For defense and intelligence systems, in particular, the United States is competing with adversaries using their own analytics while trying to defeat, disadvantage, interfere, or trick ours. The race to deploy analytics in operations to achieve mission effects and degrade and defeat adversary analytics is as critical as, and occurring at a faster pace, than any arms race we have ever confronted. The United States must set up better analytic infrastructure, employ better analytic models, update them more frequently, build better analytic systems, and train better data scientists, system operators, and end users than our adversaries in order to achieve and maintain analytic superiority.

ANALYTIC SUPERIORITY AND SOME COMMON MISCONCEPTIONS

We define analytic superiority as the operational advantage obtained through the ability to collect data

required for analytics; build useful, performant, and robust analytic models; and deploy analytic models in operational systems (from core to edge) to achieve objectives, while exploiting or denying an adversary's ability to do the same. Note that analytic superiority is not the same as information superiority. The latter is a much broader concept, while analytic superiority is specifically focused on analytics (including ML/AI, models, rules, and related concepts) and their role in warfare and other adversarial situations.

Analytic superiority is best understood in the context of adversarial analytics and thus, measured against the analytic capabilities of one's adversaries. Competitors also collect data, build models, and deploy models to achieve their own objectives and deny, degrade, or defeat the analytics of their adversaries. Those who build better models; build them over larger volumes, higher velocities, and more varieties of data; and build agile and scalable infrastructures to deploy them are more likely to benefit from the operational advantages conferred by analytic superiority.

The importance of analytic competition for military competition is not new. A good example is provided by electronic warfare in World War II. As radar provided a military advantage to detecting and shooting down planes, electronic countermeasures were developed to trick or deceive radar. These included both physical countermeasures, like deploying chaff, and electronic countermeasures, like jamming or spoofing radar systems. The development of models for detecting planes and other moving objects using radar, and electronic countermeasures for deceiving these models, is an example of adversarial analytics and shows the importance of analytic superiority for achieving military objectives—in this instance destroying targets with bombers. Clearly, the concept of analytic superiority is not unique to ML/AI; it applies to all analytic models. This dynamic may enable the basic cycle of weapons development and counter measures, but these are not synonymous because many military innovations and counter innovations do not depend on the operational deployment of analytic models.

There are several common misconceptions about analytic superiority. The first is that analytic superiority is achieved solely by building more accurate and precise analytic models. One may have much better models with much higher precision and recall, but if the force cannot deploy the models into operational systems (for military and/or intelligence purposes), or not deploy them in time, while the adversary can, even if the adversary's models are worse, one will not achieve analytic superiority. Moreover, even with better models, if one lacks an analytic infrastructure that collects the data required and efficiently builds analytic models while an adversary has these capabilities, then one will not achieve analytic superiority. In general, the higher quality and the timelier the data, the simpler the models can be and the more effective and robust they will be when deployed. It is nearly impossible to achieve analytic superiority without a sufficiently powerful computing infrastructure to manage and analyze data and to build and deploy models.

A second common misconception is that today's ML/AI models require a fundamentally new approach to analytics. The importance of analytics is not new; nor is the viewpoint, as we argue below, that analytics in competitive environments require an end-to-end process for getting and managing data, building models, deploying models, and extracting value from models (as captured in Figure 1 by the framework of the analytic diamond, which is described below). It is also not new that novel models emerge from time to time and provide disruptive changes in what is possible in analytics.⁶ Moreover, it has always been a challenge for large organizations to develop complex software systems, especially those involving large-scale or complex data.⁷ That being said, over the past forty years, the exponentially growing amount of data and computing infrastructure, enabled by Moore's Law, has created year by year ever more powerful analytic models built on ever larger amounts of data. Even experts today are surprised at the speed at which new AI/GAI algorithms, models, systems, services, and applications are being developed, and the unexpected power and capabilities that these models, systems, and services have. This is widely recognized across commercial and military sectors, and by national governments globally.

A third common misconception is that analytics is mainly about improving the workflows of intelligence analysts. Of course, AI-powered analytics will undoubtedly improve intelligence analyst workflow, but they will do much more. They will impact any military system (e.g. weapons and navigational platforms) that employs analytic models, which means virtually all military platforms.

The U.S. Department of Defense's 2023 *Data, Analytics, and Artificial Intelligence Adoption Strategy* calls for a "systematic, agile approach to data, analytics, and AI adoption that is repeatable by all DoD Components" for enduring decision advantage. U.S. Cyber Command stood up an AI task force to move from "opportunistic AI application to systematic adoption."⁸ Its focus is delivering capabilities to the force, posturing the Command to enable AI adoption, and countering AI threats. As the Department of Defense invests in new technologies, makes associated infrastructure, organizational, and force design decisions, it should do so informed by a framework for achieving and sustaining analytic superiority to provide both decision advantage for commanders and operational advantage for warfighters.

ANALYTIC SUPERIORITY IN THE COMMERCIAL WORLD

The importance of analytics for business competitiveness in general has been long recognized.⁹ More specifically, in some commercial applications, the specific role of analytic superiority has also been long recognized. We discuss two examples in some detail—payments fraud and high frequency trading—because it is easier to talk at the granular level about commercial applications than it is about sensitive military applications.

A payment network enables a group of financial institutions to transfer funds digitally between individuals, businesses, or financial institutions. The most well-known examples are the payment networks between banks and merchants that enable individuals and businesses to use credit cards to buy merchandise and pay later when billed by their banks. Payment networks also support ATM withdrawals, gift cards, mobile payments, and other types of transactions. The development of analytic models by payments networks to stop fraud and the corresponding adjustment of tactics, techniques, and procedures (TTPs) by criminals trying to perpetrate fraud, is an example of competitive or adversarial analytics.

A wide array of bad actors, including individual criminals, criminal gangs, and states conduct payments fraud. In 2022, payments fraud totaled \$11.64 billion in the U.S. and \$32.34 billion worldwide.¹⁰ Given the scale of payments fraud, a whole ecosystem has grown up around this industry. It includes the individuals using fraudulent cards or engaging in fraudulent transactions, software engineers developing software supporting fraudulent transactions, vendors selling stolen personal information that can be used to obtain fraudulent credit cards, lines of credit, etc., and electronic marketplaces where suppliers of data, software, and other services offer their goods and services to those engaging in fraud.

A payments system needs the right data at the right time to develop analytic models to quickly detect and stop fraudulent transactions. This requires a sufficiently powerful analytic infrastructure, analytic modeling capability, and, importantly, fraud systems (an example of analytic operations) to detect and stop fraudulent payments *as they are occurring* before losses are significant. At the same time, actions to stop fraudulent transactions should only minimally impact normal operations and the customer experience. By fraud systems we mean systems that not only detect putative fraud but also contact users to verify whether questionable transactions are in fact theirs, stop future use of compromised cards, and issue new credit cards to replace compromised cards. In this example, analytic operations include embedding the fraud models into operational systems for quickly stopping questionable transactions, quickly determining whether the transaction is valid or not, while minimizing impact on users. In the context of systems, there is always the balance of simpler models that work with the data immediately available versus more complex models that can leverage richer data that might not be immediately available.

Meanwhile, bad actors will adjust their tactics, techniques, and procedures to elude analytic models trying to detect and block their fraudulent transactions. The payments system is a dynamic space with sophisticated bad actors not only developing sophisticated attacks, but also quickly identifying simple new vulnerabilities that arise as systems are updated, new devices are installed, and new products are introduced. Companies tolerate a certain amount of fraud as the “cost of doing business.” As the amount of payments fraud increases and begins to impact reputation or customer experience, companies will increase their efforts to detect and stop the fraud. Building the perfect analytic model is not the goal, but

rather building a good enough model that can be deployed and updated quickly enough to keep fraud at an acceptable level.

Analytic superiority has also been a critical aspect of high frequency trading since its inception. The key to high frequency trading is not just having the best models, but rather a trading system that can beat other trading systems for the small opportunities available with each trade. This requires having the right data, the right models, and an underlying computing infrastructure that can get the latest trades a bit faster than others in the market.

A good example is provided by the efforts of different traders to reduce the latency between moving data between Chicago and the New York / New Jersey region. As described in the book *Flash Boys* by Michael Lewis,¹¹ moving data between the Chicago Mercantile Exchange, where futures and options are traded, and Carteret, New Jersey, where the Nasdaq data center is located, took about 14.5 milliseconds (ms) using commercial Internet Service Providers (ISPs). To generate a competitive advantage to high frequency traders, Spread Networks developed a dark fiber network that took a more direct route between Chicago and the Nasdaq data center in New Jersey that was able to bring the latency of sending a message down to about 13.1 ms, providing a slight advantage to traders using their networks versus commercial ISPs. Finally, since light travels slower in glass fibers than in the air, Windy Apple Technologies set up a microwave network between Chicago and the New Jersey data centers that brought down the latency of sending a message to about 9.0 ms. Of course, traders had to also develop the appropriate analytic models, trading strategies, and computing infrastructure to take advantage of these millisecond improvements in networking infrastructure.

A FRAMEWORK TO ACHIEVE ANALYTIC SUPERIORITY

Achieving analytic superiority is much easier if one adopts a framework for analytics. The analytic diamond is one such framework that has proven useful,¹² although other frameworks could be used. The analytic diamond distinguishes processes for developing an analytic strategy, for building analytic models, for deploying analytic models, and for managing all the data required to build and deploy analytic models. It further defines an associated analytic governance model as well as an analytic maturity model to measure progress towards developing the processes, infrastructure, and talent required to achieving analytic superiority.

More specifically, an analytic strategy examines analytic opportunities, decides which to pursue, ensures that efforts are most likely to result in value to the organization, and measures the value. Analytic infrastructure collects and manages the data required for analytic operations and modeling. Analytic modeling analyzes data and builds the analytic models required by the organization. Analytic operations deploy and operationalize analytic models into products, services, or internal systems to extract value. These four features, when arranged in a diamond as in Figure 1, reveal six processes required for effective analytics, which in the context of analytic superiority means effective against an adversary.

Organizations need to (1) develop an analytic strategy that prioritizes analytic projects and efforts, and (2) build the analytic infrastructure to train models over large amounts of data. Processes for (3) collecting and managing data so that it is available for modeling and operations, (4) developing appropriate analytic models, and (5) deploying the analytic models into operational systems must all be executed in the time frame required to achieve the objectives and value identified in the analytic strategy. An organization must also develop the operational systems, and their associated TTPs, to achieve the objectives and value identified in the analytic strategy. By operational systems here we mean any system that integrates, embeds, or interoperates with analytics, such as: systems for cyber defense and cyber operations; analysts’ tools and support systems; electronic warfare and electronic countermeasures; systems supporting information superiority; fire systems for weapons platforms; systems supporting the targeting cycle, etc. The next step is to (6) accurately measure the value and impact achieved through the operational systems, and update the entire process accordingly. Enabling these processes is analytic governance, analytic security and compliance, and recruiting, training, and retaining the technical staff required. The analytic diamond can be applied to a wide variety of military operational priorities, as can be seen from the list of operational systems above that leverage analytics.




Figure 1. The Analytic Diamond

The analytic diamond shows why analytic superiority is not synonymous with a data strategy (although a data strategy is a subset of analytic strategy, as are the required components of an IT strategy to support analytics). Nor is analytic superiority reducible to building models and algorithms. In practice, analytic models have little value until they are turned into a functional system through an engineering process (AI engineering) and then the system is integrated into operational processes and/or operational weapon systems to bring value to the organization (AI operations or AIOps). Analytic superiority is sometimes misleadingly viewed as an OODA loop, which is a tactical activity rather than an enterprise level construct that integrates four analytic functions—strategy, infrastructure, modeling, and operations—to achieve competitive advantage over an adversary.

Although there is no single way to achieve analytic superiority, just as there is no one way to defeat an adversary in a particular engagement, in general achieving it includes some combination of the following: analytic task advantage (for example, having better analytic models than your adversary and updating them more frequently); asymmetric advantage (for example, asymmetrically attacking your adversary’s models or infrastructure); and autonomy advantage (leveraging autonomy and semi-autonomy more effectively). Autonomy and counter-autonomy are familiar and well understood in the context of autonomous weapons

systems,¹³ but are perhaps less familiar and less well understood in the context of adversarial analytics. Autonomy and semi-autonomy provide scale and speed for analytics, both for defensive and offensive operations. Countering these operations can be done by identifying suitable weak points in adversaries' analytic workflows and then disrupting and defeating those weak points. Figure 2 shows notionally how Red and Blue attack each others' analytic infrastructure, analytic models, analytic operations, and analytic strategy. For example, Red may try to poison Blue's data and embed in Blue's analytic infrastructure, while Blue may try to weaken or poison Red's analytic models and blunt the operational impact of Red's models.

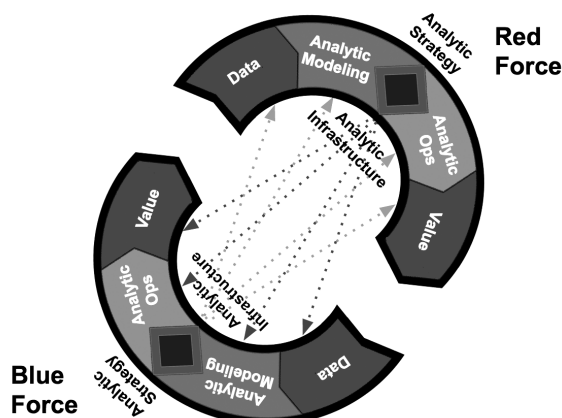


Figure 2. Achieving Analytic Superiority.

Achieving analytic superiority requires that all the analytic models needed to achieve the goals of an analytic strategy have the data they need to function, are updated as required, and are integrated with the required operational systems to achieve the desired outcomes. This is not easy to achieve. Here are two simplified examples to illustrate this point.

First, consider the role of sensors in cybersecurity. Sensors can produce a lot of data and unless there is an adequate analytic infrastructure (the bottom of the diamond) to manage the data, critical patterns that may span many days of data can be missed. Detecting behavioral patterns in log and sensor data require not only analytic models (the left-hand side of the diamond), but keeping these models up to date as TTPs change. Once models detect potential intrusions, analytic operations (the right-hand side of the diamond) determine the appropriate actions and counter measures to quickly pinpoint, contain, and mitigate the intrusions. Finally, there are always limited resources, and analytic strategy (the top of the diamond) prioritizes which problems are tackled and in what order. These factors have always been present in the management of data for analytics, the building of models for analytics, and the deployment of models to take appropriate actions. On the other hand, recent advances in ML/AI create new opportunities for both the U.S. and our adversaries, and managing and coordinating the components of the analytic diamond across the enterprise to achieve analytic superiority is critical.

As a second example, protecting a Navy ship requires object identification and tracking models to identify potential threats and track them so that appropriate fires and countermeasures can be assigned to different threats. Each of these systems requires appropriate analytics to protect the ship, and the different analytic models must work together. More precisely, the outputs of some models are the inputs to other models creating an *analytic workflow*. To achieve the

desired outcome each model must get the data it needs. This interdependence creates fragility; thus, various techniques are used to create more robust models and more robust analytic workflows. For example, models and workflows can impute data when it is missing,¹⁴ detect and remove outliers,¹⁵ and reduce the fidelity of models in order to make the models and workflows more robust and more likely to provide correct results when data is missing or delayed, or when sensors are down, for example.

The concept of a system of systems (SOS) is familiar in the context of complex weapon systems.¹⁶ In the context of analytic superiority, it is helpful to extend this to what one might call a system of systems *and models* (SOSAM). This includes not only the system of systems, but also all the analytic models and workflows needed for the system of systems to achieve the desired outcomes for the analytic strategy. This almost always requires multiple analytic models and systems. In the simplified example, an adversary need only identify and defeat the weakest link in an analytic workflow to attack the ship successfully. To achieve analytic superiority, your SOSAM must be more capable than your adversary's SOSAM in all operating conditions, including in crisis and conflict. This is a concept familiar to cyber defenders trying to keep an adversary out of their system. No matter how good the algorithms, signatures, and models in your firewalls, an adversary need only exploit any weak link in the system, whether from weak passwords, default passwords, phishing, stolen credentials, cross site scripting attacks, etc. More generally, as we discuss below in the section on analytic maturity models, it is also critical to build analytics over two or more analytic systems or workflows ("hierarchical" analytics) to detect patterns and threats that may not be apparent otherwise. Analytic superiority, it turns out, is rarely about one's analytic model simply being more accurate (measured as the percentage of correct predictions) than the adversary's model, but rather about the robustness and effectiveness of entire analytic workflows or hierarchies of analytic workflows.

MEASURING THE EFFECTIVENESS OF ANALYTIC SYSTEMS

Defense and intelligence analysts are familiar with the application of analytic models to digest intelligence and defense data, organize it, draw conclusions with confidence levels, write reports, and disseminate the reports to decision makers and policy makers. Although multiple analytic models may be used to collect and analyze the data, the output is the report; actions, if any, are the responsibility of some other part of the organization. Analytic superiority, as applied to military activities, requires that the outputs of analytic models be integrated into operational systems to achieve specific mission related objectives. The focus, therefore, must be on the effectiveness of the analytic system as a whole, not just the output of analytic models.

One way to understand this difference is to think of the output of analytic models as scores (say from 1 to 100), which are used to take *actions*, and the operational significance of these actions are evaluated with *measures*. This sequence is summarized with the acronym *SAM* for scores-actions-measures.¹⁷ Typically, the effectiveness of a single analytic model is measured

with detection and false positive rates. In contrast, from the SAM perspective, these model-specific metrics influence how often the appropriate actions are taken for the appropriate situations. The effectiveness of the actions determines the measures, which summarize the effectiveness of the analytic system as a whole.

Consider a cyber defense model that identifies users whose credentials may have been compromised. Assume that the model produces a score in the range 1-100, that credentials are revoked for scores above 90 and then manually investigated (the action), and credentials are manually investigated for scores between 80 and 90 (another action). The model producing these scores has a detection rate and false positive rate, but measures of effectiveness that incorporate action might be how quickly compromised credentials are detected (seconds, minutes, etc.), how much damage is done before they are found (say measured with dollars), and how much lost productivity (measured with person-hours) arises from improperly revoked credentials. Measures can be created at various levels, such as the analytic task level, at the analytic system level, and at the mission level.

As another example of SAM, a network or system intrusion might set off multiple alerts from multiple models, including alerts from firewalls, behavioral alerts from endpoint systems, alerts from identity and access management systems, etc. Many of these are false positives, producing alert fatigue. One solution is to build models that process the outputs of all these systems and produce alerts at the incident level. The associated actions might include automatically revoking access credentials or isolating network segments when the alert scores are high enough. Measures might reflect the amount that alerts are reduced, say from 100 alerts from the original systems that alert to 5 alerts of candidate incidents, and the percentage of true incidents detected.

Other types of analytic models produce different outputs. For example, large language models and generative AI models produce text, images and other outputs, giving rise to text-action-measures, or *TAM*. It is not the text that is evaluated, but text associated with some action, such as producing SQL code for querying a database of cyber threat data or log data, and the measure might be the productivity increase for junior and senior software developers. Measures can be devised to assess the cumulative impact of actions taken against ransomware groups or an advanced persistent threat actor such as reducing the success of intrusions or improving the targeting process. Other types of measures can be devised to assess decision advantage, such as efficiency in organizing and prioritizing information so that a commander can make decisions faster and still leverage more of the available information.

ANALYTIC GOVERNANCE

Analytic superiority requires a governance framework to ensure that the right people, processes, and frameworks are in place to identify and achieve the organization's priority analytic objectives. There are broad similarities between analytic governance and IT governance

processes designed to control and oversee risk. IT governance ensures that IT investments generate business value; that risks associated with IT are mitigated; and that the organization makes sound, long-term decisions with accountability and traceability to those funding, developing, supporting, and using IT resources.¹⁸

With this definition and the analytic diamond in mind, we argue that an analytic governance framework should accomplish the following: (1) ensure sound long-term decisions about analytics are reached and that investments in analytics generate value; (2) operate in such a way that data, derived data, and data and analytic products are protected and managed in a secure and compliant fashion; (3) ensure accountability, transparency, and traceability to those funding, supporting, and using analytic resources; and (4) ensure analytic workflow: that data is available to modelers, that analytic models can be deployed, and that impact and value of analytic models is quantified and tracked.¹⁹

Analytic governance is Commander's business and cannot be delegated because it crosses the entire enterprise. By contrast, data and sensor governance can be delegated further down in the organization. Typically, data governance is delegated to the CIO. Analytic governance, however, requires an influential champion to ensure the integration of functions—managing data, building models, exploiting models for advantage by analysts and operators—at scale. These processes must be driven by operational priorities, and by the organizational element that sets those priorities.

ANALYTIC MATURITY

Just as software maturity can be measured with a Capability Maturity Model that quantifies the level at which an organization's business processes develop software and complete a software project,²⁰ processes for developing and deploying analytics can be developed to measure an organization's analytic maturity. Depending upon the desired outcome, analytic models must be integrated across various hierarchies and across different units within an organization. For this reason, the analytic maturity of an organization is assessed along three dimensions.²¹ First, how repeatable are the processes for managing data, building models, deploying models into operations, and quantifying the effectiveness and impact of the models? Second, how widespread are these processes, both horizontally and vertically throughout the organizational structure? Third, do these processes support the organization's analytic strategy?

Each organizational unit that supports analytics can be evaluated along a maturity dimension. For example, can the unit 1) build reports based upon data; 2) build and deploy analytic models based upon data; 3) build and deploy analytic models with a repeatable process; 4) build and deploy the appropriate analytic models, as prioritized by the analytic strategy. Similarly, the entire enterprise can be assessed along a maturity dimension: 1) does the organization provide the enterprise services so that a unit can get the data, build the models,

deploy the models, and extract the value required; 2) how widespread across the organization are units that build and deploy high quality and effective analytics; 3) is there analytic governance at the enterprise level; 4) are analytic models built by different units integrated together in such a way to achieve the organization's overall analytic objectives or do they create unanticipated challenges for some analytic models?

As a simple example from the financial services world, one division of credit card issuers is responsible for acquiring new customers with acquisition models and expanding the company's "share of their wallet" with cross sell models. Another division (credit risk) is responsible for determining how likely customers are to default on the credit extended to them with credit default models. The enterprise dimension of an analytic maturity model manages whether the right customers are being acquired so that the acquisition division does not meet its yearly targets by acquiring customers who are likely to cause future problems by defaulting on their loans.

ANALYTIC SUPERIORITY AND PERSISTENT ENGAGEMENT IN CYBERSPACE

Superiority in cyberspace has come to be seen as foundational to military advantage in the physical domains of land, air, sea, and space. The United States is not alone in this assessment. The People's Republic of China (PRC) also sees superiority in cyberspace as core to its theories of victory.²² For this reason, we apply the concept of analytic superiority to cyberspace, which has itself become a major battleground in strategic competition among states.

In 2018, U.S. Cyber Command defined cyberspace superiority as "the degree of dominance in cyberspace by one force that permits the secure, reliable conduct of operations by that force, and its related land, air, maritime, and space forces at a given time and place without prohibitive interference by an adversary."²³ Cyberspace is not only a domain of military operations, however, but also a strategic environment in and through which adversaries put U.S. national security at risk.²⁴ Adversaries operate continuously in and through cyberspace below the threshold of armed conflict, extending their influence and eroding U.S. military, economic, and political power without resort to physical aggression. General Timothy Haugh, Commander of U.S. Cyber Command and Director of the National Security Agency, recently remarked that the PRC is pursuing a policy of global dominance, but hopes to achieve that without a kinetic, real-world military fight. They are using cutting-edge technologies to achieve advantage.²⁵

U.S. cyber forces adapted their operational approach to address the requirements of strategic competition outside armed conflict as well as those aimed at helping to deter armed conflict and prevail should it occur.²⁶ Recognizing that the United States was losing ground to adversaries operating in and through cyberspace below armed conflict, U.S. Cyber Command pivoted from a bias for "cyber response" to action through "cyber persistence." Empowered with new authorities, military cyber forces began proactively defending and contesting adversaries globally, continuously, and at scale. Reflecting upon his five years as commander, General Paul Nakasone

remarked, “I think we got persistent engagement completely right. ... If you’re on the sidelines watching this, you’re going to get hit. That’s why I think it’s so important for our forces worldwide to be able to be engaged, and being able to act and understand what our adversaries are doing ... Being able to continue to operate day in and day out, this is how you get really good. You operate in the domain.”²⁷ The Russia-Ukraine war has shown how persistent engagement also sets conditions for successful contingency operations by curtailing an adversary’s freedom of maneuver (e.g., precluding options, eroding confidence, generating organizational friction, and degrading capabilities), and generating options and opportunities for crisis and conflict.

Cyberspace is an interconnected, fluid space wherein those who continuously anticipate and act can set, reset, and maintain conditions in their favor. This is what it means to hold the initiative in cyberspace, where the scale and scope of change is so vast that anticipating exploitation and acting based on that insight is the key to security. A state that cedes the initiative can assume that its adversaries will set conditions to undermine its security while increasing their own.²⁸ Highly capable adversaries are conducting operations in cyberspace at a volume and pace that overtake the ability of defenders to discover and counter. Increasingly sophisticated techniques, such as “Living off the Land” (LOTL),²⁹ allow them to evade detection more easily and preposition in military and civilian critical infrastructure, setting conditions for use during crisis and conflict. Defeating complex obfuscation techniques requires timely processing of massive data and deploying analytics to detect this type of behavior into operational (in this instance defensive) systems at speed and scale to outpace the malicious actor’s ability to shift across infrastructures and preclude interdiction. With innovations in big data and AI, the cyberspace force that achieves and sustains superiority in analytic competition will have the initiative in persistent cyber engagements.

Tactically, one can still think in terms of prevailing in defensive and offensive analytics, but cyberspace is so fluid and dynamic that offensive and defensive advantage is not meaningful at a strategic level. What is meaningful is whether or not one has the initiative in setting the conditions of cyberspace in one’s favor, where “those conditions are measured as the relative balance between being cyber vulnerable to exploitation and being able to exploit the cyber vulnerabilities of others” by “anticipating the exploitation that *will come next* by either you as a defender or another State as an attacker.”³⁰ Developments in AI have made analytic superiority necessary for achieving and sustaining cyberspace superiority.

The Joint Cyber Warfighting Architecture (JCWA) is U.S. Cyber Command’s platform and associated capabilities that enable Cyber Operations Forces to conduct full-spectrum cyberspace operations, globally at-scale. The Unified Platform is the data hub of JCWA and provides the analytic infrastructure for deploying and managing data. The federated nature of this system for deploying and managing data is one reason U.S. Cyber Command was granted greater technical responsibility and authority to direct the development, integration and fielding of critical capabilities and infrastructure in the JCWA.

The framework of the analytic diamond specifies other requirements for achieving analytic superiority, foremost being a strategy that prioritizes analytic objectives and requirements. A subordinate data strategy should identify data collection, management, standards, and analytics for full-spectrum operations, which in turn should inform and drive specific computing capabilities, data management systems, data models, and analytic platforms. Designated entities should be made responsible for building and deploying analytic models. Some element of the organization should have the mission to attack all elements of the adversary's analytic diamond as a priority or objective, for example by poisoning data, countering models or infrastructure, or blunting the operational impact of the adversary's models. Just as there are well defined disciplines of counter intelligence, counter terrorism, and cyber defense to counter cyber-attacks, we need to develop a discipline of counter analytics to understand how to disrupt and defeat the analytics of our adversaries. Given that DoD is organized around weapons platforms, U.S. Cyber Command may want to think of its entire organization as a platform to acquire, manage, develop, and deploy data and analytics for full-spectrum operations at scale—and counter adversaries' ability to do the same.

CHALLENGES AND LESSONS LEARNED

Achieving analytic superiority is challenging. Many of the hurdles are familiar ones that DoD has faced when developing software, managing and analyzing data at scale, and interoperating defense systems. For example, DoD's transition to cloud computing, an important enabling technology for training and using ML/AI over large datasets, has been rocky, significantly lagging behind commercial adoption of the technology.³¹ Many software projects fail for a variety of well understood reasons,³² and software systems supporting intelligence and defense systems are particularly challenging.³³

Many analytic/AI projects, like many data warehousing projects, fail because assembling and managing teams with multiple skills is hard. Building a data warehouse requires assembling and managing a team that knows about both software and data, which are two distinct skills. Many analytic/AI projects also fail because a disproportionate amount of time, ingenuity, and effort is spent focusing on the AI algorithm or analytic model of interest, say a convolutional neural network (CNN), or fine tuning an LLM. AI projects that succeed, however, spend significant time, effort, and ingenuity to get all the data required and to deploy the model in such a way to achieve an operational advantage.³⁴ It is helpful to think of this as a three-phase process. Phase one focuses on acquiring and curating data to be used as inputs to the AI and analytic algorithms ("collecting the data required for the models"). Phase three focuses on integrating the outputs of the AI and analytic algorithms into current operational systems or developing new systems around them ("deploying the analytic models"), and then extracting operational value from the systems as a whole, once the model has been deployed. Many successful projects spend less than 20% of project's total time and effort on phase two—

finding the right algorithms and processing the data using them (“developing the analytic models”). See Figure 1.

Achieving analytic superiority also requires that the right data from the right sensor get to the right analytic model in the right time. This is especially challenging when the sensors, models, and effects cross war fighting domains. Programs such as the Joint All-Domain Command and Control (JADC2) are designed to address these issues, but face challenges of their own.³⁵ Finally, achieving and maintaining analytic superiority requires a developer, DevOps, and DevSecOps mindset, environment, and acquisition process. For many DoD software projects this has been challenging to achieve because the defense acquisition process is not agile.

AI IN THE PRC

The PRC is investing heavily in AI, cloud computing, and related technologies.³⁶ It is engaged in a campaign to attain technological superiority and place U.S. critical and national infrastructure at risk.³⁷ In some ways, the Chinese have an edge with their smart cities initiatives and extensive industrial online-to-offline (O2O) industrial base.³⁸ With these efforts, the PRC's many data engineers are gaining deep experience developing and deploying systems at scale that use big data and AI every day. Furthermore, much of the technology developed by Chinese companies for smart cities is dual use and can also be applied to military systems. Those Chinese companies work both commercially and for the government, and competition among them is usually regarded as extremely intense, yielding rapid technological advancement.

The PRC is also acquiring a wide range of data sources, to include proprietary big data, machine learning, and AI technology stolen from U.S. companies through computer intrusions.³⁹ Another source of large scale data and analysis that is particular to the PRC is provided by its social credit system, which integrates data from multiple sources, including analysis of internet traffic, monitoring of social media, analysis of mobile telemetry data from phones and cars, and pervasive video surveillance. All of this data is analyzed at scale using AI and other techniques to produce a score for each individual.⁴⁰ The scale of data provided by Baidu, Alibaba and Tencent (BAT) provides smaller companies that are part of the Baidu, Alibaba or Tencent ecosystem the data required for AI and machine learning. These qualities of the PRC's analytic ecosystem make achieving analytic superiority over China challenging, but essential.

CONCLUSION

Statistical models, machine learning, AI models and related technologies (what we call analytics in this article) have been critical enabling technologies ever since military systems integrated digital technology. A decade ago deep learning substantially increased the performance of these models. More recently, the emergence of LLM and GAI models introduced significantly new capabilities to these models. Yet powerful models alone do not necessarily provide an advantage in competition, crisis, or conflict.

We focus on the importance of analytic superiority and offer a framework—the analytic diamond—as a guide to achieving analytic superiority. The framework stresses the importance of an analytic strategy for identifying and prioritizing analytic opportunities; obtaining the required data and building an analytic infrastructure that manages and analyzes the data; analytic modeling, which builds the models; analytic operations that deploy models into operational systems; and measures that capture the operational impact of analytics against the strategic goals identified. Analytics are usually integrated into analytic workflows and adversaries can be expected to attack the weakest point in these workflows to obtain an advantage.

It is standard in many organizations today to develop a sensor strategy, a data strategy, a cloud strategy, and an AI strategy. But without an organizing principle like analytic superiority to tie these together, they are not likely to provide the advantages needed in competition, crisis and conflict.♥

DISCLAIMER

The views expressed are those of the authors; in particular, they do not reflect the official position of any U.S. Government Agency.

NOTES

1. Yoshua Bengio, Yann Lecun, and Geoffrey Hinton, "Deep Learning for AI," (2021). *Communications of the ACM* 64, no. 7: 58-65.
2. World Travel & Tourism Council, *Artificial Intelligence: Global Strategies, Policies and Regulatory Landscape* (April 17, 2024). <https://researchhub.wttc.org/product/responsible-artificial-intelligence-ai>, 4.
3. Algorithms are instructions to solve problems, perform computations, or process data that are precise enough that they can be implemented with a computer.
4. We can think of machine learning or statistical modeling as a process that takes data as an input and produces a model as an output (this is called a training model"). There are several different types of models, but some common types are models that make predictions, models that summarize data, and models that correlate one variable with another. If you fix a model, inferring, which is also called scoring, is the process that given input data produces an output. For example, in a large language model, the input is the prompt, and the output is the response. In a predictive model, the input is a feature vector, and the output is the prediction (yes/no, a numerical prediction, etc.) Rules are a very basic type of model that, in the simplest case, consists of multiple if...then...else statements, which may be nested.
5. A neural network is a type of predictive model that processes data in layers, with each layer a very simplified mathematical model of a neuron. Deep learning is a type of neural network that uses many sequential layers for processing data. Large language models are a particular type of deep learning models that take prompts asking questions as inputs and produce text responses as outputs. Generative AI are models that take prompts as inputs and can produce pictures, videos, and other objects as outputs.
6. Yash Sherry and Neil C. Thompson, "How Fast do Algorithms Improve?," (2021). Sherry, *MIT Initiative on the Digital Economy* Vol. 6. https://ide.mit.edu/wp-content/uploads/2021/09/How_Fast_Do_Algorithms_Improve.pdf, 2-4.
7. Benjamin Jensen, Christopher Whyte, and Scott Cuomo, "The Future of Algorithmic Warfare: Fragmented Development," (July 20, 2023) *War on The Rocks*, <https://warontherocks.com/2023/07/the-future-of-algorithmic-warfare-fragmented-development/>.
8. Timothy D. Haugh, *Summit on Modern Conflict and Emerging Threats*, (April 17, 2024). Remarks, Vanderbilt University, Nashville, TN.
9. Thomas H. Davenport, "Competing on Analytics," (2006). *Harvard Business Review* 84, no. 1: 98. <https://hbr.org/2006/01/competing-on-analytics>.
10. "Global Network Card Losses Worldwide," (May 2024) in *The Nilson Report* No. 1264, Credit Card Fraud, <https://nilson-report.com/the-current-issue/>, 5.
11. Michael Lewis, *Flash Boys: A Wall Street Revolt* (2014). New York: WW Norton & Company.
12. Robert L. Grossman, "A Framework for Evaluating the Analytic Maturity of an Organization," (2018) *International Journal of Information Management* 38, no. 1, <https://www.sciencedirect.com/science/article/pii/S0268401217300026>: 45-51.
13. Ruth A. David and Paul Nielsen, *Defense Science Board Summer Study on Autonomy* (Defense Science Board: 2016). Mark Maybury and James Carlini, *Defense Science Board Report on Counter Autonomy* (Defense Science Board: 2020), https://dsb.cto.mil/reports/2020s/CA_ExecutiveSummary.pdf. A. Etzioni, "Pros and Cons of Autonomous Weapons Systems," (with Oren Etzioni), in A. Etzioni, ed., *Happiness is the Wrong Metric: A Liberal Communitarian Response to Populism* (Springer International Publishing: 2018), pp. 253-263, https://doi.org/10.1007/978-3-319-69623-2_16.
14. W.C. Lin and C.F. Tsai, "Missing value imputation: A review and analysis of the literature (2006–2017)," (2020) *Artificial Intelligence Review*, 53:2, <https://doi.org/10.1007/s10462-019-09709-4>, 1487–1509.
15. V. Hodge and J. Austin, "A Survey of Outlier Detection Methodologies," (2024). *Artificial Intelligence Review* 22:2, <https://doi.org/10.1023/B:AIRE.0000045502.10941.a9>, 85-126.
16. J. Boardman and B. Sauser, "System of Systems—The Meaning of Of," (2006) *2006 IEEE/SMC International Conference on System of Systems Engineering*, 118-123. <https://doi.org/10.1109/SYSOSE.2006.1652284>, 6.
17. Robert L. Grossman, *Developing an AI Strategy: A Primer* (2020). https://analyticstrategy.com/wp-content/uploads/2020/03/Analytic_Strategy_Primer-TOC.pdf
18. Allen E. Brown and Gerald G. Grant, "Framing the frameworks: A Review of IT Governance Research," (2005). *Communications of the Association for Information Systems*, Volume 15, AIS eLibrary, <https://aisel.aisnet.org/cgi/viewcontent.cgi?article=3160&context=cais>, 696-712.

NOTES

19. Grossman (2018).
20. Watts S. Humphrey, *Managing the Software Process*, (1989). Reading, MA: Addison-Wesley, 89.
21. Grossman (2018).
22. U.S. Department of Defense, *Summary 2023 Cyber Strategy* (September 2023), p. 2, https://media.defense.gov/2023/Sep/12/2003299076/-1/-1/1/2023_DOD_Cyber_Strategy_Summary.PDF.
23. U.S. Cyber Command, *Achieve and Maintain Cyberspace Superiority: Command Vision for U.S. Cyber Command* (2018), <https://www.cybercom.mil/Portals/56/Documents/USCYBERCOM%20Vision%20April%202018.pdf>.
24. Michael P. Fischerkeller, Emily O. Goldman, and Richard J. Harknett, *Cyber Persistence Theory: Redefining National Security in Cyberspace* (2022). Oxford University Press.
25. Julian E. Barnes, "China Could Threaten Critical Infrastructure in a Conflict, NSA Chief Says," (17 April 2024). *The New York Times*, <https://www.nytimes.com/2024/04/17/us/politics/china-cyber-us-infrastructure.html>.
26. Emily O. Goldman and Michael Warner, "The Military Instrument in Cyber Strategy," (2021). *SAIS Review of International Affairs*, vol. 41, no. 2. Johns Hopkins University Press, <https://muse.jhu.edu/article/852326>, 51-60.
27. Martin Matishak, "Cyber Command, NSA usher in Haugh as new chief," (22 February 2024). *The Record*, <https://therecord.media/cyber-command-nsa-usher-in-haugh-as-new-chief>.
28. Fischerkeller, Goldman, and Harknett (2022).
29. Joint Cybersecurity Advisory (February 7, 2024), <https://media.defense.gov/2024/Feb/07/2003389935/-1/-1/0/CSA-PRC-COMPROMISE-US-CRITICAL-INFRASTRUCTURE.PDF>.
30. Joint Cybersecurity Advisory (February 7, 2024).
31. U.S. Department of Defense Defense Science Board, *Cyber Security and Reliability in a Digital Cloud*, January 2013 retrieve from <https://dsb.cto.mil/reports/> on May 10, 2024. Odell, L., Wagner, R. and Weir, T., 2015. Department of Defense use of commercial cloud computing capabilities and services. IDA P-5287, Institute for Defense Analyses, Alexandria, Virginia.
32. R. N. Charette, "Why Software Fails [software failure]," (2005) *IEEE Spectrum*, 42(9): 42–49, <https://doi.org/10.1109/MSPEC.2005.1502528>
33. William LaPlante and Robert Wisnieff, *The Design and Acquisition of Software for Defense Systems* (2018), <https://apps.dtic.mil/sti/citations/AD1048883>.
34. Grossman (2020).
35. S. Lingel, J. Hagen, E. Hastings, M. Lee, M. Sargent, M. Walsh, L. A. Zhang, and D. Blancett, D. *Joint All-Domain Command and Control for Modern Warfare* (Santa Monica, RAND Corporation, 2020) https://community.apan.org/cfs-file/__key/docpreview-s/00-00-17-19-55/8156.RAND_5F00_RR4408z1.pdf.
36. Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and The New World Order*, (2018). Boston: Houghton Mifflin, 3.
37. Robert O. Work and Greg Grant. "Beating the Americans at their Own Game: An Offset Strategy with Chinese Characteristics" (2019). *Center for New American Security*, <https://www.cnas.org/publications/reports/beating-the-americans-at-their-own-game>: 195-260.
38. Lee (2018).
39. United States Department of Justice, "Two Chinese Hackers Associated with the Ministry of State Security Charged with Global Computer Intrusion Campaigns Targeting Intellectual Property and Confidential Business Information," (December 20, 2018), <https://www.justice.gov/opa/pr/two-chinese-hackers-associated-ministry-state-security-charged-global-computer-intrusion>.
40. Rogier Creemers, "China's Social Credit System: An Evolving Practice of Control," (May 9, 2018). Leiden Institute for Area Studies, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3175792, 3.

Emerging Technologies for Data Security in Zero Trust Environments

Patrick Davis

Maj. Sean Coffey

Dr. Lubjana Beshaj

Lt. Col. Nathaniel D. Bastian, Ph.D.

ABSTRACT

Zero Trust (ZT), simply defined, is an information security framework which monitors and protects users, assets, resources, and data on a network by positively verifying all activity and never trusting anything by default. With the push to implement ZT across the public and private sectors, this transition between cybersecurity paradigms must be accomplished in a manner that is robust and enduring. This article examines emerging technologies most likely to impart the largest impact on ZT architectures (ZTAs), so that we better anticipate the pluses and minuses that will accompany those technologies. The discussion here focuses on data security, and the potential of each technology to affect security and protection across the lifecycle of data as it is generated, collected, transmitted, utilized, and stored. Technologies appraised include differential privacy, confidential computing, homomorphic encryption, quantum technology, biological technology, blockchain, and alternative computing methods.

INTRODUCTION

The timing of this assessment is critical, as the Department of Defense (DoD) and the broader Federal Government have mandates to implement a baseline ZTA by 2027. Planning and implementing ZTA is further complicated by the accelerating pace of technological change in today's operating environment. Within the

This is a work of the U.S. Government and is not subject to copyright protection in the United States. Foreign copyrights may apply.



Patrick Davis is a Research Scientist in the Data & Decision Sciences Division at the Army Cyber Institute within the Department of Electrical Engineering and Computer Science at the United States Military Academy (USMA) at West Point. He previously served as a Cavalry officer in the U.S. Army after graduating from USMA with honors with a degree in Electrical Engineering. He holds an MBA from The Wharton School and a Masters in Strategic Design from Parsons. His research and professional interests include data, decision sciences, applied computation, and innovation within strategic and leadership contexts of large enterprises and government.

past decade, in no small part due to the COVID-19 pandemic, both public and private organizations have embraced policies such as Bring Your Own Device (BYOD) and have moved computing workloads and data storage into commercial cloud-based services away from more costly, less capable, and less agile on-premises legacy architectures. At the same time, the volume, variety, and velocity of data has increased dramatically. The computational capacity and data availability driving the growth in Artificial Intelligence and Machine Learning (AI/ML) capabilities did not exist a few years ago, and those resources are expected to grow in scale for the foreseeable future. This creates new threat vectors and introduces new requirements for data security. Many enterprise-level security strategies struggle to keep up, especially in less data intensive industries. Given the pace of change and the nature of how technology shifts, prior approaches to data security using old technology will not hold up.

Data is critical to all network-reliant systems and operations. Adversarial actors are known to be operating domestically, internationally, and even within DoD entities, so the importance of data security cannot be overstated. DoD's 2023 Cyber Strategy details how the U.S. is continuously challenged by malicious cyber actors from the People's Republic of China, Russia, North Korea, Iran, violent extremist organizations, and transnational organizations; each pose threats to destabilize our democratic systems.¹ With the introduction of practical AI and more advanced emerging technologies for storing and processing data, we are now charged with countering these threats in an environment experiencing an explosion of change. The transition towards ZT will be a constantly moving target and the demand for a highly trained workforce will increase. Educating, recruiting, efficiently employing, and retaining talent who understand and can work with these emerging technologies is a matter of National Security.



Major Sean Coffey is a U.S. Army Cyber Warfare Officer trained in Cyberspace and Electromagnetic Activities (CEMA). Maj. Coffey currently serves as a Research Scientist and Machine Learning Engineer in the Data & Decision Sciences Division at the Army Cyber Institute within the Department of Electrical Engineering and Computer Science at the United States Military Academy (USMA) at West Point. Maj. Coffey earned an M.S. in Computer Science from the University of Minnesota and an M.S. in Data Analysis from the Georgia Institute of Technology. His research and engineering interests are in the areas of data science, artificial intelligence, machine learning, and scientific computing applied to CEMA, battlefield command and control systems, and privacy-enhancing emerging technologies.

Done well, ZTA reduces threat vectors and establishes a sustainable and adaptable framework that is forward compatible with future technologies. However, this modularity will not happen on its own. As such, this paper focuses on emerging technologies that are both critical to the future of ZT, but also require specific accommodations for a ZTA to truly actuate their potential benefits. In this paper, we provide an overview of ZT and discuss the relevant aspects of data security. Then, we introduce and discuss a series of emerging technologies that, as they mature and become more widely adopted, are postured to play a key role in advancing the Nation's ZT security posture and getting ahead of our adversaries. We conclude with a set of strategic recommendations for approaching these emerging technologies in terms of research, development, and innovation to better meet the needs of our future ZT environment.

ZERO TRUST

Since late 2018, National Institute of Standards and Technology (NIST) and National Cyber Center of Excellence researchers have worked closely with the Federal Chief Information Officer Council, federal agencies, and industry to address the challenges and opportunities for implementing ZTA across U.S. government networks. This resulted in publication of NIST Special Publication 800-207, which the DoD adopted for their definition of ZTA.²

ZT is the term for an “evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources.” At its core, ZT grants no implicit trust to assets or users based solely on their physical or network location or device ownership.³ This shift in philosophy is a significant change in legacy perimeter-based authentication and security mechanisms. It also represents a major cultural change that stakeholders throughout DoD's ZT Ecosystem, including the



Lubjana Beshaj, Ph.D., is a Senior Research Scientist in the Data & Decision Sciences Division at the Army Cyber Institute within the Department of Electrical Engineering and Computer Science at the United States Military Academy (USMA) at West Point. She also serves on the faculty in mathematical sciences at USMA. Dr. Beshaj earned her Ph.D. in Mathematics from Oakland University. Her research interests include both cryptography and emerging technologies, specifically in algebraic curve cryptography, post quantum cryptography, homomorphic encryption, quantum stabilizer codes, blockchain technologies, and neural network cryptography.

Defense Industrial Base (DIB), will need to embrace and execute beginning with FY2023 through FY2027 and in the future.

ZT entails a collection of information security design principles intended to replace previous perimeter-based security design principles. A ZTA is a network that adheres to ZT principles to secure its systems, data, and processes. The ZT approach is built on the realization that well-defined security perimeters can no longer be relied on to protect assets and resources, as there is no single point in the security system that is robust and capable enough that it cannot be circumvented. A ZT approach assumes the network has already been compromised and that threat actors are operating and active throughout the network. As such, ZT does not automatically trust actors, systems, or services operating within a security perimeter. Instead, rigorous analysis is conducted, and strict compliance to enterprise policies are verified before, during, and after granting access to any enterprise resource. ZT focuses on protecting critical assets, particularly data, at a granular level.

DoD's ZT Strategy is shown in Figure 1. The strategy has four top-level goals that are each supported by objectives in the areas of cultural adoption, securing

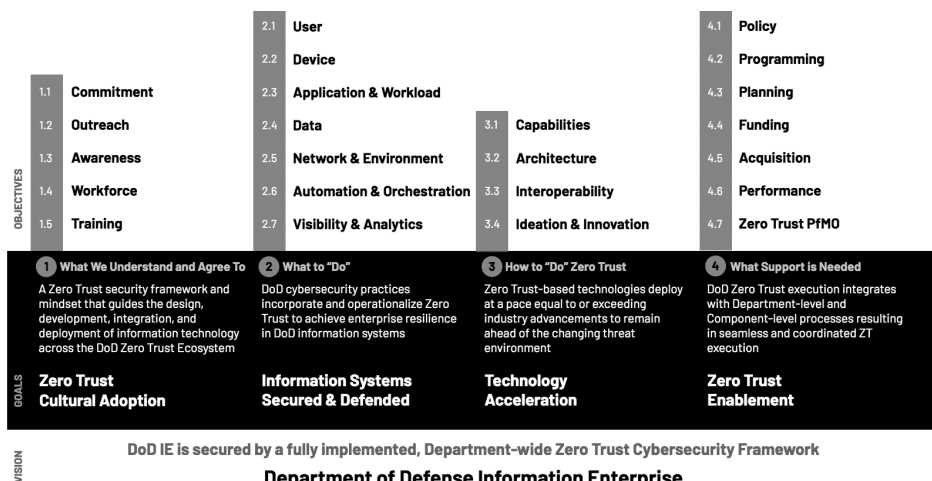


Figure 1. DoD ZT Strategy At-a-Glance



Lieutenant Colonel Nathaniel D. Bastian, Ph.D., is an Academy Professor and Cyber Warfare Officer at the United States Military Academy (USMA) at West Point. He is also concurrently serving as a Program Manager at the Defense Advanced Research Projects Agency (DARPA). At USMA, Lt. Col. Bastian serves as Division Chief, Data & Decision Sciences, Senior Research Scientist, and Chief Data Scientist at the Army Cyber Institute within the Department of Electrical Engineering and Computer Science. Lt. Col. Bastian earned a M.S. in Econometrics and Operations Research from Maastricht University and a M.Eng. in Industrial Engineering and Ph.D. in Industrial Engineering and Operations Research from Pennsylvania State University. His research interests aim to develop innovative, assured, intelligent, human-aware, data-centric, and decision-driven capabilities for cyberspace operations.

and defending information systems, technology acceleration, and execution enablement.⁴ A common theme across all guidance is that data plays a central role.⁵ In the following section, we will discuss key tenets of data security before turning to a discussion on emerging technologies that will increase our requirements for advancing and hardening data security within future ZT environments.

DATA SECURITY

Data is the lifeblood of the highly complex, interconnected socio-technical systems of societies, economies, and governments. It is the focal point of computation and conveys tactically and strategically valuable information about the behavior and intentions of individuals and organizations, and for that reason it is often targeted by adversarial actors. As such, the methods and practices used to secure data throughout its lifecycle are the cornerstone of any cybersecurity architecture.

The totality of benefits provided by modern computing networks can be reduced to data; the ability to store data for later, send data to others, and use data efficiently. Within network security, these states of data employment are known as “Data at Rest”, “Data in Transit”, and “Data in Use.” ZT requires us to broaden the perspective of data’s lifecycle by also considering “Data Generation” and “Data Collection” as two additional states.

ZT methods and practices enable data security by interweaving data encryption and secure network communication protocols with an encompassing data access paradigm based on two principles:⁶

1. Least privilege is the principle that any given user or automated process on a ZT network is only able to access the absolute minimum amount of data needed to perform the task at hand.

2. Assume breach is the principle where an organization never trusts any request or resource by default. Instead, all such requests need to be verified based on all available metrics. All security controls are conducted as if there was a known, yet to be located, malicious actor inside the network.

The employment of these principles across the data lifecycle – from generation to the point of collection, to storage, to data transmission, and while in use – make up the core of a ZTA. Each stage of the data lifecycle is given a brief section below which discusses some of the nuances of how the two guiding principles apply, all of which are relevant to specific emerging technologies discussed later in this paper.

Data at Rest

Data is stored at rest in databases and in file systems located on any storage system, which can be anything from dedicated servers to end user devices. Data at Rest is any data that is not currently actively in motion or in use, and as such it can be considered the default state of data.⁷ The word “actively” is crucial to this definition. It is not enough to secure Data at Rest until it is about to be utilized, rather, Data at Rest needs to be secured up until the very moment when it changes to one of the other states, at which point the security procedures relevant to those states should be employed. Data at Rest can be extracted from systems and stolen when malicious actors obtain direct or indirect access to an organization’s data sources or technology environments.

Critical infrastructure data is particularly vulnerable. In 2021, cybersecurity authorities in the United States, Australia, and the United Kingdom observed an increase in sophisticated, high-impact ransomware incidents against critical infrastructure organizations globally.⁸ Authorities observed incidents involving ransomware against 14 of the 16 U.S. critical infrastructure sectors, including the DIB, emergency services, food and agriculture, government facilities, and IT sectors.⁹ As threat actors become more technologically sophisticated, their tactics and techniques also evolve. They opportunistically identify weaknesses in personal and organizational data security, seeking to create an advantage or profit from exploiting sensitive data, including Personally Identifiable Information (PII) and Personal Health Information (PHI), trade secrets, intellectual property, and other private information stored in various formats, in different contexts, and across multiple devices and networks.

The General Data Protection Regulation (GDPR), drafted and passed by the European Union, includes provisions that relate to Data at Rest. The regulation specifies that only the minimum amount of data necessary should be stored, and in the case of personally identifiable data, it should be stored only for as long as it is needed for its intended purpose.¹⁰ There is no federal equivalent to GDPR in the United States, although the California Consumer Privacy Act of 2018 includes privacy protections, such as the right to be forgotten, for consumers interacting with businesses that collect personal information.¹¹

Data in Transit

Data in Transit, as the name implies, describes data currently in transit from one point in a network to another. While in transit, data is extremely difficult to secure, as most of the network traffic takes place across hardware and networks with different owners. Often even high-level administrators do not know the exact route network traffic will take from one point to another, even when operating in a semi-controlled environment.^{12,13} This makes it particularly difficult to differentiate between malicious traffic, misrouted traffic, and normal traffic. Furthermore, these data signals are physically transmitted between devices using diverse types of transmission mediums such as twisted pair cable, coaxial cable, optical fiber, high-frequency radio waves (wireless), or lasers. Each of these mediums are uniquely vulnerable to distinct types of physical attacks, electromagnetic interference, and signal interception or tapping.

For data transmitted over the Internet, Hypertext Transfer Protocol Secure (HTTPS) is the most widely used security protocol. HTTPS is a combination of the standard HTTP protocol used for web communication and the Transport Layer Security (TLS) encryption protocol. When a user connects to a website using HTTPS, their browser initiates a TLS handshake with the web server.¹⁴ During the TLS handshake, the server and client agree on a set of encryption and decryption algorithms to use for the session. Once the handshake is complete, all data transmitted between the server and client is encrypted using these agreed-upon algorithms, ensuring that unauthorized parties cannot intercept or tamper with the data. The security of HTTPS depends on the strength of the encryption and decryption algorithms used and the key management techniques employed. As such, it is crucial to use approved cryptography and regularly update the encryption algorithms and key management techniques to ensure that HTTPS remains secure against potential threats and attacks.

Peer-to-peer (P2P) security protocols are essential components of secure data transmission in cloud computing and on blockchains. P2P protocols enable secure data transmission between two or more devices, preventing unauthorized access and ensuring data confidentiality, integrity, and availability. TLS, for example, is widely used to secure web communications used in online banking, e-commerce, and email. Updates to TLS are managed by the Internet Engineering Task Force, and the most current version 1.3 was approved in 2018. Developers should use at least TLS version 1.2 or later. Secure Shell is another technique used to establish secure connections between devices for remote login and data transfer. Internet Protocol Security is used to secure Internet Protocol (IP) communication and is commonly used in enterprise networks, including virtual private networks.

Applying the second principle of ZT to Data in Transit, assume breach, it is assumed that network traffic traversing network midpoints and over different mediums is compromised. From this we arrive at the requirements for Data in Transit within a ZT security model. First, data must always be end-to-end encrypted. In addition, Digital Rights Management (DRM)

policies, which are also implemented to secure Data at Rest, should map back to relevant security policies, and ensure that only the minimum relevant data is transmitted at any given time. Finally, it is necessary to log, track, and filter network traffic to enable analytical detection of any anomalous behavior.¹⁵ Because of the sheer volume of network traffic that needs to be analyzed, organizations should deploy advanced analytical methods from AI/ML to increase the efficiency and efficacy of finding the proverbial needle in the haystack.

Data in Use

Modern digital computing largely relies on three distinct types of processing units. Central Processing Units (CPUs) are responsible for the overall performance of most computers and devices. The CPU executes instructions of programs and can perform basic arithmetical and logical operations on data inputs to generate the desired output. Sometimes CPUs include several cores, which can process data in parallel, thereby increasing capability. However, CPUs are not ideally suited to execute advanced AI/ML algorithms in a reasonable amount of time. Both Graphics Processing Units (GPUs) and Tensor Processing Units (TPUs), working in combination with CPUs, provide magnitudes more processing capability and consume less energy than CPUs alone.¹⁶ Multi-core CPU, GPU, and TPU can be deployed on local computers or servers, or ephemerally in cloud environments.

Data being processed within machine memory can be vulnerable. Attackers with physical or logical access to equipment (e.g., printers, fax machines), network appliances (e.g., routers, firewalls), workstations (e.g., desktops, laptops, tablets), servers, mainframes, or personal devices (e.g., smartphones, smartwatches), might copy or alter information stored within the device memory while it is in use. The prevalence of personally owned devices used to process organizational data have increased due to industry trends for BYOD policies where corporate and personal data is often processed on the same smartphone or laptop. Software can be reverse engineered, data can be copied, altered, deleted, or added, sensitive information such as PII or PHI might be exposed, and critical security parameters might be compromised allowing identity theft or fraud.

Additionally, the growing need for data-driven techniques and methodologies that require copious amounts of processing hardware has led to many entities outsourcing their computing needs to third parties in the form of cloud computing. This service provides the needed computation, but significantly increases the risk of sensitive data exposure, as a customer must now rely on the cloud service provider to employ adequate security techniques with limited ability to verify it for themselves.¹⁷ The data itself might not even be the most significant risk for exposure, as exposure of cloud computing techniques might reveal proprietary algorithms or techniques, negating any competitive advantage in that domain.

Data Generation and Collection

The security implications of data generation and collection should also be considered when exploring the impact of emerging technologies on ZT. Data cannot be stored, transmitted, or used until it exists, whether collected from humans via human-computer interfaces or generated by devices, sensors, or systems autonomously. Data takes many different forms including text, images, audio, video, location coordinates, or application use telemetry. For data to be considered valuable, it does not have to be structured or formatted in any set way. Nor do the originators or subjects of the data collection necessarily know that collection is taking place. Different forms of AI/ML, such as large pre-trained models, can generate extremely realistic synthetic data, and interpret, summarize, and find patterns in data regardless of its format.

Modern data technology increases the value of data, but also introduces new risks. Once data exists, its pervasive quasi-controlled and uninformed use is a serious threat to individuals and organizations. In fact, there exists a legal multi-billion dollar data brokerage ecosystem that consists of companies that collect and generate data that they then sell, license, and share with other companies who use it for their own purposes or to provide technical services. This often includes sensitive personal data, and studies have shown some data brokers do very little due diligence on customers they sell their data to.¹⁸

Historically, data collected via web-based applications introduced vulnerabilities in which threat actors could inject malicious code or scripts directly into system data storage or even directly into the memory buffer. Over time, common web development frameworks incorporated security measures such as input masking and validation to prevent these types of risks. While direct attacks can be mitigated, these measures do little to actively protect users against social engineering attacks, such as phishing, that manipulate users into disclosing sensitive information including passwords, pins, usernames. This highlights why users should be aware of their system access privileges and their responsibilities in relation to accessing, entering, and securing their organization's sensitive data.

With this foundation of data security established, we will now move on to the emerging technologies anticipated to have the largest potential impact on a ZT future.

EMERGING TECHNOLOGIES

ZT changes the data security paradigm to be more compatible with future technologies and methodologies. In the ZT Reference Architecture, the Defense Information Systems Agency and National Security Agency ZT Engineering Team view emerging technologies as technical opportunities. They state that ZT “is an evolution of technology and operational approaches which over time evolve with the threat environment it is seeking to ameliorate.”¹⁹ The following sections delve into the emerging technologies most likely to play a role in ZT as they evolve.

Privacy-Enhancing Technologies

A Privacy-Enhancing Technology (PET) is one that embodies fundamental data protection principles by minimizing personal data use, maximizing data security, and empowering individuals. PETs allow online users to protect the privacy of their PII, which is often provided to and handled by services or applications. PETs use techniques to minimize an information system's possession of personal data without losing functionality.²⁰

Differential Privacy

Differential privacy is a mathematically defined series of techniques and methods intended to allow data to be utilized and shared while reducing or eliminating the risk of sensitive exposure. Differential privacy was conceived in 2006 and is now widely used in both the public and private sectors. Companies including Microsoft, Apple, Google, Uber, Facebook, Amazon, LinkedIn, and Snapchat have embedded it into their products and continue to invest in ongoing research and development.²¹ Most notably, from a Public Sector standpoint, the U.S. Census Bureau executed one of the largest-scale implementations of differential privacy for the first time during the 2020 Decennial Census.²²

From a conceptual perspective, differential privacy adds noise to datasets to reduce the likelihood that any person or entity whose information is present within a dataset could be identified. Differential privacy also protects individuals or entities from being identified as participants in a particular data collection process.²³ This second aspect is important, as it ensures that participants in the dataset cannot be identified by simply excluding everyone who did not participate.

Differential privacy is a powerful tool that reduces risk inherent in collecting sensitive data, such as is required for ZT. Implementing and operating a ZTA relies on capturing and analyzing large amounts of network information including network traffic flow, user behavior statistics, and other data to identify suspicious activity and to adjudicate requests for system access. Collecting this data creates a vulnerability as it must be accessed, processed, and shared regularly by both users and systems. Employing differential privacy reduces both the risk of exposing sensitive information and the value of that information to malicious entities.

Confidential Computing

Confidential computing loosely defines a range of techniques that enhance the security and privacy of Data in Use and is a recognized way forward to implement ZT DoD-wide. While not an emerging technology, confidential computing should be applied to hardware designed and manufactured for sensitive computing use cases, and is a powerful framework within a ZTA.²⁴ Confidential computing seeks to reduce data vulnerability before, during, and after it is processed, and is achieved by establishing hardware-based trusted execution environments that are secured using embedded encryption keys, and embedded identity

attestation mechanisms to ensure keys are accessible solely to authorized application code. It also enhances data protection in the cloud, especially as organizations move more sensitive data and computing workloads to public cloud services.

Although somewhat new to market capability, it is fully fielded and available from several private sector vendors.²⁵ Many consider confidential computing as an overarching term that defines all Data in Use protection techniques, to include homomorphic encryption.²⁶ However, the original term, coined by Intel in 2015, referred to a hardware level encryption technique, in which modern CPUs establish physically and/or logically secure enclaves during computation run-time. This secure enclave is logically removed from other operations within the CPU, and, for specific implementations by companies like Intel and Advanced Micro Devices (AMD), it also is physically separated from the rest of the CPU, although still on the same chip.²⁷

Confidential computing has some key drawbacks, such as the absence of any mechanism to confirm data security throughout the entire process without continuously monitoring all incoming, outgoing, and internal data operations, which would be highly unfeasible, especially when relying on third party cloud computing resources. This drawback is flagged here, because vulnerabilities have been identified in confidential computing enclaves in the past.²⁸

Software Guard Extension (SGX) is Intel's first iteration of a CPU with a secure hardware enclave. In 2017, researchers found a serious potential vulnerability in its implementation when they were able to extract RSA keys from within the enclave, and the enclave itself actually hid the presence of the malware, increasing the severity of the issue.²⁹ This led to Intel deprecating support for SGX for a number of years, until its competitor, AMD, partnered with Amazon to launch a competing confidential computing cloud service, which motivated Intel to bring the product line back.

The second major drawback of confidential computing is that implementations are hardware specific and rely on how the CPU is physically constructed.³⁰ When Intel abruptly suspended its SGX series CPUs, highly significant research to develop platform specific SGX implementations was ongoing.³¹ These and other drawbacks notwithstanding, confidential computing remains a largely reliable and viable framework. The cloud based confidential computing services market is expected to multiply fourfold over the next five years, so it is highly likely all known drawbacks will be mitigated as the underlying technology and systems evolve.³²

Homomorphic Encryption

Homomorphic encryption schemes allow computations on encrypted data without needing the secret decryption key. Fully homomorphic encryption maintains the privacy of search engine queries and enables the searching and editing of data and information that stays encrypted.³³ In AI/ML, fully homomorphic encryption maintains the security of not only

model input prompts and outputs, but also details of the model itself and the contents of its potentially sensitive training data.

Homomorphic encryption was first envisioned by Rivest, Adleman, and Dertouzos in 1978, well before a system with the required processing power could be built, and decades before Craig Gentry found the first practical implementation of a fully homomorphic system in 2009.³⁴ He did this using a lattice-based encryption scheme and a process called squashing and bootstrapping. Modern quantum resistant encryption schemes rely on the ring learning with errors problem which also uses highly-dimensional lattices.³⁵

Fully homomorphic encryption allows a company or individual to have its data encrypted while other known parties run computations on that data without being able to decrypt it.³⁶ Traditionally the other party would need a private key to decrypt the data, modify it, encrypt it, and then send it back. This poses a security concern, especially when it comes to sensitive data, because if the outsourced company is subject to a cyber-attack, the attacker will gain access to all the unencrypted sensitive data. However, if homomorphic encryption is used, even if the attacker gains access to the data, it is encrypted through the entire process, yielding no clues as to what the actual data is or how it is being used. Therefore, homomorphic encryption could make a significant impact in the world of cryptography and data security.

Blockchain Technology

Blockchain technology was introduced in 2008 with the invention of Bitcoin by Satoshi Nakamoto.³⁷ Blockchain has seen significant evolution in recent years since its introduction. Bitcoin's market capitalization, the largest among all cryptocurrencies, currently exceeds \$1 trillion and it continues to be the global standard bearer for digital currency. The second largest blockchain by market capitalization, Ethereum, is worth over \$400 billion.³⁸ Ethereum has established itself as a versatile platform for various applications beyond its native cryptocurrency, ETH. Ethereum's blockchain facilitates smart contracts, decentralized applications, and decentralized autonomous organizations, showcasing its multifaceted utility. Today, there are over 200 blockchains, each with their own functional utility, with total market values over \$1 million.³⁹ The proliferation of non-fungible tokens and other digital assets further highlights blockchain's expanding influence and utility.

Recent legislative efforts in Congress to begin regulating digital assets underscore the growing recognition and integration of blockchain technologies into mainstream financial and technological ecosystems.⁴⁰ This rapid evolution and increasing regulatory attention underscore blockchain's trajectory towards widespread adoption and its potential to revolutionize various industries.

A blockchain is an ever-growing, secure, shared, record-keeping system in which each user of the data holds a copy of the records, which can only be updated if all parties involved in a transaction agree to update. It is a P2P distributed ledger that is cryptographically

secure, append-only, immutable, and updateable only via consensus or agreement among peers. Blockchain can be thought of as a layer of a distributed P2P network running on top of the internet.⁴¹

Blockchain uses cryptographic techniques and operates over a decentralized network of nodes, where each node maintains a copy of the entire blockchain. Blockchain is considered to be a novel approach to distributed computing and data storage. In the context of computing, blockchain technology enables the management and storage of data across multiple systems, ensuring data integrity, transparency, and security. Blockchains are known as distributed networks because the data and control over the network are not centralized but are instead shared among participants. Each participant (node) in a blockchain network typically maintains its own copy of the entire ledger, and updates to the ledger are achieved through a consensus mechanism among these nodes. The consensus protocols in blockchain networks are crucial as they ensure all nodes in the network agree on the validity of transactions. The most common consensus protocols used in blockchain networks are Proof-of-Work (PoW) and Proof-of-Stake.

Data stored on blockchains must be verified by nodes on the network before entering the chain. Once a certain number of nodes agree on what data can be added, each by solving and sharing results of the same computationally complex consensus algorithm (in the case of PoW), the block containing the data is added to the chain. Data stored on blockchains are immutable and transactions are fully traceable. In contrast to traditional data networks, the robust cross-validation of blocks entering a blockchain network provides extra levels of security that cannot be found elsewhere. The computational power and associated costs required for each node to execute the consensus algorithm serves as an economic disincentive for attackers who might otherwise try to execute or confirm a nefarious transaction by tricking the requisite number of nodes (usually many) into forming a consensus at the same time for the same invalid blocks. Attacks like this are prohibitively expensive due to the intentional design of the network and are not guaranteed to succeed, even if attempted at great cost.

Quantum Technology

Quantum technology is an emerging field of physics and engineering that encompasses technologies that rely on the properties of quantum mechanics.⁴² For the purposes of this analysis, we will address the three primary applications of quantum computing, quantum communications, and quantum sensing.

Quantum Computing

Quantum computing leverages quantum mechanical principles of entanglement and superposition to create quantum computers. A quantum bit, or qubit, can exist in an infinite number of states simultaneously and can also resolve to an identifiable binary state. Unlike the extremely reliable bits that underpin conventional computing, qubits are highly prone to

error and only operate in specific physical environments. Power-expensive error correction mechanisms are needed. Qubits tend to be prone to several types of environmental interference, so they can only be transmitted over limited distances and require special sensors to measure their state. Despite the challenges, there is rapid progress happening in quantum computing, in no small part a result of DoD-originated funding and research.⁴³

Because qubits can be built with a variety of technologies, including superconducting circuitry, trapped ions, and photons, there are several variations of quantum computing currently being developed and researched.⁴⁴ Circuit-based, the most conventional or mainstream type, is the focus of investment by many major technology companies. Circuit-based quantum computers can be based on transmons (the focus of Google's and IBM's efforts) or based on trapped ions (employed by IonQ). Topological quantum computing employs exotic quasiparticles called anyons that are naturally resistant to errors and perturbations.⁴⁵ This fault-tolerant approach is being developed by Microsoft but is still in the early stages of physical realization of even one qubit. Another type of quantum computing known as measurement-based, or one-way, quantum computing exists, but is still in its conceptual stage.

Quantum annealing is a non-general approach to quantum computing that applies specifically to solving combinatorial optimization problems, where the goal is to find the optimal combination or arrangement of elements from a large set. The energy usage of the system corresponds to the objective function of the optimization problem, and the annealing process efficiently finds the lowest energy configuration that is also the optimal solution. This combined hardware- and software-based approach is being developed by D-Wave, who make instances of their computing solution available in the cloud. This shows tremendous promise, particularly for efficiently solving previously intractable optimization problems encountered in AI/ML use cases.⁴⁶ As quantum computers advance, their capabilities will further accelerate analytical progress and supercharge use cases across all industries.

Quantum Communications

Communications channels secured by quantum technology differ from classical information exchange in that the two parties know when a third party has attacked. This is because measuring an unknown quantum state changes it. If a third party eavesdrops on an exchange by trying to measure the key, this creates detectable anomalies in the quantum state. In the 1970s, well before his time, Stephen Wiesner introduced the concept of quantum conjugate coding—based on a method of transmitting multiple messages such that reading one destroys the others.⁴⁷ In 1984, his theory was used as a base for Bennet and Brassard to propose a method for secure communication called BB84. The BB84 scheme is at the basis of quantum key distribution (QKD) methods which are based on the idea of one-time pad (OTP). OTPs are in theory the strongest possible algorithmic cipher: if the key is used properly, they cannot be broken, even in theory.⁴⁸

Post Quantum Cryptography

If large-scale quantum computers are ever built, they will be able to break many of the public-key cryptosystems currently in use. This would seriously compromise the confidentiality and integrity of digital communications on the Internet and elsewhere. The goal of post-quantum cryptography (also called quantum-resistant cryptography) is to develop cryptographic systems that are secure against both quantum and classical computers and with existing communications protocols and networks.

Since 2016, NIST publicly acknowledged that when large-scale quantum computers are realized, they will be able to break many of the public-key cryptosystems currently used in digital communications and for encrypting data. To counter this future vulnerability, NIST initiated a process to solicit, evaluate, and standardize quantum-resistant public-key cryptographic algorithms.⁴⁹ Following a November 2017 deadline for submission, NIST has continued to lead a public-facing process to narrow down the scope of possible candidates, and there are currently three algorithms still under consideration: CRYSTALS-Dilithium, CRYSTALS-KYBER, and SPHINCS+. The intent of NIST's effort is to specify one or more unclassified, publicly disclosed digital signatures, public-key encryption, and key-establishment algorithms and make them available for use around the world.⁵⁰

Quantum Sensing

Quantum sensors and measurement devices provide accuracy, stability, and new capabilities that offer advantages for commercial, government, and scientific applications. Examples such as atomic clocks for GPS navigation and nuclear spin control for magnetic resonance imaging are already widely used, with transformative impacts for society. Further advances in quantum information science and technology will enable a new generation of similarly transformative sensors that will increase the volume, variety, and value of data that can be generated. While quantum sensors can potentially enhance decision-making in a military context, adversarial use of the same technologies against us has serious security implications. Researching how to work with, secure, and trust data generated by quantum sensors will help drive innovation and serve as a competitive advantage against adversaries.

Biological Technology

Biological technology, or biotechnology, is expected to have a major impact across various sectors, including health, energy, and national security. Data is crucial for advancing biotechnology innovations, and ZT is an appropriate security framework for protecting sensitive health information and proprietary research data within the field. In "Bold Goals for U.S. Biotechnology and Biomanufacturing: Harnessing Research and Development to Further Societal Goals," The White House Office of Science and Technology Policy highlights the pivotal role of data and the importance of data standards in the future of biotechnology:

Global bioeconomy strategies recognize that a sustainable, safe, and secure bioeconomy is built upon standards and the availability of high-quality data. Data, particularly from genomics and multiomics, underpins advances in biotechnology, for example, by enabling the rapid design of systems to produce needed medicines, food, and materials. Standards, particularly international data standards, can help accelerate research and development (R&D) and commercialization of new, safe, and effective medicines and therapeutics; contribute to food product safety, quality, and consistency; promote international trade; and instill confidence among consumers for products in all sectors of the economy. For industry, standards can promote research and manufacturing innovation, streamline regulatory review, and enable international alignment, interoperability, and coordination. As biotechnology converges with automation, connected devices, and AI, a robust data infrastructure can also accelerate R&D and commercialization of emerging technology.²⁷

DNA Computing

By 2025, it is estimated that the entire data universe will be 175 zettabytes.⁵³ Storage systems that now hold the world's data in the form of 0's and 1's will not be able to keep up, so finding ways to store and compute biological forms of data have become increasingly relevant problems. One approach is DNA-based data storage.

DNA is genetic material in all organic organisms and is composed of two polynucleotide chains that coil around each other in a double helix formation. Each polynucleotide is composed of smaller units consisting of one of four nitrogen containing nucleobases: cytosine (C), guanine (G), adenine (A), or thymine (T).⁵⁴ Researchers are harnessing these pattern-based strands to store data. By allowing each nucleobase to represent a two-element binary value (00=A, 01=T, 10=C, 11=G), DNA can be made to represent a table for use in bitwise arithmetic operations. Encoding DNA to represent data occurs through DNA synthesizing.⁵⁵ This recent research opens the cyber world up to the possibility of harnessing DNA as a medium for larger-scale computations and data storage. DNA can be sequenced and accurately copied and is quite stable. A single gram of DNA can store a single Zettabyte (one billion terabytes).

DNA also can be very useful to encrypt and decrypt data. Nucleotides, representing binary values, allow theoretical algorithms to be derived for encrypting and decrypting data. DNA cryptography is a proposed technique of encryption and decryption in which data can be hidden within a DNA sequence. This technique is very promising due to the powerful potential of DNA computing and data storage. Hiding data within DNA alone is insufficient to keep it secret, but applying known and new ciphers to data already sequenced into DNA shows great promise.⁵⁶

DoD has a vested interest in and actively contributes towards research and development to understand, harness, and manage the implications of advances in each of these emerging technologies on society, in warfare and for diplomacy. Data, and the security of data in these

new human-technological contexts, will continue to be an important topic of discussion, debate, investment, and innovation as the opportunities and vulnerabilities of our future technological environment become more clear. Now that we have covered the preliminaries of ZT, data security, and emerging technologies, we consider their future implications.

ANALYSIS AND DISCUSSION

In the evolving landscape of ZT, quantum technology, biotechnology, blockchain, and PETs (e.g. differential privacy, confidential computing, and homomorphic encryption), present both opportunities and challenges for data security. DoD's ZT Strategy emphasizes seven key data capabilities necessary for implementing ZT: data catalog risk alignment, DoD enterprise data governance, data labeling and tagging, data monitoring and sensing, data encryption and rights management, data loss prevention, and data access control.⁵⁷ These capabilities must be examined for all new data-driven applications of technology both through the lens of data security and across the entire data lifecycle, which encompasses Data Generation, Data Collection, Data at Rest, Data in Transit, and Data in Use.

Data security implications vary significantly across different applications of technology. In ZT environments, enforcing the principle of least privilege ensures that users, devices, and systems access only the necessary data for their specific purposes. When assuming breach, every interaction involving data must be authenticated and never trusted by default. This is particularly challenging when designing secure applications in tactical environments where devices and communications are subject to denial, disconnection, intermittency, or limited bandwidth (DDIL).⁵⁸ Devices and systems that comprise the Battlefield Internet of Things, although highly capable in some cases, are often resource constrained and can only execute their main functionality, making them less capable of defending themselves against attacks. Advances in quantum computing, applied in the AI/ML development process, can reduce the size and complexity of security models, enabling them to run in resource constrained devices in DDIL environments.

In tactical environments, Electronic Warfare (EW) weapons can disrupt data generation and collection by jamming sensors and communication devices, leading to corrupted or incomplete data. For Data at Rest, electromagnetic pulses or directed energy weapons could physically damage storage infrastructure, highlighting the need for physical security and redundancy. Data in Transit is particularly vulnerable to EW, where interception and jamming can delay or compromise the transmission of sensitive information. Quantum-safe encryption and secure communication protocols, including QKD as it progresses towards operational readiness, will be vital in safeguarding Data in Transit. While emerging technologies offer potentially promising solutions, they also introduce new layers of complexity which could be infeasible given operational constraints. Regardless, in designing secure devices and applications, particularly in sensitive and tactical environments, trade-offs are inevitable,

and the focus must remain on robust data access control and continuous authentication to ensure data security.

Security considerations should be proactively integrated into the earliest phases of innovation and engineering for any application or system that touches data throughout its lifecycle. This is consistent with “Secure by Design” guidance published by the Cybersecurity Infrastructure Security Agency and global partners, which empowers designers and engineers to consider security throughout the design process.⁵⁹ Neglecting security and privacy in incipient phases will inevitably result in an ineffective reactionary security posture. At the scale and speed that modern networks operate, and as adversaries become more sophisticated, this is untenable. Therefore, it is essential for designers, engineers, and security professionals to have a shared understanding of how PETs, for instance, may apply to a data security problem as well as their unique implementation considerations.

While technologically advanced adversaries can disrupt data generation and collection, and also generate, collect, and use data for their own intelligence purposes, it is often much easier to target Data at Rest and Data in Transit in more vulnerable legacy systems. Identifying and addressing these vulnerabilities is crucial, especially as adversaries find new ways to attack us and degrade our capabilities. Frameworks such as MITRE's ATT&CK and ATLAS are invaluable for understanding the threat. The ATT&CK knowledge base compiles adversary tactics and techniques from real-world observations and aids in developing threat models and methodologies across various sectors.⁶⁰ Similarly, MITRE ATLAS describes adversary tactics and techniques in AI systems based on observed attacks and realistic demonstrations by AI red teams.⁶¹ These frameworks show how adversaries target and exploit data, and should be leveraged in the early phases of design to define an application's data security requirements; they exemplify the benefits of collaboration across industry, academia, and government, and provide a common taxonomy for threat intelligence further enabling the testing and development of improved network and system defenses.

Advances in quantum computing will revolutionize critical analytical capabilities needed for bolstering defenses and achieving ZT. Quantum machines already have the ability to solve formerly unsolvable, highly-dimensional optimization problems by processing vast amounts of data at incredible speeds and exploring extremely large solution spaces near instantaneously. Furthermore, there are known data- and model-related challenges in classical AI/ML approaches that quantum compute solutions can potentially overcome which warrant further investigation and research. As DoD increases reliance on AI/ML for security monitoring, network intrusion detection, and other ZT related use cases, quantum compute is poised to further enhance data capabilities, increase efficiency and accuracy of these systems, save time, and enable deployment further to the edge by requiring fewer compute resources.

Quantum computing can also be used to create extraordinarily realistic datasets for training AI/ML models, particularly in cases where training datasets are insufficient or not available. Data labeling and tagging is now a mission critical capability. For many cybersecurity use cases, biotechnology initiatives in particular will rely on increasingly sensitive and large scale biological data (e.g., DNA and human genes) that necessitates precise labeling and tagging. AI/ML-enhanced automation can help improve accuracy and consistency of training datasets. PETs can also be leveraged when the labeling and tagging process involves sensitive data that could be either inadvertently exposed or purposefully exploited or manipulated by adversaries intending to attack or steal information about the underlying models, their intended capabilities, or the data used to train them.

Quantum sensing capabilities will enable the generation of precise and accurate data for sensing and measuring objects and phenomenon currently invisible to us and our systems. Examples of applications include more accurate and reliable geolocation, improved medical diagnostic imaging, safer navigation of autonomous vehicles, better guidance systems, and detection, imaging, and mapping of underground environments.⁶² Quantum sensing will also significantly accelerate our ability to integrate data from several disparate sources and sensors for near real time monitoring and analysis, even in noisy environments. These are particularly valuable data generating use cases for DoD, but quantum sensing is still a capability that few people know about and has yet to reach widespread adoption.⁶³ Although quantum sensing has fewer immediate implications on data security and ZT, the intellectual property associated with ongoing research is highly sensitive and needs to be protected in accordance with ZT principles.

For systems that require many users and devices to generate, collect, store, transmit, or use sensitive data, blockchain technology will play an increasingly key role in enterprise data governance. Blockchains provide, by design, secure and transparent methods for tracking data lineage and automating compliance with governance and security policies. Its decentralized and immutable ledger guarantees that all actions and modifications to data are logged and auditable. Custom private blockchain platforms can track access to data for labeling and tagging, monitoring and sensing, and data loss prevention in a ZT environment. In addition, the rules necessary for data catalog risk alignment, rights management, and access control can be programmatically defined, managed, and automatically executed on a blockchain platform.

As emerging technologies reshape the landscape of data processing and storage, the need for robust privacy-enhancing technologies will remain critical. Homomorphic encryption allows computations directly on encrypted data, which can be used to enhance DRM and for data loss prevention without compromising privacy. It can also be used in data labeling and tagging use cases where details of the models and data used to train them need to remain secret but could

become exposed (or derived) during their usage. Engineering challenges in implementing, assuring, and verifying fully homomorphic encryption schemes necessitate advanced methods for their design and testing before they can be broadly adopted within ZTAs.

Confidential computing secures data processing at the hardware level. It therefore plays an important role in securing environments where AI/ML is used for sensitive computation and strict policies need to be enforced for Data in Use. As cloud adoption grows and global competition for semiconductors intensifies, supply chain gaps and vulnerabilities in chip design and fabrication pose serious risks. Advanced verification methods are required to ensure that sophisticated chips are free from security flaws and function as intended, as such risks may only surface during future use. Furthermore, novel confidential computing techniques will be necessary to secure the unique hardware used in quantum and biotechnological computing use cases.

Differential privacy protects individual privacy of subjects included in data sets used for model training, analysis, statistics, or informational reporting products. This will become increasingly vital as progress in quantum technology and biotechnology heighten the sensitivity and implications of Data in Use, especially in contexts which require transparency. This is the case for securely sharing data for decision-making, R&D, or innovation, as well as meeting regulatory or compliance requirements. Standards for evaluating effectiveness and providing assurance for differential privacy and other PET implementations will be necessary as ZTAs become more mature and widespread.

ZT fundamentally strengthens data security by ensuring that trusted user and device access to sensitive data is closely monitored to identify and prevent malicious exploitation. Together, the technologies discussed not only align with, but also enhance DoD's ZT posture, creating a robust, resilient, and secure data environment that can withstand evolving cybersecurity threats introduced by new technologies and persistent adversaries.

CONCLUSION

Given the immense potential of emerging technologies and their transformative applications, ZT stands as the ideal framework to guide data security practices and decision-making. Interplay among these emerging technologies and PETs underscores a complex yet promising future for ZT environments. Quantum technology enhances data generation, computation, processing, and security; biotechnology introduces new forms of data computation, storage, and relies on sensitive data that requires rigorous protection; and blockchain provides a decentralized, tamper-evident framework for data integrity and secure transactions. At the same time, PETs ensure that data produced and consumed by modern data systems remains secure and private throughout its lifecycle.

Addressing ZT security challenges at the interface between humans and machines will require innovation that capitalizes on the capabilities of today's emerging technologies. Innovation ecosystems are already emerging to support quantum technology, biotechnology, and blockchain. These communities can contribute towards addressing society-wide challenges and opportunities related to data security and privacy, but doing so will require a common vision, coordination, collaboration, and research. Prioritized and synchronized R&D is important in strategic areas such as workforce education, innovation and engineering, partnership building, and scientific research. Sustainable innovation requires careful balance between risk and reward, managing resource constraints, and making cost-benefit trade-off decisions in an extremely complex environment. The technologies described are poised to revolutionize and optimize how we do all these things in the future.

ZT is a large and complex engineering and cultural undertaking. No data technology is out of scope, whether on the bleeding edge or embedded in legacy systems. DoD's mission to accelerate decision advantage and sustained focus on data security must drive these efforts in support of broader national security objectives. DoD is uniquely suited to address complex multi-faceted problems such as ZT, where real world consequences raise the stakes above profit-oriented objectives. DoD must continue leading and expanding efforts to partner with industry, allies, and academia to drive innovation based on science and evidence-based decision-making.

Even though these emerging technologies will be cross-cutting and broadly impactful, each use case is unique. Both the magic and the devil are in the details. This requires careful prioritization and resource allocation towards current R&D efforts and for the future maintenance, sustainability, and protection of systems and infrastructure that we are increasingly dependent upon. This reinforces the absolute necessity of nurturing and expanding a diversified pipeline of educated, experienced, and motivated talent who can work with these emerging technologies and help shape the future. Failing to plan for and anticipate our future needs poses a serious risk to our technological society.🛡️

DISCLAIMER

The views expressed here are of the authors alone. They do not reflect the policy or position of the U.S. Military Academy, U.S. Army, U.S. Department of Defense, or U.S. Government.

NOTES

1. U.S. Department of Defense, *Summary of the 2023 Cyber Strategy of The Department of Defense*, (September 12, 2023). https://media.defense.gov/2023/Sep/12/2003299076/-1/-1/1/2023_DOD_Cyber_Strategy_Summary.pdf, 1-2.
2. National Institute of Standards and Technology, “Implementing Zero Trust Architecture,” NIST National Cybersecurity Center of Excellence, accessed May 20, 2008, <https://www.nccoe.nist.gov/projects/implementing-zero-trust-architecture>.
3. Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly, “NIST Special Publication 800-207: Zero Trust Architecture,” (August 2020). *Computer Security*, <https://doi.org/10.6028/NIST.SP.800-207>, ii.
4. U.S. Department of Defense (DoD) Chief Information Office (CIO), Zero Trust Portfolio Management Office, *DoD Zero Trust Strategy* (October 21, 2022). <https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf>, vi.
5. Defense Information Systems Agency (DISA) and National Security Agency (NSA), *Department of Defense Zero Trust Reference Architecture version 2.0*, (July 2022). [https://dodcio.defense.gov/Portals/0/Documents/Library/\(U\)ZT_RA_v2.0\(U\)_Sep22.pdf](https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v2.0(U)_Sep22.pdf), 21-23.
6. Barclay Osborn, Justin McWilliams, Betsy Beyer, and Max Saltonstall, “BeyondCorp: Design to Deployment at Google,” (Spring 2016). *Usenix, The Advanced Computing Systems Association Journal*, volume 41, no. 1, <https://storage.googleapis.com/gweb-research2023-media/pubtools/pdf/44860.pdf>, 28-32.
7. Iftexhar Ahmed, Tahmin Nahar, Shahina Sultana Urmi, and Kazi Abu Taher, “Protection of Sensitive Data in Zero Trust Model,” (March 2020). *ICCA 2020: Proceedings of the International Conference on Computing Advancements*, <https://doi.org/10.1145/3377049.3377114>, 1-5.
8. U.S. Department of Justice, U.S. Cybersecurity and Infrastructure Agency, U.S. National Security Agency, Australian Cyber Security Centre, and U.K. National Cyber Security Centre, “Joint Cybersecurity Advisory: 2021 Trends Show Increased Globalized Threat,” (February 9, 2022) https://www.cisa.gov/sites/default/files/publications/A_A22-040A_2021_Trends_Show_Increased_Globalized_Threat_of_Ransomware_508.pdf
9. U.S. Department of Justice et al., “Joint Cybersecurity Advisory: 2021 Trends Show Increased Globalized Threat.”
10. European Parliament and Council of the European Union, “General Data Protection Regulation- Regulation (EU) 2016/679,” (April 27, 2016). *Official Journal of the European Union*, L 119: 1-86. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>.
11. State of California Department of Justice, “The California Consumer Privacy Act (CCPA) of 2018,” (last modified April 2024), https://cippa.ca.gov/regulations/pdf/cippa_act.pdf.
12. Pascal Poupart, Zhitang Chen, Priyank Jaini, Fred Fung, Hengky Susanto, Yanhui Geng, Li Chen, Kai Chen and Hao Jin, “Online Flow Size Prediction for Improved Network Routing,” (November 8, 2016). IEEE 24th International Conference on Network Protocols (ICNP) Workshop on Machine Learning in Computer Networks, <https://doi.org/10.1109/ICNP.2016.7785324>, 1-6.
13. William Su, Sung-Ju Lee and Mario Gerla, “Mobility Prediction and Routing in Adhoc Wireless Networks,” (January 18, 2001). *International Journal of Network Management* volume 11, no. 1, <https://doi.org/10.1002/nem.386>, 3-30.
14. Su, Lee, and Gerla, “Mobility Prediction and Routing in Adhoc Wireless Networks,” 4.
15. DISA and NSA, “Zero Trust Reference Architecture,” 19.
16. Goran Nikolic, Bojan R. Dimitrijevic, Tatjana R. Nikolic, and Mile K. Stojcev, “A Survey of Three Types of Processing Units: CPU, GPU and TPU,” (June 16, 2022). *57th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST)*, <https://doi.org/10.1109/ICEST55168.2022.9828625>, 1-6
17. Osborn, McWilliams, Beyer, and Saltonstall, “BeyondCorp,” 28-34.
18. Justin Sherman et al., “Data Brokers and the Sale of Data on U.S. Military Personnel: Risks to Privacy, Safety, and National Security,” (November 2023). Sanford School of Public Policy, Duke University, <https://techpolicy.sanford.duke.edu/wp-content/uploads/sites/4/2023/11/Sherman-et-al-2023-Data-Brokers-and-the-Sale-of-Data-on-US-Military-Personnel.pdf>.
19. DISA and NSA, “Zero Trust Reference Architecture,” 63.
20. The Dutch Data Protection Authority PISA Consortium, “PET” in *Handbook of Privacy and Privacy-Enhancing Technologies: The case of Intelligent Software Agents*, (November 2003). The Hague: College Bescherming Persoonsgegevens, ed. G.W. Blarckom, J.J. Borking, and J.G.E. Olk, <https://doi.org/10.13140/2.1.4888.7688>, 33-53.

NOTES

21. Raina Gandhi, “Technology Factsheet: Differential Privacy,” (Fall 2020). Belfer Center for Science and International Affairs, Kennedy School, Harvard University, <https://www.belfercenter.org/publication/technology-factsheet-differential-privacy>.
22. Michael B. Hawes, “Implementing Differential Privacy: Seven Lessons From the 2020 United States Census,” (September 2020). *Harvard Data Science Review* vol. 2, no. 2, <https://doi.org/10.1162/99608f92.353c6f99>, 2-3.
23. Daniel L. Oberski and Frauke Kreuter, “Differential Privacy and Social Science: An Urgent Puzzle,” (January 2020). *Harvard Data Science Review* vol. 2, no. 1, <https://doi.org/10.1162/99608f92.63a22079>, 4-6.
24. DISA and NSA, “Zero Trust Reference Architecture,” 70-71.
25. Rick Merritt, “What is Confidential Computing?,” (March 1, 2023). NVIDIA, <https://blogs.nvidia.com/blog/2023/03/01/what-is-confidential-computing/>.
26. Sergei Arnautov, Bohdan Trach, Franz Gregor, Thomas Knauth, Andre Martin, Christian Priebe, Joshua Lind, Divya Muthukumar, Dan O’Keeffe, Mark L. Stillwell, David Goltzsche, David Eysers, Rüdiger Kapitza, Peter Pietzuch, and Christof Fetzer, “SCONE: secure Linux containers with Intel SGX,” (November 2, 2016). *Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation*, Savannah, GA: The Advanced Computing Systems Association, <https://www.usenix.org/system/files/conference/osdi16/osdi16-arnautov.pdf>, 689-690.
27. Do Le Quoc, Franz Gregor, Jatinder Singh, and Christof Fetzer, “SGX-PySpark: Secure Distributed Data Analytics,” (May 2019). *WWW ’19: The World Wide Web Conference*, San Francisco, CA: Association for Computing Machinery, <https://doi.org/10.1145/3308558.3314129>, 3564-3565.
28. Richard Chirgwin, “Boffins Show Intel’s SGX Can Leak Crypto Keys,” *The Register*, March 7, 2017, accessed November 1, 2023, https://www.theregister.com/2017/03/07/eggheads_slip_a_note_under_intels_door_sgx_can_leak_crypto_keys/.
29. Chirgwin, “Boffins Show Intel’s SGX Can Leak Crypto Keys”.
30. Arnautov et al., “SCONE: secure Linux containers with Intel SGX,” 689-690.
31. Arnautov et al., “SCONE: secure Linux containers with Intel SGX,” 702.
32. Merritt, “What is Confidential Computing?”.
33. Craig Gentry, “A Fully Homomorphic Encryption Scheme,” (September 2009). *A Dissertation Submitted to the Department of Computer Science and the Committee of Graduate Studies of Stanford University*, <https://crypto.stanford.edu/craig/craig-thesis.pdf>.
34. Chiara Marcolla, Victor Sucasas, Marc Manzano, Riccardo Bassoli, Frank H.P. Fitzek, and Najwa Aaraj, “Survey on Fully Homomorphic Encryption, Theory, and Applications,” (October 2022). *Proceedings of the IEEE* vol. 110, no. 10, <https://doi.org/10.1109/JPROC.2022.3205665>, 1572-1609.
35. Marcolla et al., “Survey on Fully Homomorphic Encryption, Theory, and Applications,” 1576-1584.
36. Marcolla et al., “Survey on Fully Homomorphic Encryption, Theory, and Applications,” 1572.
37. Satoshi Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” (October 31, 2008) Accessed May 28, 2024, <https://bitcoin.org/bitcoin.pdf>.
38. CoinDesk, “CoinDesk Digital Assets Data and Research, Cryptocurrency Index,” Accessed June 18, 2024, <https://www.coindesk.com/indices>
39. CoinDesk, “CoinDesk Digital Assets Data and Research, Cryptocurrency Index”
40. Alexandra Andhov, “The U.S. Financial Innovation and Technology Act: Initial Overview,” (May 29, 2024). *Forbes*, Accessed May 30, 2024, <https://www.forbes.com/sites/digital-assets/2024/05/29/the-us-financial-innovation-and-technology-act-initial-overview/>.
41. Imran Bashir, *Mastering Blockchain*, 3rd. Edition, (August 31, 2020) Birmingham, U.K.: Packt Publishing.
42. “Quantum Technology,” *Wikipedia*, accessed May 6, 2024. https://en.wikipedia.org/wiki/quantum_technology.
43. Sydney J. Feedberg Jr., “Off to the Races: DARPA, Harvard Breakthrough Brings Quantum Computing Years Closer,” (December 7, 2023). *Breaking Defense*, <https://breakingdefense.com/2023/12/off-to-the-races-darpa-harvard-breakthrough-brings-quantum-computing-years-closer/>.
44. Michael Brooks, “What’s Next for Quantum Computing,” (January 6, 2023). *MIT Technology Review*, accessed December 15, 2023, <https://www.technologyreview.com/2023/01/06/1066317/whats-next-for-quantum-computing/>.

NOTES

45. Ville Lahtinen and Jiannis K. Pachos, "A Short Introduction to Topological Quantum Computation," (September 9, 2017). *SciPost Physics* vol. 3 no. 021, 1-44, <https://doi.org/10.21468/SciPostPhys.3.3.021>.
46. Patrick Davis, Sean Coffey, Lubjana Beshaj, and Nathaniel Bastian, "Quantum Machine Learning for Feature Selection in Internet of Things Network Intrusion Detection," (June 7, 2024). *Proceedings of Quantum Information Science, Sensing, and Computation XVI*, volume 13028, 1-15, <https://doi.org/10.1117/12.3013539>.
47. Stephen Wiesner, "Conjugate Coding," (January 1, 1983). *ACM SIGACT News*, vol., 15, no. 1, Association for Computing Machinery, <https://dl.acm.org/doi/10.1145/1008908.1008920>, 78-88.
48. Steven M. Bellovin, "Frank Miller: Inventor of the One-Time Pad," (July 1, 2011). *Cryptologia* vol. 35 no. 3, <https://doi.org/10.1080/01611194.2011.583711>, 203-222.
49. National Institute of Standards and Technology (NIST), "Post Quantum Cryptography Overview," (last modified May 28, 2024). *NIST Computer Security Resource Center*, accessed May 28, 2024, <https://csrc.nist.gov/projects/post-quantum-cryptography>.
50. NIST, "Post-Quantum Cryptography Standardization," (last modified May 28, 2024). *NIST Computer Security Resource Center*, accessed May 28, 2024. <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization>.
51. Office of the President of the United States, *Bringing Quantum Sensors to Fruition*. (March 2022). Report, National Science and Technology Council, Subcommittee on Quantum Information Science, <https://www.quantum.gov/wp-content/uploads/2022/03/BringingQuantumSensorsToFruition.pdf>.
52. Office of the President of the United States, *Bold Goals for U.S. Biotechnology and Biomanufacturing: Harnessing Research and Development to Further Societal Goals*, (March 2023). Report, The White House Office of Science and Technology Policy, <https://www.whitehouse.gov/wp-content/uploads/2023/03/Bold-Goals-for-U.S.-Biotechnology-and-Biomanufacturing-Harnessing-Research-and-Development-To-Further-Societal-Goals-FINAL.pdf>, 32
53. Sang Yup Lee, "DNA Data Storage Is Closer Than You Think," (July 1, 2019). *Scientific American*, accessed May 28, 2024, <https://www.scientificamerican.com/article/dna-data-storage-is-closer-than-you-think/>.
54. Rachael Rettner, "What is DNA?," (March 17, 2023). *Live Science*, accessed November 1, 2023, <https://www.livescience.com/37247-dna.html>.
55. Federico Tavella, Alberto Giaretta, Triona Marie Dooley-Cullinane, Mauro Conti, Lee Coffey, and Sasitharan Balasubramaniam, "DNA Molecular Storage System: Transferring Digitally Encoded Information through Bacterial Nanonetworks," (September 2021). *IEEE Transactions on Emerging Topics in Computing* vol. 9 no. 3, <https://doi.org/10.1109/TETC.2019.2932685>, 1566-1580.
56. Samuel Crislip, "Operationalize DNA Technology for U.S. Defense," (April 1, 2023). *AMS Special Session on Cybersecurity and Cryptography*, *American Mathematical Society*, American Mathematical Society, <https://meetings.ams.org/math/spring2023e/meetingapp.cgi/Paper/24268>.
57. DOD CIO Zero Trust Portfolio Management Office, "DoD Zero Trust Strategy," 22.
58. Tom Gaines and Alex Suh, "Reimagining Contested Communications," (August 23, 2023). *The Modern War Institute at West Point*, <https://mwi.westpoint.edu/reimagining-contested-communications/>.
59. U.S. Cybersecurity & Infrastructure Security Agency, et al., "Secure by Design. Shifting the Balance of Cybersecurity Risk: Principles and Approaches for Secure by Design Software," (October 2023). <https://www.cisa.gov/resources-tools/resources/secure-by-design>.
60. The MITRE Corporation, "MITRE ATT&CK v15.1," (last modified April 23, 2024), accessed June 17, 2024, <https://attack.mitre.org/>.
61. The MITRE Corporation, "MITRE Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS)," accessed June 17, 2024, <https://atlas.mitre.org/>.
62. BAE Systems, "What is Quantum Sensing?," accessed June 17, 2024, <https://www.baesystems.com/en-us/definition/what-is-quantum-sensing>.
63. BAE Systems, "What is Quantum Sensing?"

An Introduction to Quantum Computing and Its Applications

1st Lt. Mandeep Singh
Dr. Albert H. Carlson

ABSTRACT

Winning is much easier if you have an edge, whether that be better personnel, strategy, and/or technology. Quantum Information Science (QIS) - which includes quantum sensing, networking, communications, and computing - provides a technology that both tactical and strategic commanders will leverage to seize the initiative and create positions of advantage. Optimizing the exploitation of quantum technology will require that senior leaders understand enough about the technology and its fast-evolving applications to outmaneuver and outthink our adversaries. This does not require expertise in all facets of QIS, any more than a computer user needs to know computer design. This article attempts to be an introduction to quantum technology and some of its potential uses in the military operational environment.

INTRODUCTION

Had the Nazis won the race to harness the power of the atom, many would agree that World War II could have been disastrously lost, leaving the maps of Europe and the Americas vastly different. But what if the Nazis had quantum technologies before the D-Day invasion?

What if the Germans had the ability to thoroughly inspect key elements of the great deception plan, called Bodyguard, which misled the Germans about the location of the

© 2024 Dr. Albert H. Carlson; 1st Lt. Singh's contribution is a work of the U.S. Government and is not subject to copyright protection in the United States. Foreign copyrights may apply.



First Lieutenant Mandeep Singh grew up in Santa Clara, CA and attended college at Santa Clara University (SCU), earning degrees in Computer Science and Electrical Engineering with a minor in Mathematics. He has worked for several San Francisco Bay Area startups with roles in software development and architecture design, and currently is pursuing his graduate degree in Computer Science & Engineering at SCU. After earning his commission into the Cyber Corps, 1st Lt. Singh was the honor graduate of graduated his Cyber Basic Officer and Electromagnetic Warfare Qualification Courses at Fort Eisenhower, GA. He was then assigned to Fort Lewis, WA, where he now serves as the Electromagnetic Warfare Detachment Commander for 1st Special Forces Group (A). 1st Lt. Singh's areas of interest for research are quantum computing, post-quantum encryption algorithms, cryptography, and artificial intelligence. He recently published two articles in the fields of mathematics and cryptography with Dr. Carlson, earning best in conference at the 2024 AI IoT World Congress.

D-Day invasion? Using data from quantum LIDAR, magnetic field, and EM sensors, the Nazis likely would have known that the camps around the city of Dover and the landing crafts assembled along the southeast coast of England were phantom camps and ghost vessels, devoid of troops and actual equipment. These cleverly designed deception elements gave the Nazis every indication that the Allied invasion would cross the English Channel into Pas-de-Calais, not Normandy, which was 150 miles to the southwest.¹

Quantum LIDAR and imaging would have enabled the Nazis to track ship movement in real-time as they approached Normandy, and identify that the large clouds of aluminum strips dropped by Allied aircraft flying toward Pas-de-Calais were not, in fact, a large fleet of Allied ships, but just a decoy.² What if, instead of Allied codebreakers deciphering Germany's secret communications, they were able to crack ours? QIS would have given the Nazis a powerful advantage over the Allied Forces on D-Day. Nazi decision makers would have data sets, devoid of deceptive data, that led them to reinforce Normandy, and this data-driven decision would have potentially changed the fate of the Allied D-Day invasion and the outcome of the war.

Today, quantum computing is mostly inaccessible. To be prepared to use quantum technology when it is available, we must have the discussion about how to employ the technology today. Unlike the classical computers in universal use today, the size, high cost, and mechanical and operational sophistication of quantum computers have rendered them highly limited for use today.³ Most quantum computers are owned by governments, large corporations, and select universities. Classes and instruction on quantum computing are currently limited to a few institutions of higher learning.⁴ While not generally available, it can be anticipated that offerings for quantum instruction will rapidly increase, similar to artificial intelligence. This does not



Dr. Albert H. Carlson was born and raised in Chicago, IL, and attended the University of Illinois at Urbana, graduating with a B.S. in Computer Engineering in 1981 and entered the U.S. Army as an Electronic Warfare Officer. After assignment to the 105th MI Bn at Fort Polk, he worked as a computer and electrical engineer in various civilian companies and earned his Ph.D. in Computer Science with a focus in Set Theory applied to Cryptography at the University of Idaho. Dr. Carlson holds nine patents on polymorphic cryptographic algorithms, and has taught at Fontbonne University in St. Louis and Austin Community College. He currently conducts research and mentors graduate students at the University of Louisiana at Lafayette. Dr. Carlson serves as the Chair for Entropy and Encryption for the Quantum Security Alliance and works with IEEE on Cybersecurity for Next Generation Privacy. His research areas include cryptography, artificial intelligence, generative adversarial network applications, quantum cybersecurity, and control systems for aquaria.

mean that it will remain hard to access for very long. The impact of quantum machines is already being felt disproportionately by governments, academia, the military, and policymakers, and as of June 2024 various nations have invested about \$55 billion,⁵ to develop a general-purpose quantum computer. Nation-states are making this investment because quantum computers are projected to be able to break what cryptographers presently see as the most secure encryption ciphers: the asymmetric Public Key Encryption (PKE) algorithms.⁶

The United States has confirmed that a cryptanalytically relevant quantum computer (CROC) could put global communications systems, critical infrastructure, and financial transactions at risk.⁷ The first country to achieve a general-purpose quantum computer, with a sufficient number of qubits, will have a significant position of advantage over others in this era of Great Power Competition. Classical computers cannot begin to achieve the same operations that quantum computers use to break PKE cipher algorithms due to their basic hardware technology.⁸ Breaking PKE algorithms will require both a technology and paradigm upgrade. We are currently in a technological cold war to achieve general-purpose quantum computing. Future leaders will need to know how to leverage each new capability to create positions of advantage over our adversaries. The time for developing the people and concepts for a quantum future is now. This article presents the basics of quantum computing to assist forward-thinking leaders and technicians in how to better operationally deploy it in the future.

BASICS OF QUANTUM COMPUTING

Quantum computers are similar in structure and function to a classical computer. From an operational viewpoint, the user interface is usually a classical computer with a backend quantum coprocessor

(see Figure 1), with a classical front end interface where users do the actual programming, and where the language, compilers, algorithms, and communications protocols are found. This computer is typically programmed using Python and stores the program locally until it is transferred to the quantum coprocessor. This transfer occurs through the cooling unit or “chandelier.”⁹ Quantum hardware and the coprocessor are placed inside a supercooled enclosure that is large enough for technicians to work on that hardware. These rooms are typically a ten foot cube consisting of air, insulating walls, refrigeration equipment, and the computer, itself. Inside the cooled unit are the quantum computing hardware and the interface to keep things cold. Then the quantum coprocessor does the computation, arrives at the answer, and transfers the solution back to the classical computer via the interface unit in the chandelier to the user.




Figure 1. Classical Front End to Quantum Coprocessor.

Classical computers represent information in bits, the smallest encoding possible for information represented as 0 or 1, that are then combined to represent larger units of information, such as characters. Quantum computers use a similar, but more powerful type of representation for information, called quantum bits, or “qubits” for short. Qubits are like classical bits on steroids. Bits are independent units, but qubits can combine to represent even more data. In the quantum processor, qubits can be in the binary state typically depicted at a 0 or 1 like classical computers, but they can occupy multiple states simultaneously. This allows for the encoding of information at multiple positions, but also multiple positions in angles from the poles, e.g. at 90-degrees, 30 degrees, or 15 degrees (in all directions, see Figure 2)¹¹ not just at the poles (0 or 1). Multiple orientations of data give rise to the concept of quantum digits, or “qudits”¹² allowing for quantum computers to encode and process at a rate much faster than a classical computer.¹³




Figure 2. Multiple Positions in All Directions.¹⁰

The interface between the two machines passes data from the classical computer to the Quantum Processing block shown in Figure 3. This block, or layer, contains the memory, registers, and logic gates for computing at the quantum level. It may also be the “edge” of the quantum machine where all the bookkeeping and storage that is needed by the quantum hardware. This layer interacts with the Classical Computing layer to set up a program and later report the results, while also working with the Quantum Hardware layer to run programs.

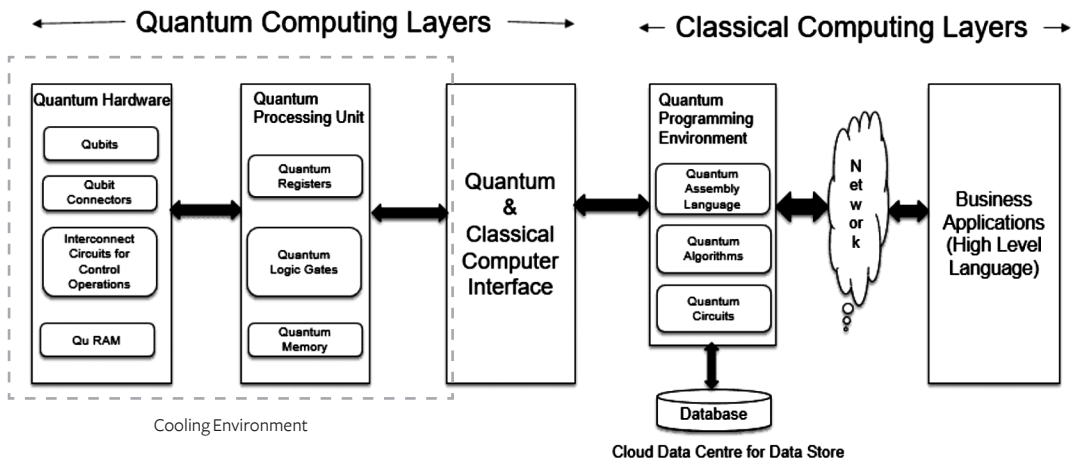


Figure 3. Detailed Structure of Quantum Computer.¹⁴

The Quantum Hardware layer consists of the storage unit for qubit memory, called the quRAM (fast local memory), and all connectors that help conduct the processing, and must be supercooled to correctly operate. As of late 2023, the time that supercooling can be maintained was an extremely short 34 ms.¹⁵ At the end of this process, the layer loses coherence (the ability to entangle qubits) and all computational activities end. Quantum computers also take a long time to cool down between runs to achieve the necessary cold temperature. In 2023, IBM quantum machines required up to 2.5 days to achieve that cooling.¹⁶ This situation should be fixed by the development of room temperature qubits.

Qubits can be created using any two-level quantum system. Some of the promising types of qubits being developed include: superconducting, trapped ions, quantum dots, and photons atoms.¹⁷

- ◆ Superconducting qubits are made from superconducting materials operating at extremely low temperatures and can be manipulated by microwave pulses. For the last decade researchers have used superconducting transmon qubits which are favored by researchers due to their high designability, scalability, and controlability.¹⁸ One problem with superconducting qubits is the coherence time, how long a qubit can run algorithms or perform operations before qubit drops its information, is relatively short. Researchers at MIT recently developed a superconducting qubit architecture using fluxonium qubits connected by transmon that can perform operations between qubits with high accuracy¹⁹ that could help make superconducting qubits the path to a fault-tolerant quantum computer.²⁰
- ◆ Trapped ions can also be used as qubits storing information in suspended free space around ions trapped using oscillating electromagnetic fields. Trapped ion qubits are noteworthy because the qubits have a relatively long lifetime, long internal state

coherence, and high-fidelity measurements²¹ but there are challenges to scaling them. Quantum computers with ion traps contain about 30 qubits each.²² However due to the nature of the oscillating electromagnetic field, combining more than multiple qubits on a single chip causes the trap to heat up significantly.²³ As of March 2024, new developments in using static magnetic fields instead of oscillating fields, called a Penning trap, show promise toward the realization of scaling trapped ion qubits and will benefit both quantum computing and sensors.²⁴

- ◆ A quantum dot is a nanoscale crystal semiconductor capable of holding a single electron or a small group of electrons that give the dot spin. These spins represent qubits of information and can exist in two states at once, both up and down, allowing computers that use them to do many calculations simultaneously.²⁵ Silicon chips, etched with quantum dots holding single electrons in their spin, appeared to be a feasible solution to bringing quantum computers to the masses but researchers were unable to manage the heat generated.²⁶ In May 2024, MIT and MITRE researchers demonstrated a “quantum-system-on-chip” architecture that integrates thousands to qubits on an integrated circuit and then connects multiple chips via optical networking.²⁷ This progress toward general-purpose quantum computing was significant “...with over 4,000 qubits that could be tuned to the same frequency while maintaining their spin and optical properties.”²⁸
- ◆ Photon qubits can be used to send quantum information through fiber optic cables by setting directional spin states of individual light particles.²⁹ Photon qubits are used in quantum communication and quantum cryptography. Recent developments at Berkeley’s Accelerator Technology & Applied Physics (ATAP) Lab include research on programmable spin-photon qubits uses a femto-second laser to create and destroy qubits on demand and at desirable locations.³⁰ Photon qubits were primarily viewed as information carriers between quantum computers to enable distributed processing. However, the development of the photonic quantum chip³¹ and a metal-ion qubit that absorbs light and emits color,³² makes a room temperature quantum computer appear to be possible.

Each type of qubit has shortcomings that must be overcome; breakthroughs cannot be accurately forecasted. However, using statistical nonlinearities it is estimated that a room-temperature quantum computer will be possible by 2031.³³ While most qubit implementations today take on the same binary values used in classical computers, many of these technologies can be used for more than representing binary values.³⁴

The ability to represent more than two values results in a qudit. The multi-level unit provides a larger space to store and process information.³⁵ Qudits can encode up to seven numbers, unlocking more computational power with fewer components, and can be a path to effective scaling.³⁶ For the purposes of this article, we will continue to describe quantum particles as qubits since most research we discuss is done using qubits.

Qubits are said to be entangled when they remain connected even when separated by large distances.³⁷ Knowing the state of one particle automatically tells you something about its entangled partners and changing one qubit will affect all the qubits entangled with it.³⁸ This connection allows qubits to perform vastly more operations per qubit than a classical computer can per bit, and computing power rises exponentially with the number of qubits. As of March 2024, quantum computers are primarily being used for simulations, optimization, and data analysis problems as most applications of quantum computing are still theoretical and, according to John Preskill, a theoretical physicist at the California Institute of Technology, it is challenging to find problem sets worthy of quantum computer “...because classical computers are pretty good at a lot of the things they do.”³⁹

Most quantum computing capabilities and research are conducted in isolated, temperature controlled, super-cooled environments connected to the outside world via an interface unit to a classical computer. As such, most current cybersecurity considerations are indirect in nature and involve indirect attack vectors on the power, cooling, or vibration abatement systems. Cyber attacks would have to target the classical computer or its user and then move to the interface unit before trying to breach the quantum computer.

A summary comparing classical and quantum computers is in Table 1.

Classical vs Quantum Computers		
Feature	Classical Computer	Quantum Computer
Computing units	Bits	Qubits
Computing Power	Linear increase with number of transistors	Exponential increase with number of qubits
Logic Type	Binary Logic	States of spin on atoms
Targets	General Problems	Optimization, Data Analysis, Simulations
Operation size	1 per N bits	2^N for N qubits
Security	Attacks directly and indirectly to multiple aspects of the system.	Currently physical. Social engineering and electronic attacks targeting the cooling system.
Temperature	Room	Super cooled

Table 1. Comparison of Classical and Quantum Computers.

ADVANTAGES AND DISADVANTAGES OF QUANTUM COMPUTING

Some of the advantages of quantum computers include:

- ◆ Speed - Quantum computers are fast. They allow solving some problems in considerably less time than presently available classical machines. In general, a quantum machine runs faster than a classical machine due to its architecture. Quantum computers can run large-scale data problems faster than classical computers. However, the problem of cooling restricts how long a quantum computer can run. Most quantum computers take a long time to cool down enough to run again.⁴⁰

- ◆ **Solving New Problems** - Quantum computers can solve problems not solvable by classical machines. Quantum computers open classes of problems that cannot be effectively represented and calculated using the binary approach. Drawing a similarity to mechanical engines, the mathematics and implementation of the V-8 and the rotary engine both generate power, but the method of generation is different. They perform the same function but in different ways. Among the complex problems is the ability to break asymmetrical key ciphers which are presently used for the protection of data.
- ◆ **Enhancing Classical Computers** - Quantum computers are complementary to classical computers. Both types of computers are useful for different problems and can be used together.
- ◆ **Memory Density** - Entangled qudits can represent problems that cannot be represented in the binary computational space. Entangled qudits have more than two states. Instead of rising linearly, like the information in classical memories, qudits rise exponentially. It only takes a relatively few qubits and even fewer qudits to perform complex algorithms that exceed the capability of conventional computers.⁴¹
- ◆ **Novel Approaches to Problems** - Quantum computers allow for a new way to view problems. Because of their new technology, new approaches to solving problems are needed and becoming available. Progress takes place when new ways to view a problem present themselves.

There are also aspects of quantum computers that are not advantages. Some of these disadvantages may be overcome with time, others are inherent to the technology. These disadvantages include:

- ◆ **Need for Supercooling** - Qubits that presently make up quantum computers require temperatures near absolute zero. Such a temperature is extremely hard to generate and maintain. Although some progress in this area was made with the photonic quantum chip and light absorbing metal-ion qubit, the need for extreme cold is still a significant problem with quantum computing today.
- ◆ **Costs** - Supercooling requires specialized equipment and power to maintain the computer's required operational environment. Supercooling coupled with the use of exotic materials to construct the various qubits, processors, and interfaces make quantum computing a very costly venture. This problem will probably be moderated as room temperature qubits are developed, but the time until that improvement is likely ten or more years from now.
- ◆ **Thermal Errors** - Just running a quantum machine can cause thermal errors to be present in the results of calculations. The problem is solved by quantum annealing,

but will require connecting multiple hardware qubits to act as a single, logical qubit and additional qubits will be needed to monitor the behavior of the ones holding the data and make the necessary corrections.⁴² Multiple runs of the same problem are required to correlate the output and decide on the correct solution, each run will also generate heat, potentially causing more thermal errors.

- ◆ **Decoherence** – Decoherence is the process in which the environment interacts with qubits, often causing uncontrollable changes to their quantum states and information to be lost.⁴³ Since qubits interact with the environment, they are affected by it. Changing magnetic and electric fields, radiation, other qubits, seismic activity, and physical movement can all cause decoherence.
- ◆ **Not Generally Available** - At this time, there are few quantum computers. Lack of supply of these units makes them harder to employ and limits access to the resource. The development of capabilities that can link multiple quantum computers together may help alleviate the shortage of materials needed to make a general purpose quantum computer.
- ◆ **Lack of Trained Users** - As with newer technologies, the number of trained users is limited. Few experts exist to design, develop, and operate the machines. The lack of specialists in this area means the cost of training and employing such trained personnel is high.
- ◆ **Standardization** - Quantum computers are early in their development timeline. Many solutions for the problems and implementation are being tried. There is little standardization for the designs or methodology, contributing to the confusion in the selection and implementation of quantum computers.

Quantum computing adds a new dimension to computing. It can speed up programming and provide new insight into information and data problems. However, it does not change information theory⁴⁴ and communications theory⁴⁵ but rather does things in a new way. It will provide a great deal of innovation in those fields as the technology matures. While there seems to be tension between classical and quantum computers, the two are complementary in functionality. An easy way to think of the differences between classical and quantum computers is that classical computers excel at discrete math (such as algebra and trigonometry) while quantum computers do a much better job of continuous mathematics (such as calculus-based problems). Many experts feel that the future of computing holds hybrid computing systems.⁴⁶ Both classical and quantum computers will exist in the same machine with a preprocessor that selects which unit is best suited to solve the problem presented to the computer. However, this solution is probably not practical in the short term due to the problems with supercooling and the need for a unified programming language solution.

THE IMPACTS OF QUANTUM COMPUTING ON DATA

The infrastructure of information and information processing in the U.S. is primarily that of classical computing. There are over 1 trillion (10^{12}) classical computing devices in use as of 2023 and more than 2 billion (2×10^9) personal computers.⁴⁷ In contrast, as of early 2023, there were less than 120 quantum computers in operation worldwide.⁴⁸ Part of the reason for the number disparity is that quantum computers are continuing to evolve. However, quantum sensors are already in widespread use and their use is expected to increase significantly as new ecosystems for data collection are formed from observing what used to be unobservable and training new AI models.⁴⁹ The numbers suggest QIS will have a significant impact on data and data analytics for our Army. Below we will discuss two areas that will have the most significant impacts to the Army: quantum sensing and quantum computing, the latter of which includes post quantum encryption.

Quantum Sensing Technologies

Quantum sensing applications are the most mature of the QIS technologies being developed for use in the DoD.⁵⁰ Quantum sensing technologies use trapped ions, solid-state spins, superconducting circuits, and quantum dots⁵¹ to provide sensing capabilities across several physical environments including magnetic and electric fields, the forces of acceleration, pressure, gravity, and time.⁵² The combinations of these sensing capabilities enable quantum sensors to collect data and conduct analysis that will enable us to see the previously unobservable, including undersea, underground, and in space.

While sonar is widely used for the detection of objects in deep water, in shallower depths sonar can become affected by echoes off of the seabed or shoreline.⁵³ Magnetic detection can discriminate magnetic objects from non-magnetic objects. Arrays of quantum sensors called atomic magnetometers have successfully detected 100% of single and multiple targets in both saline water and air, and precisely located over 93% of the objects, to include tracking of moving targets.⁵⁴ In the vast expanses of the Pacific Ocean, the application of quantum sensors to detect vessels deliberately evading satellite imagery and sonar would enable the U.S. and our allies to detect Chinese movements toward disputed waters and territories or the deployment of autonomous vessels, underwater drones, unmanned underwater vehicles,⁵⁵ or mines.

Rydberg atom electric field sensors can be used for ultra-wideband spectrum sensing and communications and high-accuracy near-field sensing.⁵⁶ Rydberg electrometry uses atoms to observe electric fields and is considered a path to accessing large operational bandwidths with one device without an array of antennas⁵⁷ paving the way for small form factor radios with autonomous hopping between multiple bands and expanding use of the electromagnetic spectrum to include frequencies higher than the 100 GHz achieved by conventional electronics.⁵⁸

Inertial navigation systems leverage quantum sensors to measure rates of acceleration and rotation of moving objects, such as a missile, airplane, or ship, to determine its position and orientation. They are widely used for navigation in GPS-denied environments and the more advanced systems use sophisticated algorithms and data fusion techniques to filter out the effects of jamming and improve resiliency.⁵⁹ Russia's Iskander short range ballistic missile, responsible for numerous deadly attacks on Ukraine in 2023 and 2024, owes its unprecedented accuracy to its inertial navigation system.⁶⁰

Quantum sensors for timing include both microwave and optical atomic clocks for use in GPS-denied timing, secure communications, and sensing requiring synchronized distributed sensor nodes such as radar and reflectometry.⁶¹ Quantum based atomic clocks do not rely on satellite connectivity to operate, mitigating the risk posed by electronic jamming or GPS signal spoofing. Microwave atomic clocks have applications in communications networks, GPS, and long baseline interferometry,⁶² used in measuring microscopic displacements and surface irregularities. Optical atomic clocks are used in GPS-denied navigation, distributed sensing (e.g. synthetic aperture radar), international time-keeping, and geodesy.⁶³

These four quantum sensing applications, magnetic fields, electric fields, force, and time come with exponential data storage demands and algorithms. Increased applications with intelligence, surveillance and reconnaissance (ISR) and precision, navigation and timing (PNT) will enable navigation in GPS-denied environments and enable us to see in maritime and underground environments.⁶⁴ Both our allies and our adversaries will use quantum sensors to collect data on all the places military forces used to hide, whether that be camouflage to blend in with the environment, a safe house for special forces missions, the tunnels under Gaza, or in the expanses of the ocean. There will be an increased need for data storage, but more importantly, we will need people who are capable of leveraging this technology, developing the algorithms to make sense of it, and designing a new way to fight in an environment where we can no longer hide.

Quantum Communications

Quantum communications will have a significant impact on data security through quantum key distribution. It also includes the ability to network quantum computers and sensors together. In this section we will discuss one of the largest threats quantum computing poses to our national defense: the ability to render current industry PKE standard encryptions useless. Currently, two types of encryptions are considered extremely strong:

1. Advanced Encryption Standard (AES), especially when wrapped in randomizing modes such as Cipher Block Chaining (CBC). AES is a block product cipher in which the plain text is encrypted 14 times with randomized input data.⁶⁵ This is meant to mix up the ciphers so no one can follow the changing data during encryption.

2. Asymmetric key public key encryption (PKE) algorithms require two keys. The idea is that if one key is hard to break then two keys are much harder and as a result data is safer. However, one of the keys is publicly revealed, so the use of the additional key is inconsequential.

This leads to the application of Shor's algorithm in using quantum computers to break PKE ciphers.

Shor's algorithm⁶⁶ demonstrates that algorithms such as the Rivest-Shamir-Adleman (RSA) cipher⁶⁷ and Diffie-Helman⁶⁸ "secure" key exchange can be broken easily. Each of these algorithms is based on "hard"⁶⁹ ciphers built on the discrete logarithm, prime and semi-prime factorization, or Euler's totient function. These problems have not been broken by a classical computer. It is believed Shor's quantum algorithm will allow large-scale quantum computers to break these asymmetric key algorithms rendering internet traffic unsecure.⁷⁰

Grover's algorithm,⁷¹ a quantum algorithm for unstructured search, shows that it is possible to reduce the effectiveness of one key ciphers, like AES-256, by reducing the time to brute force symmetric algorithms by an average time and effort of a factor of 2, meaning a 256-bit key would take roughly 2^{128} iterations to crack. Further research showing attack methods on AES-256 with bolted on randomization functions⁷² show this too is able to be cracked. AES has been shown, through isomorphic cipher reduction,⁷³ to reduce to a block substitution (S) cipher. Isomorphic cipher reduction allows mapping from one cipher to another. Even when protected by the CBC mode, a side-channel attack on parts of the hardware or surrounding software can return the original text and bypass AES.⁷⁴ Even without these outside attacks, analysis of AES in the quantum environment has resulted in an estimate that even AES-128 cannot be defeated in approximately 100 trillion years using a classical computer.⁷⁵ Smaller numbers of possible keys (known as "key spaces") than available for AES can be defeated in a matter of seconds or minutes. AES can be a component of the cryptographic solution but is not the sole answer to encryption in the post quantum environment (PQE). Symmetric key algorithms can be used in the PQE successfully and should be considered if the principles of Information Theory are properly followed.

Quantum computers will be able to break PKE-protected messages and our adversaries are preparing for that time. Both China and Russia are now implementing the Harvest Now Decrypt Later (HNDL) Attack,⁷⁶ where they collect and store encrypted data that cannot yet be decrypted and store them for a later time when quantum machines are capable of that decryption. The harvesters know that most data will not be valuable or helpful, but they hope that an occasional message will contain vital and important data that they can use.

Much of the information used in daily life is time sensitive, meaning it often loses value over time. While some users believe their data should be kept secret forever,⁷⁷ the goal of cryptography is to obscure information from an adversary until that information no longer holds value if revealed. For example, intelligence about upcoming an attack has great value,

however, once the attack begins that information loses value. It becomes part of the lessons learned and is often shared with the cyber security community. The longer the data remains encrypted and hidden, the more likely the value of that data will decrease to the point that it becomes useless.

In 2022, NIST and NSA named four algorithms they consider “quantum resistant” meaning they can run on computers today and are believed to be resistant to attacks from both classical and quantum computers: Crystals-Kyber, Crystals-Dilithium, Falcon and SPHINCS+.⁷⁸ Unfortunately, all four selections have now been proven vulnerable and cannot be considered safe. Researchers from North Carolina State University were the first to present vulnerabilities in the Falcon algorithm.⁷⁹ The SPINCS+ algorithm was shown to be vulnerable to a forgery attack in September of 2022.⁸⁰ Researchers in Stockholm used recursive trained artificial intelligence combined with a side-channel attack to crack the CRYSTALS-Kyber algorithm in February of 2023.⁸¹ In April 2024, the Dilithium based algorithms was also cracked using a side-channel attack.⁸²

In May 2024, a quantum-resistant algorithm using polymorphic encryption that incorporates sharding was presented at the IEEE Artificial Intelligence IoTConference.⁸³ In this encryption method, messages are broken into shards and each shard must be independently broken. Revealing one shard does not provide clues for other shards, effectively breaking the relationship between different portions of the message. Polymorphic ciphers change the combination of ciphers and keys used independently from each other, in frequent, but irregular intervals. Changes to the cipher/key pairs are made often so the data needed for decryption cannot accumulate.

Recent advances in using these algorithms show promise, but have not yet been sufficiently studied to allow for full acceptance. In May 2024, Anne Neuberger announced that NIST expects to release four new post quantum algorithms as early as July.⁸⁴ It will take the collective efforts of academia and the government to develop the algorithms needed to protect our nation’s data in the post quantum world.

CONCLUSION

Operationally deploying quantum capabilities in the future will provide positions of advantage in seeing formerly unseen areas of our environment undersea, underground, behind walls, and under camouflage. Enabled by quantum sensors, both our allies and adversaries will collect intelligence at rates exponentially higher than seen today. Data harvesting of military, industry, economic, and personal data will also increase, even for encrypted data, as quantum computing will be able to decrypt today’s most secure algorithms. Any data that may provide a position of advantage will be targeted.

Research in developing hybrid architectures such as the limited resources for quantum computing will likely cause independent research efforts to team together, creating powerful

architectures for general purpose quantum computers. The processing speed of these quantum computers will enable simulations of large-scale military deployments to train our forces and train with our allies. This training should include the development of new techniques, tactics, and procedures that take into account the quantum sensors that will see their every move and the AI-enabled algorithms that will try to get inside the decision cycle of our commanders. Quantum technologies will cause the way we conduct war to evolve and adapt.

Commanders will have to not just leverage our quantum capabilities to attack, but more than ever develop defensive practices to counter adversary capabilities. Deception techniques will need to be developed to frustrate sensors that can see through smoke, weather, physical obstructions, and vegetation. Breaking up the visual outline of equipment is not enough. The magnetic/gravitational outline will also have to be obscured. Therefore, alternate means of camouflage must be developed and TTPs developed for employing both at home station and while deployed.

As we develop our quantum capabilities, an equal effort needs to be applied to countering these capabilities. Once a new capability is used in the open, our adversaries and industry rapidly copy it and leverage it for their use or commercial gain. But the near-term advantage will go to the nation state which fully embraces the development and application of quantum technologies, this is why our leaders must start to understand quantum technology and start thinking about how to apply it in the future operational environment.

Leaders will have to not just maneuver troops on the battlefield. Our Soldiers and their families are in a battle for their data every day. The harvesting of personal data for future use and for cognitive warfare is real and will grow with the advent of quantum capabilities, as will the targeting of critical infrastructure around our military bases. Attention should be placed on protecting our people from these threats.

Quantum sensors will not just be used for detecting military equipment and units, quantum sensors will enable the ability to target and track individuals via their heartbeat through a crowd. Detecting and identifying individual signatures allows for the finding, fixing, isolating, and targeting of specific soldiers, equipment, and high value targets, for both sides. It will be a commander's responsibility to balance the ability to exploit the data and intelligence gained by quantum sensors and computing while, simultaneously defending friendly assets from collection and exploitation by our adversaries.

Quantum sensors, computing, and communications will provide advantages for precision fires in the near term, but it will also make our adversaries more lethal as seen in the Russian use of the Iskander missile in Ukraine. Attention should be also given to the electromagnetic spectrum. As quantum communications enable the ability to leverage more of the EMS for our secure communications via Rydberg electrometry and QKD, quantum will also enable the detection of faint changes in the EMS to pinpoint radios and the people using them for

communications. The EMS will become a battlespace that must be fully integrated into plans and operations across the continuum of military operations.

The technology race to harness the power of the qubits and qudits will not just revolutionize the battlefield, it will also render the PKE encryption that protects communications over the internet useless. This will affect everything that enables our economy and critical infrastructure to operate. Our leaders should both support and monitor the efforts of NIST and the NSA to develop secure quantum algorithms. While we recommend looking at polymorphic algorithms combined with a shard-based approach as a solution for quantum-proof encryption, there is much work to be done in this area.

Advances such as the Penning Trap, MIT's Fluxonium qubits connected by transmon, MIT and MITRE's efforts to create a "quantum-system-on-a-chip," and the development of the photonic quantum chip and light-absorbing metal-ion qubit are all steps toward overcoming the issues of scaling, coherence, and heat and show that a room temperature general purpose quantum computer will be developed in the next decade. The Army needs leaders who facilitate and support these research efforts and are prepared to leverage new capabilities in the future.♥

ACKNOWLEDGEMENT

The authors would like to thank MAJ James Armstrong, MAJ Anna White, MAJ Stephen Willson, and Deborah Karagosian for their guidance and feedback while writing this article.

DISCLAIMER

The views expressed in this work are those of the authors and do not reflect the official policy or position of the United States Military Academy, the Department of the Army, or the Department of Defense.

NOTES

1. English Heritage Website, “D-Day Deception: Operation Fortitude South,” (accessed May 22, 2024). <https://www.english-heritage.org.uk/visit/places/dover-castle/history-and-stories/d-day-deception>.
2. Klein, Christopher, “Fooling Hitler: The Elaborate Ruse Behind D-Day,” (accessed on May 20, 2024). *History Channel Online*, <https://www.history.com/news/fooling-hitler-the-elaborate-ruse-behind-d-day>.
3. N. Chandra and S. Parida. “Quantum Entanglement in Photon-induced Electron Spectroscopy of Atoms and Molecules: Its Generation, Characterization, and Applications,” (2016). *Advances in Imaging and Electron Physics*, vol. 196, Elsevier, <https://www.sciencedirect.com/science/article/abs/pii/S1076567016300404>, 1–164.
4. Matt Swaney, “20 Top Universities for Quantum Computing Research,” (May 21, 2024). *The Quantum Insider*. <https://thequantuminsider.com/2024/05/21/20-top-universities-for-quantum-computing-research/>.
5. Sylvain Duranton, “Quantum Computing Takes Off with \$55 Billion in Global Investments,” (June 26, 2024). *Forbes*, <https://www.forbes.com/sites/sylvainduranton/2024/06/26/quantum-now/>.
6. Hrithik Saini, “8 Strongest Data Encryption Algorithms in Cryptography,” (March 10, 2022). *Analytic Steps Online*, <https://www.analyticssteps.com/blogs/8-strongest-data-encryption-algorithms-cryptography>.
7. The White House, “National Security Memorandum on Promoting United States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems.” (May 4, 2022) Technical report, <https://www.whitehouse.gov/briefing-room/statements-releases/2022/05/04/national-security-memorandum-on-promoting-united-states-leadership-in-quantum-computing-while-mitigating-risks-to-vulnerable-cryptographic-systems/>, sec. 1, para. b.
8. Stephen Ornes, “Inside the Quest for Unbreakable Encryption.” (October 19, 2023). *MIT Technology Review*, www.technologyreview.com/2023/10/19/1081389/unbreakable-encryption-quantum-computers-cryptography-math-problems/.
9. Nicholas Ho, “Why Quantum Computers Resemble Chandeliers?.” (accessed June 8, 2024). *Quantum Computing Vulnerabilities are Everywhere Project*, <https://qcve.org/blog/why-quantum-computers-resemble-chandeliers>.
10. Jeff Hecht, “Quantum Science: The Quest for Quantum Information Technology Expands,” (July 14, 2020) *Laser Focus World*, <https://www.laserfocusworld.com/lasers-sources/article/14176179/quantum-science-the-quest-for-quantum-information-technology-expands>.
11. Li-Heng Chang, Shea Roccaforte, Rose, Xu, and Paul Cadden-Zimansky. “Geometric Visualizations of Single and Entangled Qubits.” (December 7, 2022). *Cornell University Quantum Physics Repository Online*, <https://arxiv.org/abs/2212.03448>, p 11.
12. Yuchen Wang, Zixuan Hu, Barry C. Sanders, Sabre and Kais, “Qudits and High-Dimensional Quantum Computing,” (2020). *Frontiers in Physics*, vol. 8, <https://www.frontiersin.org/articles/10.3389/fphy.2020.589504>.
13. R. Whitney Johnson, Kamyar Maserrat, and Jihwang Yeo, “The Basics: How Quantum Computers Work and Where the Technology is Heading,” (May 23, 2024). Wahington, D.C.: Foley and Lardner LLP, <https://www.foley.com/insights/publications/2024/05/basics-how-quantum-computers-work-where-technology-heading/>.
14. Gopala Krishna Behara, “Overview of Quantum Computer Platform,” (September 12, 2021). *Analytics Insights*, <https://www.analyticsinsight.net/latest-news/overview-of-quantum-computer-platform>.
15. Evan Brown, “The Devil in the Details: A Survey of Current Approaches to Building a Quantum Computer.” (December 18, 2023). *Center for Strategic and International Studies*, <https://www.csis.org/blogs/strategic-technologies-blog/devil-details-survey-current-approaches-building-quantum-computer>.
16. Will Fox, “New Record Length for Quantum Coherence,” (September 18, 2023). *Timeline.net*, <http://futuretimeline.net/blog/2023/09/18-record-length-for-quantum-coherence.htm>.
17. Josh Schneider and Ian Smalley, “What is a Qubit?” (February 28, 2024). *IBM Think*, <https://www.ibm.com/topics/qubit>.
18. He Liang Huang, Dacao Wu, Daojin Fan, and Xiabao Zhu, “Superconducting Quantum Computing: A Review,” (November 3, 2020), https://www.researchgate.net/publication/342380016_Superconducting_Quantum_Computing_A_Review_2.
19. Adam Zewe, “New Qubit Circuit Enables Quantum Operations with Higher Accuracy,” (September 25, 2023). *MIT News*, <https://news.mit.edu/2023/new-qubit-circuit-enables-quantum-operations-higher-accuracy-0925>.
20. James Dargan, “Fluxonium Qubits Bring the Creation of a Quantum Computer Closer,” (November 22, 2022). *The Quantum Insider*, <https://thequantuminsider.com/2022/11/22/fluxonium-qubits-bring-the-creation-of-a-quantum-computer-closer/>.

NOTES

21. Colin D. Bruzewicz, John Chiavernini, Robert McConnell, and Jeremy M. Sage, “Trapped-Ion Quantum Computing: Progress and Challenges,” (April 9, 2019). *Massachusetts Institute of Technology, Lincoln Laboratory*, <https://arxiv.org/pdf/1904.04178>, 3.
22. Oliver Morsch, “A New Ion Trap for Larger Quantum Computers,” (March 13, 2024), *Physics*. <https://phys.org/news/2024-03-ion-larger-quantum.html>.
23. Shreyans Jain, Tobias Sagesser, Pavel Hrmo, Celeste Torkzaban, Martin Stadler, et al. (March 13, 2024). *Nature*, vol. 627, <https://www.nature.com/articles/s41586-024-07111-x>, 510.
24. Jain, et al., “Penning Micro-trapo for Quantum Computing,” 513.
25. Dexter Johnson, “The Road to a Quantum Computer Begins with a Quantum Dot,” (May 25, 2020). *IEEE Spectrum*, <https://spectrum.ieee.org/the-road-to-a-quantum-computer-begins-with-a-quantum-dot>.
26. John Timmer, “Quantum Computing Progress: Higher Temps, Better Error Correction,” (March 27, 2024). *ARS Technica*, <https://arstechnica.com/science/2024/03/quantum-computing-progress-higher-temps-better-error-correction/>.
27. Linsen Li, Lorenzo DeSantis, Isaac B. W. Harris, Kevin C. Chen, et al. “Heterogeneous Integration of Spin–photon Interfaces with a CMOS Platform,” (May 29, 2024). *Nature Communications*, <https://www.nature.com/articles/s41586-024-07371-7>.
28. Adam Zewe, “Modular Scalable Hardware Architecture for a Quantum Computer,” (May 29, 2024). *MIT News*, <https://news.mit.edu/2024/modular-scalable-hardware-architecture-quantum-computer-0529>.
29. Schneider and Smalley, “What is a Qubit?”.
30. Jhuria, K., Ivanov, V., Polley, D. et al. “Programmable Quantum Emitter Formation in Silicon. (May 27, 2024) *Nature Communications*, <https://doi.org/10.1038/s41467-024-48714-2>.
31. Charles Q. Choi, “In the Race to Hundreds of Qubits, Photons May Have Quantum Advantage,” (March 5, 2021). *IEEE Spectrum*, <https://spectrum.ieee.org/race-to-hundreds-of-photonic-qubits-xanadu-scalable-photon>.
32. Kyushu University. "Generating Stable Qubits at Room Temperature." (January 11, 2024). *Science Daily*. www.science-daily.com/releases/2024/01/240111113125.htm.
33. Stefan Krastanov, Mikkel Heuck, Jeffrey H. Shapiro, and Dirk R. Englund, “Room-temperature Photonic Logical Qubits via Second-order Nonlinearities,” (January 8, 2021). *Nature Communications*, Vol. No. 12, 191. <https://doi.org/10.1038/s41467-020-20417-4>.
34. Donna Lu. “What is a Quantum Computer?” (accessed May 19, 2024). *New Scientist*, <https://www.newscientist.com/question/what-is-a-quantum-computer>.
35. Yuchen Wang, Zixuan Hu, Barry C. Sanders, and Sabre Kais, “Qudits and High-Dimensional Quantum Computing.” (November 2020). *Frontiers in Physics*, vol. 8, <https://www.frontiersin.org/articles/10.3389/fphy.2020.589504>.
36. Charles Q. Choi, “Qudit Computers Go Beyond Ones and Zeroes,” (August 1, 2022), *IEEE Spectrum*, <https://spectrum.ieee.org/qudit>.
37. California Institute of Technology, “What is Entanglement and Why is it Important?” (accessed February 24, 2024). *Caltech Science Exchange Online*, <https://scienceexchange.caltech.edu/topics/quantum-science-explained/entanglement>.
38. California Institute of Technology, “What is Entanglement and Why is it Important?” (accessed June 28, 2024)
39. Lakshmi Chandrasekaran, “Physicists Finally Find a Problem That Only Quantum Computers Can Do,” (March 12, 2024). *Quanta Magazine*, <https://www.quantamagazine.org/physicists-finally-find-a-problem-only-quantum-computers-can-do-20240312/>.
40. Lindsey Valich, “A Quantum Leap in Cooling Atoms for Better Computers.” (September 12, 2023). *University of Rochester*. <https://www.rochester.edu/newscenter/quantum-mechanics-thermoelectricity-superposition-entanglement-565852>.
41. Glenn Roberts, Jr., “Going Beyond Qubits: New Study Demonstrates Key Components for a Qudit-Based Quantum Computer,” (April 26., 2021). *University of California, Lawrence Berkeley National Laboratory*, <https://newscenter.lbl.gov/2021/04/26/going-beyond-qubits/>.
42. John Timmer, “Quantum Computing Progress: Higher Temps, Better Error Correction,” (March 27, 2024). *ARS Technica*, <https://arstechnica.com/science/2024/03/quantum-computing-progress-higher-temps-better-error-correction/>.
43. Katherine McCormick, “Decoherence Is a Problem for Quantum Computing, but...,” (March 30, 2020). *Scientific American*, <https://www.scientificamerican.com/blog/observations/decoherence-is-a-problem-for-quantum-computing-but/>.

NOTES

44. Claude E. Shannon, "A Mathematical Theory of Communication." (1948) *Bell System Technical Journal*, 27; Wang, et al., "Qudits and High-Dimensional Quantum Computing;"
45. Thomas Cover and Joy Thomas, *Elements of Information Theory*, (2005). New York: John Wiley & Sons, Inc, 2nd edition.
46. Peter Chapman and Pat Tang, "What is Hybrid Quantum Computing? (January 18, 2024) <https://ionq.com/resources/what-is-hybrid-quantum-computing>.
47. Richard Taylor, "How many computers are there in the world?" (accessed June 22, 2024). Quora, <https://www.quora.com/How-many-computers-are-there-in-the-world>.
48. Campbell, Charlie, "Quantum Computers Could Solve Countless Problems—And Create a Lot of New Ones." (January 26, 2023), *Time Magazine*, <https://time.com/6249784/quantum-computing-revolution/>.
49. Yannick Bormuth, Martina Gschwendtner, Henning Soller, Amanda Stein, and Ronald Walsworth, "Quantum Sensing Can Already Make a Difference but Where?" (2024). *Journal of Information Management*, <https://journalsojs3.fe.up.pt/index.php/jim/article/view/2578/848>, 2.
50. U.S. Department of Defense, *Applications of Quantum Technologies*, (November 2019), Report, Defense Science Board, Office of the Under Secretary of Defense for Research and Engineering, 6.
51. David L. Chandler, "Quantum Sensor Can Detect Electromagnetic Signals of Any Frequency," (June 21, 2022). *Massachusetts Institute of Technology, Lincoln Laboratory*, <https://news.mit.edu/2022/quantum-sensor-frequency-0621>.
52. Bormuth, et al., "Quantum Sensing Can Already Make a Difference but Where?," 3.
53. Tim Wogan, "Atomic Magnetometers Detect Underwater Objects," (April 18, 2018). *Physics World*, <https://physicsworld.com/a/atomic-magnetometers-detect-underwater-objects/>.
54. Cameron Deans, Luca Marmugi, and Ferruccio Renzoni, "Active Underwater Detection with an Array of Atomic Magnetometers," (March 22, 2018). *Applied Optics*, Vol. 57, No. 10, Optica Publishing, <https://opg.optica.org/ao/fulltext.cfm?uri=ao-57-10-2346&id=383870>, 2346.
55. Zubeda Anjum Niazi, "Future of Maritime Security: Navigating Complex Waters in the Indo-Pacific," (March 1, 2024). *Journal of Indo-Pacific Affairs*, <https://media.defense.gov/2024/Mar/11/2003410990/-1/-1/1/FEATURE%20-%20NIAZI.PDF/FEATURE%20-%20NIAZI.PDF>, 132.
56. David Anderson, Rachel Sapiro and Georg Raithel, "A Self-Calibrated SI-Traceable Rydberg Atom-Based Radio Frequency Electric Field Probe and Measurement Instrument," (September 2021). *IEEE Transactions on Antenna Propagation*, vol. 69, no. 9, <https://ieeexplore.ieee.org/document/9363580>, 5931-5941.
57. John Keller, "ColdQuanta Eyes Quantum Applications in Electronic Warfare (EW), Sensors, and Anti-submarine Warfare (ASW)," (April 14, 2021). *Military and Aerospace Electronics Magazine*, <https://www.militaryaerospace.com/computers/article/14201305/quantum-sensors-electronic-warfare-ew>.
58. Eun Oh, Maxwell Gregoire, Adam Black, et al., "A Perspective on Quantum Sensors from Basic Research to Commercial Applications," (June 2024), *Army Research Laboratory*, <https://arxiv.org/pdf/2407.00689>, 12.
59. James Marinero, "Inertial Navigation – What is It, How Does It Work?" (May 7, 2024). *Medium*, <https://medium.com/the-dock-on-the-bay/inertial-navigation-what-is-it-how-does-it-work-4fel1601943d>.
60. Josh Holder and Constant Meheut, "Facing and Endless Barrage, Ukraine's Air Defenses are Withering," (May 13, 2024). *The New York Times*, <https://www.nytimes.com/interactive/2024/05/13/world/europe/ukraine-missile-defenses.html>; Deagal, "Iskander Capability Overview," (August 15, 2023), <https://deagal.com/Weapons/Iskander/a001012>.
61. Nurettin Sevi, "Quantum Technology in the Defence: A Paradigm Shift," (October 23, 2023). *Defense Domain*, <https://defensedomain.com/quantum-technology-in-the-defence-a-paradigm-shift>.
62. U.S. Department of Defense, *Applications of Quantum Technologies*, 22.
63. Joe Kinast, "Prospects and Priorities for Emerging Quantum Sensors," (September 2022), *Quantum Economic Development Consortium*, <https://quantumconsortium.org/sensing22>, 10.
64. U.S. Government Accounting Office, "Defense Navigation Capabilities: DoD is Developing Position, Navigation, and Timing to Complement GPS," (May 10, 2021). Report. <https://www.gao.gov/products/gao-21-320sp>.
65. Bruce Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, (1996). New York: John Wiley and Sons Inc., 2nd Edition.

NOTES

66. Peter W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer.” (August 4, 2006) *Society for Industrial and Applied Mathematics Review*, Volume 41, <https://inspirehep.net/literature/2743879>, 303-331.
67. Richard Chirgwin. “Quantum Computers Won’t Break RSA Encryption Anytime Soon,” (January 25, 2023). *IT News Online*, <https://www.itnews.com.au/news/quantum-computers-wont-break-rsa-encryption-any-time-soon-590115>.
68. Whitfield Diffie and Martin Hellman, “New Directions in Cryptography.” (November 1976). *IEEE Transactions on Information Theory*, Vol. 22, No. 6, <https://ieeexplore.ieee.org/document/1055638>, 644 – 654.
69. Jose Balcazar, Josep Diaz, and Joachin Gabarro. *Structural Complexity I*. (1998). New York: Springer-Verlag.
70. Edward Parker, “When a Quantum Computer Is Able to Break Our Encryption, It Won’t Be a Secret,” (September 13, 2024). *Lawfare*, <https://www.lawfaremedia.org/article/when-a-quantum-computer-is-able-to-break-our-encryption-it-won-t-be-a-secret>.
71. Lov K. Grover, “A Fast Quantum Mechanical Algorithm for Database Search.” (July 1, 1996). *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, Association for Computing Machinery, New York, 212–219.
72. Albert Carlson, Bhaskar Ghosh, and Indira K. Dutta. “Using the Collision Attack for Breaking Cryptographic Modes.” (October 2022) *IEEE 13th International Congress on Computing, Communication, and Networking Technologies (ICCCNT)*. Kharagpur, India. <https://ieeexplore.ieee.org/document/9984325>, 1-7; David McGrew. “Impossible Plaintext Cryptanalysis and Probable-plaintext Collision Attacks of 64-bit Block Cipher Modes.” (March 2013). *Proceedings of the Fast Software Encryption Workshop*, Singapore, <https://eprint.iacr.org/2012/623>.
73. Bhaskar Ghosh, Indira Dutta, Shivanjali Khare, Albert Carlson, and Michael Totaro, “Isomorphic Cipher Reduction.” (October 2021). *IEEE 12th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, Vancouver, BC, Canada, <https://ieeexplore.ieee.org/document/9623135>.
74. David McGrew, “Impossible Plaintext Cryptanalysis and Probable-plaintext Collision Attacks of 64-bit Block Cipher Modes;” and Bhaskar Ghosh, et al., “Isomorphic Cipher Reduction.”
75. Ron Franklin, “AES vs. RSA Encryption: What Are the Differences?” (November 14, 2022). *Precisely*, <https://www.precisely.com/blog/data-security/aes-vs-rsa-encryption-differences>.
76. Keyfactor. “Harvest Now, Decrypt Later: A New Form of Attack,” (April 29, 2024). <https://www.keyfactor.com/blog/harvest-now-decrypt-later-a-new-form-of-attack/>.
77. Ueli Maurer, “A Universal Test for Random Bit Generators,” (March 1, 1992). *Journal of Cryptography*, Vol. 5, No. 2, <https://dl.acm.org/doi/10.5555/148544>, 89-105.
78. Alexandra Kelley, “Federal Researchers are One Step Closer to Protecting U.S. Data from Quantum Computing Decryption Capabilities,” (July 5, 2022). *NextGov*, <https://www.nextgov.com/cybersecurity/2022/07/nist-identifies-four-quantum-resistant-encryption-algorithms/368954>
79. Emre Karabulut and Aydin Aysu, “Falcon Down: Breaking Falcon Post-Quantum Signature Scheme Through Side-Channel Attacks,” (December 5, 2021), *IEEE Design Automation Conference*, <https://ieeexplore.ieee.org/document/9586131>.
80. Ray Perlner, John Kelsey, and David Cooper, “Breaking Category Five SPHINCS+ with SHA-256,” (September 21, 2022). *International Conference on Post-Quantum Cryptography*, https://link.springer.com/chapter/10.1007/978-3-031-17234-2_23.
81. Kevin Townsend, “AI Helps Crack NIST-Recommended Post-Quantum Encryption Algorithm,” (February 21, 2023). *Security Week*, <https://www.securityweek.com/ai-helps-crack-a-nist-recommended-post-quantum-encryption-algorithm>.
82. Kalle Ngo, “Single-Trace Side-Channel Attacks on CRYSTALS-Dilithium,” (April 10, 2024). Presentation, *The Fifth PQC Standardization Conference*, Rockville, MD: NIST, <https://csrc.nist.gov/Presentations/2024/single-trace-side-channel-attacks>.
83. Mandeep Singh and Albert Carlson, A., “Watermarking Using Polymorphic Algorithms for Increased Data Security,” (May 2024). *IEEE World AI IoT Congress (AllIoT)*, Seattle, WA.
84. Matt Swayne, “White House Advisor Says NIST To Release Post-Quantum Cryptographic Algorithms in Coming Weeks,” (May 24, 2024). *The Quantum Insider*, <https://thequantuminsider.com/2024/05/24/white-house-advisor-says-nist-to-release-post-quantum-cryptographic-algorithms-in-coming-weeks/>.

Data as Ammunition – A New Framework for Information Warfare

Lt. Col. Jessica Dawson, Ph.D.

Col. Katie Matthew, Ph.D.

ABSTRACT

This article presents a new framework for thinking about data and the risks posed to national security. Taking issue with the prevailing analogy of “data as oil,” this article argues that viewing data as ammunition provides a clearer understanding of the real threats and a familiar path toward risk mitigation in the information space. The “data as ammunition” analogy carries a better intuitive depiction of the risk and why, in the days of increasing storage which keeps data easily accessible seemingly forever, categorizing data through the lens of ordnance classifications can help clarify the risks to force and national security. We close this article with recommendations to adapt current privacy, security, and commercial policies to mitigate the new risks to force and personnel on and off the battlefield.

Keywords: Data, Analytics, Publicly Available Information, Micro-Targeting, Information Domain

INTRODUCTION: EXPLAINING INFORMATION WARFARE IN WARFARE TERMS

This article expands on the definition of data and data information in the Army’s newest doctrine ADP 3-13 and advances a new framework for thinking about data and the risk posed to national security. We lay out how and why thinking about how the definitions in ADP 3-13 are a vast improvement in the Army’s doctrinal understanding of data but that defining the interpretation data as “receiver centric” does not go far enough in setting up frameworks for *understanding the impact* of varying interpretations of data. If we categorize data as ammunition, we move it from the academic or doctrinal realm of “meaning of information” and into a more military focused national

This is a work of the U.S. Government and is not subject to copyright protection in the United States. Foreign copyrights may apply.



Lieutenant Colonel Jessica Dawson is a United States Military Academy professor in the Army Cyber Institute and is the Digital Force Protection Research Team Division Chief. She has written numerous articles about the risks of commercial data collection and has advised senior leaders across the federal government on these risks. She holds a Ph.D. in sociology from Duke University.

defense arena – after all, ammunition literally has impact and we argue, so does data. By thinking about data as ammunition, we argue we start to think of the impact of the data on people and society. This new framework is necessary to combat the buzzwordification of data, which prevents deeper understanding and application of the concept and uses an analogy familiar to most military personnel to describe data. The Army does not need leaders who ask for “big data” – they need to understand what problem they are asking data to solve. Ammunition, we argue, is a better framework for understanding what data can do and this shift is necessary to help the Army better understand this major shift in the character of war. Treating data as ammunition reduces the tendency of leaders to store data for data’s sake and instead recognizes the potential impacts on individuals and the force.

BACKGROUND: HOW WAR AND PERSONAL DEVICES CHANGED WEAPONS

When one of the first Americans serving in Iraq died in an explosion on May 26, 2003 – only six weeks after the initial U.S. invasion had “ended” – the U.S. military had no words to describe the type of ordnance that exploded under his vehicle and killed him.¹ The device that killed Private First-Class Jeremiah Smith came to be known as an improvised explosive device (IED) and was not a new concept in warfare. The French Resistance, for example, used IEDs extensively to derail the German supply lines during World War II, and with the massive amount of discarded and leftover ordnance lying around Iraq and the number of unemployed former soldiers who knew how to use it, Iraq became a war of IEDs.

What the U.S. military learned early on was that commercial, off-the-shelf technology could easily be repurposed to make the already deadly IEDs even more precise and more lethal. Iraqi insurgents and technology evolved together, and IEDs and the mechanisms



Colonel Katie E. Matthew, U.S. Army, is a United States Military Academy professor in the Department of Behavioral Sciences and Leadership. She has served as a logistician at home and on deployment with the 1st Theater Sustainment Command, 1st Infantry Division, 82nd Airborne Division, and Joint Special Operations Command, and most recently commanded the Brigade Special Troops Battalion and Camp Buehring, U.S. Army Central Command. She holds a Ph.D. in sociology from George Mason University.

and methods used to detonate them, grew more sophisticated. Over time, IEDs shifted from needing a vehicle or victim to detonate to having wired or remote detonation capabilities – often done by cell phone. Suddenly, once mundane activities – using a cell phone – became threatening, and allied troops had to determine if the civilian seen holding a cell phone was either a lethal threat to the force or someone simply using their phone for its intended purpose – a notion that seemed absurd before cell phones became detonators.

In many ways, the cell phones that millions carry with them continue to have dual potential: a loaded weapon and a useful device to store data. The data that each of us has to hand to aid in our daily lives also opens a vulnerability to the force as a new form of ammunition on and off the modern battlefield.² Just as IEDs were adapted to be able to target different size formations, data can be used to do wide area targeting or precision microtargeting,³ but the end state is the same: soldiers are removed from the fight. We argue that thinking about data as ammunition helps leaders to better understand and comprehend the impact on the force, viewing it as a tool in a broader arsenal.⁴

A BAD ANALOGY: THE LIMITATIONS OF DATA AS OIL

There are many analogies about data in modern society, but none is as pervasive as “data is the new oil.”⁵ Originally coined in *The Economist* in 2017, the analogy describes data as something that can be traded, bartered, owned, or stolen—as something whose value underwrites the entire digital economy. The data as oil analogy conceals many of the risks from the new data economy. ADP 3-13 defines data as “any signal or observation from the environment”⁶ but the broader societal discussion of data is more informative for defining the problem we are addressing. The “data as oil” analogy also works to prevent and limit regulation, as viewing something as a commercial

commodity makes it more difficult to quantify as dangerous and yet, considerations of risk associated with data are not new.⁷ Viewing data as a commodity also falsely leads to the idea that that more data is better; in reality, more data is simply *more*. This analogy also creates difficulties for privacy advocates who struggle to demonstrate how data can be weaponized against individuals⁸ when their arguments are outweighed by the commercial interest of those that depend on data for financial interests, which in turn influence the national conversation around the risks from data. Army leaders are not immune to absorbing these ideas and this makes it more difficult for them to conceptualize data as something that can be used against them or the force. It is difficult for everyday people to understand how something they do on their phone can be turned into data that can be used against them. This is also where one of several major myths about data and privacy from the commercial surveillance economy come into play: if you have nothing to hide, what are you worried about?

Ammunition, like data, can be used for good things like providing food or protection. Ammunition can also be destructive, just as the cell phone in Iraq could be a completely harmless device or a detonator for a mass casualty event. All military equipment and ammunition are classified according to a federal supply classification (FSC).⁹ Most ammunition, apart from nuclear associated ammunition, is classified as FSC Group 13 and it is categorized according to the size of the projectile and stored according to its potential risks. We argue that moving data from the generic conception of information to categorizing it in terms of its effects would significantly move the discussion on how to protect it more effectively both from a DoD perspective and a broader national security perspective.

From a practitioner perspective, the “data as ammunition” analogy carries a better intuitive depiction of the risk from varying interpretations of data, particularly for those who are less steeped in data expertise. In the days of increasing storage which keeps data easily accessible seemingly forever (though this is currently in doubt),¹⁰ new ways of thinking about data are necessary. For example, thinking of data as a minefield subject to algorithmic overwatch helps better define the risks, not only to individuals but to the force and national security writ large. Categorizing data through the lens of ordnance classifications (i.e. different types of ammunition) can help clarify and quantify the risks to force and national security, both for staff and commanders as well as other decision makers such as contracting officers. If the use of data comes from the interpretation of it, then decision-makers must be aware of how different interpretations of the data can be used to attack red or attack blue – just like ammunition.

A BETTER (NOT ABSOLUTE) ANALOGY: DATA AS AMMUNITION

Why do we argue that we should use an analogy of data as ammunition? Analogies can be useful in thinking about complicated topics and, there are not many more confusing topics – for a variety of reasons – than data. In the rush to obtain more data, industry leaders have focused on the potential opportunities and paid less attention to the risks. Treating data as

ammunition through its dual potential better prepares military leaders and lawmakers on how to safeguard, manage, and appropriately utilize data.

Storage of ammunition is based on the caliber and potential damage. Where data is similar is that very innocuous data can be combined and enriched to create much more powerful impacts on individuals or at scale. We currently segregate many things based on the risk they create when they are combined. We store bleach and ammonia separately. We segregate knowledge to those who have a need to know. Data should be no different – leaders need to know what data they need to answer particular questions. Simply demanding more data is like demanding more ammunition without considering what you are trying to use said ammunition for.

FROM DATA BUZZ WORDS TO DATA LITERACY

Data, like ammunition, is meaningless without the proper tools and interpretations. Analytics or interpretations of the data are necessary to get any use out of data. Information warfare and information advantage imply putting data into use as tools. It is relatively easy to understand how fuel and ammunition impact war – you can see the fuel gauge on your vehicle ticking down as you run out of gas, and you can calculate ammunition burn rates. You can see the cell phone battery drain but we lack the fundamental understanding or tracking of how much data is moving through the signals of our interconnectedness devices. Without gauges and a fundamental understanding of how much data is pushed (and pulled) from our devices, it is easy to forget or ignore the potential dangers of the free movement of data. We talk about storing data on the cloud – something fluffy and far away instead of calling it what it is: other people's computers. Putting data on someone else's computer implies a lot more risk than a cloud.

We tell commanders they must understand the operational environment but how do you train leaders to understand the types of data that are available on your unit when you don't know what questions to ask? This lack of understanding is not willful ignorance but rather a lack of literacy on both what data is, how it is generated from everyday activities such as buying medication or getting directions, and how we actually utilize data in daily and military life. It is less clear how quickly different types of data be used or exploited on the modern battlefield or arguably more importantly, prior to open conflict because its interpretation is viewed as available to the purveyors of the dark arts who can understand it and hopefully explain the impact of it. The data as ammunition analogy (re)educates leaders to examine data by potential risk and puts data management in the more familiar territory of risk management.

Physical geography influences our understanding of our equipment limitations and resupply capabilities and the mapping of both threat and advantage in the physical space is clearly understood. Situation reports and concepts of the operation both list strengths in supply and personnel and risk levels for the mission. Less clear is how to categorize and understand how different types of data influence conflict and competition. Large-scale discussions over

resource control around oil-rich areas are easier to visualize and understand than datasets residing on a server somewhere. It is precisely because of the unknown and unclassified vulnerabilities in the vast amounts of available data that a new framework is needed for conceptualizing the risk that data poses to manage that risk.

Rethinking and remapping the digital space as lines of supply/communication on the battlefield brings an awareness that there is a risk to the force if they are not secured in the same way we secure our physical lines of communication. Digital lines of communication are more porous and vulnerable, made further so by the easily accessible data carried on our personal devices. We go to great lengths to ensure ammunition is not easily accessible; how much riskier is it to ignore the easy access of personal data about our personnel? It is difficult to understand the vulnerability that the commercial and unregulated purchase of data poses when we think of it only as individual bits, but as any soldier who’s had to police the wood line for a round of ammunition will tell you, the potential threat that one unaccounted for round poses is real. Those individual rounds of data do not pose threats in and of themselves, but the potential weaponization through analytics ought to grant us pause in the same way we would comb the wood line for a round of ammunition.

UXO: UNEXPLODED ORDNANCE OR THE UNKNOWN POTENTIAL OF DATA

As shown in Figure 1, ammunition is typically classified by explosive type or family and moves from small arms for weapons such as an M4/AR15, through larger caliber weapons such as 125mm rounds to grenades, rockets through land mines, and larger explosives.¹¹ Important to this analogy, the different types of ammunition are typically used for different targets. Small arms rounds are used for smaller objects such as individual enemy soldiers whereas land mines, for example, are area obstacles meant to funnel personnel toward specific locations. Ammunition is further stored by hazard category and even smaller, less damaging rounds are segregated from other more combustible munitions. Similarly basic data, on its own, is not dangerous but when combined with analytics, can reveal powerful and potentially damaging insights.

FSC Group 13 Classes	
FSC Group 13 (Classes)	Ammunition and Explosive Type or Family
1305	Ammunition, through 30mm
1310	Ammunition, over 30mm up to 75mm
1315	Ammunition 75mm through 125mm
1320	Ammunition, over 125mm
1330	Grenades
1340	Rockets and rocket ammunition
1345	Land mines
1365	Military chemical agents
1370	Pyrotechnics
1375	Demolition materials
1376	Bulk explosives
1377	Cartridge and propellant activated devices and components
1390	Fuzes and primers
1395	Miscellaneous ammunition
1398	Specialized ammunition handling and service equipment
1410/20/25/27	Guided missiles

Note: There are other FSC groups, but they are for Class V materiel outside the U.S. Army ammunition inventory. (Look in any current copy of the DOD ammunition listing, volumes 1 through 3, for more information.)

Table 1. Table F-1, Appendix F (Department of the Army 2001a).

Data functions in a similar way to ammunition in the way that ammunition, when combined with a weapons system, creates different levels of destruction. Data is not inherently dangerous in and of itself and can be used for many different types of activities, some of which are objectively good, like using data to train machine learning algorithms to try to identify earlier means of identifying cancer. But thinking about data as ammunition can help conceptualize the potential risk it poses by tying that data back to the real people it represents rather than in an abstract form.

Consider microtargeting, a key piece of the information warfare battlefield¹² – which is the ability to develop a unique profile of an individual to ensure that advertisements are tailored to their specific interests.¹³ Advertisers used to say that only 50% of advertising worked – they just did not know which 50%. Not all microtargeting is bad; arguably it can be used to help people find the things they are already looking for. Marketers know, for example, that it takes up to seven engagements with an advertisement before someone is moved to make a purchase. Microtargeting allows tailored advertisements to probe multiple times to get those seven engagements, including how often the target simply pauses over the ad. Microtargeting promises advertisers efficiency and effectiveness because it gathers significant amounts of data to promise accuracy. Microtargeted advertising is what enabled Meta/Facebook to dominate advertising for the last decade – the ability to push ads to people who may be interested in your product based on data that Meta/Facebook collected from users.¹⁴

Like ammunition storage units that house potential devastation wrought by one round igniting the others around it, a data storage unit contains the potential for devastating effects. Data in and of itself can be inert, but when massed (enriched) or honed for a specific purpose (analytics) it can wreak havoc and render targets ineffective by taking them out of the fight. There are hundreds of stories of data being used to target people at their most vulnerable. From gambling¹⁵ to divorce¹⁶ to addiction,¹⁷ data can identify patterns or pattern changes in life that indicate opportunities for attack and exploitation. Collated data across these lived patterns further connects individuals to others, with the result of not only a precision round designed for the individual but also a cluster of these rounds that can damage or destroy important social relationships.

The same data that allows an advertising algorithm to help a shoe company find people in the market for shoes also allows a con artist to find people particularly vulnerable to scams and fraud or worse.¹⁸ Amazon recently came under fire for its algorithm recommending a common food preservative along with other products that could be used for suicide.¹⁹ Sodium nitrate is also widely used in food preservation– but the algorithm was not trained on data that enabled it to distinguish between someone looking for it for a legitimate purpose compared to one intended for self-harm. Just as we do not store dangerous products in the vicinity of each other, nor should we accept storing dangerous data in the same proverbial closet.

MODERN MINEFIELDS

If we accept the analogy that data is ammunition, then the analogy for the surveillance economy – everything from apps to internet usage to location data - is a minefield. A minefield can both injure unwitting individuals who step into it and funnel people and forces to places advantageous to the opposing force when known. In marketing, funnels are used to scan for people who might be interested in a product and use a series of gates to draw them further into the sales funnel until they complete a purchase. Think of companies that use sale items to get people into the store. Once they are in the store, the stores are designed in specific ways to increase the likelihood of additional purchases.²⁰ Minefields in real life function in a similar way, funneling people towards or away from an objective. In marketing, data provides the targeting information used to identify potential targets. In physical minefields, mines can be indiscriminate—injuring or killing anyone who happens by—or they can wait until just the right time to detonate on a specific target as happened in the Grozny Stadium bombing in 2004 where a bomb was built into the stadium during construction.²¹ Data can be gathered and analyzed using machine learning to determine the greatest vulnerabilities and, similar to a minefield, they can be used to broadly target, or they can be set to target a specific individual at a specific time.

Having a single tweet or video go viral can be similar to the effect of stepping on a mine – and not only a virtual mine but one with real-life consequences.²² Even more data is collected and aggregated on seemingly harmless and innocuous everyday life making it difficult to see the potential dangers. But if Target knows that someone is buying newborn diapers, then that data, collated with other data such as proximity to a military base can now inform an adversarial party that there is a high likelihood of an infant in a servicemember's home.²³ The physical boundaries of a kinetic fight are no longer the only battlefield concerns in the information fight but now cross war-time boundaries as well. Children of politicians have become national news stories overnight because of their social media posts or because the content of their private data lives has been exposed to the world. Targeting someone's family to influence their actions is not a new tactic – the ease with which the surveillance economy enables this targeting now is what's changed.²⁴

DEEPLY BURIED DATA IEDS

One problem with the data as ammunition analogy is that unlike physical ammunition, which is degraded over time by the elements and the bounds of physics, data does not always suffer the same limitations. Data, in all its forms, is collected and stored at an increasing rate, and through the advances in algorithms and artificial intelligence, collated on a scale unimaginable even twenty years ago.²⁵ While much of that data gathering can be used for individual and collective good—reducing shopping times and finding causes of health concerns—it can also have devastating effects. Photographs are records and digital copies of photographs can spread around the world in an instant for essentially zero cost and be collected

and stored by companies building facial recognition algorithms with no legal recourse for the person associated with said face.²⁶ Data as ammunition can be loaded into the weapon at the time and place of an adversary's choosing – meaning something you may have long forgotten about may suddenly be brought to the forefront, efficiently (re)using ammunition specifically tailored for you.²⁷ Another limitation of the data as ammunition analogy is that most of the data that would need to be categorized along these lines is not owned or controlled by the government – it is collected and stored by commercial entities.²⁸ However, there is a tendency to ignore or downplay the risks from ubiquitous technical surveillance. Dismissing the problem because “they already have everything” is like dismissing the bleeding wound. We don't dismiss risk to force or risks to mission in any other context. This, however, makes the analogy more important – we cannot ignore the risks commercial data and analytics pose to the force simply because we dismiss it as marketing data or something else harmless. Thinking about all data collection – commercial or otherwise – as potentially adversarial would go a long way to expanding how people think about data and risks to the force.

Data, like ammunition, is generally not as useful on its own as when combined with analytics. To go back to the ADP 3-13, data is interpreted by the receiver. In this case, we argue that the meaning of data is created not only through its collection but also through analytics, or for what it can be used. Data is inherently biased beginning with the decision to collect what by whom and when. The meaning it takes on is created the moment it is transcribed or captured and further modified depending on what it is used for.²⁹ If data is the ammunition, who has access to the data and what analytics they can run are the weapon system. There has been reporting for more than a decade about the outrage generated when the public becomes aware of someone having access to data they were not supposed to. The scale and scope of the data should be measured by its potential and access to it should be controlled just like we control access to ammunition. In much the same way, we have procedures when there has been a breach of ammunition storage or when ammunition is improperly accounted for, but we do not have procedures for notification when commercial data is legally shared/transferred.³⁰ We tend to pay attention to people who have large legitimate stockpiles of ammunition, yet we don't blink at companies stockpiling data.

Likewise, we ignore the fine print in the terms and conditions of apps we use for personal activities and then feel exposed when we find out that the same data is available to a third party for a price. The U.S. government is precluded from peering into citizen's lives without probable cause and yet commercial entities can collect, aggregate, sell, and transfer deeply personal data without considering the consequences to individuals.³¹ Why wouldn't America's adversaries be acquiring this data for use with their analytics?

DIGITAL IEDS REQUIRE NEW KINDS OF FORCE PROTECTION

In Army doctrine, defense sets conditions for the offense. Currently, there is limited to no defense for servicemembers and other members of the total force from commercial data

collection. Around the world, data and analytics are being used to speed up targeting while not necessarily ensuring accuracy.³² People have lost jobs, divorced, been doxed, and harassed because a piece of data hit at exactly the right moment for malicious actors to take advantage of the algorithms deciding a story was worth amplifying.³³ If these things are happening to civilians, why would we think that our nation's adversaries are not considering how to deploy these same tools during periods of heightened competition or conflict? Why wouldn't an adversary leverage the ability to rapidly employ data through analytics to target people in their homes? This capability has already been deployed against the Uighur population and people who have been critical of China abroad.³⁴ It is now possible to leverage rapid analytics to identify people's homes, families, and routines to target them individually at scale to reduce the ability of the military to muster forces.³⁵ This is the equivalent technological advancement of the machine gun in World War I, by increasing the ability to put rounds downrange. Leveraging digital ammunition in non-conflict spaces could be dramatically more effective in impacting readiness levels than waiting until forces are massed on a battlefield.³⁶ These "wounds" could render the targets just as combat ineffective as the loss of soldiers and equipment.

Using this theory, we argue that we do not have to start from scratch for policy recommendations. We argue that given the framework of data as ammunition, we can then start adapting existing privacy policies, compartmentalization procedures, and data use policies to clearly articulate who should have access to what data and for what purposes, including restricting access by commercial entities to collect on servicemembers where allowed by law. Current privacy policies should include requirements for commercial entities to notify individuals before the sale or sharing of their data with a third party, rather than a blanket statement in the terms and conditions. Buying in bulk of certain items, such as fertilizer, triggers commercial and governmental algorithms. A farmer has a right to access and need for bulk fertilizer; a similar adaptation could exist for acquiring data in bulk: demonstrated need to know. This is part of the same language we use for access to classified information: cleared to know it and need to know.

Understanding the impact of data is important for classifying that data. In April 2024, the White House expanded on Executive Order 13873 of May 15, 2019, and limited the sale of bulk personal data to countries of concern. This executive order defined sensitive data as "covered personal identifiers, geolocation, and related sensor data, biometric identifiers, human'omic data (meaning data that characterizes various elements of individual human biology), personal health data, personal financial data, or any combination thereof, as further defined in regulations issued by the Attorney General"³⁷ which is to be restricted from sale to certain countries of concern. While there is still more to be done, it is significant that the executive branch is recognizing that some of the data readily available to us companies may pose a serious risk to the U.S. in the wrong hands.³⁸ For decades, service members used their Social Security Number for nearly every aspect of their daily military lives – which also linked them to many aspects of their civilian lives such as credit reporting, medical information, and cell

phone numbers.³⁹ Understanding the legitimate risks of having one government identification number available to so many agencies caused the DOD to create a DOD ID number.⁴⁰ Using information to identify a person for more than one purpose is convenient, but it is also dangerous and opens us up to individual, or organizational, attacks. This is also why, though frustrating at times, creating unique passwords for different accounts does provide a level of protection and limit access should a breach of one account database occur. By thinking of data as ammunition, we argue that we can require both universal safety for bulk storage and practical safety for individual storage as risk mitigation.

Information warfare is a new form of contact and data is the technological advancement of this changing character of modern war—it can be purchased commercially off-the-shelf, weaponized, stored long-term, and reused in new combinations to exploit vulnerabilities both long before forces ever begin to mobilize and on the battlefield. It is efficient and can be highly effective because of the low cost, durability, and real-time customization it provides combatants. Unlike the machines of war that are the domain of state actors, data and associated analytics are available widely and have a low barrier to entry – everyday activity can inadvertently provide ammunition to adversarial actors, from pop up social media games to posting pictures of family members on significant dates. Further, even if people want to protect their data, there are limited opportunities for service members or anyone without means to hire an army of attorneys to meaningfully opt out of the surveillance economy minefield.⁴¹ Data is already ammunition purchased in broad daylight and on the black market—by friend and foe alike.

Data is the next evolution in munitions.⁴² Just as we develop Army leaders to understand how to employ specific weapons platforms and mass fires for effects, we need to start developing leaders who understand how data and analytics can be leveraged both for force protection as well as offensive actions. It is foolhardy to continue to talk about being data-centric without developing leaders who understand the risks data and analytics pose to the force. It can and is being weaponized against our forces. Our would-be adversaries are already in the market and the Army would be foolish to ignore the ways that data they cannot control can become powerful weapons systems to use against their own forces.🛡️

DISCLAIMER

The views expressed in this work are those of the authors and do not reflect the official policy or position of the United States Military Academy, the Department of the Army, or the Department of Defense.

NOTES

1. Gregg Zoroya, "How the IED Changed the U.S. Military," (December 18, 2013). USA TODAY, <https://www.usatoday.com/story/news/nation/2013/12/18/ied-10-years-blast-wounds-amputations/3803017/>.
2. Jessica Dawson, "Who Controls the Code, Controls the System: Algorithmically Amplified Bullshit, Social Inequality, and the Ubiquitous Surveillance of Everyday Life," (September 2023) *Sociological Forum*, Vol. 38, No. S1, <https://doi.org/10.1111/socf.12907>.
3. Jessica Dawson, "Microtargeting as Information Warfare," *The Cyber Defense Review* (February 2021): <https://doi.org/10.31235/osf.io/5wzuq>, 63-80.
4. Jessica Dawson and Tarah Wheeler, "How to Tackle the Data Collection behind China's AI Ambitions," (April 29, 2022) *Brookings Institute* (Online), <https://www.brookings.edu/techstream/how-to-tackle-the-data-collection-behind-chinas-ai-ambitions/>.
5. "The World's Most Valuable Resource Is No Longer Oil, but Data," (May 6, 2017). *The Economist*, <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data>.
6. U.S. Department of the Army, *Information*, ADP 3-13. (November 2023) Washington, D.C.: Department of the Army, https://armypubs.army.mil/epubs/DR_pubs/DR_a/ARN39736-ADP_3-13-000-WEB-1.pdf, 1-1.
7. Jer Thorp, *Living in Data: A Citizen's Guide to a Better Information Future* (May 4, 2021). New York: MCD; Mary F. E. Ebeling, *Healthcare and Big Data: Digital Specters and Phantom Objects*, 1st ed., (September 28, 2016). New York: Palgrave Macmillan; Mary F. E. Ebeling, *Afterlives of Data: Life and Debt Under Capitalist Surveillance*, 1st ed. (June 14, 2022). Oakland, California: University of California Press, <https://bookshop.org/p/books/afterlives-of-data-life-and-debt-under-capitalist-surveillance-mary-f-e-ebeling/17505478>.
8. Daniel J. Solove and Paul M. Schwartz, "ALI Data Privacy: Overview and Black Letter Text," (February 23, 2022). 68 *UCLA Law Review*, accessed in SSRN, <https://doi.org/10.2139/ssrn.3457563>, 1252-1300.
9. U.S. Department of the Army, Supply Bulletin SB 708-21: Federal Supply Classification (May 2005). Battle Creek, MI: Defense Logistics Information Service, https://gsaccess.gov/html/userguides/DRMS_H2_FSC_Guide.pdf, 4-7.
10. Drew Austin, "The End of Infinite Data Storage Can Set You Free," (March 16, 2022). *Wired*, <https://www.wired.com/story/cloud-computing-space-data-storage-habits/>.
11. U.S. Department of the Army, *Ammunition Handbook: Tactics, Techniques, and Procedures for Munitions Handlers*, FM 4-30.13 (FM 9-13), (2001). Washington, D.C.: Department of the Army.
12. Dawson, 2021; Christopher Wylie, *Mindf*ck: Cambridge Analytica and the Plot to Break America* (October 8, 2019). New York: Random House; Laurie Clarke, "How the Cambridge Analytica Scandal Unravelled," (October 15, 2020). *The New Statesman*, <https://www.newstatesman.com/science-tech/social-media/2020/10/how-cambridge-analytica-scandal-unravelled>.
13. Byron Tau, "How the Pentagon Learned to Use Targeted Ads to Find Its Targets—and Vladimir Putin," (February 27, 2024). *Wired*, <https://www.wired.com/story/how-pentagon-learned-targeted-ads-to-find-targets-and-vladimir-putin/>; Almog Simchon, Matthew Edwards, and Stephan Lewandowsky, "The Persuasive Effects of Political Microtargeting in the Age of Generative AI," (January 29, 2024). *PNAS Nexus* Vol. 3, Issue 2, <https://doi.org/10.1093/pnasnexus/pgae035>; Gus Lubin, "The Incredible Story of How Target Exposed a Teen Girl's Pregnancy," (February 16, 2012). *Business Insider*, <https://www.businessinsider.com/the-incredible-story-of-how-target-exposed-a-teen-girls-pregnancy-2012-2>; Kashmir Hill, "How Target Figured Out a Teen Girl Was Pregnant Before Her Father Did," (February 16, 2012). *Forbes*, <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/>.
14. Paul Hitlin, Lee Rainie, and Kenneth Olmstead, "Facebook Algorithms and Personal Data," (January 16, 2019). *Pew Research Center: Internet, Science & Tech* (Report), <https://www.pewresearch.org/internet/2019/01/16/facebook-algorithms-and-personal-data/>; Shoshana Zuboff, "Surveillance Capitalism or Democracy? The Death Match of Institutional Orders and the Politics of Knowledge in Our Information Civilization," (July-September 2022). *Organization Theory* Vol. 3, Issue 3, <https://journals.sagepub.com/doi/epub/10.1177/26317877221129290>; Francis Haugen, "Testimony of Frances Haugen to the January 6th Committee," (December 17, 2021). Transcript, U.S. House of Representatives Select Committee to Investigate the on January 6th Attack on the U.S. Capital, Washington, D.C., <https://archive.org/details/january-6th-committee-witness-testimony-20211217-frances-haugen>.
15. Wolfie Christl, "Digital Profiling in the Online Gambling Industry," (January 2022). *Institute for Critical Digital Culture*, <https://crackedlabs.org/en/gambling-data>.

NOTES

16. Rob Mudge, "US Religious Data Platform Targets Vulnerable People," (September 27, 2020). Germany: Deutsche Welle, Dw.Com, <https://www.dw.com/en/us-religious-data-platform-targets-mentally-ill-vulnerable-people/a-55062013>; Kendra Barnett, "The Repeal of Roe v Wade Will Impact Our Privacy Landscape at Large, Experts Predict," (June 24, 2022). London: *The Drum*, <https://www.thedrum.com/news/2022/06/24/the-repeal-roe-v-wade-will-impact-our-privacy-landscape-large-experts-predict>.
17. Juliana Gruenwald Henderson, "Alcohol Addiction Treatment Firm Will Be Banned from Disclosing Health Data for Advertising to Settle FTC Charges That It Shared Data Without Consent," (April 11, 2024). U.S. Federal Trade Commission, Press Release, <https://www.ftc.gov/news-events/news/press-releases/2024/04/alcohol-addiction-treatment-firm-will-be-banned-disclosing-health-data-advertising-settle-ftc>.
18. Alistair Simmons and Justin Sherman, "Data Brokers, Elder Fraud, and Justice Department Investigations," (July 25, 2022). *The Lawfare Institute*, <https://www.lawfaremedia.org/article/data-brokers-elder-fraud-and-justice-department-investigations>.
19. Megan Twohey and Gabriel J. X. Dance, "Lawmakers Press Amazon on Sales of Chemical Used in Suicides," (February 4, 2022). *The New York Times*, Technology Section, <https://www.nytimes.com/2022/02/04/technology/amazon-sui-cide-poison-preservative.html>.
20. Romain Dillet, "Google and Mastercard Reportedly Partner to Track Offline Purchases," (August 31, 2018). *TechCrunch Online*, <https://techcrunch.com/2018/08/31/google-and-mastercard-reportedly-partner-to-track-offline-purchases/>; Minna Ruckenstein and Julia Granroth, "Algorithms, Advertising and the Intimacy of Surveillance," (February 11, 2019). *Journal of Cultural Economy Vol. 13, No. 1*, 12–24. <https://doi.org/10.1080/17530350.2019.1574866>; Hirsh Chitkara, "How Lax Social Media Policies Help Fuel a Prescription Drug Boom," (July 6, 2022). Protocol, <https://www.protocol.com/policy/meta-tiktok-prescription-drug-ads>; Eduardo Schnadower Mustri, Idris Adjerid, and Alessandro Acquisti, "Behavioral Advertising and Consumer Welfare: An Empirical Investigation," (April 3, 2024). SSRN Scholarly Paper, Rochester, NY, <https://doi.org/10.2139/ssrn.4398428>.
21. "2004 Grozny Stadium Bombing," NGTC Group Training (blog), United Kingdom, accessed April 1, 2022, <https://www.ngtc.co.uk/2004-grozny-stadium-bombing/>.
22. Davis Winkie, "Pat Donahoe, Civilian, Wants a Word with the Army," (January 5, 2022). *Army Times*, <https://www.armytimes.com/news/your-army/2023/01/05/pat-donahoe-civilian-wants-a-word-with-the-army/>.
23. Lubin, 2021.; Jon Keegan and Joel Eastwood, "From 'Heavy Purchasers' of Pregnancy Tests to the Depression-Prone: We Found 650,000 Ways Advertisers Label You," (June 8, 2023). The Markup, *The Markup*, <https://themarkup.org/privacy/2023/06/08/from-heavy-purchasers-of-pregnancy-tests-to-the-depression-prone-we-found-650000-ways-advertisers-label-you>; Anya E. R. Prince, "I Tried to Keep My Pregnancy Secret," (October 10, 2022). *The Atlantic*, <https://www.theatlantic.com/ideas/archive/2022/10/can-you-hide-your-pregnancy-era-big-data/671692/>.
24. Yohannes Lowe, Kevin Rawlinson, and Warren Murray, "Wife of Ukrainian Military's Top Intelligence Official Poisoned; Stoltenberg Urges Nato to 'Stay the Course' – as It Happened," *The Guardian*, November 28, 2023, sec. World news, <https://www.theguardian.com/world/live/2023/nov/28/russia-ukraine-war-live-enemy-attempts-to-storm-av-diiivka-from-all-directions-says-ukrainian-official>.
25. Justin Brookman, "Understanding the Scope of Data Collection by Major Technology Platforms," (May 2020). *Consumer Reports*, https://innovation.consumerreports.org/wp-content/uploads/2023/05/Understanding-the-scope-of-data-collection-by-major-platforms_2020_FINAL.pdf
26. Kashmir Hill, "A Face Search Engine Anyone Can Use Is Alarmingly Accurate," (May 26, 2022). *The New York Times*, Technology Section, <https://www.nytimes.com/2022/05/26/technology/pimeyes-facial-recognition-search.html>.
27. Kashmir Hill, "A Vast Web of Vengeance," (January 30, 2021). *The New York Times*, Technology Section, <https://www.nytimes.com/2021/01/30/technology/change-my-google-results.html>.
28. Ronan Farrow, "How Democracies Spy on Their Citizens," (April 18, 2022). *The New Yorker*, <https://www.newyorker.com/magazine/2022/04/25/how-democracies-spy-on-their-citizens>; Finn Myrstad and Ingvar Tjostheim, "Time to Ban Surveillance-Based Advertising: The Case Against Commercial Surveillance Online," (June 22, 2021). Oslo, Norway: Forbruker Rådet, <https://www.forbrukerradet.no/wp-content/uploads/2021/06/20210622-final-report-time-to-ban-surveillance-based-advertising.pdf>; Byron Tau, Andrew Mollica, Patience Haggin, and Dustin Volz, "How Ads on Your Phone Can Aid Government Surveillance," (October 13, 2023). *The Wall Street Journal*, <https://www.wsj.com/tech/cybersecurity/how-ads-on-your-phone-can-aid-government-surveillance-943bde04>; Dell Cameron, "The US Is Openly Stockpiling Dirt on All Its Citizens," (June 12, 2023). *Wired*, <https://www.wired.com/story/odni-commercially-available-information-report/>.

NOTES

29. Ebeling, *Afterlives of Data: Life and Debt Under Capitalist Surveillance*.
30. Hill, "A Face Search Engine Anyone Can Use Is Alarming Accurate."
31. Tau "How the Pentagon Learned to Use Targeted Ads to Find Its Targets—and Vladimir Putin;" Byron Tau, *Means of Control: How the Hidden Alliance of Tech and Government Is Creating a New American Surveillance State*, (February 27, 2024). New York: Crown; Hill, "A Face Search Engine Anyone Can Use Is Alarming Accurate;" Kashmir Hill and Corey Kilgannon, "Madison Square Garden Uses Facial Recognition to Ban Its Owner's Enemies," (December 22, 2022). *The New York Times*, New York Section, <https://www.nytimes.com/2022/12/22/nyregion/madison-square-garden-facial-recognition.html>.
32. Melissa Heikkila, "Dutch Scandal Serves as a Warning for Europe over Risks of Using Algorithms," (March 30, 2022). *POLITICO, Europe Edition*: <https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/>.
33. Kelly Kennedy, "Short Changed: No Justice at Fort Hood for Another Woman Soldier and a Commander Who Tried to Help," (December 14, 2023). *The War Horse* (blog), <https://thewarhorse.org/fort-hood-commander-punished-after-reporting-sexual-harasser/>.
34. Elias Groll, "Chinese Hackers Target Family Members to Surveil Hard Targets," (March 26, 2024). *CyberScoop Online*, <https://cyberscoop.com/china-hacking-family-members/>; National Counterintelligence and Security Center, *China's Collection of Genomic and Other Healthcare Data from America: Risks to Privacy and U.S. Economic and National Security*, (February 2021). Washington, D.C.: National Counterintelligence and Security Center, https://www.dni.gov/files/NCSC/documents/SafeguardingOurFuture/NCSC_China_Genomics_Fact_Sheet_2021revision20210203.pdf.; Albert Zhang and Danielle Cave, "Smart Asian Women Are the New Targets of CCP Global Online Repression," (June 3, 2022). *The Strategist*, Australian Strategic Policy Institute, <https://www.aspistrategist.org.au/smart-asian-women-are-the-new-targets-of-ccp-global-online-repression/>.
35. Madhumita Murgia and Max Harlow, "Who's Using Your Face? The Ugly Truth about Facial Recognition," (September 18, 2019). United Kingdom: *Financial Times*, <https://www.ft.com/content/cf19b956-60a2-11e9-b285-3acd5d43599e>; Paul Lewis, "'I Was Shocked It Was So Easy': Meet the Professor Who Says Facial Recognition Can Tell If You're Gay," (July 7, 2018). *The Guardian*, <https://www.theguardian.com/technology/2018/jul/07/artificial-intelligence-can-tell-your-sexuality-politics-surveillance-paul-lewis>.; Ryan Mac, Caroline Haskins, and Logan McDonald, "Clearview's Facial Recognition App Has Been Used By The Justice Department, ICE, Macy's, Walmart, And The NBA," (February 28, 2020). BuzzFeed News, <https://www.buzzfeednews.com/article/ryanmac/clearview-ai-fbi-ice-global-law-enforcement>; Hill, "A Face Search Engine Anyone Can Use Is Alarming Accurate;" Hill and Kilgannon, "Madison Square Garden Uses Facial Recognition to Ban Its Owner's Enemies."
36. Tim Kane. *Bleeding Talent: How the US Military Mismanages Great Leaders and Why It's Time for a Revolution*. (January 4, 2013). New York: Palgrave Macmillan; Gareth Corfield, "Why LinkedIn Is a Snooper's Paradise," (August 24, 2023). *The Telegraph*, <https://www.telegraph.co.uk/business/2023/08/24/how-linkedin-became-a-hotbed-for-spying/>.
37. President Joseph R. Biden Jr., Executive Order 13873, *Executive Order on Preventing Access to Americans' Bulk Sensitive Personal Data and United States Government-Related Data by Countries of Concern*, (February 28, 2024). The White House, <https://www.whitehouse.gov/briefing-room/presidential-actions/2024/02/28/executive-order-on-preventing-access-to-americans-bulk-sensitive-personal-data-and-united-states-government-related-data-by-countries-of-concern/>.
38. Dan Lamothe, "China's Military Seeks to Exploit U.S. Troops, Veterans, General Warns," (September 9, 2023). *The Washington Post*, <https://www.washingtonpost.com/national-security/2023/09/08/china-military-exploit-us-troops-veterans/>.
39. Jim Routh, "The Emergence and Implications of Unconventional Security Controls," (July 1, 2017). *The Cyber Defense Review Vol. 2, No. 2*, https://cyberdefensereview.army.mil/Portals/6/Documents/CDR%20Journal%20Articles/The%20Emergence%20and%20Implications%20of%20Unconventional%20Security%20Controls_Routh.pdf?ver=2018-07-31-093712-843,35-44.
40. U.S. Army Training and Doctrine Command, "Information Paper: Department of Defense Identification Number," (October 2020). <https://www.tradoc.army.mil/wp-content/uploads/2020/10/DoD-ID-Number-PII-Policy-2.pdf>.

NOTES

41. Steven Melendez and Alex Pasternack, "Here Are the Data Brokers Quietly Buying and Selling Your Personal Information," (March 2, 2019). *Fast Company*, <https://www.fastcompany.com/90310803/here-are-the-data-brokers-quietly-buying-and-selling-your-personal-information.>; Kashmir Hill, "General Motors Quits Sharing Driving Behavior with Data Brokers," (March 22, 2024). *The New York Times*, Technology Section, <https://www.nytimes.com/2024/03/22/technology/gm-onstar-driver-data.html>; Justin Sherman, Hayley Barton, Aden Klein, Brady Kruse, and Anushka Srinivasan, "Data Brokers and the Sale of Data on U.S. Military Personnel," (November 2023), Duke University, Sanford School of Public Policy, <https://techpolicy.sanford.duke.edu/wp-content/uploads/sites/4/2023/11/Sherman-et-al-2023-Data-Brokers-and-the-Sale-of-Data-on-US-Military-Personnel.pdf>.; Henrik Twetman and Gundars Bergmanis-Korats, "Data Brokers and Security," (January 14, 2021). NATO Strategic Communications Centre of Excellence, <https://stratcomcoe.org/publications/data-brokers-and-security/17>; Ari B. Friedman, Raina M. Merchant, Amey Maley, Karim Farhat, Kristen Smith, Jackson Felkins, Rachel E. Gonzales, Lujo Bauer, and Matthew S. McCoy, "Widespread Third-Party Tracking On Hospital Websites Poses Privacy Risks For Patients And Legal Liability For Hospitals," (April 2023). *Health Affairs Vol. 42, No. 4*, 508–15, <https://doi.org/10.1377/hlthaff.2022.01205>; Gennie Gebhart and Josh Richman, "Science Shouldn't Give Data Brokers Cover for Stealing Your Privacy," (June 12, 2023). *Scientific American*, <https://www.scientificamerican.com/article/science-shouldnt-give-data-brokers-cover-for-stealing-your-privacy/>; U.S. House of Representatives Committee on Energy and Commerce, "Letter to Mr. Chat Engelgau, CEO Acxiom LLC," (May 10, 2023), https://d1dth6e84htgma.cloudfront.net/05_10_2023_Acxiom_Data_Brokers_Letter_lcb81da32.pdf?updated_at=2023-05-10T16:19:56.031Z.; Lauren Feiner, "Lawmakers Press Companies That Collect U.S. Consumer Data to Reveal How They Buy and Sell It," (May 10, 2023). CNBC, <https://www.cnbc.com/2023/05/10/lawmakers-press-data-brokers-to-reveal-how-they-buy-sell-information.html>.
42. Kehrt, Sonner, "Troops, Veterans Are Targets in the Disinformation War, Even If They Don't Know It Yet." (September 1, 2022). *The War Horse* (blog), <https://thewarhorse.org/military-veterans-learn-to-fight-disinformation-campaigns/>.

Overcoming the Labeled Training Data Bottleneck: A Route to Specialized AI

Dr. Ravi Starzl

ABSTRACT

This article explores the potential of large language models (LLMs) to transform dataset creation and analysis in cybersecurity. The proposed method leverages LLMs to overcome the labeled data bottleneck by generating high-quality, task-specific datasets for AI model tuning. Existing network intrusion analysis datasets are synthesized with domain knowledge extracted from cybersecurity literature to create a new dataset tailored for supervised training of zero-day exploit detection systems. LLMs interpret the semantic content of relevant literature to identify crucial characteristics and values of zero-day exploit signatures in network traffic. The resulting synthesized dataset is primarily based on 'organic' data collected by genuine sensors, with key feature characteristics intelligently interpolated by LLMs. This approach enables the creation of suitable training data for high-performance ML models. This article demonstrates the effectiveness of this method by utilizing advanced AI techniques to generate a dataset for zero-day exploit detection, illustrating the potential for accelerated progress in specialized AI for cybersecurity. The proposed solution offers a promising approach to address the challenge of labeled data scarcity in developing specialized AI for cybersecurity, facilitating more efficient and effective protection against emerging threats.

© 2024 Dr. Ravi Starzl



Dr. Ravi Starzl, on faculty at Carnegie Mellon University (CMU) and the University of Colorado, was recently awarded the 2024 Award for Artificial Intelligence (AI) Excellence by the Business Intelligence Group. His current area of focus is the application of AI, Machine Learning (ML), and Complex Systems Science to radio frequency (RF), operational technology (OT), and Cyber Operations for asymmetric advantage. Dr. Starzl earned his M.S. and Ph.D. in ML and AI from Carnegie Mellon University

INTRODUCTION

The advent of LLMs heralds a transformative era in dataset creation and analysis.¹ By generating high-quality, task-specific labeled datasets, LLMs are setting the stage for unprecedented advancements in data variance identification and system analysis, enabling the development of specialized AI with superior performance and precision.² This article explores the potential of LLMs to revolutionize large-scale dataset processing in cybersecurity,³ focusing on their ability to overcome labeled data bottlenecks and refine AI model tuning through tailored dataset generation.

The Role of Information in AI and ML for Cybersecurity

Information is the foundation of inference tasks in ML and AI. Extracting and processing meaningful patterns and insights from cybersecurity data enable AI systems to identify threats, predict attacks, and perform complex security analyses.⁴ However, accessing and leveraging information for cybersecurity AI present major challenges, such as the lack of labeled training data specific to cybersecurity tasks, which often hinder development of high-performing AI models.

Overcoming the Labeled Data Bottleneck in Cybersecurity

To overcome the labeled data bottleneck in cybersecurity, this work proposes an innovative way to leverage the power of LLMs. By utilizing existing datasets, even if not designed to address the specific task at hand, and leveraging the descriptions of systems and semantic relationships within the context of the objective task, LLMs can generate high-quality, task-specific labeled datasets. This approach harnesses the expressive power of current LLMs and other transformer-based learning systems to create new, synthesized datasets rooted in organic examples but synthetically

fused together and interpolated to create suitable training data for high-performance ML models on objective tasks that have insufficient traditionally gathered training data.

Demonstrating the Method: Zero-Day Exploit Detection

To ground the proposed method for use by cybersecurity practitioners, this work demonstrates how advanced AI techniques can overcome a data label bottleneck in developing an AI for zero-day exploit detection. By leveraging traditional intrusion detection datasets and LLM capabilities, the method creates a new dataset capable of detecting network traffic patterns that may be associated with zero-day exploits, showcasing the potential for rapid advancements in specialized AI for cybersecurity.⁵

The Future of Specialized AI in Cybersecurity

As the sophistication of contextual inference capabilities and the scale of data access and processing grows, we anticipate further rapid growth of methods like the one presented here. These advances will enable swift scaling of specialized AI in cybersecurity, revolutionizing threat detection, attack prediction, and overall network security.

The proposed method, which leverages LLMs to synthesize task-specific labeled datasets, offers a promising solution to overcome labeled data bottlenecks in specializing AI for cybersecurity. By demonstrating the method's effectiveness in creating a dataset for zero-day exploit detection, this work highlights the potential for rapidly advancing AI-driven cybersecurity solutions, paving the way to better protect against evolving threats.

Information, Error, and Synthesizing Specialized Datasets

At the heart of developing effective cybersecurity AI are the fundamental concepts of information and error. Information, in this context, refers to meaningful patterns and relationships within cybersecurity data that enable AI systems to detect threats and anomalies. Error represents discrepancies between the AI's predictions and ground truth, which can arise from various sources such as data inconsistencies, model limitations, or the inherent complexity of the cybersecurity domain.

Decomposing error into its constituent components—bias, variance, and irreducible error—provides a mathematical and philosophical framework for understanding the challenges and opportunities in creating specialized datasets for cybersecurity AI. Bias, which arises from oversimplified assumptions or limited data representation, can be mitigated by incorporating a wider range of cybersecurity scenarios and data sources. Variance, reflecting the model's sensitivity to fluctuations in the training data, can be addressed through techniques like regularization or ensemble learning. Irreducible error represents the inherent uncertainty or “noise” within the data, which cannot be eliminated entirely. This is a byproduct of assumptions and methods by which data are collected or sampled.

To overcome these challenges and create specialized datasets for tasks like zero-day exploit detection, we can leverage the power of information theory and ML techniques to create systems that implicitly capture the information structures that drive the observed variations associated with cyber security events.⁶ This information structure is necessarily reflected in the use of language and code that are the substance of cyber security or cyber operations. By analyzing the information content and relationships within existing datasets, such as intrusion detection data, we can identify key patterns and features most relevant to the target task. This analysis can guide the synthesis of new, specialized datasets that capture the essential characteristics of the target domain while minimizing the impact of irreducible error.⁷

The approach we explore here leverages the expressive power of current LLMs to generate realistic and diverse cybersecurity scenarios from related datasets designed to address different objective tasks. By also training these models on a wide range of cybersecurity literature, threat reports, and technical documentation, we can create a rich semantic understanding of the domain. This understanding can then be used to generate synthetic dataset designs, as well as data points that mimic the patterns and relationships observed in real-world cybersecurity events, while introducing controlled and semantically viable variations to enhance the model's ability to generalize to new threats.

This can further be extended with techniques from transfer learning and domain adaptation⁸ to leverage knowledge gained from adjacent cybersecurity tasks. By identifying common patterns and features shared between related tasks, like intrusion detection and malware analysis, insights and representations learned from previous tasks can be transferred to a new task.

Synthesis of specialized datasets for cybersecurity AI requires deep insight into the concepts of information and error. Decomposing error into its three components and leveraging the power of information theory and ML techniques allows us to create rich and diverse datasets that capture the essential characteristics of the target domain while ‘borrowing’ insight from existing problem-adjacent datasets and literature. Approaches like language model-based data generation, transfer learning, and adversarial training allows us to overcome the limitations of manual labeling and enhance our ability to generalize to new threats.⁹

TRANSITION FROM INTRUSION DETECTION TO ZERO-DAY EXPLOIT DETECTION

The 1999 KDD cybersecurity dataset, while influential in developing ML models for network intrusion detection, falls short in addressing the unique challenges posed by zero-day exploits. These exploits, which target unknown vulnerabilities, require a novel approach to detection that goes beyond the scope of the KDD dataset.

The method presented here utilizes LLMs to create a new objective-task training data matrix specifically designed for detecting zero-day exploits.¹⁰ The process begins by granting the LLM access to a vast array of cybersecurity literature, including research papers, technical reports, and industry publications. This extensive knowledge base enables the LLM to develop a deep understanding of the characteristics, patterns, and indicators associated with zero-day exploits.¹¹

Armed with this domain-specific knowledge, the LLM is then tasked with inferring the essential features and attributes that are crucial for effective zero-day exploit detection. This inference process involves analyzing the cybersecurity literature to identify the unique signatures, anomalies, and behavioral patterns that distinguish zero-day exploits from known threats and normal network traffic.^{12,13}

Based on these inferred characteristics, the LLM designs a comprehensive objective-task training data matrix that encapsulates the key elements necessary for training ML models to detect zero-day exploits. This matrix serves as a blueprint for the synthesized dataset, ensuring that it includes the most relevant and informative features for the task at hand.

To populate the objective-task training data matrix, the LLM analyzes the KDD dataset and a diverse array of other cybersecurity and network traffic datasets, extracting pertinent features and patterns that align with the designed matrix. The LLM then generates specific scripts and transformation functions to integrate and adapt the data from these datasets into the new synthesized dataset.

The resulting synthesized dataset is not a mere amalgam of existing datasets, but rather a crafted resource tailored to the unique challenges of zero-day exploit detection. By leveraging the LLM's knowledge of cybersecurity literature and its ability to infer the critical characteristics of the objective task, the synthesized dataset provides a rich and comprehensive training ground for advanced ML models.

New models for anomaly detection, unsupervised learning, and more, can then be trained on the synthesized dataset to learn the subtle nuances and patterns associated with zero-day exploits.^{14,15} Exposing the models to a wide range of realistic and diverse scenarios encompassed within the synthesized dataset enables them to develop the ability to identify and flag potential zero-day exploits in real-world network traffic.

This method marks a step forward for cybersecurity. By harnessing the power of LLMs and their access to extensive cybersecurity literature, this approach enables the creation of highly specialized and effective training datasets for zero-day exploit detection. The resulting models, trained on these synthesized datasets, have the potential to shift the way organizations detect and respond to previously unknown vulnerabilities, bolstering their overall security posture.

PROMPTING PROCEDURE

Utilizing a Llama2-70B model embedded in a Python 3.11 program and the transformers library from huggingface, an initial prompt asked the LLM how it would transform the existing network dataset into a dataset useful for zero-day anomaly exploit detection. A table was also provided with the prompt showing a reduced set of features from the 1999 KDD cybersecurity dataset:

Please provide suggestions on how to modify the data matrix and labels to enable zero-day exploit detection. Think creatively and provide a detailed explanation of the changes you are making and why they are important. Here are the features of the data matrix and some example values:

Protocol Type	Service	Flag	Src Bytes	Dst Bytes	Land	Wrong Fragment	Urgent	Count	Srv Count	Error Rate	Srv Error Rate	Error Rate	Srv Error Rate	Same Srv Rate	Diff Srv Rate	SrvDiff Host Rate	Label
tcp	http	SF	215	45076	0	0	0	8	8	0.00	0.00	0.00	0.00	1.00	0.00	0.00	good
udp	private	SF	105	146	0	0	0	15	15	0.00	0.00	0.00	0.00	1.00	0.00	0.00	good
tcp	telnet	RSTO	0	0	0	0	0	23	10	0.00	0.00	1.00	1.00	0.17	0.83	0.00	bad
icmp	eco_i	SF	8	0	0	0	0	13	13	0.00	0.00	0.00	0.00	1.00	0.00	0.00	good
tcp	ftp_data	So	0	0	0	0	0	1	1	1.00	1.00	0.00	0.00	0.00	1.00	0.00	bad
tcp	http	SF	337	981	0	0	0	2	2	0.00	0.00	0.00	0.00	1.00	0.00	1.00	good
tcp	http	SF	0	0	0	0	0	50	25	0.20	0.20	0.00	0.00	0.50	0.50	0.00	bad
udp	domain	SF	44	133	0	0	0	8	8	0.00	0.00	0.00	0.00	1.00	0.00	0.00	good
tcp	smtp	SF	789	334	0	0	0	4	4	0.00	0.00	0.00	0.00	1.00	0.00	0.50	good
tcp	ftp	SF	0	0	0	1	0	16	8	0.75	0.75	0.25	0.25	0.50	0.50	0.00	bad

Table 1: A reduced set of features from the 1999 KDD cybersecurity dataset, provided in the prompt.

The response received outlined a suggested set of modifications illustrated in Figure 1.

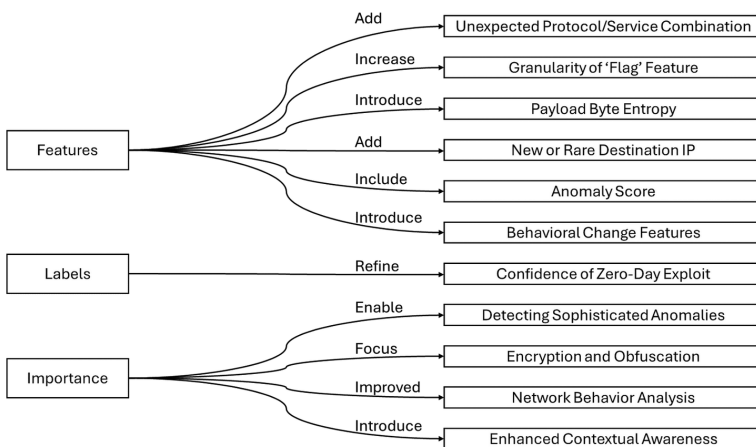


Figure 1: LLM suggested data set modifications.

The next prompt elicited a specific set of features most likely to provide powerful predictive performance in detecting network traffic associated with a zero-day exploit, utilizing network traffic data only. This was a two-prompt process.

Prompt 1: Ask the LLM for assistance in crafting a more comprehensive prompt for the objective task:

Imagine you are entering a data science competition aimed at developing a model to classify network traffic into three categories: normal, suspicious, and highly suspicious. You will use well-known network intrusion detection datasets for this task. Your objective is to identify the most informative features within those datasets that can be combined to help distinguish between these traffic types effectively.

Draft a strategy for selecting key features from basic traffic descriptors, behavioral indicators, and advanced metrics that could potentially highlight network anomalies. Describe how these features will be utilized to enhance the accuracy and effectiveness of your classification model.

Prompt 2: Ask the LLM to directly specify the set of features that would be most informative to the objective task:

Imagine you are competing in a data science challenge with the goal of developing a cutting-edge model that classifies network traffic into three distinct categories: normal, suspicious, and malicious. This task is vital for boosting network security measures by pinpointing potential threats through a detailed analysis of network traffic.

To undertake this challenge, you plan to leverage datasets from publicly available network intrusion detection resources, like the NSL-KDD dataset. Your task is to devise a model that uses critical network traffic features, meticulously selected for their ability to shed light on network behavior and potential threats.

Your first step is to identify and articulate the most informative features within the data that could distinguish effectively between the three traffic types. Consider what basic traffic features (like protocol types and data volume), behavioral features (such as connection patterns and error rates), advanced indicators (including unusual protocol-service combinations and payload characteristics), and anomaly detection metrics (which measure deviations from normal behavior) might be crucial for this task.

Craft your feature selection strategy to include both well-established and novel metrics, aiming to build a robust model that excels in accuracy and reliability for classifying network traffic. This approach will ensure that your model not only identifies but also understands the subtleties and complexities of network traffic, enhancing your ability to detect and classify potential security threats efficiently.

The features identified through this process as being most useful in detecting network traffic associated with zero-day exploits by the LLM are:

Protocol Type Service, Flag, Src Bytes, Dst Bytes, Land Wrong Fragment, Urgent, Count, Srv Count, Serror Rate Srv Serror Rate, Rerror Rate, Srv Rerror Rate, Same Srv Rate, Diff Srv Rate, Srv Diff Host Rate, Unexpected Protocol/Service, Payload Byte Entropy, New/Rare Destination IP, Anomaly Score, Behavioral Change Score.

The next prompt provided an interstitial reasoning exercise for the LLM to gain more traction on how to modify the dataset in detail, as well as to provide a scaffold for where to discover the needed data to make reasonable interpolations for synthetic data generation.

Think creatively. How would you generate the actual data for the new features and add them to the new table. Provide a detailed explanation. Use all the publicly available knowledge about these features and how they can be estimated.

The response received outlined suggested process and method illustrated in Figures 2 and 3.

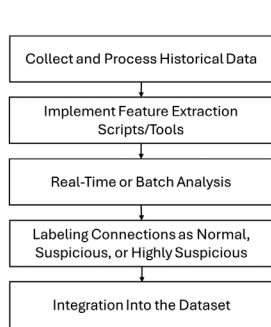


Figure 2. LLM suggested pipeline overview for synthesized dataset creation.

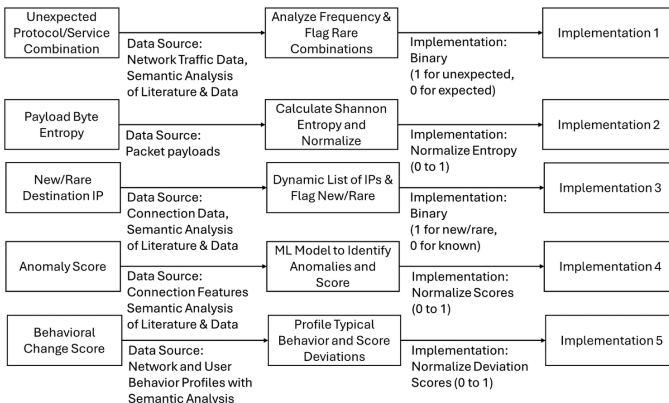


Figure 3. LLM suggested features, implementations, and method for designing a dataset aimed at detecting network anomalies and zero-day exploits through network traffic analysis.

DATA ALIGNMENT AND DATASET IDENTIFICATION THROUGH VECTOR EMBEDDING

The data were aligned to enhance the detection of zero-day exploits and improve cybersecurity network traffic analysis. A custom Python program, leveraging the Ray parallel processing library, was developed to traverse the May 2023 Common Crawl dataset. Utilizing BERT base uncased, this program generated vector embeddings for the elements of the dataset that were pre-identified as potentially relevant with phrases indicative of cybersecurity interests such as "zero-day exploit data set," "cyber security network traffic analysis," "comprehensive cybersecurity dataset," and "network intrusion anomaly dataset." This dataset was then further mined using a vector-database (Weaviate) and retrieval-augmented-generation (RAG) to identify other datasets that contained the type of information the LLM had already identified as potentially helpful to the prediction task. Identifying and including these additional datasets provided more information for synthesizing values for missing features in the newly designed dataset for the objective task. This generally has the effect both of increasing accuracy of the synthesized values and providing a broader but still realistic range of variations in those values.

This process unearthed several key datasets for cybersecurity analysis:

- ◆ **KDD Cup 99 Dataset:** Despite its age, this dataset remains a cornerstone in network intrusion detection, offering a broad spectrum of network connection features.
- ◆ **NSL-KDD Dataset:** An evolution of the KDD Cup 99 dataset, the NSL-KDD addresses previous limitations by eliminating redundant records, thereby enhancing the dataset's utility for training and testing intrusion detection models.

- ◆ **CICIDS2017 (Canadian Institute for Cybersecurity Intrusion Detection System 2017):** Featuring both malicious and benign attacks, this dataset reflects contemporary attack scenarios with detailed network traffic features and, for some attack types, payload data.
- ◆ **UNSW-NB15 Dataset:** Provided by the Australian Cyber Security Centre, this dataset mixes real normal activities with synthetic attack behaviors, offering a diverse array of network traffic analysis features.
- ◆ **ISCX VPN-nonVPN Traffic Dataset (ISCXVPN2016):** Containing labeled network traffic that differentiates between VPN and non-VPN traffic, this dataset is invaluable for studying encrypted traffic patterns and potentially high-entropy payloads.
- ◆ **CTU-13 Dataset:** This dataset includes botnet traffic alongside normal and background traffic, facilitating the study of botnet behaviors which may share similarities with zero-day exploit traffic patterns, especially in command and control communications and lateral movements.
- ◆ **MAWI Working Group Traffic Archive:** A compilation of real-world internet backbone traffic datasets from the Wide project, capturing a variety of internet activities over extended periods. While not cybersecurity-specific, it offers a rich baseline for normal traffic pattern analysis.
- ◆ **The CAIDA UCSD Datasets:** Provided by the Center for Applied Internet Data Analysis (CAIDA), these datasets consist of anonymized internet traces, aiding in the modeling of both normal and anomalous network behaviors.
- ◆ **ToN IoT Telemetry Dataset:** A contemporary dataset focusing on Internet of Things (IoT) telemetry, including network traffic, logs, and attack data. It is particularly suited for exploring IoT-specific threats and anomalies.

The next prompt asked the LLM to derive the best way to align features and labels across these diverse datasets:

Please provide a detailed feature alignment plan to align the features of these datasets with these features: Protocol Type Service, Flag, Src Bytes, Dst Bytes, Land Wrong Fragment, Urgent, Count, Srv Count, Serror Rate Srv Serror Rate, Rerror Rate, Srv Rerror Rate, Same Srv Rate, Diff Srv Rate, Srv Diff Host Rate, Unexpected Protocol/Service, Payload Byte Entropy, New/Rare Destination IP, Anomaly Score, Behavioral Change Score.

The five methodologies the LLM defined are shown in Figure 4. This alignment between existing dataset features (“organic” features derived from actual sensor observations), and designed dataset features (ideal features for the objective task that the LLM suggests) minimizes the distance between the two feature sets, helping the inferred values to draw the maximum amount of information from the actually observed values, and helping to minimize the attenuation of the underlying signal in the process of synthesizing the new dataset.

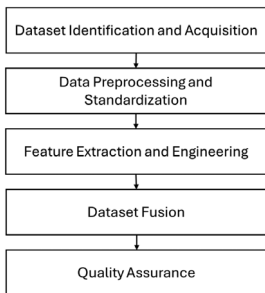


Figure 4. Overview of the implementation process for creating a synthesized dataset using LLMs and existing cybersecurity datasets. The process consists of five steps that synthesize diverse cybersecurity datasets into a unified dataset.




Figure 5. Dataset Identification and Acquisition - The first step in the dataset fusion process involves identifying relevant cybersecurity datasets.

The “closer” the new synthesized dataset features are to the existing dataset features, the greater the real-world performance and generalization of the models trained on that data.

IMPLEMENTATION STEPS

This method for fusing diverse cybersecurity datasets, guided by an LLM and automated through Python scripts, creates a unified, analysis-ready dataset.

Dataset Identification and Acquisition

The first step in the dataset fusion process involved identifying relevant cybersecurity datasets, as described in Figure 5. The LLM, implicitly trained on a vast corpus of cybersecurity literature, provided guidance on selecting datasets that encompass a wide range of attack types, network environments, and data formats. The identified datasets were downloaded and stored in a centralized repository.

Data Preprocessing and Standardization

The acquired datasets underwent a preprocessing phase shown in Figure 6. The LLM generated dynamic prompts, tailored to each dataset's specific attributes and requirements, which were then used to develop Python scripts for automated data preprocessing. These scripts performed tasks such as decoding features into plain text, handling missing values, and converting the datasets into a uniform tabular format, typically CSV files.




Figure 6. Data Preprocessing and Standardization - The acquired datasets undergo a preprocessing phase.

Feature Extraction and Engineering

The LLM defined a detailed set of features for the fused dataset shown in Figure 7. By analyzing the characteristics of each dataset and leveraging its knowledge of cybersecurity domain expertise, the LLM generated a feature alignment plan. This plan outlined the necessary features to be extracted from each dataset, as well as advanced features that could be derived through feature engineering techniques.




Figure 7. Feature Extraction and Engineering - The LLM defines a detailed set of features for the fused dataset. By analyzing the characteristics of each dataset and leveraging its knowledge of cybersecurity, the LLM generates a feature alignment plan.

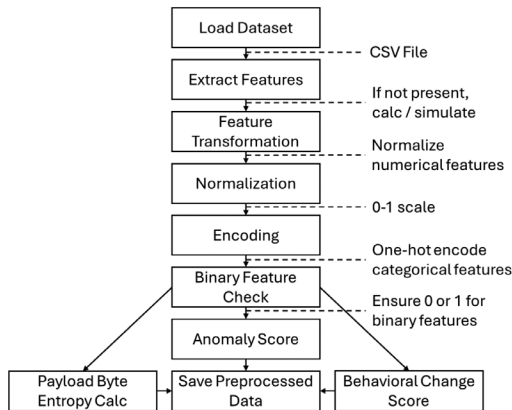


Figure 8. Dataset Fusion – The final stage involves fusing the preprocessed and feature-engineered datasets into a unified dataset. The LLM provides guidance on the techniques used, considering factors such as data format, feature compatibility, and scalability.

Python scripts, guided by the LLM's feature alignment plan, were developed to automate the feature extraction and engineering process. These scripts utilized various data manipulation libraries and analytical tools to calculate complex features, such as entropy measures and behavioral change scores, based on the available data. The extracted and engineered features were then integrated into the preprocessed datasets, resulting in a set of feature-rich, analysis-ready datasets.

Dataset Fusion and Quality Assurance

The final stage of the method involved fusing the preprocessed and feature-engineered datasets into a unified dataset as shown in Figure 8. The LLM provided guidance on the appropriate data fusion techniques, considering factors such as data format, feature compatibility, and scalability. Python scripts were developed to automate the dataset fusion process using pandas and numpy.

To ensure the integrity and reliability of the fused dataset, a quality assurance process was implemented. The LLM generated prompts for automated data validation scripts, which checked for anomalies, inconsistencies, and outliers in the fused dataset. Additionally, the scripts verified the correct application of labels and analyzed feature distributions to ensure the coherence and representativeness of the fused dataset. Manual checking of samples of the results was also undertaken.

By leveraging the knowledge and reasoning capabilities of the LLM, the method draws from a wide range of sources to ensure the fused dataset is well-aligned, feature-rich, and suitable for the objective task within the limits of current semantic analysis. The automated nature of the process, facilitated by Python scripts, enables efficient and scalable dataset fusion, reducing manual effort.

Individual scripts were produced with the assistance of LLM crafted prompts. The prompts were generated with a seed prompt such as:

Please provide me with a prompt I can give to code llama to normalize the features of one of these datasets with the target set of features. Include the target features as a template in the prompt.

For brevity, the logical structure of two such resulting prompts, one to normalize data and the other to conduct literature review and incorporate resulting insights, are shown in Figure 9.

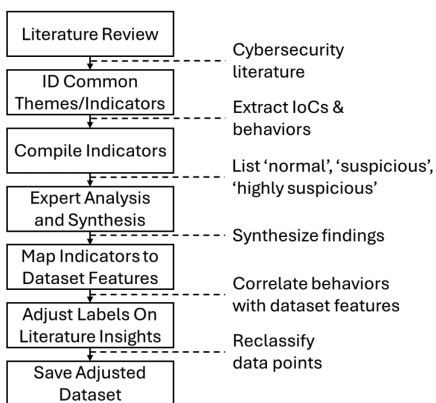


Figure 9. Quality Assurance - To ensure the integrity and reliability of the fused dataset, a quality assurance process is implemented. The LLM generates prompts for automated data validation scripts, which check for anomalies, inconsistencies, and outliers.

The code generated by code llama required manual editing to ensure correct paths; manual installation of required libraries was also needed. Full automation of the integration of the prompt and output of the two LLMs is eminently achievable but left for future work.

As an additional step to ensure synthetic data were properly aligned with credible feature values to enable reasonable real-world performance, an additional prompt was submitted asking Llama2 to ensure the labels 'normal,' 'suspicious,' and 'highly suspicious' were all properly associated with the feature values that were credible representations:

Please explain how you will ensure that the labeling of 'normal', 'suspicious', or 'highly suspicious' will be associated with the appropriate feature values that are representative of real-world cases. Incorporate expert knowledge to mine cybersecurity papers, publications, and other written material without interviewing live personnel.

The LLM response was a near duplicate of that already shown in Figure 9, so the process defined there was implemented as a procedure run twice, back-to-back.

Manual editing of the generated code to incorporate necessary API keys for the online services queried was necessary. Some online repositories of cyber security literature did not have an API, so access was provided in the form of web-scraping the search results pages directly. These steps can be automated but are left to future work.

STREAMLINED APPROACH TO NETWORK BEHAVIOR CLASSIFICATION

Introduction

This section presents a streamlined approach to classifying network behaviors using the Llama2-70B model, a versatile language model. The objective is to categorize network traffic into three classes: 'normal', 'suspicious', and 'highly suspicious'. This classification task is important for identifying potential security threats and ensuring network infrastructure safety.

To achieve this goal, the Llama2-70B model is fine-tuned using a training method designed to address the challenges posed by limited computational resources in a home setup. The fine-tuning process includes appending a classification head to the pre-trained model, optimizing the model's performance using appropriate loss functions and metrics, and employing techniques such as gradient accumulation and data streaming to manage memory constraints.

Post-training evaluations are conducted to assess the model's performance across metrics including accuracy, precision, recall, and F1 score. Based on these evaluations, iterative refinements are made to the dataset, feature engineering, and model architecture to improve the model's ability to accurately classify network behaviors.

To validate the effectiveness of the fine-tuned model, normal and zero-day exploit traffic are simulated in a controlled environment using Python and the Scapy library. This testing

method generates realistic network traffic patterns and assesses the model's ability to detect and classify potential security threats.

The following subsections detail the fine-tuning process, training method, evaluation and refinement, and the testing environment.

Setup for Fine-tuning

The task at hand involved categorizing network behaviors into 'normal', 'suspicious', and 'highly suspicious' classes. For this purpose, the Llama2-70B model, known for its versatility, was fine-tuned. A classification head with a dense layer outputting three categories was appended, using a softmax function to generate a class probability distribution. The model's optimization employed the categorical cross-entropy loss function, ideal for multi-class tasks, with accuracy, precision, recall, and F1 score as performance metrics.

Fine-tuning Execution

The fine-tuning initiated with a learning rate of $1e-5$ to adjust the pre-trained model subtly. Given the data and model's extensive memory requirements, a batch size of 8 was chosen, targeting a training span of 3 to 4 epochs. The Adam optimizer facilitated adaptive learning rate adjustments, while a 0.1 dropout rate in the classification head aimed to prevent overfitting.

Training Method

Training adapted to a 128GB RAM home setup required strategic planning. The training leveraged PyTorch's `DataLoader` for data streaming, allowing efficient batch processing from disk, thus bypassing RAM constraints. Gradient accumulation was employed to mimic larger batch effects, enhancing training efficacy. Frequent model checkpointing safeguarded against data loss.

Evaluation and Refinement

Post-training evaluations on a separate test set provided crucial feedback on the model's performance across metrics. Scikit-learn was used to compute precision, recall, and F1 scores, offering a detailed performance overview. Based on these insights, iterative refinements addressed dataset balance, feature engineering, and model adjustments to rectify misclassifications and improve metrics.

Preparing for Deployment

The deployment phase involved model quantization via PyTorch, enhancing the model's efficiency for use in limited-resource settings. Dynamic quantization was chosen for its balance of efficiency and simplicity, with an eye on static quantization and quantization-aware training for future enhancements. Inference testing on a reduced data subset confirmed the model's performance and responsiveness, despite the hardware limitations of a home com-

puting setup.

The Testing Environment

To simulate both normal and zero-day exploit traffic in a controlled environment using Python, we leveraged popular cybersecurity libraries like Scapy for network traffic manipulation and generation. Scapy enables creation, sending, and capturing of network packets in a detailed manner, making it ideal for both benign and malicious traffic simulations.

Generating Normal Traffic with Python and Scapy

For generating normal traffic, the script simulates common activities such as web browsing, email communication, and file transfers. The aim is to create a realistic background traffic pattern that a typical home network would experience.

```
```python
from scapy.all import *

def generate_normal_traffic():
 # Simulate HTTP web browsing
 ip_layer = IP(dst="www.example.com")
 tcp_layer = TCP(sport=RandShort(), dport=80)
 http_get = "GET / HTTP/1.1\r\nHost: www.example.com\r\n\r\n"
 packet = ip_layer / tcp_layer / http_get
 send(packet, verbose=0)

 # Simulate email traffic (SMTP)
 ip_layer = IP(dst="mail.example.com")
 tcp_layer = TCP(sport=RandShort(), dport=25)
 smtp_hello = "HELO mail.example.com\r\n"
 packet = ip_layer / tcp_layer / smtp_hello
 send(packet, verbose=0)

 # Simulate FTP file transfer
 ip_layer = IP(dst="ftp.example.com")
 tcp_layer = TCP(sport=RandShort(), dport=21)
 ftp_hello = "USER anonymous\r\n"
 packet = ip_layer / tcp_layer / ftp_hello
 send(packet, verbose=0)

generate_normal_traffic()
```
```

Generating Simulated Zero-Day Exploit Traffic

For the "WindowsGate" zero-day exploit, the script simulates the network behavior characteristic of the exploit's activity, such as scanning, exploitation attempts, and data exfiltration. It's important to note that in a real-world scenario, executing such a script should be done with utmost caution and strictly within a controlled environment to prevent unintended harm.

```

```python
from scapy.all import *

def generate_zero_day_traffic(target_ip="192.168.1.10"):
 # Simulate scanning activity
 for port in range(20, 25):
 ip_layer = IP(dst=target_ip)
 tcp_layer = TCP(sport=RandShort(), dport=port, flags="S")
 send(ip_layer / tcp_layer, verbose=0)

 # Simulate exploit payload delivery
 payload = "WindowsGate simulated exploit payload here [redacted code]"
 ip_layer = IP(dst=target_ip)
 tcp_layer = TCP(sport=RandShort(), dport=445, flags="PA")
 packet = ip_layer / tcp_layer / Raw(load=payload)
 send(packet, verbose=0)

 # Simulate data exfiltration activity
 exfiltrated_data = "Exfiltrated data here"
 ip_layer = IP(dst="malicious.server.com")
 tcp_layer = TCP(sport=RandShort(), dport=80)
 packet = ip_layer / tcp_layer / Raw(load=exfiltrated_data)
 send(packet, verbose=0)

generate_zero_day_traffic()
```

```

Results

The study employed a simulated 10-fold cross-validation method to assess the performance of the Llama2-70B language model fine-tuned for cybersecurity threat detection. The analysis focused on the model's ability to classify network behaviors into three categories: 'normal', 'suspicious,' and 'highly suspicious.' The evaluation shows confusion matrices for each fold, providing insights into model classification accuracy and the influence of synthetic data on performance.

The confusion matrices revealed that the model demonstrated a high true positive rate for 'normal' network behaviors, with the number of correctly identified 'normal' instances ranging from 950 to 980 across the folds. This indicates the model's capability to recognize legitimate network activities, which is an important foundation for effective threat detection.

The classification of 'suspicious' and 'highly suspicious' behaviors showed encouraging results, with the model correctly identifying 'suspicious' behaviors in a range of 780 to 850 instances and 'highly suspicious' behaviors in 704 to 760 instances. While there were instances of misclassification, primarily in the form of false negatives, the overall performance suggests that the model has the potential to detect and classify malicious activities effectively.

The confusion matrices also highlighted areas for model improvement, such as enhancing the model's ability to differentiate between varying degrees of threat severity and fine-tuning

its sensitivity to nuanced network behaviors. These insights provide valuable guidance for future refinements and optimizations.

The cross-validation results demonstrate the potential of using synthetic data to train models for improving the performance of specialized AI in complex objective tasks, particularly when some amount of prior data is available. The model's high accuracy in identifying 'normal' behaviors and its effectiveness in recognizing 'suspicious' and 'highly suspicious' activities indicate its potential for real-world applications.




Figure 10. 10-fold cross validation confusion matrices showing performance on simulated data, using model trained on synthetic/synthetically enhanced training data.

However, it is important to acknowledge that there is still significant work to be done to refine this approach. The observed misclassifications underscore the need for model refinement and optimization to enhance its ability to discern between complex threat behaviors. Future research should focus on improving data synthesis techniques, incorporating more diverse datasets, and exploring advanced model architectures to further improve the model's performance and generalizability.

IMPLICATIONS, DISCUSSION, AND FUTURE WORK

The cross-validation results, as demonstrated by the confusion matrices, affirm the potential of using synthetic data to train models for improving the performance of specialized AI in complex objective tasks, particularly when some degree of prior data is available. In the example of cybersecurity threat detection, the model's consistently high true positive rate for

'normal' network behaviors, ranging from 950 to 980 correctly identified instances across the ten folds, indicates its strong capability to recognize legitimate network activities.

However, the classification of 'suspicious' and 'highly suspicious' behaviors encountered more variability. While the model maintained a promising accuracy, with 'suspicious' behaviors correctly identified in a range of 780 to 850 instances and 'highly suspicious' behaviors in 704 to 760 instances, there were notable instances of misclassification. These errors were primarily false negatives, where 'suspicious' or 'highly suspicious' activities were incorrectly labeled as 'normal,' and to a lesser extent, as false positives within the 'suspicious' and 'highly suspicious' categories. The confusion matrices revealed a pattern of errors suggesting areas for model improvement, e.g., challenges in differentiating between varying degrees of threat severity and the need for fine-tuning the model's sensitivity to nuanced network behaviors.

This work highlights the potential of using synthetic data to create extensive, labeled training datasets by merging real data with synthetic extensions, integrating both numerical and semantic insights from existing data sources. The promising results demonstrate the practicality of this novel approach, involving a complex integration of real-world data attributes with synthetically produced data points to enhance the dataset's variety and representativeness.

The relationship between the complexity of the synthesized datasets and the expressiveness of the learning system being trained is crucial. The synthesized dataset must accurately portray how the system under analysis would behave naturally. An overly simplistic synthesized dataset could lead to rapid overfitting by a powerful deep learning system. The use of LLMs and the integration of semantic information from the literature help mitigate this risk by ensuring the dataset's complexity matches the learning system's expressiveness.

The challenges in overcoming the labeled training data bottleneck are particularly evident when mimicking the complexity of continually changing systems, such as cybersecurity threats. These threats are characterized by their intricacy and the adaptiveness of their perpetrators, making them difficult targets for predictive modeling. Accurately mirroring the subtleties of such systems with synthetic data requires continuous advancement in data generation methods.

Despite these challenges, the use of synthetic data represents progress in creating customized readily available training sets for specific analytical goals. The automated amalgam of real and synthetic data enables highly flexible and scalable training environments, suitable for a wide range of informational needs. The implications of this method extend beyond cybersecurity, and provide an example for developing on-the-spot training sets across diverse fields. By reducing dependence on extensive real datasets, which are often restricted by availability and privacy issues, this strategy promotes a more exploitable approach for the progression of ML and AI.

To address the limitations and build on the current findings, future work should explore:

1. Developing advanced data synthesis techniques to better mimic the complexity and variability of real-world cyber threats.
2. Incorporating more granular features and exploring additional training strategies.
3. Integrating synthetic data with curated real-world data to provide a richer training dataset.
4. Implementing incremental learning techniques continuously to update the model with new data.
5. Testing the model's performance across different network environments and against various types of cyber threats.
6. Developing explainability mechanisms to understand the model's decision-making processes and increase trust in its predictions.
7. Automating the integration of API keys and web-scraping for seamless data access.

The goal is to create systems that can automatically generate substantial and credible labeled training datasets for any complex objective task or information need.

While the fine-tuned Llama2-70B model using this data synthesis approach shows promise for cybersecurity threat detection, this is an initial step in developing a comprehensive and reliable threat detection system. With ongoing research and refinement, this approach will contribute significantly to the field of cybersecurity, enabling more effective and efficient detection of evolving threats in real-world network environments.

This work offers promising insights into the potential of using synthetic data for creating labeled training sets by fusing organic information with synthetic extrapolation. While challenges remain, particularly in simulating complex and evolving systems, this approach represents a step forward in developing accessible, effective, and scalable approaches for creating in-situ training sets for arbitrary objective tasks or information needs in a largely automated manner. 🍷

DISCLAIMER

The views expressed in this paper are those of the author and do not reflect the official policy or position of the U.S. Military Academy, U.S. Army, U.S. Department of Defense, or U.S. Government.

NOTES

1. Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan et al. "Language Models are Few-Shot Learners." (July 22, 2020). Conference Paper, Vancouver: *34th Conference on Neural Information Processing Systems*, <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac-142f64a-Paper.pdf>, 33.
2. Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Hongzhi Wen, Xiaochi Wei, Shuaiqiang Wang et al. "Exploring the Potential of Large Language Models (LLMs) in Learning on Graphs." (March 28, 2024) *Association for Computing Machinery (ACM) Special Interest Group on Knowledge Discovery in Data (SIGKDD) Explorations Newsletter Vol 25, No. 2*, <https://dl.acm.org/doi/10.1145/3655103.3655110>, 42-61.
3. Jie Zhang, Haoyu Bu, Hui Wen, Yu Chen, Lun Li, and Hongsong Zhu. "When LLMs Meet Cybersecurity: A Systematic Literature Review." (May 6, 2024) <https://arxiv.org/abs/2405.03644>.
4. HanXiang Xu, ShenAo Wang, Ningke Li, Yanjie Zhao, Kai Chen, Kailong Wang, Yang Liu, Ting Yu, and HaoYu Wang. "Large Language Models for Cyber Security: A Systematic Literature Review," (May 8, 2024). <https://arxiv.org/abs/2405.04760>.
5. Diana Levshun and Igor Kotenko. "A Survey on Artificial Intelligence Techniques for Security Event Correlation: Models, Challenges, and Opportunities." (January 7, 2023) *Artificial Intelligence Review* 56, no. 8, 8547-8590. <https://link.springer.com/article/10.1007/s10462-022-10381-4>.
6. Yoshua Bengio, Aaron Courville, and Pascal Vincent. "Representation Learning: A Review and New Perspectives." (June 24, 2013) *IEEE transactions on pattern analysis and machine intelligence* 35, no. 8, 1798-1828. <https://arxiv.org/abs/1206.5538>.
7. Alexander Ratner, Stephen H. Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. "Snorkel: Rapid Training Data Creation with Weak Supervision." (2018) In *Proceedings of the VLDB endowment. International Conference on Very Large Data Bases, vol. 11, no. 3*, p. 269; Eman T. Hassan, Xin Chen, and David Crandall. "Unsupervised Domain Adaptation Using Generative Models and Self-Ensembling." (2018) <https://arxiv.org/abs/1812.00479>, 8; Eman T. Hassan, Xin Chen, and David Crandall. "Unsupervised Domain Adaptation Using Generative Models and Self-Ensembling." (December 2, 2018) <https://arxiv.org/abs/1812.00479>, 5.
8. Hassan, Eman T., Xin Chen, and David Crandall. "Unsupervised Domain Adaptation Using Generative Models and Self-Ensembling." (2018) *arXiv preprint arXiv:1812.00479*.
9. Mostofa Ahsan, Kendall E. Nygard, Rahul Gomes, Md Minhaz Chowdhury, Nafiz Rifat, and Jayden F. Connolly. "Cybersecurity Threats and Their Mitigation Approaches Using Machine Learning - A Review." (July 2022) *Journal of Cybersecurity and Privacy* 2, no. 3, https://www.researchgate.net/publication/361906485_Cybersecurity_Threats_and_Their_Mitigation_Approaches_Using_Machine_Learning-A_Review, 527-555.
10. Hanan Hindy, Robert Atkinson, Christos Tachtatzis, Jean-Noël Colin, Ethan Bayne, and Xavier Bellekens. "Utilising Deep Learning Techniques for Effective Zero-day Attack Detection." (October 14, 2020) *Electronics* 9, no. 10, <https://www.mdpi.com/2079-9292/9/10/1684>.
11. Iqbal H. Sarker, A. S. M. Kayes, Shahriar Badsha, Hamed Alqahtani, Paul Watters, and Alex Ng. "Cybersecurity Data Science: An Overview from Machine Learning Perspective." (2020) *Journal of Big Data* Vol. 7, No. 1, <https://journalof-bigdata.springeropen.com/counter/pdf/10.1186/s40537-020-00318-5.pdf>, 1-29.
12. Yang Guo. "A Review of Machine Learning-based Zero-day Attack Detection: Challenges and Future Directions." (January 15, 2023) *Computer Communications* Vol. 198, <https://www.sciencedirect.com/science/article/abs/pii/S0140366422004248>, 175-185.
13. Pierre Parrend, Julio Navarro, Fabio Guigou, Aline Deruyver, and Pierre Collet. "Foundations and Applications of Artificial Intelligence for Zero-day and Multi-step Attack Detection." (April 24, 2018) *EURASIP Journal on Information Security*, <https://jis-urasipjournals.springeropen.com/articles/10.1186/s13635-018-0074-y>, 1-21.
14. Jin Chen and Lei Xu. "Zero-Day Exploit Detection Using Machine Learning." (November 8, 2022). Unit 42, Palo Alto Networks. <https://unit42.paloaltonetworks.com/injection-detection-machine-learning/>.
15. Dhruv Nandakumar, Robert Schiller, Christopher Redino, Kevin Choi, Abdul Rahman, Edward Bowen, Marc Vucovich, Joe Nehila, Matthew Weeks, and Aaron Shaha. "Zero-day Threat Detection Using Metric Learning Autoencoders" (2022) in *2022 1st IEEE International Conference on Machine Learning and Applications (ICMLA)*, <https://arxiv.org/abs/2211.00441>, 1318-1325.

THE CYBER DEFENSE REVIEW

◆ SENIOR LEADER PERSPECTIVE ◆

Leadership Matters

Lieutenant General (Retired) Edward C. Cardon

Looking back to 2010 and the creation of Army Cyber Command, I am impressed by the tremendous progress in developing, employing, and sustaining those capabilities required to operate in the cyber domain, including the electromagnetic spectrum, at scale. Advances in cyber policy, force design, force development, and operational deployment are stunning with the creation of a branch, units, and operating concepts and practices. Yet this progress remains fragile, and we must continually adapt because change is an enduring characteristic of war and in this domain, change happens at warp speed. There are significant opportunities and challenges ahead for our people and teams, as well as for our future operations and technologies. Future cyber forces will neither resemble nor operate like today's forces. Operational concepts are changing. We are experiencing revolutionary changes in technology, and much of that change applies both to the government and commercial sectors and can be used for good or evil. Thus, this note is a call for every cyber leader to exercise the fundamentals of leadership to ensure that our cyber force into the future will remain unequaled.

The Army's adaptation to build, sustain, and employ the initial cyber capabilities as envisioned is largely complete – but it is also insufficient. General Nakasone called this out at the Intelligence and National Security Alliance Leadership Breakfast in December 2023 saying, “I think all options are on the table except status quo.”¹ To engage this challenge will require fully and perpetually engaged leadership, willing to imagine and doing the hard work of leading change within the Army's cyber community.

© 2024 Lieutenant General (Retired) Edward C. Cardon



Lieutenant General (Retired) Edward C. Cardon is the United States Military Academy Academic Chair for Cyber and has served the Nation for over 36 years, including assignments in Germany, Bosnia-Herzegovina, Iraq, and Korea. From September 3, 2013 to October 14, 2016, he transformed and scaled U.S. Army Cyber Command into the peerless cyber force we know today, with new organizations (including Task Force Ares), operational constructs, headquarters, and talent models that served as the foundation for the ARCYBER of today. Since retirement he has created a portfolio approach focused on helping individuals and teams tackle and solve intractable problems.

The cyber domain is a man-made, contested, and competitive domain that is continuously evolving and adapting with the ever-changing convergence and divergence of people, technologies, and processes. It is characterized by disruptive technologies and applications. Time is a critical component – software changes at the speed of new code, hardware at the speed of new chips, and people at the speed of new ideas. This domain will always be somewhat unstructured; hence, today’s established law, authorities, regulations, processes, structure, and concepts may be inappropriate, or even dysfunctional, tomorrow. This dynamic renders even more important the question: why?, which in turn forces continuously revisiting the purpose behind our missions and operations, especially when the current structures, processes and tools increasingly become constraining rather than enabling.

As the character of war in competition, crisis, and conflict evolves, so too must our strategies and capabilities with constant invention, innovation, and integration of cutting-edge technologies and operations. Our Soldiers, Civilians, Industry and Academic partners from Army Cyber Command and the supporting/supported Army and Joint organizations will lead this effort to increase the effectiveness of our Army and further the broader missions of our Nation. Our storied history of past achievement was built by the dedication of those who went before us. Going forward, the sheer pace of change renders that pioneering spirit more important than ever. As leaders, we are responsible for inspiring and motivating those we lead not only to accomplish today’s missions, but even more important, do the hard work of understanding, visualizing, describing, and directing tomorrow’s cyber force.

People thrive when driven by a shared sense of purpose, impactful mission, and a strong sense of community. Cyber is brimming with opportunities for growth and contribution. Embracing the traits and norms of dynamic, innovative cultures is a collective endeavor

that requires internalization and adaptation. Effective communication and idea-sharing are paramount in this endeavor. It aligns seamlessly with the Army's renewed emphasis on professional writing, exemplified by the Harding Project.² The power of writing extends beyond mere documentation; it directly shapes policy, disseminates valuable lessons learned, and informs critical debates within the Army, particularly regarding realms like cyber and electronic warfare. As such, this journal must evolve into one of the primary platforms facilitating such exchanges within our ranks. We must all scrutinize ourselves with a critical eye, seeking avenues for improvement and sharing insights that elevate our collective community. While introspection is essential, we must also cast our gaze outward, exploring novel approaches to capability development and operational enhancement. Only by synthesizing external perspectives with our internal realities can we maintain a perpetual combat edge in effectiveness and resilience. Effective leaders improve the community every day, always making it better than they found it.

Good leaders must have a competitive mindset – to be and stay the best. Harvard Business School Professor Michael E. Porter, the author of *The Competitive Advantage of Nations*,³ wrote extensively on competition, encouraging a focus on gaining, maintaining, and sustaining a competitive advantage through a culture of continuous improvement, innovation, and inventions. This mindset is essential to success in our domain. For example, artificial intelligence and large language models are causing sweeping changes today across the commercial sector. How is the Army moving at speed and scale in harnessing these technologies for cyber and electronic warfare? Competitive edge is not limited to technologies – the same exists for concepts. The concept “hunt forward” has been outstanding in delivering advantage for our partners and allies. What is the next big concept? Given advances in digitization, mobile, cloud adoption, advanced connectivity, and artificial intelligence, new concepts are needed at scale for both offense and defense.

Technology and concepts alone are inadequate – our real competitive advantage remains our people. The way we recruit, train, educate, deploy, and retain our force also must adapt. Again, the warp speed pace of change today – some highlighted above – requires everyone in our force to be life-long learners, perpetually looking outside the box. I also believe in permeability – people need to move seamlessly in and out of our cyber forces. For example, the advances in large language models are moving much faster in commercial industry. Large language models may have the highest probability of causing disruptive levels of change within our force in the near term. How best can we leverage our commercial sector capabilities?

Aviation efforts to build, develop, and test new aircraft and new technologies are commonly referred to as “pushing the edge of the envelope.” The flight envelope of an aircraft is commonly referred to as a Vg diagram, the linear representation of speed and altitude, representing the upper and lower limits of speed, power, maneuverability, altitude, etc of a given

aircraft.⁴ Pushing the edge of the envelope entails looking “out there,” to and even beyond the outer limits, to see where we can go, what we can do, and what we can learn. As cyber leaders, we must call upon ourselves, and each other, to push the envelope’s edge. Just like aviation, each specialty in cyber has its own Vg diagram with limits of what is technologically possible today. Pushing the limits of AI, machine learning, the electromagnetic spectrum, and quantum computing will provide us, or our adversaries, the leverage to shape the future to our ends. If we don’t do it, it will be shaped by an adversary that gets there first. We should strive always to know where the edge of the envelope is and push out past that edge.




Figure 1. The Vg Diagram represents the velocity vs G Load for an aircraft. Each aircraft has its own Vg diagram that is valid at a certain weight and altitude. The envelope of flight is in the center, labeled ‘Normal operating range.’ Pushing the envelope involves pushing the limits of the boundaries of this box in terms of acceleration, maneuvering speed, load airspeed, etc.⁵

Our recent past has brought us to this point, but not all our histories are successes. Hence, leaders at every level must direct a sharp eye to everything we do and strive to be the best. Our near and distant future is where we build tomorrow’s capabilities and weapons. Moreover, across whole of government, academia, industry, partners, and particularly from our adversaries, there are lessons to be learned. Be curious about the world outside your lane. See what others are doing, how they are competing for tomorrow’s high ground in cyberspace, synthesize it with what you know, and, whenever appropriate, share with your teammates in forums like this journal.

Our margin for error is often paper thin, and our teammates in other domains count on us to get it right and get it right faster than any adversary. In ancient times, warriors with sword and shield were limited by muscle power in their ability to impact the battlefield. Thousands of years later, dozens of warfighters in several tanks, or a squadron of aircraft, or a submarine, could leverage technology and defeat thousands of combatants. Eighty years ago, a small number of cryptologists broke codes and ciphers and changed the course of history. Tomorrow, several warfighters on keyboard may disrupt, defeat and maybe even destroy nation-states as we know them today. The future is there to be seized by those who constantly hone the skill and the will to execute. We must be hard on ourselves, as we burn the midnight oil and strive to master our professions, to be the best we can be at what we do. We can, and must, both deter and, whenever necessary, defeat those future adversaries that are training, experimenting, learning, and competing to defeat us. We build strength through demonstrated capability that gives our competitors and adversaries pause. For our leaders, that competition is already here.

Leaders lead change. Change and innovation can be top down but the best new ideas often come from those in the fight who can see problems first-hand, what succeeds and fails, and possible solutions that work. The cyber and electronic warfare community in particular needs a free exchange of ideas across the force to be more agile than our adversaries. I call on our force, especially our leaders at all levels, to embrace change and make our community better. Write and publish: in this journal, in our schoolhouses, and at the edge where our warfighters are in contact with our adversaries every day. Leadership matters. 🛡️

DISCLAIMER

The views expressed here are those of the author and do not reflect the official policy or position of the U.S. Military Academy, U.S. Army, U.S. Department of Defense, or U.S. Government.

NOTES

1. General Paul Nakasone, Commander, U.S. Cyber Command, (December 8, 2023). Remarks, Intelligence and National Security Alliance Leadership Breakfast.
2. Harding Project, <https://mwi.westpoint.edu/introducing-the-harding-project-renewing-professional-military-writing/>
3. Michael E. Porter, *The Competitive Advantage of Nations*. (New York: The Free Press, 1999).
4. The term “flight envelope” comes from the Journal of the Royal Aeronautical Society in 1944. It was used to cover all probable conditions of flying and maneuvering an aircraft. In 1978, *Aviation Week & Space Technology* used the term “push the envelope” to describe the way NASA pilots would test the boundaries of altitude for aircrafts. In 1979, Tom Wolfe used the term in his book *The Right Stuff* and it became common language for getting out of your comfort zone. <https://www.phrases.org.uk/meanings/push-the-envelope.html>.
5. Federal Aviation Administration, “Aerodynamics of Flight” in *Pilot’s Handbook of Aeronautical Knowledge, FAA-H-8083-25C*. (November 3, 2023). https://www.faa.gov/sites/faa.gov/files/07_phak_ch5_0.pdf, p. 5-38.


THE CYBER DEFENSE REVIEW


CONTINUE THE CONVERSATION ONLINE

 CyberDefenseReview.Army.mil

AND THROUGH SOCIAL MEDIA

 Facebook [@ArmyCyberInstitute](https://www.facebook.com/ArmyCyberInstitute)

 LinkedIn [@LinkedInGroup](https://www.linkedin.com/company/ArmyCyberInstitute)

 Twitter [@ArmyCyberInst](https://twitter.com/ArmyCyberInst)
[@CyberDefReview](https://twitter.com/CyberDefReview)



ARMY CYBER INSTITUTE ♦ WEST POINT PRESS



WEST POINT PRESS



THE CYBER DEFENSE REVIEW IS A NATIONAL RESOURCE THAT CONTRIBUTES
TO THE CYBER DOMAIN, ADVANCES THE BODY OF KNOWLEDGE,
AND CAPTURES A UNIQUE SPACE THAT COMBINES MILITARY THOUGHT
AND PERSPECTIVE FROM OTHER DISCIPLINES.