NAVAL SCIENCE AND TECHNOLOGY
FUTURE FORCE™
VOL. 5, NO. 3, 2018

WHAT WILL AI MEAN FOR WARFIGHTERS?

THE RED TEAM: THE "GOOD" BAD GUYS

A DECADE OF VIRTUAL WORLDS

ARTIFICIAL INTELLIGENCE

VOLUME 5
#03
2018

FUTURE FORCE
NAVAL SCIENCE AND TECHNOLOGY MAGAZINE




IN THIS ISSUE

COLUMNS

To submit an article or subscribe to *Future Force*, please visit our website or contact the managing editor.

Future Force Magazine
Office of Naval Research
875 N. Randolph Street, Suite 1425
Arlington, VA 22203-1995

Email: futureforce@navy.mil
Phone: (703) 696-5031
Web: http://futureforce.navylive.dodlive.mil
Facebook: http://www.facebook.com/navalfutureforce

BEYOND FUTURE FORCE ▼

Interactive features are enabled with the digital copy of Future Force:

futureforce.navylive.dodlive.mil

Mobile Download




Future Force is a professional magazine of the naval science and technology community. Published quarterly by the Office of Naval Research, its purpose is to inform readers about basic and applied research and advanced technology development efforts funded by the Department of the Navy. The mission of this publication is to enhance awareness of the decisive naval capabilities that are being discovered, developed, and demonstrated by scientists and engineers for the Navy, Marine Corps, and nation.

This magazine is an authorized publication for members of the Department of Defense and the public. The use of a name of any specific manufacturer, commercial product, commodity, or service in this publication does not imply endorsement by the Department of the Navy. Any opinions herein are those of the authors and do not necessarily represent the views of the US government, the Department of the Navy, or the Department of Defense.

Future Force is an unclassified, online publication. All submissions must be precleared through your command's public release process before being sent to our staff. To subscribe to *Future Force*, contact the managing editor at futureforce@navy.mil, (703) 696-5031, or *Future Force* Magazine, Office of Naval Research, 875 N. Randolph Street, Ste. 1425, Arlington, VA 22203.

All photos are credited to the US Navy unless otherwise noted.

Front Cover: Designed by Jeff Wright

By Capt. Robert Palisin, USN (Ret.)

Photo by John F. Williams

Today, we can create, capture, and process data in ways that were once unimaginable. We can envision a thousand new ways to utilize artificial intelligence (AI)—and indeed, that's part of our task here at the Naval Research Enterprise. AI is not a new field, but what is new are the ways we can employ its different components in support of America's Navy and Marine Corps, and how it will transform the business of research and development.

The chief of naval research, Rear Adm. David Hahn, has tasked the naval research community not only to envision the future of AI for the Department of the Navy (DoN), but also to create better ways to get AI into the hands of our Sailors and Marines. We need a practical implementation model for AI in real-world systems—a new business model to get from ideation to implementation much faster than our current naval research-and-development system allows. And not over decades, or years, but now.

We must figure out how to rapidly take AI "to scale"—the more pervasive it is, the more effective, efficient, and useful it will become. Our men and women defending freedom around the globe deserve no less.

We're in a new era for this field that is exciting and challenging. With topics ranging from data management to machine learning and artificial neural networks, the sky is the limit (actually, there is no limit) to what we can do to advance AI for our nation's security. The naval environment is highly complex, and naval AI is a chance for our researchers and partners to be involved with some of the toughest AI challenges possible, as we improve our multidomain warfare effectiveness.

AI will make our Navy and Marine Corps more effective because it will make our Sailors and Marines more responsive as we sense, think, decide, and act faster than our adversaries. In fact, "augmented intelligence" may be a better way to define AI and how it will affect naval forces.

Naval AI also is a chance to effect important new policy that will guide the development and use of AI-enabled capabilities. The mission, ethics, and shared values of the DoN bound our work on AI, while the accountability framework of naval command ensures human judgment has the final say in the use of force.

This issue offers examples of exciting work in naval AI research. The authors are not only helping the Navy solve some of its most vexing challenges, but their expertise contributes to policy that guides future development and use of AI-enabled capabilities. ONR recently hosted the A2I Summit in San Diego to learn from AI industry/academia leaders, and much of what was discussed is relevant to this issue of *Future Force.* In fact, we started a website, A2IQ.com, at the summit to collect feedback/inputs about AI investments ONR should consider and how we can do the business of AI research, better. I invite your thoughts.

So how do we get from theory to reality? We must focus on the right AI applications and figure out how we can work better together in the naval research community, and with leading researchers and companies to accelerate new AI-enabled capabilities.

Capt. Palisin is the artificial intelligence portfolio director for the Office of Naval Research.

Artificial Intelligence

Rear Adm. David Hahn, chief of naval research (center), and Dr. David Walker, research and development portfolio director at the Office of Naval Research (second from left), tour Prof. Howie Choset's Biorobotics Laboratory at Carnegie Mellon University while attending an artificial intelligence and autonomy for humanitarian assistance and disaster relief workshop co-hosted by the two organizations.

Photo by MC2 Jess Lewis

WHAT WILL ARTIFICIAL INTELLIGENCE MEAN FOR WARFIGHTERS?

By Dr. Stuart H. Rubin, Dr. Matthew Stafford, Dr. Sukarno Mertoguno, and Dr. Jamie R. Lukos

LAUDED BY SOME, FEARED BY OTHERS, ARTIFICIAL INTELLIGENCE IS AN IMPORTANT PART OF THE NAVY AND MARINE CORPS' FUTURE. FROM WARGAMING TO PHOTO RECOGNITION, ARTIFICIAL INTELLIGENCE IS BECOMING AN IMPORTANT AND INTEGRAL TOOL IN THE DAILY WORK OF SAILORS AND MARINES.

Artificial intelligence (AI) has recently risen in prominence to become a part of the public psyche. It promises, and more or less delivers, everything from keeping warfighters out of harm's way to enabling scientific discoveries and engineering achievements that formerly were the exclusive domain of humans. Detractors have gone so far as to attempt to block AI's progress by claiming that its advancement is inimical to human interests. This exposé attempts to put the promise and the pitfalls of AI into perspective by explaining what it can now do, what it will likely do in the future, and why the development of AI serves the interests of the free world.

Introduction

AI was born at Massachusetts Institute of Technology (MIT), Stanford, Carnegie Mellon University, and the Department of Defense (DoD) in the 1950s. Stanley Kubrick featured it in the movie *2001: A Space Odyssey*, under the tutelage of MIT professor Marvin Minsky. Many other Hollywood movies have been released since, casting AI as a villain that enslaves or threatens humanity. As is so often the case, however, Hollywood has the story backwards.

AI is simply the next step in human evolution. It has been and will continue to be shaped in the form of human associates. AI is an evolving collection of computational techniques for solving problems. It started by learning to play games such as tic-tac-toe, then continued with checkers, chess, Go, and virtually everything in between.

One of the first things researchers discovered in working with AI was that problems assumed to be difficult (e.g., playing chess) were much easier to solve than was thought; and problems, thought to be fairly easy (e.g., distinguishing a photo of a cat from a photo of a dog) were more complicated.

As AI progressed from merely learning and playing games, it moved into applications involving game theory. This was attractive to the national defense establishment. The Defense Advanced Research Projects Agency (DARPA) was one of the earliest sponsors of AI. DARPA understood that simple battle management—concepts such as red-versus-blue teaming—could involve computations of extraordinary complexities, much in the same way a competitive game of chess does. Game theory gave rise to the subfield of AI known as heuristics—a Greek word meaning "serving to discover." Heuristics are rules of thumb that save computer processors enormous amounts of time allowing them to solve far more complex problems than would otherwise be possible. An example of a heuristic within a war game might be: "The enemy will seek the simplest approach to causing disruption and successful strategies are more likely to be repeated".

Neat vs. Scruffy AI

In the 1990s, two AI camps evolved: the symbolic folk, or neats and the neural folks, or scruffies. The symbolic approach to AI is to define one or more types of symbolic logics (e.g., the predicate calculus) and use those logics to deduce, induce, and/or abduce knowledge about the world, using whatever knowledge is supplied (e.g., the axiomatic set). For example, I may be supplied with the knowledge that the grass is wet; one can use that knowledge to abduce that it must have rained last night. Neats maintain that if one builds a society of such reasoning engines, then a very useful AI will emerge from the collective. Neats build their AI by leveraging and combining many forms of symbolic reasoning. An important part of the neat approach is the advantage of being able to explain conclusions and recommendations. Such explanations are vital to building user trust, but symbolic logics are difficult to apply to conceptual entities that defy symbolic representation (e.g., an engine that sounds different).

These contentions gave rise to "scruffy AI." Scruffies maintain that much of what we sense in our environments and reason about cannot be symbolically captured—at least not easily. Scruffies point to problems in computer vision (e.g., being able to distinguish a dog from a cat) and hold that these are not problems best addressed through symbolic reasoning. Scruffies prefer to address such problems through the use of neural networks, which run the gamut from the early perceptrons to the recent deep learning neural networks. Such scruffy approaches, however, can be extraordinarily costly and time consuming to train and, unlike the neat approaches, do not lend themselves to trust-building explanations of proposed actions.

Which approach is preferable? The answer is sometimes one, and sometimes the other, but more often the answer is both. That is, neural and symbolic approaches are compatible, and complex systems can be evolved to use the best features of each to solve specific problems. For example, in simulating wargames, not unlike the game of chess, we can use symbolic reasoning to reason out our next move. The possibility space is so large, however, that we need some form of pruning to be able to search out the best alternatives. Neural nets can and have served as heuristics, which learn to guide the search based on what has worked best in the past.

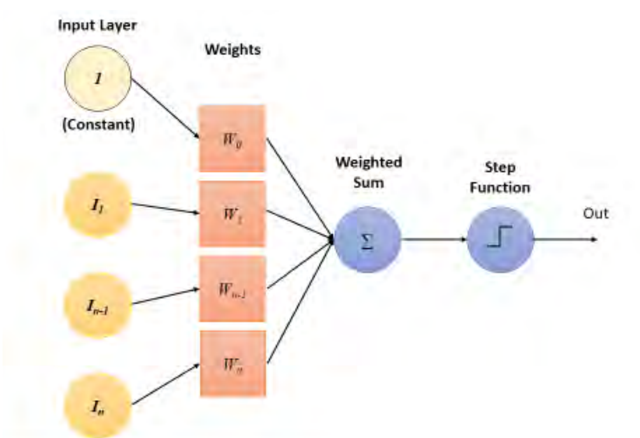

Figure 1: Perceptron. Here, a set of weights is adjusted so that various sets of inputs (e.g., visual pixels) are properly mapped to similar trained inputs (e.g., similar learned images). The step function ensures that these mappings fall into the "best well" before being output. Inputs can be excitatory (i.e., positive) or inhibitory (i.e., negative), just as in the natural brain. The difficulty lies in adjusting all of the weights so that a newly learned pattern does not unlearn one or more previously learned patterns.

Instead of being unnecessarily constrained by human-invented nomenclatures, we need to take an engineering point of view. This means we need to build what works and iteratively optimize it using both symbolic and neural AI to do so. One of the fundamental tenets of the most-advanced AI systems is the use of self-reference. Not only is this the foundation of computability theory, it also is what separates intelligent lifeforms from lifeless machines and even defines our personalities.

Unfortunately, these devices did not always converge so they could not properly map objects unless they had been exposed to previous training on the category. Without the training, these AI engines often were unable to recognize images. Unfortunately, these engines did not always notify users of their inability to recognize images, resulting in erroneous guesses that reduced confidence in the capability.

Perceptrons evolved into neural networks consisting of hidden layers that distinguished images with greater accuracy. There was, however, a drawback. As noted in a classic paper by Lin and Vitter, training these algorithms—specifically using backpropagation as invented by Prof. Geoffrey Hinton at the University of Toronto in 1986—ultimately resulted in failure.[1] No matter how well they worked, they ground to a halt once the size of the training data set grew to an unmanageable point.

About this time, the first papers on deep learning appeared.[2] The purpose of hidden layers in neural networks was to generalize inputs into outputs. Deep learning had more than one hidden layer—in some cases, there were hundreds of hidden layers. The debate on just how many hidden layers is sufficient continues today. In theory, two hidden layers are sufficient to recognize any object. Indeed, the concept of building two hidden layers led to the discovery of support vector machines (SVMs) for linear discrimination. SVMs are supervised learning models that analyze data to support classification into categories within a domain model. SVMs remain competitors for hidden-layer neural networks, although the latter have proven more popular because these networks do not obligate users to build working domain models.

The neats have their preferred AI, most notably expert systems and case-based reasoning (CBR). Expert systems work by acquiring the rules of a domain expert and replaying them into the solution of a domain problem.

CBR works by recording a case base of what has worked in past situations and matching the current situation against this experience base to find the best-matching solution, which is then practically adapted. CBR has been applied to problems in battle management, corrosion-resistant alloy designs, and the like. CBR systems are more manageable (because they scale) than expert systems because, unlike expert rules, cases are mostly independent of one another.

Knowledge Amplification Employing Structured Expert Randomization (KASER) was invented at the turn of the

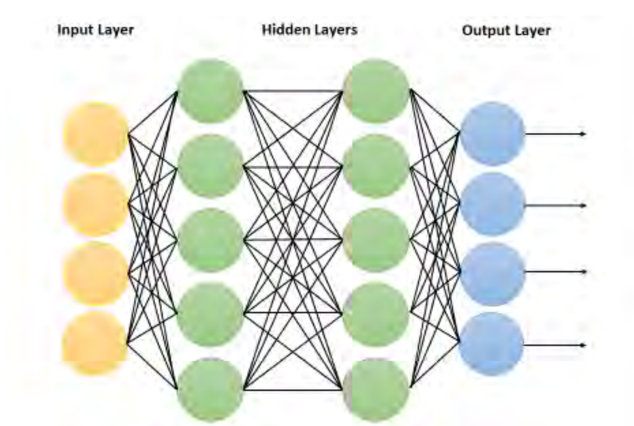

Figure 2: Hidden-Layer Neural Network (Deep Learning). In theory, at most two hidden layers are sufficient to distinguish any complex pattern, but hundreds of hidden layers have been used in practice. Conceptually, each successive hidden layer learns an increasingly abstract representation of its input. This is useful for processing natural language and embedded images. Unfortunately, as the training time for each hidden layer is intractable, deep learning is not the final word.

millennium.[3–5] These symbolic (neat) systems create new knowledge through the use of analogical reasoning, and, in so doing, ameliorate the knowledge acquisition bottleneck—the empirically sluggish process of programming knowledge.[6] The knowledge acquisition bottleneck is what makes knowledge-based systems costly—not the hardware platforms or the software codes. We note that neural (scruffy) systems are incapable of modus ponens (i.e., deductive reasoning). Thus, while a strictly neural KASER is not possible, it should be possible to conceptually achieve the neural KASER using symbolic operations. This is an area for future research and development.

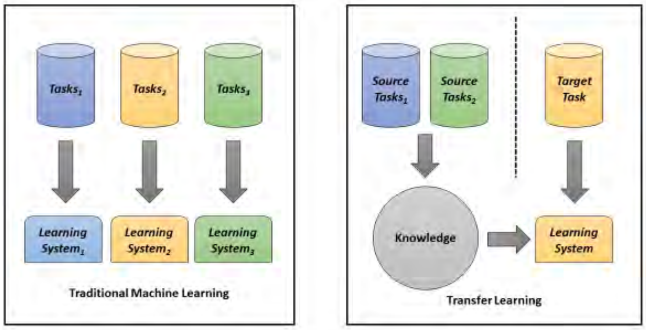

Figure 3: Traditional Machine Learning vs. Transfer Learning. On the left we see that traditional machine learning systems are brittle, using separate learning systems for each task with little or no integration. Conversely, on the right, we see that transfer learning applies acquired knowledge to the acquisition of derived knowledge. Transfer learning acquires some, but not all, knowledge by learning maps for existing knowledge, which are also knowledge, and so on.

KASERs have been applied to military problems ranging from predicting multiple faults in the Air Force F-16 ANG-68 radar system to helping automate naval record keeping and medical report generation.[5] As it turns out, future AI systems may become tractable (and, therefore, run far faster) simply by having a basis in symbolic AI instead of scruffy AI. This would enable AI systems to anticipate, learn, and compute previously unknown enemy countermeasures in real time—again, something that hidden-layer neural networks cannot do.[1] When it comes to information dominance, AI is king—i.e., the right kind of AI. This is where the AI contest with China is currently taking place.

The term transfer learning refers to knowledge acquisition, which is facilitated by having acquired knowledge in a similar domain. For example, humans having learned checkers will transfer more or less of that knowledge into helping them learn chess. As a result of transfer learning, they will learn chess faster and better than otherwise. Current computer checkers and chess programs are incapable of transfer learning. Were they so capable, it would be a most-worthy breakthrough in machine learning. Transfer learning holds the promise of making much, but certainly not all, of intractable learning times tractable.

Scruffy AI cannot do transfer learning. This was demonstrated in a classic experiment by Hosseini and Poovendran.[7] They showed that the deep neural networks that are currently finding their way into hundreds of military subsystems, once trained on flashcards containing integer images, cannot do any better than random chance in recognizing photographic negative images—something children can do even before they can talk. This is significant from a military perspective because of the knowledge acquisition bottleneck.[6] Humans understand the world through transfer learning. If deep learning is incapable of this, then there is no hope of it being scalable. This means that neural networks will need to be reinvented if the United States is to achieve global information dominance. Some of the obvious consequences of Hosseini and Poovendran's classic experiment follow below.

Without transfer learning, the fact that parallel AI first beat the world chess champion and later the world Go champion is relatively meaningless. It was simply a result of blind search and will, not scale. It means that if autonomous vehicle navigation systems are based on deep learning they can never learn to interpret complex sensor inputs properly in conjunction with dynamic mission parameters. Instead, as it stands now, heuristics will need to be built in or more powerful processors will need to be built. Their requisite training time, however, will always outstrip even quantum computers as they scale. Power demands must also be considered, beginning with midscale applications.

The sine qua non of AI is that it leverages knowledge and additional processing power in the solution of a particular nontrivial problem. Yet it is how one leverages processor cycles and even how one designs the (hybrid) hardware to process the AI (e.g., analog photonics vs. digital electronics) that can lead to superiority across the full range of military applications. Knowledge is king, but its representation is problem-specific. One size does not fit all. Military systems must be designed to be cost-effectively modifiable (the so-called open architecture), their designs must be knowledge based, and this knowledge need be captured for reuse and transfer learning. Taken together, all of this enables information dominance—a central requirement of modern warfare.

On Formal vs. Informal Data Mining

A few years ago the term data mining emerged to describe the capability of knowledge-based systems to extract hidden knowledge from formatted and/or free-text databases. As additional data sources were targeted by system queries, the capability became known as big data mining. Big data mining provides a heretofore unprecedented ability to conduct statistical analysis. For example, one can mine the acoustic signatures of helicopter engines to determine a preventative-maintenance schedule. Big data mining is analogous to scruffy AI—it is to be used when we don't know exactly how to represent causality.

Although the benefits of big data mining are trumpeted across the globe, why is it that one never hears of small data mining? Here, one applies the data to a more or less generic template and the fit determines the equivalent of a formal equation that is precisely aligned with the data. Many laws of physics can be discovered through small data mining (e.g., Kepler's second law of planetary motion)—except, of course, quantum physics, as it appears to only fit big data mining. Small data mining is analogous to neat AI—it is to be preferentially used when we know how to represent the information or at least the causal variables. For example, the distance, direction, and/or speed of a sonar or radar contact reflection is not statistical; we know the equations governing the return reflections. Conversely, in determining the type of object returning a reflection, we know equations for determining its size, but this does not necessarily allow us to determine the exact nature of the object returning the reflection (e.g., a C-17 or a B-52?).

The Air Force is already leveraging AI. Here, Secretary of the Air Force Heather Wilson visited the AFWERX-Austin, a new innovation laboratory, in June 2018. US Air Force photo by Senior Airman Gwendalyn Smith

In such cases, a probabilistic determination is the best we can do in the absence of additional data.

Akin to neat vs. scruffy AI, military systems will, more often than not, benefit from a combination of small and big data mining. That is, if you have a formal equation or symbolic codes to define some object, use it. Otherwise, use probabilistic inference. For example, if one senses an aircraft exhaust burst and an increase in its altitude, one can apply small data mining to determine the aircraft's short-term positions looking forward. Conversely, if the aircraft banks left and then right, the best that one can do (in the absence of additional information) is apply big data mining for a statistical prediction of its positions looking forward. Big data mining can be applied to the statistical selection of candidate models for small data mining, and, in this manner, combines the two approaches synergistically.

In complex military systems, there is a need to use small data mining insofar as possible and big data mining otherwise. Taken together, these AI-based data mining technologies can serve as an engineering associate for the constraint-based design of all manner of military hardware as well as a scientific associate aiding in the discovery processes. These technologies are great economic enablers for the military-industrial complex, offering a range of potential applications. From more efficient logistics to the design of more efficient jet engines to the discovery of projectile-resistant alloys, AI is inherently at the forefront of military planning, programming, production, purchasing, and employment.

On Genetic Algorithms, Evolutionary Programming, and Cyber-Security

Lower-level lifeforms evolve through chance search. Higher-level lifeforms evolve through chance search as guided by evolved heuristics. Genetic algorithms represent an attempt to mimic natural evolution in problem solving through chance search (e.g., antenna design) or heuristic search (e.g., logistic problem solving). Genetic algorithms were invented by the late John Holland. It is noteworthy to realize that just as evolution changes the genotype to change the phenotype, genetic algorithms are most effective when they operate on the parent algorithms in preference to the child data. Dr. Rubin illustrated this concept with the solution of problems taken from the domain of linear programming. Here, evolution operates on problem knowledge vice problem data so that every solution is reusable—minimizing the overall cost of solutions.

This approach is termed evolutionary programming (EP). It is defined as AI because knowledge is required to make sense of more-complex knowledge. This includes, of course, addressing the knowledge acquisition bottleneck.[6] EP has been used by the Air Force to successfully evolve a rule set to control a fighter jet engaged in autonomous combat. It also finds use in cybersecurity. More recent techniques such as semantic randomization are more resilient, as they can better detect and recover from cyberattacks.[8]

Semantic randomization is similar to its predecessor technique of syntactic randomization, entailing the compilation of codes using distinct compilers, but it applies AI to functional programming to create distinct semantically equivalent source functions.[8] Cyberattacks cannot be easily designed to infiltrate more than one distinct function at a time and thus are detectable. Syntactic randomization applies multiple algorithms (compilers) to create distinct object codes, which will vary in their behavior under cyberattack, which is then detectable. Semantic randomization applies multiple knowledge bases (AI) to create distinct, but equivalent algorithms, which will vary in their behavior under cyberattack, and is likewise then detectable. Syntactic and semantic randomization may be used together for increased cybersecurity. Semantic randomization is very effective, but it has been costly to manually create similar algorithms. Automatic programming techniques are under development, which promise to mitigate this cost.

Semantic randomization is a consequence of knowledge-based automatic programming.[8] Much progress has been made in this area by the Air Force through its annual knowledge-based software engineering conferences. Impressive Lisp programs (one of the earliest programming languages), which converted specifications to code were synthesized by Biermann using knowledge bases.[9] Similarly, Cordell Green wrote a compiler for the Navy, which converted one line of Refine code to an average of 50 lines of Fortran for use in submarine programs. These programs were impressive because they realized fifth-generation programming languages through AI insertion.

It may well be that nature put the cybersecurity problem into the AI evolutionary path to ensure that the software automation problem was concomitantly solved. AI, in the form of expert compilers, was successfully applied decades ago by Hindin to automate the generation of software.[10] Objects Inc. did likewise with their Layout language.

On Education and Training

An often-overlooked component of military readiness is education and training. The high cost of education and the ever-decreasing availability of financial aid is forcing some of the nation's students to embark on less-costly career paths. AI is not only cost effective, but when properly employed in tutoring systems, offers one-on-one training that is tailored to the needs of the individual and improves with each lesson and each additional student. AI enables question-and-answer learning, learning by doing, and self-directed discovery. When truly intelligent, distributed AI becomes ubiquitous; it will revolutionize the quality and capabilities of workers in every facet of society.

One new area will be the use of AI-based education and training to provide instruction on all manner of AI. Humans will focus on novel work—that which only humans can do—and delegate repetitive tasks to machines. The breadth of what defines repetitive tasks will increase over time. Although initially AI will have the potential to displace some human workers, it simultaneously has the potential to retrain them—all at significantly reduced costs compared to current training models.

Faster Hardware to Run AI

Given all the constraints that must be satisfied for an AI project to serve the joint-service needs properly, our group, based out of the Space and Naval Warfare Systems Center Pacific, has embarked on patenting and building a new type of AI. It is not EP nor is it a neural net, yet it incorporates the advantages of both. The goal is to develop a natural-language processing engine that enables access to the knowledge bases of a conventional helpdesk. It will map user queries onto known semantically equivalent helpdesk queries and similarly map helpdesk outputs to semantically equivalent outputs. Moreover, it will learn to create new maps through user interaction and through the use of a novel generalized AI. This will be a symbolic AI that learns through processor-intensive search operations.

Unlike existing hidden-layer neural networks this approach ensures that the AI remains tractable as scaled.[1] As a result, this new technology will be applied to the development of Airmen, mirroring one-on-one tutoring. It improves with use as its knowledge bases (including the conventional helpdesks) are augmented over time.

As it develops, this training capability may be expected to support multiple functions within the armed services

The Summit Computer at Oak Ridge National Laboratory, the current world's most powerful super computer.

and industrial and academic institutions as well. Last year, some of our group completed an AI-based, natural language-controlled unmanned aerial vehicle, termed a biologically inspired robotic drone (BIRD). We subsequently built a demo for an AI-based, natural language-controlled ship—the USS *John F. Kennedy*. These efforts demonstrated that the current help-desk project has both merit and a good chance of success because new neural networks, capable of tractable learning and domain transference, are indeed possible. Again, deep learning is not the end-all result.

The United States recently regained the distinction of building the fastest computer in the world—the Summit machine, built by IBM and running at Oak Ridge National Laboratory. The helpdesk AI experiment promises a capability that can use extreme processing power more fully. Thus, the experiment involves the use of one or more teraflop machines to run it. It will be scalable to use the speed of the Summit, and designed using an open architecture to facilitate the insertion of faster, more energy-efficient photonic hardware as that becomes ready.

In order to exercise the hardware and test related software, the experimental team plan on using a much smaller supercomputer—again, in the teraflop range, and leveraging this prototype for the larger computer, if and when needed. The researchers agree that the key to making all this work is a novel AI—the first of a class of AI technologies heralded by this article, that once proven, will require nurturing through additional research and development.

How can photonic hardware be used to support a faster AI? Photonic hardware replaces or augments traditional VLSI (very-large-scale integration), which switches using wires and electrons, with optical processing, which switches using light beams and photons. Unlike electronic crosstalk, light beams do not interfere with each other, can better support some types of massive parallelism, and can result in more energy-efficient architectures.

Rather than go into technical details here, consider the following example: Suppose AI is required to recognize an image of a certain person for whom it has been trained. Humans would recognize friends with or without a hat, but conventional AI cannot do this with any degree of reliability.

Thus, the question arises, why use time-consuming, power-hungry, digital hardware, when the extra accuracy does not contribute to the solution of the overall visual recognition problem? If one can design a more general AI, then one can select for faster, more energy-efficient types of hardware, on which it can be executed. This will result in a thousand-fold increase in speed with virtually no degradation in performance, but only if the AI can represent the problem at a more general level through the use of features. Features represent the core salient details of a problem minus computational irrelevancies. This is not a trivial task, but the potential payoff is enormous.

The Space and Naval Warfare Systems Center Pacific in San Diego is working with the Office of Naval Research and the Air Force to realize a more capable AI for interfacing helpdesks with natural language. The resultant helpdesks will be used for training Airmen, but are not so limited. The Air Force also plans to incorporate the technology with their vSED systems and other functions, which could benefit from off-loading redundant questions to the machine. Here, the AI determines if a question is known and if not, who to best route it to for a solution. It also learns from human-solved problems so that no human need solve the same or closely related problem twice—except to offer (conditional) alternative solutions. Government-owned, high-performance computers will be incorporated to expand the knowledge base to improve helpdesk solution capability as well as to enlarge the capability for natural language understanding. Finally, this project will attempt to achieve tractable scruffy transfer learning and otherwise fill in the great divide between symbolic and neural learning approaches by way of realizing actual brain functionality, in preference to the limited type discussed above.

On AI and National Defense

Defense undersecretary for research and engineering Michael Griffin recently created the Joint Artificial

Intelligence Center (JAIC). The JAIC's goal is to evolve US national defense strategy through the incorporation of AI applications. It should be added that it is vital that research and development into AI not necessarily be encumbered by applicative needs. Beautiful AI, like a flower, comes about for its own reasons. It needs to be unencumbered to support that genesis. Otherwise, AI will evolve through incremental innovations, the type of innovation that rarely leads to revolutionary improvements.

Just as the national strategy has moved to rapid prototyping to save on developmental costs, leaders in AI need to ensure that no more than a reasonable fraction of AI development funding is spent on incremental changes, chasing specific applicative needs. Otherwise, the United States runs an increased risk of being left behind by revolutionary foreign AI developments. Conversely, a reasonable fraction of AI development funding needs to be spent on those incremental, applicative needs. Otherwise, we run the risk of being outmanned and outgunned by conventional AI systems and capabilities. Ultimately, it appears that we are today exactly where Marvin Minsky, the "father" of AI, said we'd be, some 20 years ago:

Many AI researchers seem to have lost their way in recent years and should be investigating core problems such as how a person reasons. Today, artificial intelligence is going in a lot of directions, but it hasn't advanced in depth for quite a long time. There's no machine today that can look around a room and tell you what it sees. What bothers me is that; although people thought about such problems in the 1960s and 1970s, no one is thinking about them now except a handful of brave people—whereas 10,000 people are doing genetic algorithms and neural networks, not understanding where those methods might work and for what problems they were intended. People are trying to avoid the things that need to be done.

AI presents a discontinuity in human evolution. Humans are the first species to reach a point whereby our nonviolent evolution is not only possible but nearly assured. As AI develops, it will play an increasingly fundamental role in its own evolution—not only technically but socio-politically as well. It is clear that for the foreseeable future, a more capable US military force will emerge as a result of advancing AI but will remain simultaneously dependent on that advancement to retain its premiere place among the world's fighting forces.

References

1 J.H. Lin and J.S. Vitter, "Complexity Results on Learning by Neural Nets," *Mach. Learn.*, vol. 6, no. 3, 1991, pp. 211-30.

2 S.H. Rubin, ACM 18th Annual Computer Science Conference, Washington, DC, 22 February 1990. GMM-PAM: A Genetic Multilevel Multicategory Perceptron Associative Memory, pp. 366-72.

3 S.H. Rubin, S.N.J. Murthy, M.H. Smith, and L. Trajkovic, "KASER: Knowledge Amplification by Structured Expert Randomization," *IEEE Transactions on Systems, Man, and Cybernetics*—Part B: Cybernetics, vol. 34, no. 6, December 2004, pp. 2317-29.

4 S.H. Rubin, "A System and Method for Knowledge Amplification Employing Structured Expert Randomization (KASER)," US Patent No. 7,047,226, issued 16 May 2006.

5 S.H. Rubin, Chapter 17, "Knowledge amplification by structured expert randomization—KASERs in SoS design," in *System of Systems—Principles and Applications* (Boca Raton, FL: Taylor & Francis Group, 2009), pp. 421-50.

6 E.A. Feigenbaum and P. McCorduck, *The Fifth Generation* (Reading, MA: Addison-Wesley Publishing, 1983).

7 H. Hosseini and R. Poovendran, "Deep Neural Networks Do Not Recognize Negative Images," work supported by ONR grants N00014-14-1-0029 and N00014-16-1-2710, ARO grant W911NF-16-1-0485, and NSF grant CNS-1446866, 2017, pp. 1-3.

8 S.H. Rubin, W.K. Grefe, T. Bouabana-Tebibel, S.C. Chen, M.L. Shyu, and K.S. Simonsen, "Cyber-Secure UAV Communications using Grammatical Inference Amplified by Adaptive Waveform Synthesis and Evolution," *Advances in Intelligent Systems and Computing,* (Switzerland: Springer International Publishing), to appear December 2019.

9 A.W. Biermann, "Automatic Programming: A Tutorial on Formal Methodologies," *J. Symbolic Computation*, vol. 1, no. 2, 1985.

10 J. Hindin, "Intelligent Tools Automate High-Level Language Programming," *Computer Design*, vol. 25, 1986, pp. 45-56.

About the authors:

Dr. Rubin is a scientist in the Space and Naval Warfare System Center Pacific's intelligent sensing branch.

Dr. Stafford is the chief learning officer for Air Education and Training Command at Randolph Air Force Base, Texas.

Dr. Mertoguno is a program manager at the Office of Naval Research specializing in cybersecurity, cyber-physical systems, and higher-level secure software.

Dr. Lukos is a neuroscientist and branch head at Space and Naval Warfare System Center Pacific.

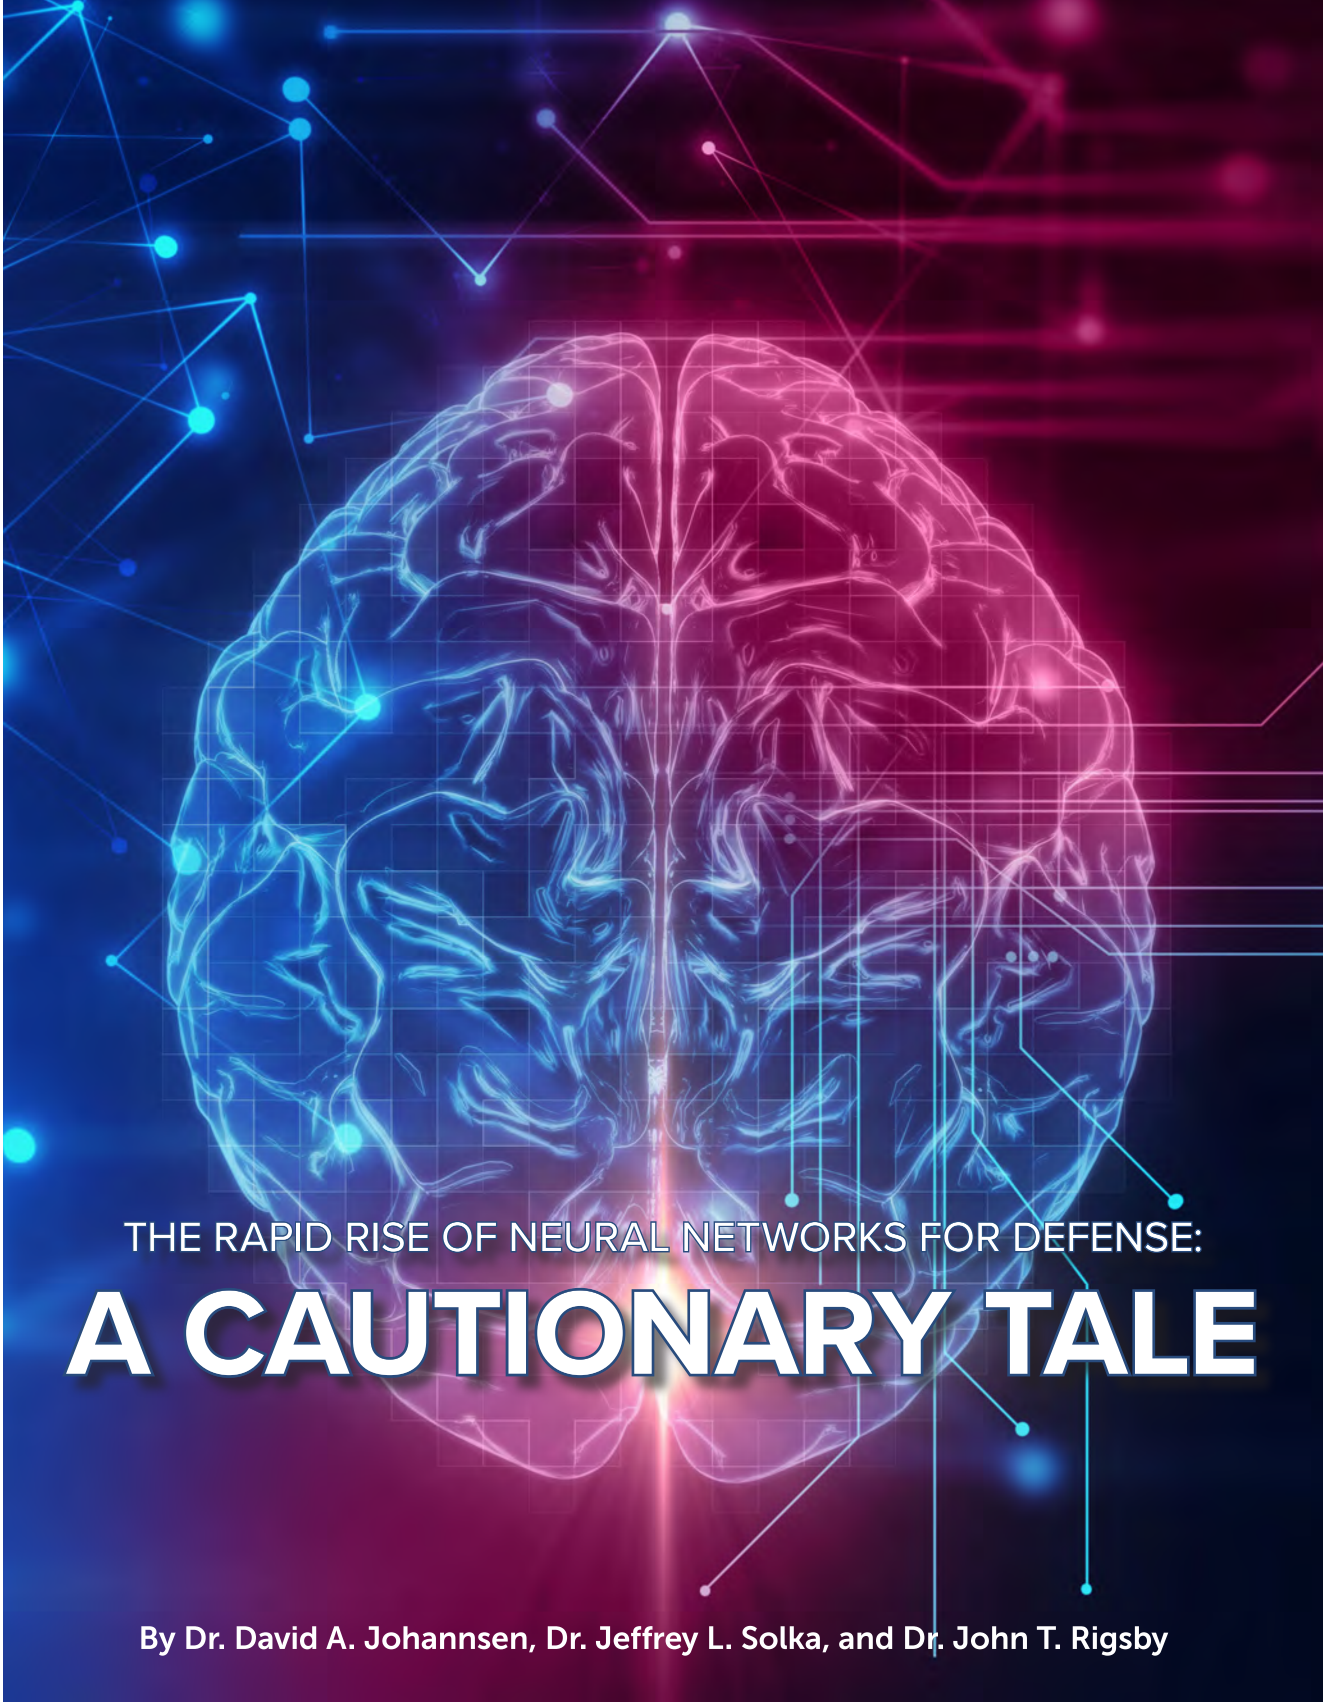

THE RAPID RISE OF NEURAL NETWORKS FOR DEFENSE:

A CAUTIONARY TALE

By Dr. David A. Johannsen, Dr. Jeffrey L. Solka, and Dr. John T. Rigsby

NO LONGER IN THE REALM OF SCIENCE FICTION, ARTIFICIAL INTELLIGENCE IS BECOMING UBIQUITOUS, IN THE CIVILIAN WORLD AS WELL AS THE MILITARY. BUT ARE WE MOVING TOO QUICKLY WITH A TECHNOLOGY WE STILL DON'T FULLY UNDERSTAND?

Fueled by rapid increases in computer storage capacity and processing power (principally through the use of graphical processing units) and the widespread availability of powerful software for designing and implementing neural networks (such as Google's TensorFlow), the application of artificial neural networks to significant problems in the real world has seen tremendous growth over the past several years. During this time, neural networks have demonstrated successes on problems as varied as automatic recognition of handwritten digits, automated image captioning and indexing, and have even beaten human masters in the game of Go. In fact, the field has reached the state of maturity that a person with only casual knowledge of computer programming and equipped with a modest PC can implement a neural network for whatever problem he or she might have at hand.

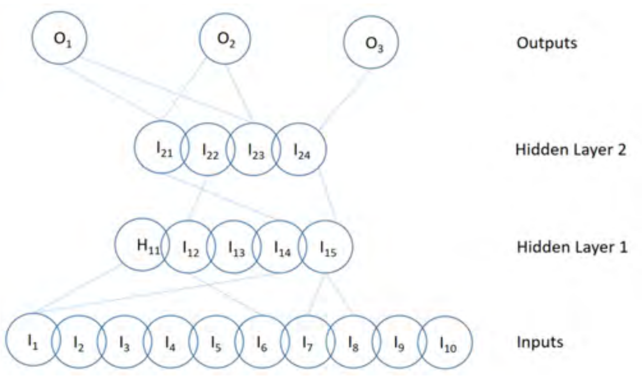

Figure 1: An example neural network with four layers.

Given this climate of success, there is growing interest in fielding neural networks in Department of Defense (DoD) systems. In this brief note, we will discuss the nature of neural networks in a language that can be broadly understood, especially in the context of the unique environment within DoD, so that Navy leaders can be better informed about the strengths and limitations of this technology as it impinges ever more frequently on the DoD. We will first attempt to explain to nonspecialists what an artificial neural network is. We will then discuss some of the inherent limitations of this class of machine learning tools and some of the ways that we and other members of the DoD are studying these limitations. Finally, we will discuss the consequences of these limitations.

What Is a Neural Network?

Rather than giving a precise mathematical definition of a neural network, we will begin by giving a functionally oriented definition. Thus, we describe a neural network as a nonlinear function from the space of inputs to outputs. The particular function is chosen from a broad class of nonlinear functions through a process known as training. Often, in current practice, the choice of nonlinear function is underdetermined; that is, the function contains more parameters to be learned than the number of observations that one has at hand for training the algorithm.

For example, in the context of image captioning, the space of inputs is the collection of all possible digital images, and the output space is the collection of all meaningful captions. The neural network accepts a digital image as input, and produces a caption. Along the way, hidden from the end-user, the computer treats the image as a mathematical object, performs mathematical operations on it, and then produces a numerical output (which is often a vector of probabilities of membership in the various classes). This vector of probabilities is then converted to a caption that is presented to the user.

We will now be a bit more precise in defining a neural network. The formulation of neural networks as a method of machine learning was motivated by analogy with the functioning of neurons in the human brain. Thus, neural networks consist of a set of nodes (neurons) with edges between them. Outputs of the nodes are multiplied by the weights associated with the edges and fed forward to the nodes in the next layer of the network. This information is adjusted by a threshold function associated with the node and then propagated through the neural network. In Figure 1 we present a four-layered neural network with an input layer, two hidden layers, and an output layer. Following our discussions above, one might imagine that each node in the output layer provides the probability of the input belonging to one of three classes.

A Virtue Is a Vice

Statistical pattern recognition has historically involved a somewhat standard pipeline, which is illustrated in Figure 2.

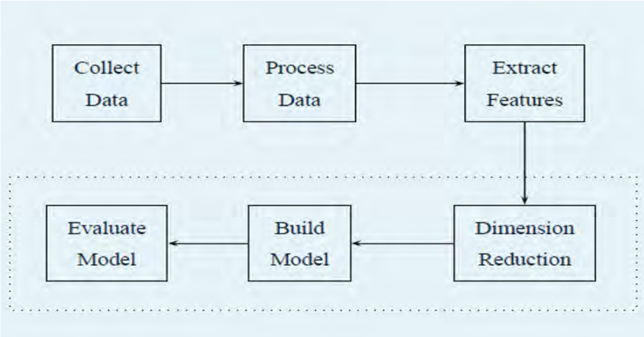

Figure 2: Pattern Recognition Pipeline.

The first three steps of this process (i.e., data collection, data processing/cleaning, and feature extraction) are often time consuming. If possible, a practicing statistician is well served to spend his or her time on the steps contained in the dotted box. The "extract features" and "dimension reduction" steps can be particularly daunting. The process of feature extraction and selection is usually the domain of experts and often will require significant time and experimentation to determine what features should be selected and utilized for the task at hand. Neural networks promise to revolutionize this pipeline by incorporating these two steps directly into the model building step without the need for experts. The virtue of artificial neural networks is that one can train a network to perform complex machine learning tasks such as interpolation, classification, regression, etc., without having to go through the process of feature generation and dimensionality reduction. In our discussions below we will focus on the consequences of automating these two steps of the pattern recognition process.

Before getting too far along, we should note that in some settings it may be impossible to generate features specifically tailored for each of the possible classes. For example, in automatic image captioning, as the image may be of anything in the world, there are virtually a limitless number of different classes. It is therefore not possible to specify optimal features to be extracted for each of the classification tasks. Thus, it is indeed a benefit of neural networks that they free scientists and engineers from the necessity of crafting specific features for each possible class. In many problem domains, however, neglecting to fully understand the processes that gave rise to the data (i.e., the training data) and then not actively participating in feature selection yields a model of which we have no understanding. In the following paragraphs we will briefly describe some of the implications of this aspect of neural networks.

The "Black Box"

If one reads the scientific literature on neural networks, one will quickly see that they are often described as "black boxes." What is meant by this? The complexity of a network trained to tackle nontrivial "real-world problems" is very high. That is, the internals of the network hide an incredible mathematical complexity. In fact, the complexity is so great that one cannot interpret how the input features provide a basis for the output. This is known in the statistics field as a nonattributable model. In the context of an application such as image captioning, this is of some concern. For example, if the neural network errs and gives the label "dog" to an image of a cat, the designer is troubled by not knowing exactly what features in the image caused the misclassification and therefore is unable to alter the algorithm to prevent future misclassifications.

This lack of insight into the relation between input and output is much more troubling for DoD applications, where the consequences of misclassification are often much more serious. For example, if one is designing an autonomous vehicle, one would like to be able to predict how the vehicle will react to a given input from its sensors. With a neural network, it is generally impossible to know the output of the network prior to presenting the input to the system and observing the output.

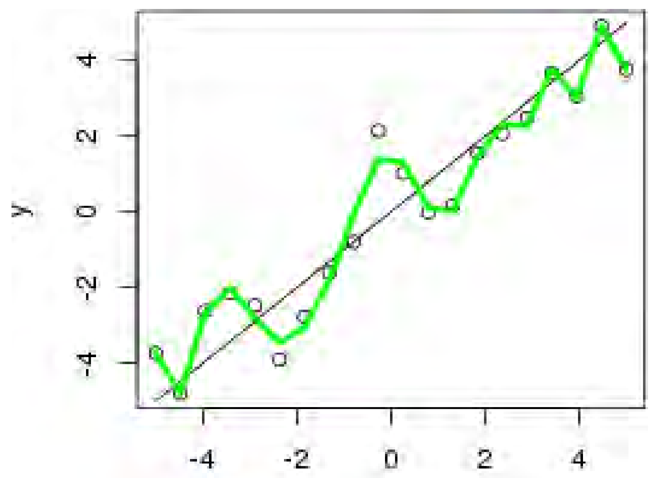

Figure 3: Simple linear model with white noise.

Generalizability

Generalizability is the ability of the model to produce reasonable output when presented with an input that is different from the data used to train the network. The issue of generalizability is a central concern in DoD applications: how will a fielded system perform with subtle changes to the environment? As we do not know much about either the processes that gave rise to the training data nor much about the "black box" nature of neural networks, it is impossible to predict the output. There are no guarantees of the behavior of the neural network, even when presented with inputs that do not vary substantially from the training data. In fact, there is currently no body of theory that governs the behavior of neural network outputs.

The class of functions that a neural network model is chosen from has tremendous richness and great power to approximate highly nonlinear behavior. Though this expressive power is useful for learning from training, model richness is a potential problem when one wants to predict/classify a new observation. The figure below represents how a model (green line) created with far too much complexity suffers from overfitting and does not generalize to the simple linear model (black line) that generates the data points (black circles).

Sufficient Training for Realistic Tasks?

As neural network models are potentially highly nonlinear, their predictive power can degrade quickly when presented with an input that differs significantly from the training data. In fact, because of the highly nonlinear dependence of network outputs on the inputs, significant degradation can sometimes occur when presented with an input that is nearly indistinguishable from one used in training the network. This problem is a significant one when one considers military applicability—a system that may be required to perform highly consequential tasks in an environment that may be at any location in the world. There is no way that one can accumulate enough training data to ensure that one is never presented with an input that differs from the training data.

In fact, in any real-world application, the complexity of the environment guarantees that the sensor will encounter inputs that differ significantly from the data used to train the network. As evidence of this phenomenon, consider the difficulties encountered with self-driving cars. Staying on a reasonably well-delineated road and interacting with objects that are frequently obeying a fairly rigid rule-set still contains enough "surprises" that we have recently witnessed several catastrophic failings. How much more difficult is the problem of deciding friend or foe in a jungle at night and in the rain?

Mitigating These Problems

Numerous organizations have begun programs to better understand these limitations. The Defense Advanced Research Projects Agency (DARPA) has started the Explainable Artificial Intelligence (XAI) program, which seeks to develop methods to better understand the decisions made by AI systems. DARPA also has begun the Lifelong Learning Machines (L2M) program, which seeks to develop machine-learning-based systems that provide the capability to train themselves in the field in the face of new environmental or mission-based conditions.

Our own organization, the Naval Surface Warfare Center Dahlgren Division (NSWCDD), also has begun efforts to help better understand these shortcomings. Our ongoing effort, "Neural Networks for Manifold Discovery," seeks to apply advanced mathematical methodologies to better characterize the fragility of neural networks and other machine-learning methodologies. Our new start effort, "Adversarial Learning for Robust AI," seeks to use recent research in "adversarial examples" to better understand how we can make neural-network-based systems more robust to environmental or enemy precipitated changes to operational environments. Both of these efforts were funded under the Naval Innovative Science and Engineering program, which is designed to serve as a major innovation catalyst for the naval surface warfare centers.

Final Comments

We hope that we have presented a fairly objective overview of artificial neural networks. We have tried to describe both the strengths of this class of machine learning algorithms, as well as illustrating some of their current limitations, especially in the unique environment of the DoD. We acknowledge the demonstrated successes of neural networks and believe that there are settings where the technology works very well; for example, developing AI for wargaming, planning, or training seems a very good use of the technology. In situations of complex environments where system performance errors have the potential for tremendous fiscal cost and potential to endanger lives, we need to be very cautious. We remain optimistic that programs at DARPA, NSWCDD, and other organizations can help better understand and ultimately mitigate these shortcomings. Until theory can catch up with practice, is a system whose outputs we can neither predict nor explain really all that desirable?

About the authors:

Dr. Solka is the chief scientist of Naval Surface Warfare Center Dahlgren Division.

Dr. Johannsen and **Dr. Rigsby** are members the applied mathematics and data analytics group at the Naval Surface Warfare Center Dahlgren Division.

A NEW METHODOLOGY FOR NONCOOPERATIVE TARGET IDENTIFICATION

By Thach Nguyen

IDENTIFYING AIRCRAFT THAT DO NOT CARRY IDENTIFICATION FRIEND OR FOE (IFF) TRANSPONDERS HAS ALWAYS BEEN A CHALLENGE. WHETHER IN PEACE OR WAR, GETTING IT WRONG CAN HAVE BIG CONSEQUENCES. NEW RESEARCH IS IMPROVING THE ALGORITHMS USED TO PROCESS SIGNALS OF INTEREST.

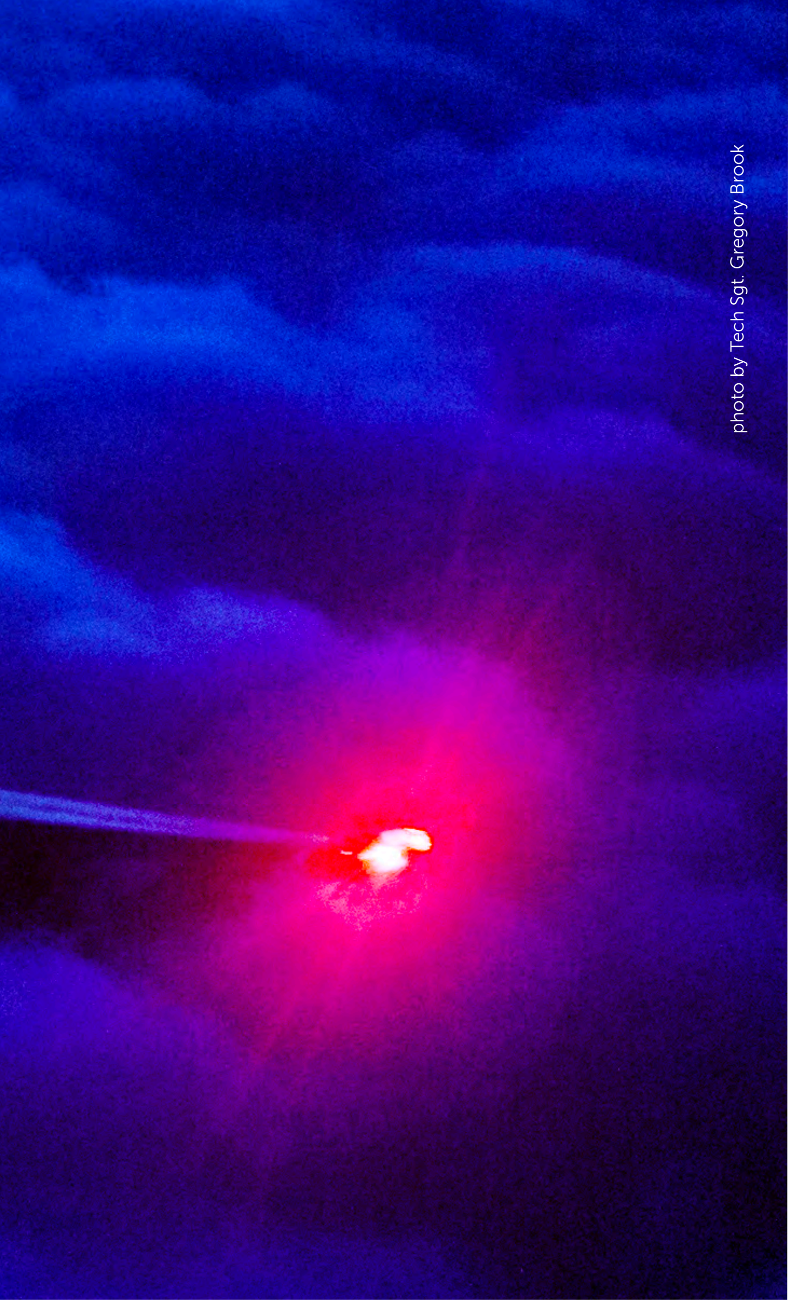

photo by Tech Sgt. Gregory Brook

In today's warfighting environment, noncooperative target identification (NCTI) systems are a valuable piece of the air combat identification puzzle. NCTI exploits the physical attributes of targets of interest by using the principles of physics, exploiting differences in reflected sensor signatures due to the physical attributes of targets not otherwise identified or carrying an Identification Friend or Foe (IFF) transponder. While current NCTI techniques can provide an effective means of general target identification, some targets differ only in their low-order NCTI features, which are not easily discernable using existing technology because of a low signal-to-noise ratio. Misidentifying these targets in battle scenarios could lead to serious errors causing unnecessary loss of life or systems. (Note that the word "target" is used here to mean any item of identification interest, including both friend and foe.)

In this paper, we present a technique [Ref 1] designated by the hidden Markov model (HMM) to improve algorithms for NCTI. There are many uncertainty parameters relating to signals of interest. These signals are randomly determined with respect to the amplitudes of the feature vectors; therefore, a statistical method of characterizing the spectral properties of the feature vector is required. In other words, given an observation of data, presumably representing an unknown signal, a statistical model of historical data is chosen that can most likely explain the sample. The technique is appropriate for the identification problem because of its ability to classify patterns based on a number of features that have certain types of underlying structures, especially if that structure results in a stationary feature distribution over some spatial period.

Our hypothesis of the technique is, if frequencies of the sensor data from the targets vary with time or differ only in their low-order features, or both, then cepstral coefficients can be used for identification purpose of the targets. The cepstral coefficients are in the cepstrum domain, which is a frequency of a frequency. The name "cepstrum" was created by reversing the letters of the first syllable of "spectrum." Cepstral coefficients have been used successfully in voice recognition [Ref 2], and the voice recognition problem is similar to our target identification problem in many ways.

Detailed Description of the Preferred Embodiments

The technique uses digital signal processing to identify unknown targets. In general, the unknown target is not "entirely" unknown because it is assumed to be a member of a group of one or more known targets. Thus, identifying the unknown target involves choosing its identity from one of the known identities in the group. The target may be an inanimate object that creates acoustic radiation, such as any mechanized vehicle with engine or other noise.

The target need not actively emit any measurable radiation, either electromagnetic or acoustic. In the passive target case, electromagnetic or acoustic radiation is directed to the target and reflected from the target to a receiver. The apparatus used to implement the technique may comprise known transmitters, receivers, and computers. The known apparatus is modified by using computer software that implements the method.

The target identification method we envision begins with creating a density function of cepstral coefficients for known targets. Once the density functions of known targets are identified, the identification system can be deployed. The new system (cepstral-based) receives a signal from the unknown target, transforms the signal from a time spectrum to a frequency spectrum using a Fourier transform, transforms the frequency spectrum to a cepstrum, creates a density function of cepstral coefficients for the unknown target, and compares the density function of the unknown target with the density function of the known target. If the two density functions match, the unknown target is identified as an instance of the known target.

The time domain signal is transformed to the frequency domain and the frequency domain signal is then transformed to the cepstrum domain. The technique extracts the signal of interest in the cepstrum domain. One may think of the cepstrum domain as the frequency of a frequency. Each sample vector is represented by an n-dimensional vector of low-order cepstral components (or coefficients). For a given signal vector Y, cepstral coefficients are defined as: Y=(y1, y2, y3, . . . , yN). The discrete Fourier transform of Y is:

$$Y_m = \sum_{n=1}^{N} Y_n \exp(-j\frac{2\pi}{M}mn)$$

Where $M \geq 2N$ and m=1, 2, 3, . . . , M

The cepstral components are obtained from [Ref 3]

$$C_y(n) = \frac{1}{M}\sum_{m=1}^{M} \log\{\frac{1}{N}|Y_m|^2\}\exp(j\frac{2\pi}{M}mn)$$

Where n=1, 2, 3, . . . , M

The cepstrum transform capitalizes on the physical phenomenon of harmonics of a frequency to pull critical target identification data from low signal-to-noise-ratio situations. Likewise, the coefficients of the ensuing discrete frequency values allow a cross correlation between a priori data of possible targets.

The cepstral coefficients may be chosen by a variety of methods. Preferably, the cepstral coefficients are chosen from a group beginning with a smallest period (or a window side) and then a next smallest period and so on. For example, the cepstral coefficients would be chosen starting with the smallest period (i.e., just to the right of 0 on the horizontal axis) and then the next smallest period, etc. The cepstral coefficient for any period is the amplitude of the cepstrum at that period. Any number of cepstral coefficients may be chosen to create the density function. In one instance, the number of cepstral coefficients is in the range from 2 to about 100, preferably from 2 to about 50, more preferably from 2 to about 30, and most preferably about 12.

As discussed above, the unknown target is a member of a group of one or more known targets. Therefore, the first step is to create templates or signatures for the known targets. A known target signal (active or reflected) is received. If the signal is an analog signal, the analog signal is digitized. The digital signal is transformed from the time domain to the frequency domain using a Fourier transform. The frequency domain is then transformed to the cepstrum domain using a cepstrum transform. The cepstral coefficients are extracted from the cepstrum domain. Then, a density function of the chosen cepstral coefficients is created. The density function of the cepstral coefficients of the known target is the template or signature that is compared to the density function of the cepstral coefficients of an unknown target. After all the applicable templates are created, the unknown target signal "y" is compared against the known templates.

Training for the Methodology

The density functions created with the cepstral coefficients may be a Gaussian, a Non-Gaussian, or a Hidden Markov Model (HMM). In one embodiment of the technique, the density functions are HMMs. The signal of interest may be characterized as a parametric random process, and the parameters of the stochastic process can be determined in a precise, well-defined manner using the Markovian principal. Training comprises an estimation of Initial state, State transition, and Mean and Covariance matrices.

Below are three basic steps that must be solved for the HMM to be useful in the technique:

1) Given the observation sequence O={o1,o2,o3,o4, . . . } and the model λ=(A,B,π),

Where:
A=State transition probabilities.
B=Observation symbol probabilities in states.
π=Initial state probabilities.

Compute $P(O|\lambda)$, the probability of the observation sequence, given the model.

2) Given the observation sequence O={o1,o2,o3,o4, . . . } and the model λ, choose a corresponding state sequence q=(q1q2q3 . . . qT) that is optimal in some sense (i.e., best "explains" the observation).

3) And lastly, adjust the model parameters $\lambda=(A,B,\pi)$, to maximize $P(O|\lambda)$. An expectation algorithm may be used to perform maximum likelihood estimation of these parameters.

Testing of the Methodology

The minimum probability of recognition error is achieved by choosing the template target x, which makes the given signal, y, most probable:

$$\arg\max p(y/x)$$

Where:
y = Received or observation signal.
x = Signature of the known targets in library.

Results

The new methodology was verified by Matlab codes and was tested on real RF (radar frequency) signatures in in-phase and quadrature value format from six targets (previously-collected data). A total of 360 files were randomly selected from the six targets (i.e., 60 different files from each target: 30 files for training and 30 files for testing). There were several parameters of the HMM, (i.e., number of cepstral coefficients, number of state, window types, etc.). By using the first 12 cepstral coefficients, number of state 7 and window type Kaiser (3) results of performance metrics of the methodology are probabilities of the correct identification (83.89 percent), misidentification (00.00 percent), and no identification (16.11 percent).

We have presented a new methodology for NCTI. The preliminary results are full of promise. This technique appears to have the potential to increase the safety and effectiveness of our military. More importantly, it will reduce fratricide and civilian casualties.

Most important, the technique will mature innovative new algorithms that can be added to current NCTI to enhance system performance and will encourage the Naval Air Warfare Center Aircraft Division (NAWCAD) to initiate new ideas, rapidly develop new identification technology to support warfighters, and enhance technology expertise in-house. The methodology will help NAWCAD engineers to build capability that meets customer needs.

References

1 Thach Nguyen, US patent No. 7,180,442, "Target Identification Method Using Cepstral Coefficients," 20 February 2007.

2 Lawrence R. Rabiner, "A Tutorial On Hidden Markov Models and Selective Applications in Speech Recognition," *Proceedings of the IEEE*, Vol. 77, No. 2, February 1989.

3 Yariv Ephraim, Mazin Rahim, "On Second-Order Statistics and Linear Estimation of Cepstral Coefficients," *IEEE Transactions on Speech and Audio Processing*, Vol. 7, No. 2, March 1999.

About the author:

Thach Nguyen is an electronics engineer with the Naval Air Warfare Center Aircraft Division at St. Inigoes, Maryland.

The author would like to thank the NAWCAD Section 219 program for providing funding and supporting the project. He is also pleased to express his sincere gratitude to Christian Utara, Dr. Kristin Farry, Peter Zawada, Paul Wainwright, Brian O'Dell, and Colin Hoylman whose gracious encouragement, support, and valued constructive input helped bring this in-house research project to fruition.

Photo by Joe Bullinger

FROM VIRTUAL SCHOOLHOUSE TO TACTICAL ENVIRONMENTS: A DECADE OF VIRTUAL WORLDS

By Susan Farley

INSPIRED BY THE SUCCESS OF ONLINE GAMES AND VIRTUAL ENVIRONMENTS SUCH AS *THE SIMS* AND *SECOND LIFE,* A TEAM AT NAVAL UNDERSEA WARFARE CENTER DIVISION NEWPORT HAS BEEN BUILDING A VIRTUAL WORLD OF ITS OWN.

This year marked a milestone for the Naval Undersea Warfare Center's (NUWC) Undersea Warfare Combat Systems Department. It has been 10 years since its Virtual Worlds (VW) program began with an idea from Don McCormack, then NUWC's technical director, and now the surface and undersea warfare centers' executive director.

"Virtual environments hold tremendous potential, and NUWC's exploration of them as possible future venues for system innovation, collaboration, and rapid-prototyping reflects a commitment to 'open our apertures' to new, ground-breaking means to sustain today's fleet efficiently and effectively while working to build an affordable future fleet," McCormack said in 2008. "Virtual world technologies reflect the convergence of two trends—an Internet evolution in which the community of users has also become the content creators such as Facebook and Wikipedia and a shift in the gaming industry to massively multiplayer settings with wildly popular online games such as Blizzard's 'World of Warcraft' and Maxis' 'The Sims Online.'"

Division Newport's Undersea Warfare Combat Systems Department set out to research the ability of a virtual environment to train and educate Sailors as well as its potential for modeling and simulation, integration and interoperability, and rapid prototyping. Over the past 10 years the department has executed numerous projects using VW technology.

Since Division Newport's initial $2-million strategic investment, VW exploration and application research has generated a significant number of prototypes and capabilities in many different domains. The success of VW can be attributed to the program's ability to enable innovation while fostering cross-community collaboration through the creation of a powerful yet intuitive simulation environment that meets both social and technical objectives.

The VW team has included many individuals and interns over the past 10 years, but currently it is supported by Steven Aguiar, technical program manager; Diana Nolan, lead systems architect; Elise Augustine, research and development lead; and several computer scientists, including Derrick Huang, Makia Powell, Nicole Vaillancourt, Ken Figuerado, and Maya Umehara.

Members of Division Newport's Virtual Worlds team include (standing, from left) Makia Powell, Jack O'Sullivan, "Virtual Steve" (a life-sized version of Steven Aguiar), Diana Nolan, Steven Aguiar, and Derek Huang; (seated, from left) Nicole Vaillancourt, Ken Figuerado, Maya Umehara, and Elise Augustine. Photo by Rich Allen

Applying Software

Early on, the VW team coordinated with the Army and Air Force on how to best use VW software, though those services focused mainly on training applications with products such as the Air Force's MyBase and the Army's Military Open Simulator Enterprise Strategy.

In 2011, the Submarine Learning Center in Groton, Connecticut, identified the need for a distance learning environment that would support delivery of live tactical displays over significant distances with minimal latency. The training application known as Virtual Schoolhouse (VSH) is one example of a VW-created concept that resulted from early collaboration with the training community. With shrinking budget lines, the learning center estimated the annual cost for temporary duty alone for convening of the sonar technician courses at more than $100,000 each. If VSH could deliver the same experience and content through a high-fidelity virtual training environment accessible from existing distributed electronic classrooms at Sailors' home stations, then the potential capability could show significant results.

In 2013, Division Newport developed the first VSH prototype. From a practical standpoint, this meant that Sailors and instructors, for example, at a fleet training center in San Diego, California, could stay in home port while training on a live tactical system located at Patuxent River, Maryland. As avatars, the Sailors using the virtual trainer performed as if they were in the same room as the physical trainer and not in a room located 3,000 miles away.

An important component of VSH that enabled the project to move forward was an evaluation that virtual students scored as well as their counterparts in actual physical classrooms. In 2015, VSH met its technical and security requirements and was deployed with access to Sailors anywhere in the country. But beyond replicating locations and facilities, VW can create any training world necessary to present the learning material.

"VW is about the visualization of data. VW let us open the aperture and explore how we could use virtual reality where we were not using it before," Aguiar said. "This kind of technology is great for collaboration and innovation."

TIMELINE OF VIRTUAL WORLDS

The idea of a dedicated naval virtual world came about in 2008. Today, it is deployed across multiple military networks, providing real capability for distance learning, rapid prototyping, modeling and simulation, and collaboration. It is the culmination of 10 years of work and development.

2008

VW was about the exploration of new technologies in Second Life. The Code 25 team jumped in to build a virtual environment and create the public-facing "Virtual NUWC" in Second Life that represented Division Newport domains. Applications included virtual training, rapid prototyping of combat systems, visualization of acoustic data, virtual ranges, and collaboration events.

The VW team educated themselves along with the rest of Navy and Defense Department on the state of the technology.

"We did a lot of demos that year," Aguiar said.

2009

The team worked with many program sponsors including the Submarine Learning Center in Groton, Connecticut, educating them on relevant VW capabilities. They also completed the first virtual Lean event.

"In 2009 we were validating to ourselves what the technology could do and what were the limitations," Aguiar said. "We had to figure out what direction we take."

2010

Division Newport's Second Life presence: The Army, Navy, and Air Force came together and founded "the Military Islands of Second Life," a central hub for all military organizations through a collaborative public-facing, joint-service presence.
This year also was pivotal because the team decided to go exclusively behind the firewall through research in enterprise VW—it was the only way they could ensure the capability reached the fleet since there was no commercial capability that met the Navy's needs.

From 2010-11 the team assessed 10 different VW and open simulator-type environments before formally adopting OpenSimulator to host behind the firewall.

2011

Early sponsors such as the Office of Naval Research and PMS-450 (Virginia-class submarines) used VW to prototype 12 variations of future attack spaces with real-time inputs by the fleet and subject matter experts.

"It really demonstrated the power of VW for rapid prototyping," Aguiar said.

For the training aspect, the team explored with the Submarine Learning Center what could be done with immersive learning. Students could "step into" the information space and experience information in different ways than typical classrooms.

They created the immersive target motion analysis trail, which is the process of generating a fire-control solution from acoustic bearings. A highlight of this immersive experience is the ability for Sailors to step into tactical plots to better visualize data. Tactical plots are usually two-dimensional; VW makes them 3D.

2012

The team worked on the training side streets, and completed a project with Defense Finance and Accounting Service by creating a training trail for ERP and its various functions. For example, the ERP training trail allowed users to visualize what happens to a document as it travels through the system allowing the user to interact with information through a 3D diorama.

2013

VW hit an important milestone in the area of live virtual construct (LVC) modeling and simulation, working with the ASN RDA and the Navy modeling and simulation community to connect VW to distributed LVC events over the secret Defense Research and Engineering Network. This allowed users in the VW not only to visualize tactical exercises but also to collaborate with each other. Highlights of that effort included connectivity to Division Newport's Weapons Analysis Facility, which provided a heavyweight torpedo, and Naval Air Systems Command's MH-60R helicopter flight simulator. It was a great example of hardware and people in the loop connected to a virtual world.

2014

The Submarine Learning Center had a capstone event to demonstrate the VSH distance learning capability. Students and instructors from the Fleet Anti-Submarine Warfare Training Center in San Diego, California, held several courses for sonobuoy deployment using a remote SQQ-89 sonar training system in Patuxent River, Maryland. The courses demonstrated that VW could provide full fidelity access to trainers and instructors over distance. The successful completion of that event led to the deployment and production capability of VSH at the Submarine Learning Center. Subsequently, the team identified additional classroom courses. They also evaluated the most difficult instruction delivery: trainer-based tactical training. That initial success was the beginning of VSH, which is still evolving today.

2015

In 2015, Secretary of the Navy Ray Mabus stated that the Navy would be scaling up the use virtual environments. He requested the development of a roadmap for the Department of the Navy that identifies opportunities, risks, and barriers.

"Virtual and simulated environments offer an unprecedented opportunity for the Department of Navy to transform how it connects people, ideas and information," Mabus said. "They also provide a risk-friendly environment to experiment with innovative ideas and have the potential to reduce cost. The DON has been exploring the use of virtual environments across four primary domains: Training and education, modeling and simulation, integration and interoperability, and rapid prototyping."

Division Newport achieved the only accredited virtual world in the government when the VW Enterprise Solution was deployed, which continues to be the foundation for the only military-accredited VW available at the enterprise level.

In 2015, the team also created and deployed VSH at the Submarine Learning Center, which was accessible from all six remote learning sites. An initial set of courseware was transitioned to submarine training, successfully teaching two specific courses.

2016

A collaboration began between Code 25 and PMW-770 (undersea communications) to develop a VW-based set of behavior agents that allow communications, sensors capabilities, command and control, power, weapon and vehicle launching, platform, size, weight, and power parameters. The implications were that what was previously an animated VW model could now be a fully functioning platform within a complex kill chain.

Known as "Links of the Chain," this joint effort with Codes 25 and 34 took off in 2016 as the team saw how constructive agents complement the ability of VW to support the kill chain continuum. This latest VW application has taken the simplicity of an operational view static diagram to animations of scenarios to externally driven, high-fidelity LVC event support.

"VW was turning out to be one of the most innovative things we were bringing to the M&S table," Aguiar said. "War gamers, analysts and engineers can now emulate any behavior, feasible or infeasible, to support the evolution of complex future battle spaces. The LOTC agents allow us to describe any behavior and then exercise that in a system-of-systems kill chain."

2017

One of the latest innovative applications of VW is combining collaborative capabilities with access external tools (e.g., web systems, system architecture tools, MagicDraw) and the ability to provide kill chain performance analysis in support of Enterprise Set-Based Design engineering. VWES is now hosted at the Maui High-Performance Computing Center in Hawaii, which provides remote tool access for government supercomputing capabilities supporting high-fidelity, large-scale analysis.

VW supports model-based systems engineering by a distributed workforce while linking to real-time performance analysis and engineering design.

2018

NAVSTA VIRTUAL debuted in addition to numerous other projects. It is the culmination of the 2008 vision of Virtual NUWC but in a government-owned, fully secure, military environment. NAVSTA VIRTUAL continues to support NUWC Division Newport's mission for the exploration of VW technology for the sake of Navy innovation and collaboration but in a fully government-owned military environment. Expansion to include Air Force and U.S. Marine Corps in the next year is underway.

"We're growing the project and helping more people adopt the technology and become knowledgeable and become SMEs," said Diana Nolan, distance learning lead, who has been on the VW team since the beginning.

The VW team continued to explore different applications for current uses and identified potential applications for future uses. Specifically, the team was interested in looking across the undersea warfare domain, implementing human factors engineering as well as rapid prototyping and support for acquisition systems.

Making It Secure

The biggest hurdle for the VW team was ensuring the necessary level of security needed for Navy applications. The team pursued Second Life, a commercial virtual world hosted by Linden Lab, but as information assurance restrictions evolved, the benefits were limited with many agencies reducing their Second Life interests to educational outreach only.

In 2010, the VW team made a strategic decision to collaborate with Linden Lab to see if they could make an enterprise product that would work behind military firewalls; a year later, Linden Lab discontinued that effort because security issues posed too great a challenge. When they abandoned the enterprise product, the VW team moved to OpenSimulator (opensimulator.org), an open-source server for hosting virtual worlds.

The team spent a significant amount of engineering resources to make the OpenSimulator VW fully accredited for government networks and this product is now known as the VW Enterprise Solution (VWES).

"That may be our biggest accomplishment," Aguiar said. "The other services gave up, industry went in other directions, but NUWC prevailed. Even though it took six years to achieve an ATO [authority to operate], we did it. We just weren't going to quit because we believe in its potential. Everyday we're building and creating something that's never been done before."

Sponsorship

As it proved its success, VW acquired two main sponsors, one is their training community sponsor, the Submarine Learning Center, which has been exploring VW for the past six years with the VSH capability—a deployed distance learning product.

"It's been a capability that is being used by the fleet and we're very proud of that. That's a success story where NUWC innovation is now a fleet capability," Aguiar said.

The other sponsor is the assistant secretary of the Navy for research, development, and acquisition under Amy Markowich, the director of integrated battlespace simulation and test at Naval Air Systems Command. Markowich has championed the use of VW as a virtual collaborative environment for Naval Enterprise Modeling and Simulation (NEMS).

The project grew from initially exploring VW as another visualization tool for modeling and simulation to being a

Secretary of Defense Ashton Carter (standing center) was briefed on Virtual Worlds during a visit to NUWC Division Newport in 2016. Photo by Rich Allen

NUWC technical director Ron Vien (right) attended a 10-year anniversary celebration for Virtual Worlds in May 2018. Photo by Rich Allen

collaborative tool to bring the modeling and simulation community together. The ability to plan, execute, and analyze complex events adds innovation and powerful visualization tools to represent more comprehensive kill chains. Under NEMS, VW has been identified as an enterprise solution.

Taking It to the Next Level

The next phase for VW will be with a capability called Naval Station Virtual (NAVSTA VIRTUAL). The team has always used VW to create interesting projects and capabilities but until now they have not been able to provide as many interactions among warfare centers and within systems commands as necessary because of accessibility limitations. With a recent increase in demand, the VW team is working with other warfare centers to grow VW expertise at sites including the naval surface warfare centers at Panama City, Dahlgren, and Crane.

"Now we're trying to help other organizations use VW effectively so they can access VW and can create their own innovation and collaboration capabilities," Aguiar said. "That's our main focus this year."

"These are new ways of visualizing data," Augustine said. "Every department could do their own exploring. It does a good job of facilitating human interaction."

In support of this effort, the team established NAVSTA VIRTUAL on the Defense Research and Engineering Network, with dedicated spaces for different projects, system commands, and mission areas. The team is exploring alternative immersive learning spaces, support for engineering processes such as set-based design, and even rapid prototyping of future submarine command-and-control systems.

There is also a classified version of VW for experimentation and analysis applications such as integration and interoperability.

According to Aguiar, the next wave will be driven by a more immersive experience. VW is adding Oculus Rift and HTC Vive headsets to increase presence and immersion when necessary. Previously, a user was limited to looking into VW; now, the user can step into the VW.

VW also supports rapid prototyping of ideas, in innovation cells. For example, VW serves as a tool that captures ideas (e.g., technologies, concepts of operation, processes) while also allowing their visualization and simulation from within a tactically relevant environment.

As the Navy continues to push VW technology across the Department of Defense, Division Newport is leading the way. The command is working with its defense counterparts as well as its industry and academic partners.

"There were waves of interest in virtual reality in the late 1990s and again from 2006 to 2009. VW like Second Life captured news headlines. That's when we jumped in," Aguiar said. "It was revolutionary in its social and collaborative aspects. But many other organizations didn't have broad enough interest, sufficient investment or the perseverance to go beyond simple training applications."

As early adopters, VW team is in a good position and ready to step into that next environment.

"To see it transition from a concept to receiving consistent external funding to a fleet capability has been rewarding," Aguiar said. "NUWC is now identified as a center of excellence for VW and a Department of the Navy enterprise solution."

About the author:

Susan Farley , with the McLaughlin Research Corp., is a writer with public affairs at the Naval Undersea Warfare Center Division Newport.

NAVY MANTECH PROGRAM REDUCES COSTS AND ACCELERATES DEVELOPMENT

By Bobby Cummings

BUILDING NAVAL VESSELS—SUCH AS THE NEWEST AIRCRAFT CARRIER USS *JOHN F. KENNEDY* (CVN 79) SEEN HERE—AND OTHER PLATFORMS IS EXPENSIVE AND TIME CONSUMING. THE MANUFACTURING TECHNOLOGY PROGRAM PARTNERS WITH INDUSTRY TO HELP MAKE IT MORE ECONOMICAL AND FASTER.

Technological advancement—changing everything we do, from how we communicate to how we fly—is moving at a pace like never before in history. The same can be said in the realm of national defense, where the United States has been the world leader in developing emerging technologies for decades. This nation's competitors, however, are catching up.

Though it is losing ground on its competitive advantage, the United States remains the world's leader in the advancement of military technology. The key to remaining the leader is to invest in affordable technologies and accelerate their implementation into the battlespace. Enter the US Navy Manufacturing Technology (ManTech) program.

The purpose of the ManTech program is to support affordable manufacturing technologies and processes that will benefit warfighters around the globe, with investments focused on those areas of greatest need and potential benefit for national defense. Along with an emphasis on affordability, the program is designed to accelerate the Department of the Navy's transition of technology to the fleet through partnerships with industry.

Investment Strategy

The ManTech investment strategy concentrates on reducing the acquisition and life-cycle costs of vital Navy

Huntington Ingalls photo by Matt Hildreth

acquisition programs. ManTech transitions manufacturing technologies that, upon implementation, result in a cost reduction or cost avoidance and save taxpayers money.

These transitions provide affordable technology to the fleet by concentrating resources on high-priority acquisition platforms, focusing on reducing costs for platforms, developing new solutions for manufacturing and repair and sustainment obstacles, and engaging with vested industry leaders and companies throughout a platform's production cycle.

There are two key measures of success for ManTech's programs: transition, which is the point at which the project is completed and the technology meets customer criteria and goals for implementation; and implementation, the point at which the actual use and application on the factory floor results. The development process of ManTech programs involves multiple players, almost always a combination of government, academia, and industry.

ManTech is managed by the Sea Warfare and Weapons Department within the Office of Naval Research (ONR), with direct oversight from the chief of naval research.

"ManTech is an industrial preparedness program—it improves the manufacturing affordability of key platforms included in the naval investment strategy," said John Carney, ManTech director. "These platforms are selected by the chief of naval research."

Execution

ManTech projects are executed through centers of excellence (COEs). The COEs were established as focal points for the development and transition of new manufacturing processes and equipment in a cooperative environment with industry, academia, and the Naval Research Enterprise, which comprises ONR, ONR Global and the Naval Research Laboratory.

These COEs are located at various locations throughout the United States and play an integral role in the definition and execution of the program. COEs execute projects; manage project teams; facilitate transfer of developed technologies; collaborate with acquisition program offices and industry to identify and resolve manufacturing issues; and develop and demo manufacturing technology solutions for identified US Navy requirements.

For more than 10 years, the ManTech program, in partnership with the COEs, has focused on affordability improvements for key acquisition platforms such as *Arleigh Burke* (DDG 51)-class guided-missile destroyers, *Gerald R. Ford* (CVN 78)-class aircraft carriers, the F-35 Lightning II, *Virginia* (SSN 774)-class submarines, *Columbia* (SSBN 826)-class ballistic missile submarines, and the CH-53K King Stallion heavy-lift helicopter.

Arleigh Burke-Class Guided-Missile Destroyer

ManTech's *Arleigh Burke* affordability initiative has 23 projects implemented or in the process of implementation, with savings to date of $18.2 million per hull. Among these projects is a partnership with Ingalls Shipbuilding in Pascagoula, Mississippi, together developing the Enhanced Task Assignment and Progressing (eTAP) tool in 2017.

The eTAP toolset improved shipbuilding productivity by coordinating daily ship construction tasks for thousands of craftsmen. It also streamlines daily work activities for craftsmen by assigning line item work to crew members; projecting progress and work volumes; assessing, calculating, and reporting progress; and recording workforce data.

This technology will result in the reduction of 8,500 hours of labor and $731,000 in cost savings per hull attributed to hourly productivity from foreman availability.

Also in 2017, ManTech partnered with Bath Iron Works (the other builder of *Arleigh Burke*-class destroyers) to enhance shipyard capacity through the development of a shipyard-wide capacity planning system. Ship delivery requirements were met more effectively and there was a seven-percent reduction in overtime costs equaling $2.92 million in savings.

The future USS *Thomas Hudner* (DDG 116), the latest *Arleigh Burke*-class guided missile destroyer, returns after successfully completing acceptance trials in May 2018. ManTech has 23 projects implemented or under way with Arleigh Burke destroyers.Photo courtesy of Bath Iron Works

Gerald R. Ford-Class Aircraft Carrier

USS *Gerald R. Ford* is the U.S. Navy's newest aircraft carrier. Ford-class ships will deliver greater lethality, survivability, and joint interoperability than their predecessors. Technologies are advancing at rapid pace, and the versatility of the *Ford* class will allow these vessels to adapt more easily to the future.

ManTech's CVN 78 affordability initiative encompasses 12 projects implemented or in process, with savings of $22.1 million per hull to date.

Recently, ManTech developed a simulation tool that was implemented in 2017. The tool is used to analyze material impact on build strategy decisions. It reduces construction costs by forecasting and reducing logistical delays, linking proposed maintenance strategies to associated material delivery logistics, analyzing material effects on build strategy decisions, and reducing the number of movements of large units and material.

Huntingdon Ingalls estimates a labor reduction of more than 25,000 hours and associated cost savings of $3.08 million per CVN 78 hull.

F-35 Lightning II

The multirole F-35 was designed to dominate airspace for decades. Each variant of the airframe is equipped with the newest technologies, including a state-of–the-art helmet worn by pilots. Each helmet is equipped with a Helmet Mounted Display System (HMDS), which provides a multitude of advantages, including 360-degree visibility, visual targeting, increased comfort, a display of flight information, night vision, and more.

ManTech helped with the manufacturing processes for the improved Helmet Mounted Display System, built by Rockwell Collins for the F-35 Lightning II aircraft. Photo courtesy of Rockwell Collins

ManTech is working with the HMDS designers (Rockwell Collins ESA Vision Systems) on the manufacturing processes for a new version of the F-35 helmet with the goals of saving power, weight, and cost.

An issue pilots have pointed out in reference to the current helmet is a distinctive green glow from the cockpit display that could potentially distract pilots during night landings. The new ManTech and Rockwell Collins helmet has been tested and evaluated in a carrier environment for night landings, and feedback from pilots has been encouraging. This project is on pace for a new helmet to be delivered to the Navy by the end of 2018, and is expected to result in $6,000 in cost savings per aircraft.

In total, all ManTech projects associated with the F-35 program save an estimated $430,000 per aircraft, which equates to a projected $800 million in savings for the Department of Defense.

Virginia-Class Submarine

ManTech's *Virginia*-class submarine affordability initiative has drastically reduced expenses through various production improvements. To date, ManTech has implemented 46 projects per vessel, with a recognized cost savings of $37.4 million per hull.

"The affordability initiative has been a major success for both ManTech and PMS 450 [the *Virginia*-class program office] and was a key contributor to the Navy's '2 for 4 in 12' cost reduction initiative," said Carney, referring to the Navy's desire to purchase two boats for $4 billion in fiscal year 2012.

As a result of the two-per-year build rate, cost savings are greater than $70 million per year. The annual ManTech budget is returned with Virginia-class cost savings alone.

ManTech enhanced the production line capability of the upgraded acoustic sensor designed initially for USS *South Dakota* (SSN 790) and all subsequent hulls in the class. The demand for the sensors—developed by Naval Underwater Weapons Center Newport—required a higher production capacity, and ManTech demonstrated a production line capable of manufacturing 10,000 units per year—resulting in a cost reduction of $4.5 million per hull. The project began in 2017. The sensor also will be used by the *Columbia* (SSBN 826)-class ballistic missile submarine program.

In addition, ManTech partnered with General Dynamics-Electric Boat to develop a new thermal spray coating solution that prevents damaging buildup on the *Virginia* class' retractable bow-planes. Calcareous deposits led to premature failure of bow-plane equipment, and the continuous repairs hindered operational readiness.

The thermal spray coating solution prevents the damaging buildup, extending the service life of the affected parts and reducing the need for unscheduled maintenance. The refined coating was first applied to USS *Vermont* (SSN 792), with a projected savings of $9.2 million per hull over the life of the platform. The project team was awarded the 2017 Defense Manufacturing Technology Achievement Award for Readiness Improvement.

A CH-53K King Stallion helicopter demonstrates its capabilities at the 2018 Berlin Air Show. The CH-53K King Stallion is a replacement for the CH-53E, the Marine Corps' main heavy-lift helicopter since the 1980s. The first aircraft was delivered in May 2018. Photo by Cpl. Hailey D. Clay

Columbia-Class Ballistic-Missile Submarine

The Navy plans to build 12 *Columbia*-class ballistic-missile submarines to replace the aging fleet of *Ohio*-class submarines. Since 2014 ManTech has incorporated the *Columbia* class into its investment strategy.

As a cost saving measure, *Columbia*-class designs will incorporate technology and components from both the *Ohio* and *Virginia* classes.

ManTech has developed the Self-Locating/Self-Fixtured (SLSF) structures project. The development uses notched beams that interlock allowing ship deck structures to conjoin. The SLSF project will simplify ship deck construction, reduce labor costs and save an estimated $2.19 million on the first *Columbia* hull.

Thus far ManTech has completed seven projects for Columbia subs, with $10.9 million in savings per hull.

CH-53K King Stallion

The Sikorsky CH-53K helicopter is the Marine Corps' heavy lift replacement for the CH-53E. The helicopter is slated to provide an initial operating capability in 2019.

ManTech and the CH-53K program office (PMA-261) have been working with industry on a 3D data exchange project that was selected in March 2018 as the Naval Air System Command Commander's Award Business Innovation runner-up.

The combined government and industry team is working on simplifying a process that will reduce the amount of reverse engineering required for the creation, verification, and validation of 3D technical data management for PMA-261. This program will create a secure 3D data exchange system for non-engineers working on the CH-53K platform. In addition, the system will result in $9 million in annual cost savings for the CH-53K program.

The Time to Accelerate Is Now

In conjunction with affordability is a focus on acceleration. Many of the manufacturing technology programs ManTech is involved in are increasing the pace of implementation. Various countries around the world, such as China and Russia, have modernized their military and manufacturing processes. Competitors of the United States are gaining ground. But ManTech and similar programs are here to reestablish technological separation and deliver the assets required for Sailors and Marines to accomplish their missions.

"Leaders within the DoD, US Navy, and US Marine Corps have adamantly made clear their priority to accelerate the development of capabilities intended for warfighters," said Carney. "The most important asset within the DoD, is its men and women in uniform defending the United States and its allies' interests abroad. Navy ManTech takes great pride in hastening delivery of technologies that will help them accomplish their mission."

About the author:

Bobby Cummings is a writer with the Office of Naval Research.

Photo by Joe Bullinger

THE BEST KEPT SECRET IN THE NAVY: DIGITAL SLEUTHING WITH CYBER FORENSICS

By Susan Piedfort

IN TODAY'S WORLD, WHERE SO MUCH CRIME NOW TAKES PLACE ONLINE, CATCHING CROOKS RED-HANDED TAKES THE SPECIAL SKILLS OF A MODERN-DAY DIGITAL SHERLOCK HOLMES. ONE NAVY TEAM IS SETTING A VERY HIGH STANDARD FOR THIS NEW REALITY.

Critical data that could prove a suspect is defrauding the government is located on a hard drive that has been submerged in water. Naval Criminal Investigative Service (NCIS) agents need access to information on a suspect's cellphone, which was intentionally destroyed.

There are two challenges, but one solution: Space and Naval Warfare Systems Center (SSC) Atlantic's Digital Media Criminal Forensics Investigations (CFIX) Laboratory. The lab, the Navy's first and only of its kind to earn the prestigious American Society of Crime Laboratory Directors international accreditation, serves as a base of operations from which SSC Atlantic's cyber forensics team performs its mission of helping the Navy and other federal agencies recover data and solve criminal cases using its unique operational cyber forensics capabilities. Their customers include NCIS, the Department of Justice, Department of Veterans Affairs, Department of Homeland Security and the Marine Corps.

Making full use of the CFIX lab capabilities, the team conducts data recovery, responds to cyber forensics incidents, performs malware analysis and reverse engineering of malware, and helps protect the Navy's network infrastructure.

The CFIX lab also contains a data recovery laboratory with physical rebuilding and submerged hard drive recovery capabilities, which the team uses to perform reconstruction, mobile device data extraction, chip-off and bad sector recovery. They also can perform advanced memory analysis, extract hidden data with steganography and other processes, analyze firmware, and recover data from solid state drives and Flash media.

SSC Atlantic's cyber forensics capability started in 2008 with one machine and two people: Robin Corkill, now cyber forensics competency lead, and Bill Littleton, cyber forensics integrated product team lead. From these meager beginnings the team has grown to 15 computer engineers, computer scientists, forensics analysts, and information technology specialists, as well as more equipment, a new building on the SSC Atlantic campus, and greatly increased capabilities. The demand for the work they do has increased over the past 10 years and continues to grow as customers learn of SSC Atlantic's cyber forensics capabilities.

More Department of Defense (DoD) and Navy emphasis on cybersecurity, incident response and data recovery has led to the growth and recognition of the team, according to Corkill. "We can grow as needed to respond to our caseload," he said, adding that the data recovery lab averages more than 100 hard drive recoveries a year for SSC Atlantic employees, and has served multiple organizations with a more than 95 percent success rate.

Two members of the team have master's degrees in digital forensic science, six others have masters in computer science, cyber or math, and many have advanced certifications. All told there are more than 20 certifications, some in specialized areas such as smart phones and mobile apps, Windows, Mac, Linux or memory forensics, and even vehicle forensics.

"We also do lots of research, and go to advanced schools to stay current on data recovery, malware analysis and digital forensics criminal lab processes," added Littleton. The research they do while solving casework often leads into new research areas with different artifacts.

"With technology growing exponentially, there is constant learning," Corkill said.

When not helping to solve crimes, recovering data and responding to cyber incidents, the team provides cyber training to SSC Atlantic and other commands on how to prevent breaches, minimize system failure and prevent catastrophic loss of data. That training includes cyber defense tracks in response, forensics, intrusion and detection to show how to prepare for attacks and how to respond to an incident and secure evidence. The team also covers how criminals get access, hide data on a system and cover their tracks to help students better understand, prepare for, respond to and even prevent cyberattacks."

"The fact that we are an all-government team, all Sensitive Compartmented Information-cleared, with strict standards on how we handle data helps us maintain trust with our DoD customers and other agencies," Corkill

SSC Atlantic intern Sage Glidewell looks on as employee Alisha Slebodnik examines the internal components of a computer hard drive on a "clean bench." (Photo by Joe Bullinger) said. "Our Air Force, Navy and other DoD customers have specifically mentioned the value of working with all-government team." The SSC Atlantic team's technical and analytical expertise, problem solving and out-of-the-box-thinking also are often lauded. Both local and federal law enforcement agencies have consistently provided positive feedback about the CFIX laboratory capabilities, and have commented on the team's efficiency, depth of information provided and the outstanding quality of the final product.

That praise is the result of a team effort, Corkill emphasized. Having a drive to dig deeper and look at a problem from a number of angles, to use an analytical approach and follow the "cyber bread crumbs" makes for a sharp and effective cyber forensics team.

"It's easy to wake up every day and be excited about what we're doing here," Corkill said. "You never know what you might be facing. Finding information, solving a problem—it's always a challenge."

"That's probably why most of us on the team enjoy participating in STEM outreach events such as the Palmetto Cyber Defense Challenge [PCDC] and cyber summer camps," Littleton said. "It stretches us to be creative and generate scenarios, then analyze and develop solutions for them."

The lab's outreach activities not only increase interest and proficiency in the cyber forensics domain for the next generation, but their hugely successful cyber summer camp is being used as a model by the Office of the Secretary of Defense.

The team's outreach has also attracted students who want to work alongside Corkill and Littleton as career cyber forensics professionals. Sage Glidewell was a PCDC and cyber camp participant, then a mentor, and now an intern with Corkill's team as she pursues a degree in computer science at the College of Charleston. Two interns who were hired this summer were hand-picked by Corkill and Littleton through SSC Atlantic's involvement in the

SSC Atlantic employee Alisha Slebodnik performs an examination of the head stack from a computer hard drive on a "clean bench." Slebodnik works in the Digital Media Criminal Forensics Investigations Laboratory. Photo by Joe Bullinger

SSC Atlantic employee Jason Staker conducts a forensic triage on a cellular phone inside a Faraday enclosure.
Photo by Joe Bullinger

Office of Naval Research's cybersecurity outreach with historically black colleges and universities. "They are computer science and cybersecurity majors with a focus on forensics," Littleton said, "and they brought great new ideas to the team."

"What we offer DoD and other federal agencies is a full service and advanced digital forensics capability that has matured over the past 10 years into a cutting-edge answer to the most critical cyber defense needs of today. Our customer feedback is that our quality is unmatched," Corkill said. "World-class data recovery, advanced digital forensics, a top quality criminal media forensics laboratory, a highly skilled malware analysis and incident response team arguably provide one of the most advanced overall cyber forensics capabilities in the DoD. Customers often identify us as the 'Best kept secret in the Navy.'"

SSC Atlantic commanding officer Capt. Scott Heller praised the team as a valuable asset to the Navy's cyber defense capabilities. "We all know too well that the threat is real. The urgency to answer these threats is real," he said. "I'm excited and proud that we have been able to apply the intellect, energy and ideas found here at SSC Atlantic in new and exciting ways to respond to these ever-evolving threats."

About the author:

Susan Piedfort is a writer with SSC Atlantic public affairs.

Photo by Joe Bullinger

THE RED TEAM: The Good "Bad Guys"

By Susan Piedfort

TO PROTECT CYBER SYSTEMS FROM HACKERS, YOU HAVE TO THINK LIKE HACKERS. THAT'S ALL IN A DAY'S WORK FOR SPACE AND NAVAL WARFARE SYSTEMS CENTER (SSC) ATLANTIC'S RED TEAM, AN EXPERT COLLECTIVE OF GOOD "BAD GUYS" WHO CONDUCT ADVERSARIAL ASSESSMENTS ON DEPARTMENT OF DEFENSE NETWORKS.

SSC Atlantic's Red Team, certified by the National Security Agency and accredited by the United States Cyber Command, is one of nine certified Department of Defense (DoD) Red Teams and one of only two in the Navy. The SSC Atlantic Red Team assesses DoD cyber security service providers (CSSP), provides adversarial and aggressor support to cyber exercises, and supports cyber developmental and operational testing to acquisition programs with information technology (IT) components. Their customers include the Defense Health Agency (DHA), Defense Contract Management Agency, U.S. Marine Corps, Special Operations Command, and the Naval Enterprise Networks program office (PMW 205).

The Red Team's real-world attack simulations are designed to assess and improve the effectiveness of an entire information security program, including those controlling weapons systems, platforms, sensors, and networks.

"The thinking is, if you simulate bad guys and put network defenders and system owners under controlled stress in a controlled environment, you get a better sense of how they will perform," said Jason Jurand, the Red Team's director.

"If you wait long enough, the real-world adversaries will tell you what's wrong with your system, usually at the worst time," Jurand said. "Our first rule is 'do no harm.' Our adversaries don't have that rule."

Jurand emphasizes that the Red Team better positions customers to deal with these vulnerabilities on their terms rather than those of their adversaries.

The Red Team's functional capabilities were developed when SSC Atlantic's CSSP was created and certified. The CSSP's mission is to protect, detect, respond, and sustain IT systems, and, as part of the "protect" service, the Red Team also assesses the defense capabilities of CSSPs across the DoD Information Network.

SSC Atlantic's Red Team has 13 government employees and can adjust their size according to need through the use of their contracting strategy. They are technically skilled with backgrounds in computer science, computer engineering, software development, test and evaluation, networking, and system administration.

Red Team director Jason Jurand (right) and deputy director Scott West discuss the details of a remote assessment. Photo by Joseph Bullinger.

According to Jurand, a knowledge of how things work—and an understanding of how to degrade, disrupt, or deny a customer's cyber environment while actually doing no harm—requires a deep technical background.

"From a temperament point of view, you have to be naturally curious and think unconventionally. Red Team people are tinkerers," he said, "with maybe a little bit of a dark side."

SSC Atlantic's Red Team is certified to perform a variety of assessments, including local ones, where they are invited in by a customer and work collaboratively and cooperatively to help identify and mitigate known vulnerabilities—and often to discover new ones.

They also perform remote assessments, which are more covert in nature. The Red Team tries to gain access to the customer's network without the knowledge of the customer's CSSP or "Blue Team." Persistence missions involve the Red Team staying in the network as the customer's Blue Team is actively pursuing it.

"They are trying to pry us out of network, and we are trying to burrow in and stay in," Jurand said.

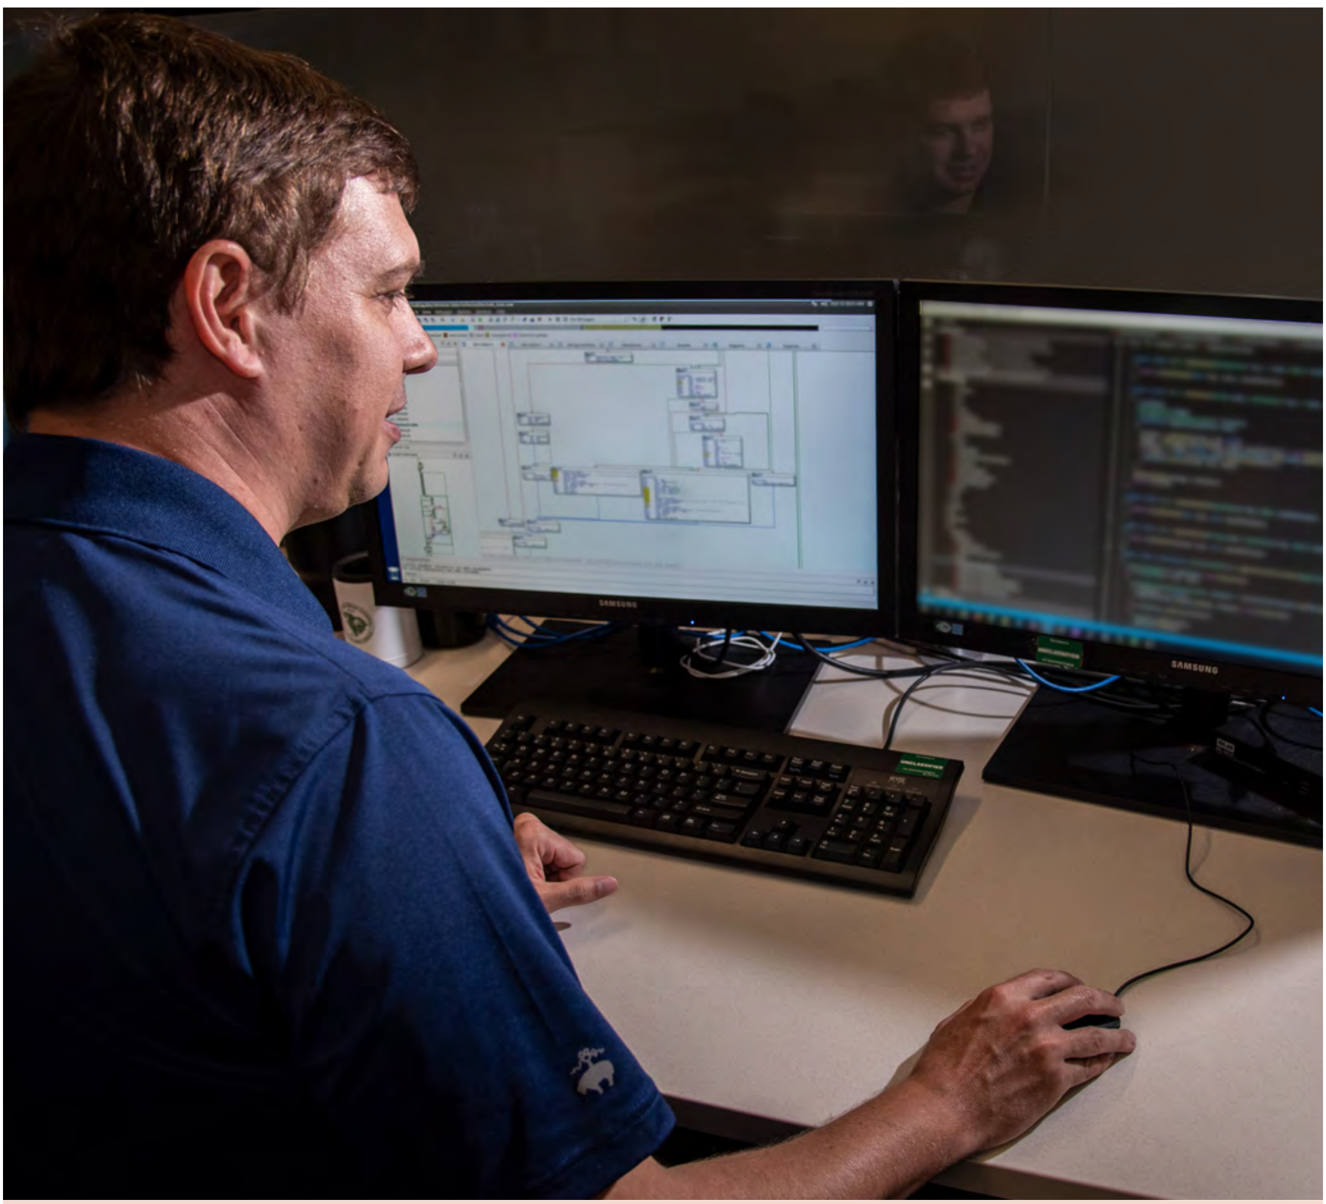

Red Team member Bryan Rhodes reviews software code for vulnerabilities. Photo by Joseph Bullinger.

The Red Team assesses wireless security, which ranges from systems as innocuous as a home Wi-Fi to anything in the RF spectrum, such as shipboard or aircraft wireless systems.

The Red Team is very effective with user-driven attacks, which Jurand describes as complicated but usually the most successful.

"Most cyberattacks are user-driven, where you manipulate the user into doing something that gets them in trouble," Jurand said. "For a Red Team, it's the easiest to get at and yields the most reliable results. We've never had a phishing campaign that failed."

Jurand explained that cybersecurity deficiencies found by the Red Team fall into the categories of people, processes and technology, with people being most common deficiency found.

"Insider threats are real. It's not just about getting past the guy at the front gate or tailgating into a building; it's user attacks and social engineering," he said. "And even though everyone gets cybersecurity training every year, invariably we'll find some kind of shortcoming."

Something as simple as going into a hospital or military health clinic can pose cybersecurity challenges that can actually risk lives. Those going in for outpatient appointments or visiting patients admitted to a hospital may want to use their phones or tablets on the facility's Wi-Fi. In a worse-case scenario, these devices could pose a threat to IT systems that connect patients to life-saving equipment. To combat this threat, SSC Atlantic's Health Systems Security Engineering integrated product team, headed by Cal Stephens, provides full scope network/cybersecurity services to DHA, including network protection suite design and development, accreditation, deployment and operations fused with Cyber Command-accredited Tier 2 CSSP services.

"Cal was part of developing a secure intranet for DHA, engineering the design, deploying it, doing network operations and sustainment of that infrastructure, and we were serving in an information assurance capacity," Jurand said.

This series of events provided SSC Atlantic a unique operational cyber perspective within the Navy. Given their capability, it made sense for SSC Atlantic to provide CSSP and Red Team services for other customers. The CSSP team was originally certified by the Defense Information Systems Agency and accredited by US Strategic Command in 2012.

Today, SSC Atlantic's Red Team is more and more in demand.

"Once we got certified, the phone starting ringing off the hook and it hasn't stopped since," Jurand said. "It has really led to a great capability for SSC Atlantic."

"There is so much complexity in cybersecurity threats; new ones pop up every day. We make folks take training and we do checkups to try to keep networks and systems healthy, but invariably, when Red Teams do assessments we always find shortcomings," Jurand said.

"We are looking for stuff that is unusual," he said, spending lots of time and energy looking through the assessment data to find what he describes as a "horrifying collection of success event audit records" that may indicate compromise.

For example, why is someone logged in at 2 a.m. on Christmas morning? Why is an administrator surfing the Internet and downloading data to the server? Are detections being made the way they are expected even when there are no failure or deny event audit records?

While the Red Team's mission is to help and protect customers, they are not always welcomed with open arms.

"People are often taken out of their comfort zones or feel violated when the Red Team shows up," Jurand said. "That's a healthy reaction to have," he said, since some people think they could get fired or that the network is actually being compromised.

"We are not the bad guys, we are trying to teach them about threats and how to mitigate them," Jurand said. "Red Team operations really represent an investment in a customer's cybersecurity infrastructure and in the people who use it. We are teaching them to be more aware of their vulnerabilities.

"In the end they realize that a real adversary would probably teach the same, but on much worse terms."

About the author:

Susan Piedfort is a writer with SSC Atlantic public affairs.

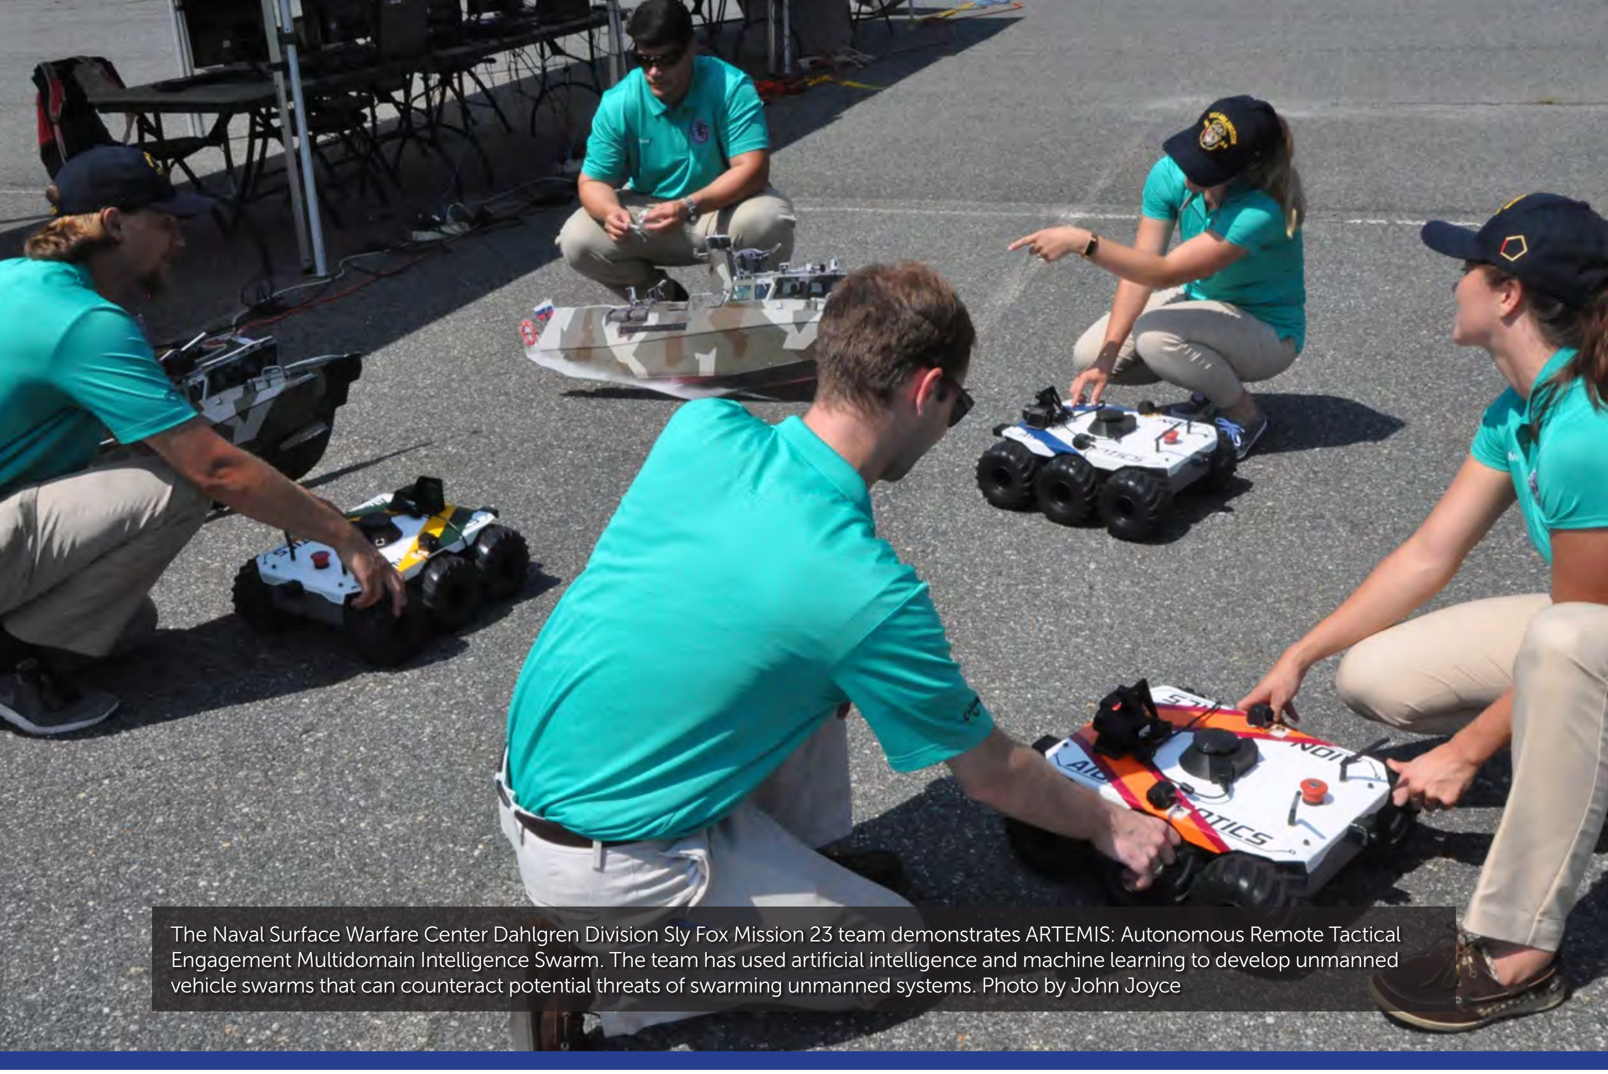

The Naval Surface Warfare Center Dahlgren Division Sly Fox Mission 23 team demonstrates ARTEMIS: Autonomous Remote Tactical Engagement Multidomain Intelligence Swarm. The team has used artificial intelligence and machine learning to develop unmanned vehicle swarms that can counteract potential threats of swarming unmanned systems. Photo by John Joyce

FUTURE FORCE is a professional magazine of the naval science and technology community published quarterly by the Office of Naval Research.

Future Force
Office of Naval Research
875 N. Randolph Street, Suite 1425
Arlington, VA 22203-1995

Email: futureforce@navy.mil
Phone: (703) 696-5031
Web: http://futureforce.navylive.dodlive.mil
Facebook: http://www.facebook.com/navalfutureforce