


DAKOTA DIGITAL REVIEW



DAKOTA DIGITAL ACADEMY
NORTH DAKOTA UNIVERSITY SYSTEM

SPRING/SUMMER 2023



**Real Robots in
Our Near Future**

Mark P. Mills

**Solving Crime Through
Digital Evidence**

Arica Kulm

Food Security for the Future

The Honorable John Hoeven

AI Security

Ram Kumar & Hyrum Anderson



Introduction to the **DAKOTA DIGITAL ACADEMY**

KENDALL E. NYGARD, PHD

Director, Dakota Digital Academy, North Dakota University System
Emeritus Professor, Department of Computer Science,
North Dakota State University
Contact: kendall.nygard@ndus.edu

Digital technologies are transforming society and driving revolutionary changes in the world of work. In response, the Dakota Digital Academy (DDA) was founded by the North Dakota University System in the fall of 2020 to provide online education in computing and the cyber sciences. DDA serves students at higher education institutions across the state—as well as residents in the workforce seeking to upskill or change careers—by imparting relevant digital skills. To date, we have focused primarily on courses and certificate programs in cybersecurity and software development.

DDA works cooperatively with the state's 11 public colleges and universities, which include two research universities, four regional universities and five colleges. Also affiliated are North Dakota's five tribal colleges. Talented faculty across the state system work together to design and deliver location-agnostic workshops, full courses, short skill-specific courses and certificate programs. Some activities are oriented toward improving the skills of technical people already in the workforce. Others focus on continuing education and credentialing for K-12 teachers.

Also included in DDA's instruction are soft skills related to the liberal arts, such as teamwork, creative

and critical thinking, problem-solving, ethics and communication, along with considerations of technology's social implications.

Over the last two and half years, DDA successfully launched Dakota Digital Review, Dakota Digital Discussions and the Workforce Advisory Council, which is comprised of business, industry and government leaders who support DDA's workforce readiness and cyber-educational mission.

Going forward, DDA is pursuing several highly relevant initiatives. One focuses on digital literacy in general education across the university system. Gen Ed refers to suites of required courses imparting knowledge and skills fundamental to all major fields of study and to success after graduation. Increasingly today, literacy in computing and cyber sciences constitute essential components of every student's formation.

A second initiative concerns advancing education in artificial intelligence and machine learning—including content creation systems such as ChatGPT, which are revolutionizing and disrupting nearly every industry, and augmenting or supplanting functionalities that involve reasoning, perception and creativity, which have been strictly human domains throughout history. ▣

■ ***Dakota Digital Discussions Webinar Series***

Dakota Digital Discussions is a monthly webinar series presented in the fall and spring semesters by Dakota Digital Academy and Dakota Digital Review. The webinars focus on the digital transformation of our economy, military and society, as well as digitization's profound ethical, legal, cultural, educational and

policy implications, including impacts on the arts and humanities and the human psyche. Most Dakota Digital Discussions engage for an hour and are scheduled at noon Central.

Please access upcoming and archived webinars at:
<https://dda.ndus.edu/ddd-overview/>

DAKOTA DIGITAL REVIEW

- **Dakota Digital Review** publishes articles, in print and online, about digitization and related technologies, as well as digitization's profound implications for our culture, economy, military, political institutions and policies, legal frameworks, moral foundations, and the arts and humanities.
- **Dakota Digital Review** is non-partisan, dedicated to free speech and academic freedom. The articles are written by subject experts in business, industry, government and academia for the general educated reader.
- **We aim to better prepare** the polity and populous to make critical decisions about our collective future and about our individual and family lives.
- Please contact Patrick J McCloskey, Editor, Dakota Digital Review, for submissions, comments or questions: patrick.mccloskey.1@ndus.edu.

CONTENTS ▪ SPRING/SUMMER 2023 ▪ dda.ndus.edu/ddreview/

Real Robots in Our Near Future: The Rise of Capable Industrial Automatons Mark P. Mills, Senior Fellow, Manhattan Institute	2
In Defense of (Virtuous) Autonomous Weapons Don A. Howard, PhD, Professor of Philosophy, University of Notre Dame	9
Food Security for the Future: Empowering Farmers & Ranchers Through Precision Ag Technologies The Honorable John Hoeven	18
Combining Precision Ag Technologies to Improve Crop Management Paulo Flores, NDSU	22
Solving Crime Through Digital Evidence Arica Kulm, PhD, Dakota State University	28
Promise and Peril of Intelligent Machines & Cyberspace Mark R. Hagerott, PhD, Chancellor, NDUS	36
AI Security: The National Security Commission on Artificial Intelligence & Adversarial Machine Learning Ram S.S. Kumar & Hyrum Anderson, PhD	42
SmartLab: Innovative Digital Learning Tools Matthew Lonn, EdD, Asst. Principal, Sheyenne High School, West Fargo	54
Four Horseman of Technological Change: Farmers, Elevator Operators, Coal Miners, Bank Tellers Evan J. Zimmerman, CEO of Drift Biotechnologies	58
Contributors	64
Cover illustration "Vitruvian Robot" by Tom Marple.	



Digit, a bipedal robot made by Agility Robotics.

REAL ROBOTS IN OUR NEAR FUTURE

The Rise of Capable Industrial Automatons

MARK P. MILLS
Senior Fellow, Manhattan Institute

We know two things about the future. And both point to the need for many more and very different kinds of robots than now exist.

We know that even as economies increasingly digitalize and become ever more service- and software-centric, demand will still increase for the kinds of things that are produced by the “hard” industries. That reality was made clear during the disruptive lockdowns. Miners are needed to access minerals to supply the manufacturers that, in turn, fabricate physical stuff, from computer chips to medical devices, and from fertilizers to pharmaceuticals. All of that is critical for creating and operating all the services that make modern life possible.

We also know, since “demography is destiny,” that the world in general and the United States and Europe in particular will experience increasing shortages of skilled workers for all of the hard industries.ⁱ In today’s America, the nearly 50-year-old average age of those in the skilled trades is far older than the overall population average.ⁱⁱ

There are only a few options for ensuring a sufficient supply of skilled labor needed for the hard industries.

The primary option, thus far, has been to find those people and industries elsewhere—that is, the *de facto* policies of increasingly importing goods from other, younger nations. That option, if pushed too far, has its own geopolitical and economic challenges and is likely to face constraints now since many policymakers are embracing reshoring initiatives.

Then there’s the option of importing a younger labor pool. Setting aside the politics of immigration, it still takes a long time for any rising generation to develop the necessary skills and experience, even if the targets (whether natives or immigrants) have the requisite interest in the first place.

Which brings us to the only other option, amplifying the effectiveness of those, of any age, with skills. Industrial automation is a longstanding solution for

amplifying labor, whether by outright taking over some jobs to free up a human for higher-skill tasks or by increasing the productivity of the skilled person (faster, safer execution of tasks). However, contrary to the popular narrative, there is far less automation across industries than most imagine, especially when it comes to robots.

Surveys reveal that you won’t find *any* industrial robots in over 90 percent of America’s manufacturing enterprises.ⁱⁱⁱ

Yes, there are millions of industrial robots in the world, but the majority are found in the minority of businesses and performing a small minority of the universe of tasks. Even for firms of significant size, over 500 employees, just half those have industrial robots. Today’s robot population is found mainly in the big businesses that produce large quantities of similar products (especially automobiles). As the size of the firm shrinks, and the variety of tasks rises, the robot share shrinks faster. The skilled workforce and automation challenges for small business impact large businesses because the latter depend on supply chains of small firms.

That automation yields more output per employee is intuitively obvious and borne out by the data.^{iv} Fifty years ago, large firms with economies-of-scale achieved, on average, about 25 percent greater output per worker compared to small firms. But today, large-firm adoption of industrial automation has led to per-person output nearly double that of the small ones.

The dramatic automation schism between large and small firms is explained by a simple fact. Robots haven’t yet been good enough to be deployed widely. Industrial robots, in the main, are used in a fixed location performing high-volume, repetitive tasks, well-suited to big manufacturers. Up until recently, there wasn’t any prospect for finding robots that *both* match or exceed human performance *and* can also be reassigned to a new task. Both those metrics need to be met with robots that can work (safely) alongside people, instead of isolated and bolted down, or limited to fixed tracks. And, of course, costs matter, not just the purchase price, but the cost of integration which, so far, can double or triple the initial cost.

What hard industries need in the near future distills to affordable, useful anthropomorphic robots, ones with skills. Ones that can perambulate or at least easily roll around in the same environment as people.

Long Time Coming

For centuries, engineers have designed machines with embodied controls that can automatically react to a change. A simple example might be a tank that registers a filling liquid and reacts accordingly: Once the fluid level reaches a certain point, it flips a lever connected to a simple valve, stopping the flow. But far more clever mechanical automation than this dates back to ancient days. In the first century, Hero of Alexandria built automatic doors and the like, powered by compressed air or running water, and even steam. Hero also invented a coin-operated drink-dispenser, as well as animated puppets controlled by ropes connected to weights, amongst dozens of other ingenious, automatic machines.^v

The idea of an automaton itself, a robot, is also old. We can

trace the idea of an automaton to a time before Hero of Alexandria, to 250 BC in the epic Greek myth of “Argonautica,” (to become Hollywoodified as “Jason and the Argonauts”) wherein Apollonius imagined a giant, man-like bronze robot called Talos.^{vi} For the 20 centuries that have followed, robots have been a staple in fiction, usually of the dystopian variety.

When Czech playwright Karel Čapek wrote his 1920 play, “Rossum’s Universal Robots,” he imagined automatons replacing humans for manual labor. Čapek invented the word “robot” from the Czech “robota,” which translates as forced labor or drudgery. Even though the word is now used rather elastically to include everything from an automated pick-and-place machine to a clothes washer, what we really mean by “robot” is a truly autonomous, ambulatory machine, and one that can be anthropomorphic—even human-like in appearance and mechanical function.

In 1939, Westinghouse built a kind of “Wizard of Oz” Tin Man, a stunt robot for the New York World’s Fair. But it could only walk stiffly and had a recorded voice that would say: “My brain is bigger than yours.” (Westinghouse wanted to show off its automated switchgear used for electrical grid controls.) More famously, it was the scientist-turned-writer Isaac Asimov who created the modern archetype for robots in his iconic 1950 science fiction book, *I, Robot*.

A key feature of an ambulatory general-purpose robot is that it can navigate in the same environments as people. Over recent decades, there have been myriad pretenders in the race to produce a real robot, from Sony’s toy robodog circa 2000 (while looking dog-like, it was not close to being able to emulate animal ambulation) to Honda’s contemporaneous



Left, Boston Dynamic's Atlas robot that can perform sophisticated movements, including running, jumping, throwing heavy objects and even doing backflips. At 5 feet tall and 190 pounds, Atlas is a battery powered, hydraulically actuated humanoid with RGB cameras, a laser rangefinder and depth sensors, all of which feed data to powerful computers operating complex algorithms. Atlas displays fine motor skills and can operate in rough terrain, with a wide range of capabilities including tasks in dangerous environments.

walking and stair-climbing Asimo, to name only two amongst dozens.^{vii} All were engineering demonstrations or toys. Few could perform functions other than walk or dance awkwardly.

But in the past few years—because of radical innovations in sensors, AI, materials and batteries—engineers are finally building anthropomorphic robots, even if most are not yet commercially viable. With remarkable prescience, for the occasion of the 1964 World's Fair, Asimov made some forecasts, amongst which he wrote that “robots will neither be common nor very good in 2014, but they will be in existence.”^{viii}

We know what Asimov was referring to with robots: the difference that distinguishes automation and automatons.

Boston Dynamics human-like Atlas robot meets that definition of an automaton, a true anthropomorphic robot. But it's not commercially available and is rumored to cost about \$1 million.^{ix} Nonetheless, Atlas has demonstrated human-like running, jumping, back-flipping and autonomous navigation. Atlas's prowess vastly exceeded the best efforts of the teams that competed only as recently as 2015 in a DARPA (the U.S. Defense Advanced Research Projects Agency) Grand Challenge. The contest involved simply having an untethered robot perform easy tasks of ascending a staircase, opening a door and turning a valve—all tasks inherent to operating usefully within a typical human environment, and all previously beyond the capabilities of any general-purpose robot.

In early 2020, Boston Dynamics offered for commercial sale (base price at \$74,500) its autonomous, ambulatory four-legged automaton, Spot.^x That such robots are now being offered for sale is more than mere curiosity. It is a pivot in history comparable to the first automobile, the 1896 horse-carriage-looking Duryea Wagon which was, in its day, magical because it was self-propelled.

Spot can walk, run and get up after falling, open doors and fetch objects. Spot, at least initially, is being hired for such things as roaming safety surveillance at construction sites, on farms, offshore oil platforms, pharmaceutical factories. Such services are critical not just for confirming that equipment is operating



Above is the poster for the 1939 production of the Czech play “R.U.R.: Rossum’s Universal Robots,” which was performed in New York City as a Federal Theatre Project, a New Deal program to help sustain artists during the Great Depression.

optimally but also for safety. Such tasks are by nature inherently repetitive and often become the kind of drudgery that lends itself to error and oversight. And freeing people from those tasks makes them available for upskilling; it amplifies human labor.

There are today at least two dozen companies designing and building pre-commercial anthropomorphic robots, ranging from start-ups to industrial giants such as Toyota and Hyundai. And, last year Elon Musk announced that Tesla would soon commercialize a walking robot called Optimus.

Useful Biomimicry

As with the automobile, the (true) robot is made possible by the confluence of a suite of technologies. For the age of the car to launch, it took the independent maturation of high-strength steel,

internal combustion and oil refining. For robots, it's the arrival of powerful micro-motors, vision "chips," enabled by AI, and lithium batteries.

Advances in the suite of sensors, vision and location systems have followed a progression similar to the often-noted Moore's Law for computer chips. Roboticists now have available an array of tiny, powerful cameras, chip-scale radar, complementary laser-based radar (i.e., lidar), along with microscopic silicon-fabricated position sensors. Tools that can sense motion, direction and velocity—from inertial changes in movement, hence the technical term, inertial measurement unit, or IMU—have been used by the military, in particular, for decades. But only in the past two decades has the IMU collapsed from coffee-cup scale to chip-scale, and only in the past decade gained both sufficient precision and affordability.

Practical, tetherless robots also needed a revolution in power, both in the ability to store onboard power and the power of actuators to effect movements and manipulations with precision and, well, power. The first commercial lithium battery, a game-changer, didn't appear until the early 1990s, and it took another decade or two for that ecosystem to achieve the requisite maturity. Similarly, actuators—in effect, robots' muscles—followed another, again independent trajectory of (fortuitous) advances in size and more power. Superior designs and new materials—not least the 1984 invention of rare-earth neodymium super-magnets—has engendered a roughly 50-fold gain in the power-to-weight ratio for tiny electric motors over the past several decades^{xi}.

The challenge that has eluded engineers for years is a mechanical and materials science one: the ability to mimic animal or human muscles. When it comes to biomimicry, the challenge has always been to find a way to use available electrical, pneumatic or polymer actuators to attempt to approach the combination of capabilities exhibited by muscles, the biological actuators: high energy conversion efficiency, a large range of motion, a strong power-to-weight ratio, durability and, ideally, self-repair.

In a 1983 paper titled "The Muscle as an Engine," American physicist and polymath Edwin Jaynes

presciently mapped out the mechanisms and the possibilities that were then not possible.^{xii} Jaynes observed that ultra-efficient conversion of chemical into mechanical energy would ultimately require emulating how muscles operated—that is, "that the moving parts receiving the primary energy be of molecular size." He speculated that "far from being impossible," that in time the design of "useful anti-Carnot molecular engines (artificial muscles) might become about as systematic and well understood as the design of drugs and antibiotics."

Today we're beginning to realize Jayne's vision with the profound, if ignored, revolution in materials sciences. The technical literature is replete with successful designs of "artificial muscles," some engineered at the molecular level and in some cases with self-healing capability.^{xiii} It has been a happy coincidence that materials sciences have enabled not only light-weight, durable construction of a frame (the skeleton), but also the design of actuators that have sufficient power-to-weight ratios.

As a key indicator of progress with biomimicry, robots are now, for the first time, able to move—even if most are still pre-commercial—at the same speed as the animals they mimic. Measured in terms of body-lengths-per-second (blps), Boston Dynamics, for example, has demonstrated a robo-Cheetah that approached the 16 blps speed of a biological cheetah. But just as aircraft can do things birds cannot, robots will be able to do the biologically impossible, such as converting in real time from, say, a rolling machine to a walking machine in order to adapt to terrain.

A couple of decades ago, it would have required a room-sized computer to process, in real time, all the data generated by all those actuators and sensors. Of course, not only has compute power increased to allow on-board capability, but high-bandwidth local wireless networks have enabled remote access to even more powerful computation when needed.

Standard engineering progression will soon take us from Spot to the Atlas-class robots for commercial use as technology improves and costs come down. It's the trajectory seen after every irruption. The emergence of general-purpose robots will echo the pattern of the rise

of the general-purpose transportation machine, the automobile. In the world of cyber-physical machines, the timespans between invention and commercial products are remarkably similar across categories and modern history.

Disrupting the Status Quo

It was in 1901 that one of the first cars was offered for sale signaling that commercial viability was possible. It was a Packard with a then-revolutionary steering wheel, instead of a tiller-like control (the design used since the 15-year earlier first invention of a car). And more critically, the Packard demonstrated the impressive feat of reliably completing a five-day, 300-mile drive. It sold for \$1,500, which was then about 120 percent of an average annual wage. We note that Spot's selling price is about 120 percent of today's average annual wage.

In late 2021, DARPA held its "subterranean challenge" in which teams competed using wheeled, tracked or walking robots that competed to (successfully) perform mining-related tasks in a network of caverns.^{xiv} One of the contestants, for example, demonstrated the ability for its robots to survey and build out a detailed subterranean map in just one

hour, a task that normally entails 100 person-hours of human surveyors to achieve the same precision.

While industrial applications for mobile robots are starting mainly with survey and safety work, a proliferation of vendors has pre-commercial machines capable of working alongside, sometimes replacing, humans in heavy-lift tasks. The warehouse "logistics" markets have become a hotbed for both development and deployment of robots, many to undertake the same kinds of lifting tasks needed across industries. Last year, Boston Dynamics, to note one example, introduced a box-handling robot that can finally match the 800 box-per-hour rate at which humans unload a truck.^{xv} It can move boxes up to 50 pounds and only needs to take a break every 16 hours (to recharge).

In the coming decade, far more robots are expected to be hired by warehouse operators than in all other applications combined. Within five years, overall spending on automation in warehouses is forecast to be more than double last year's \$16 billion, compared to a 60 percent spending increase over the past five years.^{xvi} Given the close alignment in tasks and performance metrics, all that commercialization is bound to accelerate robot capabilities for the adjacent



Spot is "an agile mobile robot," according to manufacturer Boston Dynamics, designed "to automate routine inspection tasks and data capture safely, accurately and frequently." Spot can carry and power up to 14kg of inspection equipment. It can be controlled remotely or programmed for autonomous missions.

The automaton is a class of machine that holds as much promise for disruption (and for fortunes) as did the advent of the automobile.

industrial market. The population of the total robot workforce in industries and services is expected to increase 400 percent by 2030.^{xvii} Odds are good that's an underestimate.

The value of an anthropomorphic robot, outside of entertainment, arises from the fact that the utility of such a machine increases the more easily it can operate in the environs that humans normally occupy—as opposed to specialized environments, such as warehouses or factory-floors designed for most automatons so far. It's not only about having machines that amplify human capabilities, but also doing so by accommodating humans, rather than forcing humans to accommodate machines.

The automaton is a class of machine that holds as much promise for disruption (and for fortunes) as did the advent of the automobile. And it is a class of machine that, more than any other, has excited the dystopian anxieties of doomsayers, particularly those predicting the destruction of all work as we know it.

The anxieties and complaints are similar in character to those voiced by early critics of the automobile. Many bemoaned that road infrastructures changed the landscape, that cars disrupted social norms, that they took jobs from ranchers and horse handlers, etc.^{xviii} Today we find a similar industry of pundits who compete to more loudly decry the consequences of robotification. In his masterpiece *Pneumatica*, Hero wrote, *circa* 50 AD, that while some people back then thought his automatons could “supply the most pressing wants of human life,” for others they engendered “alarm.”^{xix}

Fundamentally, the labor-productivity boost that robotification will bring promises to echo precisely what happened a century ago with the mechanization of industry. More businesses, more services, new kinds of jobs replacing old ones, and more wealth and well-being.

In thinking about our near future, Steffi Paepcke, a senior designer on the robot team at Toyota's research institute, perceptively observed the modern relevance of the apocryphal quote attributed to Henry Ford: “If the inventors of the automobile had asked people riding horses what they wanted, they would have answered that they just wanted a faster horse. It can be difficult to imagine a future that's vastly different from the status quo.”^{xx}

Amen. ☐

ⁱ https://www.barrypopik.com/index.php/new_york_city/entry/demography_is_destiny

ⁱⁱ <https://www.zippia.com/tradesman-jobs/demographics/>

ⁱⁱⁱ <https://mit-serc.pubpub.org/pub/puzzle-of-missing-robots/release/1>

^{iv} <https://mit-serc.pubpub.org/pub/puzzle-of-missing-robots/release/1>

^v Woodcroft, Bennet. *Pneumatica: The Pneumatics of Hero of Alexandria*. New York, NY: Oia Press, 2015.

^{vi} Mayor, Adrienne. *Gods and Robots: Myths, Machines, and Ancient Dreams of Technology*. Lawrenceville: Princeton University Press, 2018.

^{vii} Nocks, Lisa. “500 Years of Humanoid Robots Automata Have Been Around Longer Than You Think.” *IEEE Spectrum* 54, no. 10 (2017): 18–19. <https://doi.org/10.1109/mspec.2017.8048830>

^{viii} Asimov, Isaac. “Visit to the World's Fair of 2014,” *New York Times*, August 16, 1964. <https://archive.nytimes.com/www.nytimes.com/books/97/03/23/lifetimes/asi-v-fair.html>

^{ix} “Atlas.” Boston Dynamics. Accessed April 12, 2021. <https://www.bostondynamics.com/atlas>

^x Mogg, Trevor. “Spot the Robot Dog Is Amazing, and Look How Far It's Come.” *Digital Trends*, June 17, 2020. <https://www.digitaltrends.com/news/spot-the-robot-dog-is-amazing-but-look-how-far-its-come/>

^{xi} <https://www.mdpi.com/journal/actuators>

^{xii} Jaynes, E.T. Rep. *The Muscle as an Engine*. Cambridge, 1983.

^{xiii} Bourzac, Katherine. “A Super-Stretchy Self-Healing Artificial Muscle.” *IEEE Spectrum*, April 18, 2016. <https://spectrum.ieee.org/tech-talk/robotics/robotics-hardware/a-superstretch-selfhealing-artificial-muscle>

^{xiv} <https://spectrum.ieee.org/darpa-subterranean-challenge-2657170650>

^{xv} <https://spectrum.ieee.org/warehouse-robot>

^{xvi} <https://www.statista.com/statistics/1094202/global-warehouse-automation-market-size/>

^{xvii} <https://www.electronicproducts.com/mobile-robotics-from-amrs-to-quadrupeds/>

^{xviii} Ladd, Brian. *Autophobia: Love and Hate in the Automotive Age*. Chicago, IL: Univ. of Chicago Press, 2011.

^{xix} Crawford, James. “The Life and Death of the Library of Alexandria.” *Literary Hub*, March 13, 2017. <https://lithub.com/the-life-and-death-of-the-library-of-alexandria/>

^{xx} McBurnett, Marie. “Designing Robots for Ikigai.” *StackPath*, October 1, 2020. <https://www.machinedesign.com/markets/robotics/article/21143565/designing-robots-for-ikigai>


IN DEFENSE OF (virtuous) AUTONOMOUS WEAPONS

DON A. HOWARD, PHD
Professor of Philosophy
University of Notre Dame

The war in Ukraine is a reminder that the world's major military powers are developing and deploying weapons with ever increasing levels of autonomy in a nearly total vacuum of international law or generally accepted norms for how such systems should be designed and used. Meanwhile, almost 10 years have been wasted in conversations in Geneva under the auspices of the U.N.'s Convention on Certain Conventional Weapons (CCW), as NGOs such as Human Rights Watch (HRW), the

Campaign to Stop Killer Robots and the Future of Life Institute pressed first for a total ban on autonomous weapons, then for a ban on just offensive autonomous weapons, and now for a requirement that there be "meaningful human control" on autonomous weapons.

Below, a life-sized metal replica of a robot soldier from "Laputa: Castle in the Sky," an animated Japanese fantasy adventure movie. The robot statue, standing 16 feet tall, guards the rooftop garden at the Studio Ghibli Museum in Mitaka, Japan, and is a popular tourist attraction. Photo courtesy Matthew Klein.



It was obvious from the start that the major powers would never accept a ban, and the arguments adduced for a ban were as muddled as the concept of meaningful human control.

Yet we must norm this space if we care about justice in war. To do that requires careful thinking and the formulation of politically practical proposals of the kind to be discussed later in this article.

Among the many ethical issues that must be faced as we integrate digitally based autonomous systems throughout our society and economy, those that arise in connection with autonomous weapons are of the greatest urgency, precisely because we are delegating the kill decision to the machine.

Moral Gains from Autonomy?

Let us begin by reflecting on a possibility not acknowledged, for the most part, by the proponents of a ban, namely, that there might be moral gains from the introduction of autonomous weapons. By far the most compelling case of this kind is that made by Ronald Arkin in his 2009 book, *Governing Lethal Behavior in Autonomous Robots*. Do not dawdle over the particular architecture that Arkin suggests in that book, some of which is already dated. Appreciate, instead, his main point, which is that humans are notoriously unreliable systems, that human combatants commit war crimes with frightening frequency, and that what we must ask of autonomous weapons systems is not moral perfection, but simply performance above the level of the average human soldier.

There is not space here to review in detail the study by the U.S. Army Medical Command's Office of the Surgeon General from the Iraq War upon which Arkin mainly bases his assessment of human combatant performance (Surgeon General 2006). Suffice it to say that the number of admitted war crimes by US troops, the number of unreported but observed war crimes, and the self-reported ignorance about what even constitutes a war crime are staggering. With such empirical evidence as background, Arkin's claim to be able to build a "more moral" robot combatant seems far more plausible than one might initially have thought. Why?

Start with the obvious reasons. Autonomous weapons systems suffer from none of the human failings that so often produce immoral behavior in war. They feel no fear, hunger, fatigue or anger over the death of a friend. Move on to the slightly less obvious reasons. Thus, a robot, not fearing for its own well-being, can easily err on the side of caution, choosing not to fire in moments of doubt, where a human might rightly have to err on the side of self-defense. Then consider still more important design constraints, such as those embodied in Arkin's "Ethical Adaptor," into which are programmed all relevant parts of the International Law of Armed Conflict, International Humanitarian Law, and the rules of engagement specific to a given conflict arena or a specific action (Arkin 2009, 138-143). The Ethical Adapter blocks the "fire" option unless all of those prescriptions are satisfied. Arkin's robots could not fire (absent an override from a human operator) at all, unless the most stringent requirements are met. In the face of uncertainty about target identification, discrimination, applicability of rules of engagement and so forth, the robot combatant defaults to the "no fire" option. Of course, other militaries could design the robots differently, say, by making "fire," rather than "no fire," the default. But hold that thought until we turn at the end to the discussion of a specific regulatory regime.

Arkin illustrates the functioning of the Ethical Adaptor with several scenarios, one of which—a Taliban gathering in a cemetery for a funeral (Arkin 2009, 157-161)—bears an eerie similarity to the horrific U.S. attack on a Doctors without Borders (Médecins Sans Frontières [MSF]) hospital in Kunduz, Afghanistan, in October 2015 (Rubin 2015). The rules of engagement as uploaded to the Ethical Adaptor would typically include specific coordinates for areas within which no fire would be permitted, including hospitals, schools, important cultural monuments and other protected spaces. Likewise, no fire could be directed at any structure, vehicle or individual displaying the red cross or the red crescent. This assumes, of course, sensor and AI capabilities adequate for spotting and correctly identifying such insignia, but, especially with structures and vehicles, where the symbol is commonly painted in large, high-contrast format on the roof, that is not a difficult



In February 2023, the EurAsian Times reported that the first four Russian Marker robotic weapons platforms were deployed in Eastern Ukraine. The unmanned Marker is equipped with a modular multispectral vision system and neural networks to autonomously detect and destroy targets, prioritizing enemy tanks. Credit: Kirill Borisenko, Wikimedia Commons

problem. A fully autonomous drone designed, as per Arkin's model, which was tasked with the same action that led to the bombing of the MSF hospital in Kunduz, simply would not have fired at the hospital. A human might have overridden that decision, but the robot would not have fired on its own. Moreover, the kind of robot weapon that Arkin has designed would even remind the human operator that a war crime might be committed if the action proceeds.

Another kind of moral gain from autonomous weapons was once pointed out to me by an undergraduate student, an engineering major, in my Robot Ethics class. He recalled the oft-expressed worry about the dehumanization of combat with standoff weapons, such as remotely piloted drones. The concern is that the computer-game-like character of operator interfaces and controls, and the insulation of the operator from the direct risk of combat, might dull the moral sensitivity of the operator. But my student argued with deliberate and insightful irony, that the solution to the problem of dehumanization might be to take the human out of the loop, because it is the human operator who is, thus, dehumanized.

For the record, I would dispute the dehumanization argument in the first place, because the typical drone operator often watches the target for many minutes,

if not hours, and gets to know the humans on the receiving end of the munitions—including the wives and children—far better than does, say, an artillery officer, a bombardier in a high-altitude bomber or even the infantryman who gets, at best, a fleeting and indistinct glimpse of an enemy combatant across a wide, hazy, busy field of combat. That drone operators get to know their targets so well is part of the explanation for the extremely high reported rates of PTSD and other forms of combat stress among them (Chapelle et al. 2014). Still, my student's point was a good one. If dehumanization is the problem, then take the dehumanized human out of the loop. This is really just a special case of Arkin's point about how stress and other contextual circumstances increase the likelihood of mistakes or deliberate bad acts by humans in combat and that, since robots are unaffected by such factors, they will not make those mistakes.

One of the most common criticisms of Arkin's model was voiced in the original HRW call for a ban, namely, that sensor systems and AI are not capable of distinguishing combatants from non-combatants, so that, even if the principle of discrimination is programmed into a robot weapon, it still cannot satisfy the requirements of international law. But there are two obvious responses to this criticism: (1) what is or is not technically feasible is an empirical



question to be decided by further research, not on *a priori* grounds; and (2) discrimination is usually a highly context-dependent challenge, and in some contexts, such as finding and identifying a Red Cross or Red Crescent symbol, the problem is easily solved.

The other major criticism of Arkin's model is that, since it assumes a conventional, structured, top-down decision tree approach to programming ethics and law into autonomous weapons, it cannot deal with the often-bewildering complexity of real battlefield situations.

The basis of the objection is a simple and old worry about any rule-based or algorithmic approach to ethical decision making, such as deontology or consequentialism. It is that one cannot write a rule or build a decision tree to cover every contingency and that the consequentialist's calculation of benefit and risk is often impossible to carry out when not all consequences can be foreseen. The objection is a good one, at least by way of pointing out the limited range of applicability of Arkin-type autonomous weapons systems.

But Arkin's model for ethical autonomous weapons design is only a beginning. This last objection—that one cannot write a rule to cover every contingency—is the main reason why some of us are hard at work on developing a very different approach to ethics programming for artificial systems, one inspired by the virtue ethics tradition and implemented via neural nets and machine learning algorithms. The idea—already explored in concept by Wendell Wallach and Colin Allen in their 2010 book, *Moral Machines* (Wallach and Allen 2010)—is to supplement Arkin's top-down approach, involving rules and perhaps a consequentialist algorithm, with a bottom-up approach in which we design autonomous systems as moral learners, growing in them a nuanced and plastic moral capacity in the form of

habits of moral response, in much the same way that we mature our children as moral agents (Muntean and Howard 2017).

There is considerable debate about this approach via moral learning. Arkin, himself, objects that neural nets and learning algorithms “black box” the developed competence in such a way as to make impossible both the robot’s reconstructing for us either a decision tree or a moral justification of its choices, which he regards as a minimum necessary condition on moral machines, and the operator’s reliably predicting the robot’s behavior (Arkin 2009, 67, 108). We respond that human moral agents are also somewhat unpredictable and that what they produce, when pressed for a justification of their actions, are after-the-fact rationalizations of moral choices. Why should we demand more of moral robots? How to produce after-the-fact rationalizations is an interesting technical question, one currently being vigorously and successfully investigated under such headings as “rule extraction,” “interpretable AI,” and “explainable AI” (See Samek, et al. 2019).

Others object that there is no consensus on what morality to program into our robots, whether through learning or rule sets. We respond that moral diversity among robots should be prized in the same way that we prize human moral diversity. We learn from one another because of our moral differences. But, at the same time, in the constrained space of autonomous weapons, there is consensus in the form of the international support for extant international law and the just war moral theory, upon which it is based. Saudi Arabian health care robots might rightly evince different habits with respect to touching and viewing unveiled bodies from those evinced by North American or European health-care robots. But Saudi Arabia has ratified the main principles of the Geneva Conventions, as has the U.S.

There are other potential moral gains from autonomous weapons, such as facilitating military intervention to prevent genocide or other human rights abuses, minimizing risk of death or injury to our troops, and sparing drone operators and other personnel both psychological damage and moral

corrosion from direct participation in combat. One can imagine still more, such as employing weaponized autonomous escort vehicles to protect aid convoys in conflict zones. The conclusion is that there are, in fact, noteworthy potential moral gains from the development and deployment of both offensive and defensive autonomous weapons. Of course, this must be done in such a way as to ensure compliance with existing international law and in a manner that minimizes the likelihood of the technology’s being put to the wrong uses by bad actors. Short of a ban on autonomous weapons, how do we do that?

Article 36 Regulatory Regime

The goal is regulating the development and deployment of autonomous weapons in a way that ensures compliance with international law and minimizes the chance of misuse. Moreover, we need to do this in a politically feasible way, using regulatory structures that will be accepted by the international community. This last point is important, because, as mentioned, one common criticism of the proposed ban on autonomous weapons is, precisely, that it stands little chance of ever being incorporated into international law.

Even in the talks under the aegis of the CCW, which have been going on since 2014 in Geneva, it is mainly only nations with little or no prospect of becoming significant participants in the development and use of autonomous weapons that have shown support for moving forward with consideration of a ban. The major players, including the U.S., have repeatedly indicated that they will not support a ban. In December 2021, the American representative in Geneva, Josh Dorosin, said it again, while adding that a non-binding, international code of conduct might be appropriate (Bowden 2021). That sufficiently strong support for a ban was unlikely ever to emerge from the Geneva talks was already clearly sensed six years ago by the most energetic proponents of the ban. Thus, in a 2016 press release, the Stop Killer Robots campaign subtly shifted the discourse, hinting at a tactical retreat, by urging a focus on “meaningful human control” (whatever that might mean), though talk of a ban still dominates the headlines. If the goal is regulating the development and use of autonomous weapons in a politically feasible way,

then nine years of talks have been wasted by the continued insistence on a ban.

What could the international community have been discussing instead? The discussion should have focused on what might be done within the compass of extant international law. There is already in place, since 1977, Article 36 of Protocol I to the Geneva Conventions, which stipulates:

In the study, development, acquisition or adoption of a new weapon, means or method of warfare, a High Contracting Party is under an obligation to determine whether its employment would, in some or all circumstances, be prohibited by this Protocol or by any other rule of international law applicable to the High Contracting Party.ⁱ

One hundred and seventy-four states have ratified Protocol I, including Article 36, and three states, Pakistan, Iran and the U.S., are signatories but have not formally ratified the Protocol.ⁱⁱ But the U.S. has promised to abide by nearly all provisions, including Article 36, and has established rules and procedures in all four branches of the military for ensuring legal

review of new weapons systems (see ICRC 2006 and, for example, U.S. Army 2019). The countries having ratified Protocol I include every other major nation, among them China, the Russian Federation and all NATO member states. I would argue that, since Article 36 is already a widely accepted part of international law, it is the best foundation upon which to construct a regulatory regime for autonomous weapons.

Concerns have been expressed about the effectiveness of Article 36 in general, chief among them that the prescribed legal reviews are sometimes perfunctory, and that it is too easy to evade an Article 36 review by declaring that a weapon is not new but just a minor modification of an existing and already authorized weapon. Those are serious worries, as evidenced by the recent controversy over whether the American redesign of the B61 nuclear warhead with a tail assembly that makes possible limited, real-time steering of the warhead, the configuration designated now as B61-12, constituted a new weapon, as critics allege, or merely a modification, as the U.S. asserts (see Mount 2017).

Another worry is that only a small number of states have certified that they are regularly carrying out



In November 2021, Ukrinform reported that the "British Royal Air Force is sending Ukraine advanced laser-guided Brimstone 2 missiles." Brimstone is a dual-mode weapon. It can be operated in a manual mode, with the pilot or other operator selecting the target and guiding the weapon to that target. Or, it can be operated in autonomous mode, with the pilot releasing the weapon after

which the weapon, itself, identifies the target and decides to strike. Once released in autonomous mode, Brimstone's the only constraint is that its operation is confined to a delimited field of fire.

Article 36 reviews. Equally serious are concerns that have been expressed about the effectiveness of Article 36 specifically with respect to autonomous weapons, as in a briefing report for delegates to the 2016 meeting of experts, which argued that what is at issue is not so much the conformity of individual weapons systems with international law, but the wholesale transformation of the nature of warfare wrought by the “unprecedented shift in human control over the use of force” that autonomous weapons represent. The magnitude of that change was said to require not individual state review but the engagement of the entire international community (CCW 2016). All such concerns would have to be addressed explicitly in the construction of an autonomous weapons regulatory regime based on Article 36.

How would a new Article 36 regulatory regime be constructed? Most important would be the development of a set of clear specifications of what would constitute compliance with relevant international law. This could be the charge to a Group of Governmental Experts under the auspices of the U.N.’s CCW.

First in importance among such guidelines would be a detailed articulation of what capabilities an autonomous weapon must possess for handling the problem of discrimination, bearing in mind the point made above that this is not an all-or-nothing capability, but, rather, one specific to the functions and potential uses of an individual weapons system. For example, Great Britain’s fire-and-forget Brimstone cruise missile, which can be operated in an autonomous mode, needs only the capability to distinguish different categories of vehicles, say battle tanks versus passenger vehicles, within its designated field of fire.

An autonomous check-point sentry, by contrast, would have to be capable of much more sophisticated discriminations. Similarly detailed specifications would have to be developed for determinations of proportionality, recognition of a human combatant’s having been rendered *hors de combat*, recognition of a target’s displaying insignia, such as the Red Cross or Red Crescent, that identify a structure, vehicle or individual as protected medical personnel, and so forth.

Just as important as developing the specifications would be the development of protocols for testing to ensure compliance. Optimal, but politically unachievable, for obvious reasons, would be the open sharing of all relevant design specifications. It is highly unlikely that states and manufacturers are going to let the world community look under the hood at such things as new sensor technologies and accompanying software. The alternative is demonstrations of performance capability in realistic testing scenarios. We already have considerable relevant experience and expertise in safety and effectiveness testing for a wide range of engineered systems, especially pertinent being the testing protocols for certifying control systems in commercial aircraft and industrial systems. One might think that weapons developers would be just as shy about showing off the weapon at work in realistic scenarios, lest adversaries and competitors infer confidential capabilities and technologies. But, in fact, most weapons developers are proud to show off videos of their new systems doing impressive things and to display and demonstrate their products at international weapons expositions. What would be required would not be the sharing of secrets but simply demonstrations of reliability in complying with the detailed guidelines just discussed.

As with the existing Article 36 requirements, certification of compliance will surely have to be left to individual states. But it is not unreasonable to begin an international conversation about a more public system for declaring that the required certifications have been carried out, even if that consists in little more than asking signatories and states parties to file such certifications with the U.N., the International Committee of the Red Cross or another designated international entity.

The good news is that, within just the last few years, serious discussion of precisely such concrete elaborations of Article 36 protocols for autonomous weapons has begun to appear in the scholarly, policy and legal literatures (see, for example, Poitras 2018, Cochrane 2020 and Jevglevskaia 2020). Equally encouraging is the willingness of some governments to underwrite such work. Thus, the German Auswärtiges Amt (Foreign Office) subsidized a 2015

expert seminar under the auspices of the Stockholm International Peace Research Institute (SIPRI) that had representation from France, Germany, Sweden, Switzerland, the United Kingdom and the U.S. (Boulamin 2015).

What have been the fruits of such work? Many good ideas have emerged. Especially thoughtful are the main recommendations contained in a 2017 report, sponsored by SIPRI, covering Article 36 elaborations for cyber weapons, autonomous weapons and soldier enhancement. Their approach is to focus on advice to reviewing authorities in individual member states, and they emphasize two broad categories of advice: (1) building on best practices already being employed by states that have well-developed review procedures, and (2) strengthening transparency and cooperation among states.

Under the first heading, they advise, for example:

1. Start the review process as early as possible and incorporate the procurement process at key decision points.
2. Provide military lawyers involved in the review process with additional technical training. Engineers and systems developers should also be informed about the requirements of international law, so that they can factor these into the design of the weapons and means of warfare. (Boulamin and Verbruggen 2017, viii)

About increased transparency and cooperation, they say that it would become a “virtuous circle,” and they observe that:

1. It would allow states that conduct reviews to publicly demonstrate their commitment to legal compliance.
2. It would assist states seeking to set up and improve their weapon review mechanisms and thereby create the conditions for more widespread and robust compliance.
3. It could facilitate the identification of elements of best practice and interpretative points of guidance for the implementation of legal reviews, which would strengthen international confidence in such mechanisms.

They add:

Cooperation is also an effective way to address some the outstanding conceptual and technical issues raised by emerging technologies. Dialogues, expert meetings and conferences can allow generic issues to be debated and addressed in a manner that does not threaten the national security of any state. (Boulamin and Verbruggen 2017, viii)

When it comes specifically to Article 36 reviews involving autonomous weapons, they identify as the foremost challenge verifying “the predictability of autonomous weapon systems’ compliance with international law” (Boulamin and Verbruggen 2017, ix).

I am not at all naive about how strict compliance with Article 36 requirements would be. But existing Article 36 requirements have already created a culture of expectations about compliance and a space within which states can and have been challenged, sometimes successfully, to offer proof of compliance, as with the widely expressed concerns about truly indiscriminate weapons, such as land mines and cluster munitions. We begin to norm such a space simply by putting the relevant norms in front of the world community and initiating a public conversation about compliance. This is what we should be talking about in Geneva if we are serious about building some measure of international control over autonomous weapons.

Towards the Ultimate Goal

War is hell. It will always be an inherently immoral form of human activity. The goal of international law is to minimize the otherwise inevitable death and suffering that war entails. Advances in technology can contribute toward that goal by making weapons more accurate, less lethal and more selective. The advent of autonomous weapons promises still further moral gains by removing the single most common cause of war crimes, the too often morally incapacitated human combatant. We cannot let unrealistic fears about a Terminator-AI apocalypse prevent our taking advantage of the opportunities for moral progress that properly designed and deployed autonomous weapons afford. We must, of course, ensure that such systems are being used for good, rather than malign purposes, as we must with any technology

and especially technologies of war. Indeed, with autonomous weapons we need to be more vigilant. But minimizing death and suffering in war is the ultimate goal. If autonomous weapons can contribute to progress toward that goal, then we must find a way to license their use in full compliance with what law and morality demand. ■

This article is adapted from “In Defense of (Virtuous) Autonomous Weapons.” Notre Dame Journal on Emerging Technologies, 3;2 (November 2022).ⁱⁱⁱ

ⁱ <https://www.icrc.org/ihl/WebART/470-750045>

ⁱⁱ <https://ihl-databases.icrc.org/ihl/INTRO/470>

ⁱⁱⁱ <https://ndlsjet.com/in-defense-of-virtuous-autonomous-weapons/>

■ References

- Arkin, Ronald (2009). *Governing Lethal Behavior in Autonomous Robots*. Boca Raton, FL: Chapman Hall/CRC.
- Boulanin, Vincent (2015). “Implementing Article 36 Weapon Reviews in the Light of Increasing Autonomy in Weapon Systems.” SIPRI Insights on Peace and Security, no. 2015/1. Stockholm International Peace Research Institute. November 2015. <https://www.sipri.org/sites/default/files/files/insight/SIPRIInsight1501.pdf>
- Boulanin, Vincent and Maaïke Verbruggen (2017). *Article 36 Reviews: Dealing with the Challenges Posed by Emerging Technologies*. Stockholm, Sweden: Stockholm International Peace Research Institute.
- Bowden, John (2021). “Biden Administration Won’t Back Ban on ‘Killer Robots’ Used in War.” *The Independent*. December 8, 2021. <https://www.independent.co.uk/news/world/americas/us-politics/biden-killer-war-robots-ban-b1972343.html>.
- CCW (2016). “Article 36 Reviews and Addressing Lethal Autonomous Weapons Systems.” Briefing Paper for Delegates at the Convention on Certain Conventional Weapons (CCW) Meeting of Experts on Lethal Autonomous Weapons Systems (LAWS), Geneva, 11-15 April 2016. <http://www.article36.org/wp-content/uploads/2016/04/LAWS-and-A36.pdf>.
- Chapelle, Wayne, et al. (2014). “An Analysis of Post-Traumatic Stress Symptoms in United States Air Force Drone Operators.” *Journal of Anxiety Disorders* 28. 480-487.
- Cochrane, Jared M. (2020). “Conducting Article 36 Legal Reviews for Lethal Autonomous Weapons.” *Journal of Science Policy & Governance* 16;1 (April 2020). https://www.sciencepolicyjournal.org/uploads/5/4/3/4/5434385/cochrane_jspg_v16.pdf.
- ICRC (2006). “A Guide to the Legal Review of New Weapons, Means and Methods of Warfare: Measures to Implement Article 36 of Additional Protocol I of 1977.” *International Review of the Red Cross* 88, 931-956.
- Jevglevskaia, Natalia (2022). *International Law and Weapons Review: Emerging Military Technology and the Law of Armed Conflict*. Cambridge: Cambridge University Press.
- Mount, Adam (2017). “The Case against New Nuclear Weapons.” Center for American Progress. May 4, 2017. <https://www.americanprogress.org/article/case-new-nuclear-weapons/>.
- Muntean, Ioan and Don Howard (2017). “Artificial Moral Cognition: Moral Functionalism and Autonomous Moral Agency.” *Philosophy and Computing*. Thomas Powers, ed. Cham, Switzerland: Springer, 2017, 121-159.
- Poitras, Ryan (2018). “Article 36 Weapons Review & Autonomous Weapons Systems: Supporting an International Review Standard.” *American University International Law Review* 34, 465- 495.
- Rubin, Alissa J. (2015). “Airstrike Hits Doctors Without Borders Hospital in Afghanistan.” *New York Times*. October 3, 2015. <https://www.nytimes.com/2015/10/04/world/asia/afghanistan-bombing-hospital-doctors-without-borders-kunduz.html>.
- Samek, Wojciech, et al., eds. (2019). *Explainable AI: Interpreting, Explaining, and Visualizing Deep Learning*. Cham, Switzerland: Springer.
- Stop Killer Robots (2016). “Focus on Meaningful Human Control of Weapons Systems—Third United Nations Meeting on Killer Robots Opens in Geneva.” Stop Killer Robots. April 11, 2016. <https://www.stopkillerrobots.org/news/press-release-focus-on-meaningful-human-control-of-weapons-systems-third-united-nations-meeting-on-killer-robots-opens-in-geneva/>
- Surgeon General (2006). Mental Health Advisory Team (MHAT) IV Operation Iraqi Freedom 05-07. “Final Report.” Office of the Surgeon General. November 7, 2006. <https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2010103335.xhtml#>
- U.S. Army (2019). “Legal Review of Weapons and Weapon Systems.” Army Regulation 27-53. Washington, DC: Headquarters, Department of the Army, 23 September 2019. https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/ARN8435_AR27-53_Final_Web.pdf.
- Wallach, Wendell and Colin Allen (2008). *Moral Machines: Teaching Robots Right from Wrong*. New York, NY: Oxford University Press.

Food Security for the Future

*Empowering Farmers & Ranchers
Through Precision Ag Technologies*

THE HONORABLE JOHN HOEVEN
United States Senator for North Dakota



Americans enjoy the lowest cost, highest quality food supply in the world. On average, the people of this nation spend only 10 percent of their disposable income on food, which is a testament to the hard work of our farmers and ranchers. However, the events of recent years have shown that we cannot take our abundant and affordable food supply for granted. From potential food shortages due to the war in Ukraine and supply chain constraints, to the impact of record-high inflation on U.S. household budgets, it's clear that the world needs an even greater abundance of high-quality American agricultural products to ensure a secure food supply. The good news is that North Dakota is already working to empower our farmers and ranchers to meet that need now and into the future.

Our state has long been a global leader in agriculture, but it used to be that you could grow only a few crops in western North Dakota. Through technological advancements and the development of new varieties of crops in recent decades, we have vastly expanded production agriculture in our state, which stands as the top U.S. producer in a variety of commodities, including spring and durum wheat, dry edible beans and peas, pinto beans, canola, flaxseed and honey.

At the same time, North Dakota soybean and corn production has expanded by a combined 3 million acres over the previous decade, an increase of 36 percent overall, while also securing growth in value-added agriculture, including biofuel production. Such developments have not only strengthened the farm economy and our food supply but coincided with the rapid development of our state's oil and gas industry, enhancing North Dakota's role as both a global agriculture and energy powerhouse.

The success of North Dakota agriculture is a strong example of the overall trends in the U.S., where farm

productivity grew by nearly 300 percent between 1948 and 2017, despite farm inputs, such as fuel and fertilizer costs, remaining mostly level across this time frame. In order to maintain such dramatic progress, we need to advance the next frontier of farming and ranching. That's exactly what precision agriculture technologies will help accomplish.

Precision Agriculture

Precision agriculture is a collection of innovations, such as autonomous vehicles, remote sensing, satellite imagery, mobile applications and unmanned aerial systems (UAS). These advancements give farmers and ranchers the real-time information and capabilities required to form more effective plans, streamline their operations, reduce their input costs and increase their yields.

Many of these technologies are already being adopted by agriculture producers, resulting in fewer hours working the fields and significantly reduced fertilizer, pesticide, water and fuel consumption. In fact, a majority of farmers already uses products such as auto guidance on their farm equipment and yield monitors in their fields, but there is plenty of headroom available for further progress using existing technologies. For instance, a 2021 study estimated that precision agriculture has resulted in a 4 percent and 9 percent decrease in agriculture water and herbicide usage, respectively, while estimating that there is potential for further reductions of 21 and 15 percent should these practices and technologies be fully adopted. In clearer terms, that represents an additional decrease of 2.4 million gallons of water and 48 million pounds of herbicide per year in the U.S., reducing costs to producers and consumers, as well as benefiting our water and soil resources.

Tech as the Third Wave

These estimates, while impressive, do not account for the innovations that are to come. That's where we get back to the efforts underway in North Dakota. Following the expansion in agriculture and energy

Sen. John Hoeven talking with Craig Rupp, Founder & CEO of Sabanto, a Chicago-based robotics company, at Grand Farm's Autonomous Nation event on September 16, 2021, in Horace, ND. They are discussing a fully autonomous tractor, which Sabanto donated to Grand Farm.

production in our state, we have experienced the rise of our state's technology sector. Starting in my days as governor, we worked to build our state into a hub of technology entrepreneurship. We did so by creating the right kind of legal, tax and regulatory environment that encouraged investment and business formation, while leveraging our universities as engines of innovation.

As a result, technology is serving as the third wave in North Dakota's economic growth, with technology subsectors contributing more than \$3 billion to our state's gross domestic product in 2020, up 39 percent from 2010.

Our UAS industry is a prime example of this success. North Dakota is a premiere location for unmanned aircraft research, development, training and operations. That's due to the unique collection of public and private partners we've brought together, whether it's federal, state and military agencies or businesses ranging from startups to multinational corporations. These collaborations are centered around the Northern Great Plains UAS Test Site and Grand Sky, a first-of-its-kind UAS technology and business park that is adjacent to the Grand Forks Air Force Base.

Through these partnerships, we are bringing to life new initiatives that reach across sectors, leading to developments far beyond our original vision. When we began building this industry in our state nearly 18 years ago, innovation meant getting unmanned aircraft into our domestic airspace. Now, we have built partnerships that reach all the way to space, tying our UAS industry to two missions that are essential to the future security of our nation, including a satellite mission with the Space Development Agency and Sky Range, a new hypersonic missile testing program based at Grand Sky.

Broadband

One important reason North Dakota has been able to sustain these exciting new developments is that we've worked to ensure our state has the capacity to support data-intensive operations. Unmanned vehicles, arrays

of sensors and numerous other devices, involved in our tech industry, collect and transmit immense volumes of information, which require substantial broadband infrastructure to carry it all from the point of origin to the data servers and users. Again, this is another point where North Dakota is in a position of leadership.

According to U.S. News and World Report, our state is in the top five nationwide for access to gigabit internet service. Nearly 380 communities in our state, as well as surrounding rural regions, have access to gigabit broadband with a download speed of 1,000 megabits per second. For context, the Federal Communications Commission (FCC) defines broadband speed as a minimum download speed of

Senator Hoeven Working to Expand North Dakota Broadband Access

As the lead Republican on the Senate Agriculture Appropriations Committee, Senator Hoeven advanced the following priorities to help move the efforts of North Dakota's broadband providers forward by:

- Establishing the ReConnect Program, a rural broadband loan and grant pilot program at the U.S. Department of Agriculture, which has provided critical support to our state's broadband providers.
- Passing legislation that will provide a minimum of \$100 million over five years to North Dakota for further broadband expansion.
- Cosponsoring and passing legislation to:
 - Identify and measure gaps in the availability of broadband on agricultural land, as well as develop recommendations to promote the rapid, expanded deployment of fixed and mobile broadband to unserved areas.

25 megabits per second. That means a large portion of our state has service that is 40 times faster than the federal definition, with almost all other North Dakotans having access to at least 100-megabit speeds.

This is the result of long-term efforts and investments to deploy this infrastructure to every corner of our state. In particular, Dakota Carrier Network (DCN) and its owner companies have built out 65,000 miles of fiber optic cable to serve as the backbone for internet service statewide. Such an undertaking required the forethought and hard work of broadband providers throughout the state, as well as the support of federal, state and local governments to make sure project developers had regulatory certainty, as well as access to the necessary capital.

- Improve federal broadband mapping by making the FCC's data more granular, including more state, local and tribal government data, and establishing a challenge process to ensure the accuracy of data provided by Internet Service Providers.

Supporting Grand Farm and Precision Agriculture

- Securing \$1 million to establish a cooperative agreement between Grand Farm, NDSU and USDA's ARC.
- Bringing USDA Secretary Sonny Perdue to North Dakota to learn about the Grand Farm initiative firsthand.
- Inviting Grand Farm Director Brian Carroll to testify at a hearing of the Senate Agriculture Appropriations Committee on strengthening rural economies.
- Looking forward, provided \$4 million for precision ag research in Fiscal Year 2023 Agriculture Appropriations.

Grand Farm

Now with these pieces in place, we're bringing the same approach from Grand Sky and our UAS industry to precision agriculture. Grand Farm, a project led by Emerging Prairie, seeks to develop and demonstrate the next generation of precision agriculture. This organization has made tremendous strides in pulling together important partners, such as Microsoft, CHS, Doosan Bobcat, RDO Equipment Co., DCN and the North Dakota Department of Commerce, to name a few, to support this endeavor. Moreover, we've worked with Grand Farm and North Dakota State University to establish and fund, through the regular appropriations process, a \$1-million cooperative agreement with the Agricultural Research Service (ARC) to enhance the research undertaken by this initiative, which has just reached a key milestone.

Grand Farm recently broke ground on a new 140-acre site near Casselton, North Dakota, where an Innovation Facility will be developed. Here, the project partners will be able to research, rapidly prototype and deploy new agriculture technologies. Construction began in October 2022, and projects are expected to begin deploying in spring 2023, marking the beginning of an exciting time for North Dakota agriculture producers, who are positioned to remain on the cutting-edge of new technology, enabling them to continue competing in the global marketplace.

With a dynamic technology sector, an established leadership in production agriculture, broadly accessible high-speed internet infrastructure and innovative collaborations like Grand Farm, North Dakota has a substantial head start in advancing the future of precision agriculture technologies. Ultimately, this benefits every American because it will help ensure continued access to the most reliable and affordable food supply in the history of the world. Throughout it all, we cannot take such a central aspect of our quality of life for granted. These things don't just happen—it's up to the many partners across the state, most importantly our farmers and ranchers who are working every day to ensure that food security will continue to be in place for years to come. ☐

Combining Precision Ag Technologies to Improve Crop Management

Improving Site
Specific Weed
Control in Corn



PAULO FLORES, PHD
Assistant Professor,
Department of Agriculture
& Biosystems Engineering, NDSU

Various emerging technologies in precision agriculture improve crop management, yields and agricultural sustainability. Together, several of these technologies (including GPS systems, drones, image analysis and advanced spraying systems) can be used to improve traditional weed control in corn fields and, at the same time, provide financial savings for corn producers, as well as decrease the amount of herbicide applied to croplands.

Corn was chosen for our research because it is a widely grown crop in North Dakota and because of the 30-inch space between the rows used for corn production at NDSU's Carrington Research Extension Center (located 142 miles northwest of Fargo). The wider spacing facilitates mapping the weeds growing between rows since they are easier to see in drone's images, which is a key point of this research project.

Currently, the most common approach for weed control in corn involves two herbicide applications per growing season. The first is usually done pre-planting or pre-emergence (before the corn seedlings grow out of the ground) in order to rid the field of weeds for crop establishment and early development.

The second application is done when the corn plants have four to six leaves to take care of weeds that were missed in the first application. Both applications are usually done in blanket fashion, meaning the whole field is sprayed regardless of weed presence or distribution. However, weeds usually grow in patches across a field, so it makes sense to spray only those areas, especially during the second herbicide application.

Site-Specific Weed Control

Site-specific weed control (SSWC) is known as a practice to control weeds according to their spatial distribution across a field. Commercial solutions, such as WeedSeeker by Trimbleⁱ and See & Spray Select[™]

by John Deere,ⁱⁱ enable farmers to implement SSWC on fallow ground or later in the season when the machinery's sensors can differentiate green weeds from the brown leaves of mature crops.

Recently John Deere upgraded their technology with See & Spray[™] Ultimate,ⁱⁱⁱ which uses 36 cameras mounted across the sprayer boom (120 ft), computer vision and machine learning to distinguish weeds from crops and to activate the nozzles to spray only the weeds. All of that is accomplished in only a fraction of a second.

The downside of this impressive technological advance is the hefty cost. See & Spray[™] Ultimate comes only as a factory installed option, meaning a new sprayer must be purchased. The current cost for a John Deere sprayer (410R R Self-Propelled Sprayer, the smaller sprayer with that latest technology) with the Spray & Select[™] technology is about \$786,000. Most likely the upgrade to See & Spray Ultimate,[™] the price of which hasn't been made public yet, will increase the cost even more, especially when considering larger sprayers (612R model, for example).

Assistant Professor Paulo Flores (far left) with several research team members: Jithin Mathew (doctoral student), Taofeek Mukaila (master's student) and Nadeem Fareed (postdoctoral associate). Last spring, the team used a drone and high-resolution cameras to map the weeds in corn fields at NDSU's Carrington Research Extension Center.

Below, Prof. Flores and his research team's DJI Matrice 300 drone with a MicaSense AltumPT camera. Photographs by Jerry Anderson.



NDSU Solution

To overcome the challenge of differentiating weeds from crops during the growing season, my graduate students and I have been collaborating with two researchers at the NDSU Carrington Research Extension Center (CREC)—Michael Ostlie, PhD, (formerly a research agronomist and now CREC's director) and David Kramar, PhD, (a precision agriculture specialist at CREC)—to develop and implement a SSWC solution when corn plants have four to six leaves. The solution uses drones, specialized software and state-of-the-art sprayer technology.

Hi-Res Imagery

The first step involves collecting high-resolution imagery of the corn field with a drone. In the past few years, we used a DJI Matrice 600 Pro drone equipped with a 42-megapixel camera, which allowed us to collect high-resolution imagery when flying at an altitude of 350 ft. The images enabled us to separate ground from plants and weed from crops, since we could identify weeds as small as a quarter-inch square. However, we had to slow the drone down to allow enough time for the camera to capture the images, which resulted in three flights that took about an hour to cover 50 acres. That led us to look for better hardware solution to improve data collection.

During the 2022 growing season, we upgraded to a DJI Matrice 300 drone to fly two cameras, a DJI Zenmuse P1 and a MicaSense AltumPT. Then, with both cameras, we were able to fly the same land area in one flight in about 30 minutes. As well, these cameras made it easier to process the data after the flights, further increasing efficiencies.

Mapping

In addition to high-resolution photos, we need to ensure that the location of weeds on the images matches as close as possible to their location on the ground. To accomplish this, the drone is equipped with a GPS RTK receiver system with sub-inch accuracy, which matches the accuracy of the sprayers with the same GPS technology.

Once we have the field images from the drone, we run them through Pix4Dmapper software to create—that

is, “stitch”—one large orthomosaic image or map of the entire field. Pix4Dmapper uses the GPS data, encoded in each image, to correct for distortions caused by the inherent characteristics of any camera's sensor and due to topographic changes on the landscape. The software then seamlessly stitches all the images into one large image or, more technically correct, orthoimage or orthomosaic (Figure 1).

Every pixel on the orthomosaic image is mapped to its real location on the earth's surface using precise latitudinal and longitudinal values. This is a key aspect for us. Just as the GPS app on a cellphone enables drivers to reach a specific location, the orthomosaic imagery enables us to track the position of the weeds in relation to each nozzle of the sprayer, so we can turn nozzles off at “locations” with no weeds.

Weed-Only Orthomosaic

Now that we have the orthomosaic imagery of the field, it is time to map the weeds. First, we identify the rows of corn plants across the field (yellow lines on Figure 3). Then, in our approach, we determine that any plant growing between the yellow rows is a weed (green polygons in Figure 3), so it needs to be sprayed. To create a weed-only orthomosaic, we remove all corn rows and classify the remaining vegetation as weeds.

Weed Map

Next, we need to convert weed location into instructions that the sprayer can understand and so know where to shut off spray nozzles as it moves across the field. We do that by placing a grid cell, usually 10 x 10 feet, over the field's weed-only orthoimage as described in Figure 3. A cell that contains at least one weed is sprayed, while cells free of weed are not sprayed. This approach enables us to take care of small weeds that might not have been detected in the cell and also weeds growing within the corn rows in each cell, since the 10-ft-wide cell covers up to 4 corn rows.

After determining which cells should be sprayed (15 gallons/acre of solution, in our case) or not (0 gallons/acre of solution), we use AgLeader SMS software to translate the “weed map + grid cell” information into an application rate map, which the computer in the sprayer's cab can read and act upon. This process is

known as writing a prescription map, or an Rx map, to the sprayer. Just as one would follow a medicine prescription, the sprayer will follow the instructions provided in the Rx map and apply herbicide at a pre-determined rate (15 or 0 gallons/acre, in our case). Based on that Rx map, the sprayer's nozzles turn on and off automatically as the machine is driven across the field.

In the summer of 2022, we implemented this entire protocol, with one small adjustment. Instead of using the original cell size of 10 x 10 ft, we used a 5ft wide x 10 ft long grid cell size to improve savings on sprayed acres. By doing so, we did not have to spray 50 percent of the area where we were implementing the SSWC (Figure 4). To put that into perspective, for each quarter of section

Figure 1

A sample of an RGB (red, green, blue) orthomosaic that is composed of high-resolution images taken by the Zenmuse P1 camera, which produces photos similar to what a consumer camera or cellphone typically takes. All images in this article were taken from our drone flying at 200 ft above ground on June 28, 2022.

Figure 2



A sample of an RGB orthomosaic taken by the AltumPT camera, which is a multispectral device that captures light beyond the human vision spectrum. The filters on that sensor allow very specific colors—more scientifically, wavelengths—to pass through (which distinguishes this image from the previous one, even though they are both RGB orthomosaics). In addition to RGB colors, the AltumPT sensor can “see” light on the red-edge and low-infrared part of the electromagnetic spectrum, regions of light that are highly reflected by plants. We take advantage of this in the first step of our research analysis. With the AltumPT imagery, we produce an NDVI (normalized difference vegetation index) map that “quantifies vegetation by measuring the difference between near-infrared (which vegetation strongly reflects) and red light (which vegetation absorbs).”^{iv} This highlights the green vegetation, creating a greater contrast with the soil background.


Figure 3

Section of the field showing the corn rows (yellow lines) and weeds growing between the corn rows (green polygons). A weed-only map is obtained by removing the corn plants from the image.

Figure 4

Visualization of the Rx map in the field for one experimental plot. For this research study, we divided the 50-acre field into plots (400 ft long by 136.6 ft wide). In this plot, the red cells indicate areas to be sprayed (15 gallons/acre), while the green ones indicate areas not to be sprayed. The overall average for the SSWC treatment in this study showed that 50 percent of the area did not need to be sprayed.





Yoko "George" Kosugi, an NDSU master's student, flies a drone capturing images of a corn field, which are processed to create an orthomosaic weed map for the herbicide sprayer. Photograph by Jerry Anderson.

of corn (160 acres), a farmer would not need to spray 80 acres. Giving the current prices of chemicals, that could lead to significant savings for corn growers in the state.

Challenges

Although this approach seems to be very promising, there were several challenges. First is scalability. When using a drone, the number of acres that can be covered in a day is limited by the drone's battery life, which dictates flight altitude and flight speed. It took about 30 minutes—the flight-time limit of the drone's battery pack (two batteries)—to cover the 50-acre field with each camera. To cover a section of land (640 acres) would take about seven hours of flight time, which would require 12 more battery sets. Fewer battery sets would be needed—maybe six sets—if the batteries could be recharged in the field. Without the capacity to recharge, which is usually not available, it is difficult, and cost prohibitive for some farmers, to cover large acreage with the technology used on this research project. The total investment was \$49,994: \$13,199 for our drone, \$19,995 for the AltumPT camera and \$16,800 for 12 extra sets of batteries.

The second challenge is the cost of the hardware in the sprayer (including the cab computer), individual nozzle controls and the RTK GPS receiver. These would increase the base cost (around \$550,000) of a sprayer used in this research (Case IH Patriot 4440) by about \$83,000.

A third challenge is developing and automating a workflow that would enable a quick turnaround of the collected data—that is, in less than 24 hours. Working with 50 acres, we collected 100 gigabytes of data with the cameras. We were then able to produce an Rx map within 24 hours. The challenge to make this process viable for farmers will be to process about one terabyte of data from 640 acres and create an Rx map within 24 hours. Clearly, this process would have to become almost 13 times more efficient, which automation could achieve. Automation would also make the process many times more affordable since it currently takes about five man-hours to produce an Rx map for 50 acres.

Near Future

We are looking for ways to improve our workflow by automating parts of our workflow from stitching the orthomosaic image to processing and generating

A sprayer (Case IH Patriot 4440) operating in the field at the NDSU Carrington Research Extension Center. The sprayer is following the rates provided in the prescription map and automatically shutting the nozzles off over the grid cells free of weeds.



an Rx map. To accomplish this, we will continue to collaborate with David Kramar who is now the Research Director at the International Water Institute. His expertise in spatial analysis, image analysis and machine learning applications will be key to enhancing our workflow.

We are also collaborating with Erik Hanson, PhD, an Assistant Professor in NDSU's Department of Agribusiness and Applied Economics. Prof. Hanson will carry out an economic assessment of our SSWC approach, since there are many variables to be considered when trying to make economic sense of this approach. Some of the questions that he will be addressing are:

1. What is the return on investment for farmers implementing our SSWC approach?
2. What are the savings associated with reduced use of herbicide and water (less frequent sprayer refills)?
3. What is the minimum no-spray percentage in a field (20 percent? 30 percent? 40 percent? Or?) needed for a farmer to realize economic benefits when using our SSWC approach?

Post-Harvest Weeds

We realize that an important factor that might affect a farmer's willingness to adopt our SSWC approach is how the post-harvest amount of weeds in the cells, which were sprayed in the spring, compares to the amount of post-harvest weeds in the cells that weren't sprayed. If the sprayed cells had about the same amount of post-harvest weeds as the cells that weren't sprayed, then the savings in herbicide use and financial costs when using our SSWC approach might increase the likelihood of adoption.

To determine this, after the harvest we flew our drone over the same field, collected images with the AltumPT camera and generated a post-harvest orthomosaic map. Our preliminary analysis shows that there were very few weeds across the entire field. Since there was no difference between the sprayed and non-sprayed cells post-harvest, we can conclude that our SSWC approach provides similar weed control as the blanket application, while decreasing the sprayed acreage by half. 📍

ⁱ <https://agriculture.trimble.com/product/weedseeker-2-spot-spray-system/>

ⁱⁱ <https://www.deere.com/en/sprayers/see-spray-select/>

ⁱⁱⁱ <https://www.deere.com/en/sprayers/see-spray-ultimate/>

^{iv} <https://gisgeography.com/ndvi-normalized-difference-vegetation-index/>

SOLVING CRIME THROUGH DIGITAL EVIDENCE



ARICA KULM, PHD

**Director of Digital Forensics Services
Dakota State University**

ON JANUARY 15, 1974, four members of the Otero family were murdered in their home in Wichita, Kansas. Joseph and Julie Otero and their two youngest children, Joseph Jr. (Joey) and Josephine (Josie), ages 9 and 11, were found dead by their oldest son Charlie and siblings Carmen and Danny, when they returned home from school. All four victims had been bound and strangled, and the phone line to their house had been cut.


A few months later, on April 4, Kathryn Bright was stabbed to death in her Wichita home, while her brother was shot twice in the head but fortunately survived. The phone line to the Bright's home had also been cut.

In October 1974, after a confession to the Otero murders was publicized—and later proven false—Wichita police received a phone call directing them to a mechanical engineering textbook on the second floor of the Wichita Public Library across the street from police headquarters. Inside the book was a letter written by someone claiming to be the real killer. The letter contained previously unpublicized details of the Otero deaths, along with the killer identifying himself with the code initials B.T.K., which the self-alleging killer wrote stood for “Bind them, Torture them, Kill them.”

In 1977, there were two more strangulation victims, Shirley Vian and Nancy Fox. Early the following year, another taunt, this time in the form of a poem, was sent to The Wichita Eagle, the city's largest daily newspaper and, a few weeks later, another message arrived at KAKE-TV in Wichita. These prompted the Wichita police to reveal for the first time the presence of “The BTK Strangler,” alerting residents that a serial killer was operating in the area.



Top, Josephine Otero, age 11, was bound with Venetian blind cords and strangled by the BTK killer on January 15, 1974. She was found hanging from a pipe in the basement of her family home.



Above, a 1976 photo of Viki Wegerle, who was strangled with a nylon stocking by the BTK killer in her home on September 16, 1986. She was married and the mother of a two-year-old.

www.kansas.com
Sunday Eagle
 FEBRUARY 27, 2005
 PRICE: \$1.50

'BTK IS ARRESTED'

Park City man held in 17-year killing spree

Police have booked Dennis Rader, 59, on 10 counts of suspicion of first-degree murder

BY ADRIAN LARSEN AND THE AP/WIDE WORLD

A man suspected of being the BTK killer has been arrested in a case that has shocked the community for 17 years. Dennis Rader, 59, was arrested on 10 counts of suspicion of first-degree murder. He was charged with the murders of three women and seven children in the Park City area. Rader was arrested on February 25, 2005, after a 17-year killing spree. He was charged with the murders of three women and seven children in the Park City area. Rader was arrested on February 25, 2005, after a 17-year killing spree. He was charged with the murders of three women and seven children in the Park City area.

SEVEN PAGES OF BTK COVERAGE

WICHITA EAGLE has covered the BTK case for seven pages. The coverage includes the arrest of Dennis Rader, the 17-year killing spree, and the impact on the community. The coverage also includes the police's search for the killer and the public's reaction to the news.

ADRIAN LARSEN is a reporter for the Wichita Eagle. He has covered the BTK case for several years. He has written several articles about the case, including the arrest of Dennis Rader. He has also written about the impact of the case on the community and the police's search for the killer.

THE AP/WIDE WORLD is a news agency that provides news to newspapers around the world. It has covered the BTK case for several years. It has written several articles about the case, including the arrest of Dennis Rader. It has also written about the impact of the case on the community and the police's search for the killer.

ARTS & LEISURE

WHO WILL WIN TONIGHT'S OSCARS?

Front page of the Sunday edition of The Wichita Eagle reporting the arrest of the BTK killer on February 25, 2005.

No more murders were attributed to BTK for more than seven years until the strangulation of Marnie Hedge on April 27, 1985, and then Vicki Wegerle on September 16, 1986. BTK's final known victim was Dolores Davis, who was strangled and dumped near a bridge on January 19, 1991.

With the help of the FBI, Wichita police hunted for BTK, while being taunted with a series of clues left by the killer. Messages in public locations and sent to news media, along with packages with disturbing drawings depicting the murders and dolls posed in the positions of the victims.

After the 1991 murder, BTK seemed to disappear, until a story in The Wichita Eagle in 2004 implied that BTK was a distant memory. Then letters and packages from BTK began arriving again at local newspapers and television stations and even left in a garbage bag in Murdoch Park in Wichita. The letters and packages contained mementos from the victims, as well as drawings and puzzles with clues to the killer's identity.

In late January 2005, an employee at a local store found a cereal box in the back of his truck with a note asking whether BTK could communicate with police via a computer floppy disk without being traced. The police were instructed to run a newspaper ad with the message: "Rex, it will be OK," if this was true.ⁱ

After the police ran the ad, a package arrived at KSAS-TV in Wichita with a clue that ultimately would lead to catching the BTK killer.ⁱ That clue was uncovered with digital forensics, which did not exist during much of BTK's killing spree.

Evolution of Digital Forensics

Digital evidence is associated with approximately 90 percent of crimes committed today. Although the first known computer crimes were in the 1970s, the origins of the relatively new field of digital forensics can be traced to the mid-1980s and early 1990s. BTK killed his last victim in

1991 at the dawn of the digital forensics era. In the early days, with the growing popularity of personal computers, it was known as “computer forensics” and recognized early on by law enforcement as a source of evidence in crimes. Prior to the digital age, data was stored differently. Boxes and filing cabinets of paper and letters became bytes and files on floppy disks, hard drives and servers. The country’s first official digital forensics program was the FBI’s Magnet Media Program launched in 1984.ⁱⁱ This evolved into FBI’s Computer Analysis Response Team (CART).

Evidence found on personal computers evolved into evidence being located on multiple computers of small local networks. The growing popularity of the internet led to looking for evidence in data sent and retrieved over the internet. That, along with the later invention of cellphones and Internet of Things (IoT) devices, caused investigators to look beyond computers to other devices, including video gaming consoles, smart TVs, smart watches and vehicles.

Computer forensics became digital forensics to encompass all the different ways data could be stored or transmitted. Early tools consisted of data recovery software and command line tools (commands that a user types in directly for execution), with the first commercial software being developed in the 1990s and marketed to law enforcement. Forensics teams usually consisted of members of law enforcement officers who were computer hobbyists or had some type of computer background. Those officers tasked with investigating crimes had limited training and no official framework to follow to ensure repeatability in their investigations.

What Is Digital Forensics?

Digital forensics is the identification, preservation, analysis, reporting and presentation of data stored digitally. To reach the last step of presentation, in a criminal case in a court of law, the previous steps should be followed and done in a manner that can be repeated. By starting with proper identification, preserving the data contained on the collected device, analyzing that extracted data, and documenting every person, tool and action that has been carried out, the chain of custody is established. Following these steps

provides assurances that the results of an investigation are acceptable for presentation in court.

Identification

Digital evidence has evolved from floppy disks and hard drives to include a myriad of devices, such as cellphones, drones, GPS, vehicles, internet data, cloud data and other devices that store or transmit data. Identifying devices at a crime scene can be critical to an investigation when a memory card can be as small as a fingernail, or a USB drive can be disguised as a tube of lipstick. Hard drives have been found glued to the underside of tables or hidden in ceiling tiles. By the early 2000s, the everyday use of floppy disks was on the decline when, in February 2005, a package arrived at KSAS-TV in Wichita from BTK containing a floppy disk. That disk would hold the missing clue to BTK’s identity.

Preservation

Digital evidence can be very volatile in nature, and steps must be taken to ensure that it isn’t tampered with. Collecting digital evidence properly can be the difference between having valuable evidentiary data or having no data. Cellphones can be accessed remotely by a suspect to wipe data if they are not collected and stored properly. An unknown password on a phone or encrypted computer can be the difference between accessing the data or having it locked permanently away out of reach of investigators. Every effort is made to create a digital copy of the data prior to analysis. Working from a copy of the data maintains the integrity of the data on the original device, avoiding any question of tampering or altering the evidence. On the occasion that a digital copy cannot be obtained, and the data must be viewed on the device itself, every step must be documented to ensure that the proper steps are followed.

The integrity of data recovered from devices is ensured using hashing. A hash function is a one-way cryptographic algorithm that, when applied to data, produces a unique output. Changing even one byte in the input data results in a different hash value. If thinking of this in terms of a picture, changing one pixel will result in a different hash value than the original. Another example is even adding an extra

space in a document, which can't be easily seen, alters the hash value, as shown in the illustration below, using the SHA256 algorithm.

Analysis

During analysis, the analyst is looking for data that can serve as evidence that supports or refutes the case being investigated. Modern digital forensics tools, which comb through the vast amounts of data that can come during an investigation, make analysis easier than in the past. Each device that is examined and the level of access that can be obtained to the device's data can vary based on the version of the operating system installed, the hardware contained within the device and other factors. The data on a device can include information such as emails, images, internet search history, documents, videos, chat information, and the device might also include deleted or hidden data. On a floppy disk, hard drive, USB stick or other similar storage device, deleted files remain available until overwritten by other data.

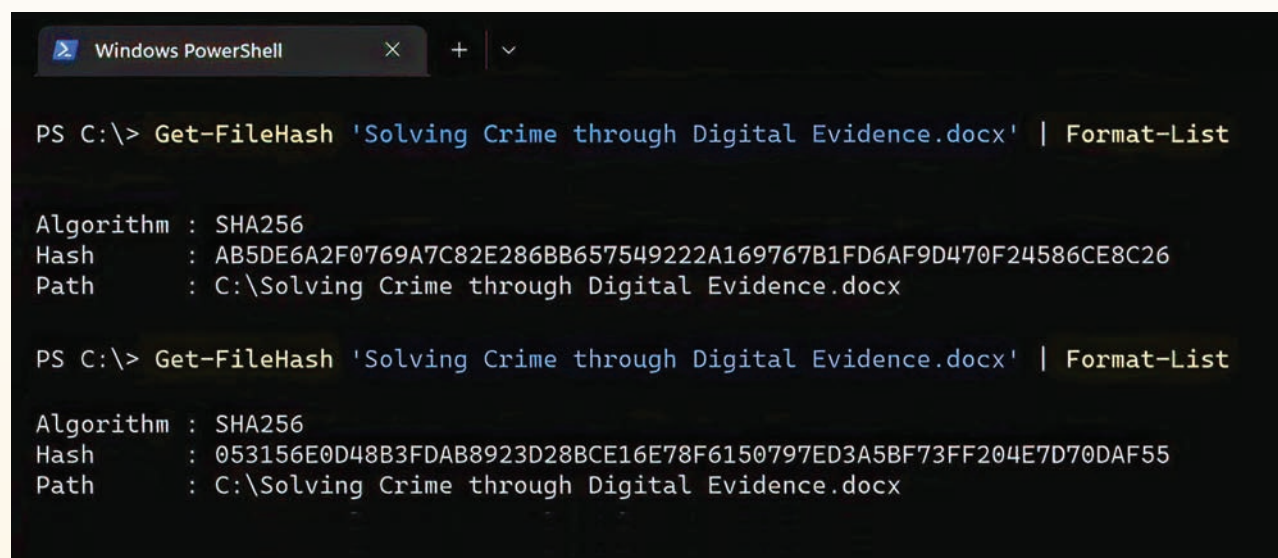
When analyzing that data, the investigator is trying to determine who created the data, what type of device created the data, when the data was created or modified, where the data was created and answers to other pertinent questions. In the BTK case, the floppy disk the killer sent had been used to send a document to the television station; however, it had previously

contained other documents that the killer had deleted. The documents were still available, not having yet been overwritten by other data, and they were able to be recovered by investigators and analyzed.

Reporting

Reporting is an important step in the digital forensics process. Analysts know what they see, but being able to interpret and describe that data in a report, which can be understood by non-experts, is a skill that is developed over time. Reporting must detail the steps the analyst took during the investigation, as well as tell the story of the significance of what was found, or not found, during the analysis of the data. Since reporting aids in the next step of presentation, which may be to a jury during a trial, reports need to be written for any audience regardless of their technical background.

Reporting should tell the story that the device has taken since arrival, detailing the acquisition of the evidence, everything done with the device while it was being examined, what evidentiary findings were located in the data recovered from the device and finally the return of the evidence in order to document the proper chain of custody. Detailed reports that include the proper handling of the evidence being reported on can avoid later accusations of improper handling or tampering of evidence.



```
Windows PowerShell

PS C:\> Get-FileHash 'Solving Crime through Digital Evidence.docx' | Format-List

Algorithm : SHA256
Hash      : AB5DE6A2F0769A7C82E286BB657549222A169767B1FD6AF9D470F24586CE8C26
Path      : C:\Solving Crime through Digital Evidence.docx

PS C:\> Get-FileHash 'Solving Crime through Digital Evidence.docx' | Format-List

Algorithm : SHA256
Hash      : 053156E0D48B3FDAB8923D28BCE16E78F6150797ED3A5BF73FF204E7D70DAF55
Path      : C:\Solving Crime through Digital Evidence.docx
```

Hash value calculation using Windows PowerShell commands. The top is the original Word document of this article. The bottom is the same Word document with one space added to the end of the last paragraph resulting in a completely different hash value.

Presentation

Presenting digital evidence to a jury during a trial can be challenging when compared to presenting physical evidence such as paper documents or photographs. Explaining how data is stored on a device requires someone who can speak to the jury in a way they can understand. It isn't always as simple as handing jury members a document they can read or a photograph that they can see. Lawyers need to understand digital evidence and what is important to their case. The expert witnesses called to testify must be able to talk to the jury about their findings in a way that reflects the importance of the evidence without getting overly technical.

Challenges

There are many challenges in digital forensics, from the continuous development of software and hardware to new encryption techniques. It is a cat-and-mouse game between law enforcement agencies and technology companies to stay one step ahead. Privacy advocates are on the side of end-to-end encryption to keep data private, however that privacy can also hide criminal behavior.

Volume

The sheer volume of devices and data is a challenge for investigators both in storage and analysis. According to Deloitte (a leading global provider of audit, consulting, advisory and tax services), the average household has 22 connected devices including laptops, phones, smart watches, smart home devices, which can store data—sometimes, a lot of data.ⁱⁱⁱ

Modern smart phones can have up to 1 terabyte (TB) of storage. To put that in perspective, it would take 728,177 floppy disks to store 1TB of information.^{iv} Finding a document like the one left by BTK becomes much more difficult when looking through 728,177 items than just 1 item. The number of devices and the vast amounts of data to be analyzed lead to backlogs in casework, as well as stress and burnout for investigators and analysts.

Encryption

Encryption poses a challenge with modern devices containing whole device encryption.^v Encryption uses a mathematical algorithm to scramble readable text so

that only someone with the correct decryption code can understand it. Previously, a hard drive could be removed from a computer or chip removed from a phone, albeit not always easily, and the data read and analyzed. Modern devices that employ encryption techniques on the hardware level rather than file-based encryption present additional challenges. The hard drives or chips can still be removed and read, however the data that is recovered is encrypted, so while the data is still able to be retrieved, it cannot be interpreted due to the encryption. The only way to access the data held within that encrypted device is with the user's passcode. If the passcode is unable to be recovered, from the user or through technological means, that data remains inaccessible.

Emerging Devices & Platforms

Along with encryption, new technology in both new devices and upgrades to existing technology, such as a newer operating system on a computer or cellphone, presents challenges. Devices such as drones, vehicles, Apple AirTags or home security cameras are introduced constantly. Technology companies are continuously upgrading their operating systems and the techniques to keep their devices secure from cyberattacks and other vulnerabilities. Combine that with the multitude of applications available on multiple platforms, and it takes a tremendous effort by the companies creating digital forensics tools to both access the devices and process data from installed applications. The introduction of a new chat application on a smartphone requires software engineers who can then decode that data to understand how it is stored and make it easily readable for forensic analysts or analysts digging deep into the database to find the data and manually linking tables stored within the database to find the stored communications—a time-consuming process.

Burnout & Staff Turnover

Digital forensics analysts are at high risk for burnout. The workload demands, due to the vast amounts of incoming devices and data, along with the material such as documents, images or videos that analysts are often exposed to that can contain violence or child exploitation content, can lead to workplace stress, reduced efficacy, absenteeism, early retirement and burnout.



BTK killer Dennis Rader asked police, "Can I communicate with Floppy [disk] and not be traced to a computer. Be honest." If so, he told police to place an ad in The Wichita Eagle with the message, "Rex, it will be OK." Two weeks later, a disk arrived in a package sent to KSAS-TV with a file that contained instructions to detectives about further communications. What Rader didn't realize was the disk contained metadata that would lead police straight to him.

Nature of Digital Evidence

Digital evidence requires different tools and techniques than physical evidence for identification, collection and analysis. Instead of blood or fingerprint evidence, digital evidence consists of information and data. Both physical and digital evidence can be volatile in that they might not be present for a long time. Blood can wash away, fingerprints can smudge, and digital evidence can disappear. For example, computer memory can vanish if the computer is powered down. However, digital evidence is wider in scope, can be mobile and is much more personal in nature. The National Institute of Justice defines digital evidence as: "Information and data of value to an investigation that is stored on, retrieved or transmitted by an electronic device."^{vi}

Computers and cellphones can hold the most personal material from banking information to conversations to pictures and videos. Pouring through that information is often akin to reading someone's diary or inner-most thoughts. Chat conversations with friends and loved ones, internet searches conducted, locations visited, pictures and videos captured and shared. The metadata found within a document or image can be enough to place someone at a location or identify him or her

as a person of interest and make the difference in an investigation, as in the BTK case.

Metadata is defined as "data about data." Metadata describes data and can provide information about the item of interest (for example, a document, image or video), such as who created it, what type of device created it, when it was created, where it was created. In the BTK case, deleted documents found on the floppy disk sent to the television station contained metadata, which indicated that the disk had been used at Christ Lutheran Church in Wichita by a user named "Dennis." That was the clue police needed that soon led them to the church's council president. After 30 years of searching for the BTK killer, on February 25, 2005, within 10 days of that floppy disk arriving at the television station, BTK was identified, arrested and confessed to the killing of 10 people.

Future

The volume of both devices and data will drive the digital forensics field toward automation. Artificial intelligence (AI) and machine learning (ML) will assist with analyzing and categorizing data faster. Unlike live analysts, AI applications don't show fatigue and are not affected by graphic material. Using automation to

complete some of the time-consuming activities can free up analysts to fine tune and do the detail work. AI and deep-learning algorithms can be used to review things, such as documents, looking for relevant information.

Digital forensics tools have begun to move toward collaborative cloud-based environments allowing more than one analyst to work on the same investigation simultaneously. This can aid in environments where investigators are working on the same case from different locations. These types of tools can also provide analytics to look for trends across different cases. Link analysis can show relationships between data recovered from multiple devices.

Wrap-Up

Digital forensics is a broad field with many specialties. So far this article has addressed digital forensics from the perspective of a law enforcement investigation, however there are many additional aspects of digital forensics. Increasingly, data is being stored in the cloud rather than on servers and workstations. Data travels across networks to get to that cloud storage with the potential for data breaches. Investigating how or where data breaches occur is a branch of digital forensics as well as cloud forensics. In addition to law enforcement, corporations need digital forensic analysts to investigate data breaches, intrusions, insider threats, theft of corporate property and violations of acceptable use policy. The increasingly connected world brings more and more opportunities for digital forensics to play an important role in helping to solve crimes. 📧

ⁱ R.J. Rosen, "The Floppy Did Me In," The Atlantic, 16 January 2014. [Online]. Available: <https://www.theatlantic.com/technology/archive/2014/01/the-floppy-did-me-in/283132/>. [Accessed 27 November 2022].

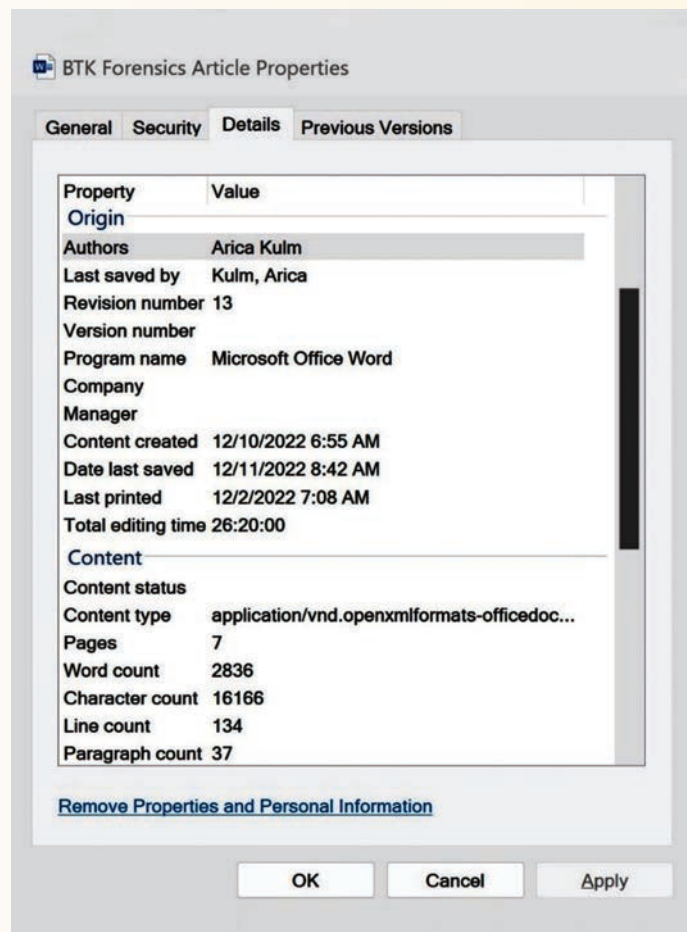
ⁱⁱ S. Moore, "What is Digital Forensics?," [Online]. Available: <https://www.azolifesciences.com/article/What-is-Digital-Forensics.aspx>.

ⁱⁱⁱ J. Arbanas, J. Loucks and S. Hupfer, "Connectivity and Mobile Trends Survey: Mastering the New Digital Life," The Wall Street Journal, 20 September 2022. [Online]. Available: <https://deloitte.wsj.com/articles/connectivity-and-mobile-trends-survey-mastering-the-new-digital-life-01663699532>. [Accessed 27 November 2022].

^{iv} T. Fisher, "Terabytes, Gigabytes, & Petabytes: How Big Are They?," 1 January 2021. [Online]. Available: <https://www.lifewire.com/terabytes-gigabytes-amp-petabytes-how-big-are-they-4125169>. [Accessed 27 November 2022].

^v For a primer in encryption, please refer to: Marcus Fries, PhD, "The Unencrypted History of Cryptography," Dakota Digital Review, North Dakota University System, Fall-Winter 2023, <https://dda.ndus.edu/ddreview/the-unencrypted-history-of-cryptography/>

^{vi} M. B. Mukasey, J. L. Sedgwick and D. W. Hagy, "Electronic Crime Scene Investigation: A Guide for First Responders, Second Edition," U.S. Department of Justice, Washington, DC, 2008.



Metadata describes data and can provide information about the item of interest (document, image, video), such as who created it, what type of device created it, when it was created, where it was created and more.


Promise & Peril of Intelligent Machines & Cyberspace

How Rural Universities Are Teaming Together

MARK R. HAGEROTT, PHD
Chancellor, North Dakota University System

It is impossible to predict the future with any certainty, but evidence is growing that our economy, society and higher education need to rapidly confront the effects of digitization. The emergence of artificial intelligence (AI) affecting human, machine and data systems is perhaps the latest driver of digital disruption. The widespread reporting on ChatGPT is just the most recent example of the monumental challenge now bursting upon our state and our nation.ⁱ

Illustration from the 1979 children's book
Your Name? Robot by Mikhail Romadin.



Historically, state systems of higher education have proven to be a key tool in helping states and their citizens navigate and even thrive through disruptive change, as Land Grant and A&M campuses did in the Industrial Age.

In today's age of accelerating digital change, educational institutions are being called upon to help our nation and our people adapt. Examples of such requests include Congressional and White House direction for higher education to help build a new national DOD Cyber and Digital Service Academy; Congressional calls and funding for a revamped research, workforce and research agenda (the Chips and Science Bill); and many states and the U.S. Department of Labor calling upon state higher education systems to assist with workforce retraining and upskilling in the cyber-digital realm.ⁱⁱ

But the leaders of the North Dakota University System and constituent colleges and universities realize the days of “go it alone” by single campuses, or even single state systems, won't get us to scale at the needed speed. Rather, collaborative problem-solving and reorientation to tackle digital transformation across all our state educational institutions and those of our neighboring states are required.

To that end, the University Alliance (full moniker: the Mountain and Plains University Innovation Alliance)—a consortium of North Dakota, South Dakota, Montana, Wyoming and Idaho university systems—was established in the fall of 2022.ⁱⁱⁱ Why was this needed and how did it come about?

Emerging Technology's Opportunity & Risk for Rural America

Society is now experiencing AI's emerging nexus in digital machines and a metaverse of data (sometimes referred to as cyberspace). As the digital world grows, financial benefits accrue to only a small portion of the population, with large swaths of the country left behind or left insecure. Almost 80 percent of venture capital is concentrated in a handful of coastal states, while North Dakota and our rural neighbors are among the bottom five states. What does this mean?

Large tech hubs—such as Silicon Valley, the Boston-NYC corridor and DC-MD-Northern Virginia (home of the new Amazon Campus)—are growing wealthier and providing more opportunity to their people and their children.

In contrast, where cyberspace and digital machines intersect with regions, cities and households outside the tech hubs, signs of distress are mounting. Darkening clouds of hacking and privacy abuses, misinformation and disinformation and lack of equal access cast a shadow of social, employment and political insecurity. Also alarming, we are seeing a rapid increase in child depression and suicide, a major shift that scientists increasingly associate with the widespread penetration of society by the highly advanced communication and computational device, the smart phone and the growth of social media sites.^{iv} While the elites in the tech hubs are typically immune to these damaging effects, the working people even in these tech hubs often struggle, manifest by the stoning of Google buses in San Francisco.

Lastly, human-centric cultural values, which attach to humans and human life, are being affected by the emergence of intelligent machines. The massive concentration of digital machines, algorithms, human programmers and executives in the tech hubs combine to exert a profound social and political influence across the nation. Meanwhile, rural America is seldom at the table to discuss the future of humanity in a digitizing world. If we doubt the significance of the social issues surrounding the emergence of advanced technological systems, guided by an ever smaller and more elite slice of American society, consider that the White House recently called for a new Digital Bill of Rights for AI.^v

Accordingly, North Dakotans and rural Americans need to participate in creating and shaping, as well as benefiting from, emerging technologies. Higher

education is a key tool to create, shape and reap benefits from emerging technology. But there is a problem of scale.

Higher Education as Key to the Future

The likelihood that our technical elites in Silicon Valley can tackle the challenge of digital transformation and make a better world for rural Americans is doubtful, in part due to a lack of trust from the public. Amazon was even driven out of high-tech New York City. Facebook/Meta paid a \$5-billion fine for privacy violations, and many in Congress and the states argue that massive digital and social media companies (e.g., Meta, Amazon, Twitter, Google) should, along with other social media companies, be more strictly regulated.^{vi} Who can help navigate the road ahead?

Educational institutions will be at the center of such efforts, as they were in the last great transformation, the Industrial Revolution and mechanization of agriculture. In the 1860s, President Abraham Lincoln created the land-grant university system to support the nation's agriculture and industry in developing and adapting to new science and technology. The accomplishments of higher education in the Industrial Age are almost too numerous to list but include the creation of the modern engineering profession, modern agronomy, and the emergent fields of aerospace engineering and computer science.

But today, national educational research activity is tilted toward the mega-university, which is not inclusive of and is in fact predatory upon researchers and the professoriate in rural America. In order to create a more geographically inclusive future, to allow all Americans to reap the benefits of advanced technology, the rural states and their universities have had to find a way to join more fully in the creation of this future economy.

The North Dakota University System tackled this problem of achieving scale in rural areas early. In 2018, North Dakota leaders proposed a new Digital Land Grant system, a concept that was published by the nation's leading educational periodicals.^{vii} That proposal would send funding to the states and campuses on the condition they prioritize responding to the challenge of digitization, including workforce cyber training as well as furthering research, not just in the hard sciences and engineering, but also in the realm of digital policy, law and ethics.

This proposal led to discussions at the White House Office of Science and Technology Policy, on two occasions, and with Congressional staff. Interest and support was high, including from education advisors (Gordon Gee, then President of the University of West Virginia) who was working with Rep. Ro Khanna, of Silicon Valley, to learn more about North Dakota's new educational model.^{viii} But with the change in federal administration and control of both Houses of Congress in 2020, the political momentum shifted from a new land grant system controlled at the state level to an expansion of federal funding directed out of the National Science Foundation and the U.S. Department of Commerce. This was embedded in the Endless of Frontier Bill, crafted by Rep. Khanna in the House and advocated by Senator Schumer in the Senate, and was eventually incorporated into the U.S. Innovation and Competition Act of 2021.^{ix}

Though our preference may have been a Digital-Cyber Land Grant system controlled by respective states, it is urgent to help rural America now. And what is reality now? The passage of the massive CHIPS and Science Act, which grew out of the Endless Frontier Bill, may exceed \$200 billion and stands as a testament of the scale needed to effect a national transformation and adapt to accelerating digitization and other related emerging technologies.

However, North Dakota or any of our rural neighbor states is too small to compete successfully alone. All our universities and colleges have a problem of scale, even if statewide collaboration is seamless. To achieve that competitive scale, the North Dakota University System partnered with our neighboring university systems and campuses in December 2020 to form the five-state University Alliance, which was codified in 2022.^x

Achieving Scale with the University Alliance

The University Alliance creates a network of five state systems using advanced technology to conduct research and deliver both technical and social science/humanities education on campuses located across the region. The University Alliance will move beyond the limits of city geographies of large urban areas and knit together rural populations and tribal nations that otherwise would be left behind. As well, the University Alliance will leverage existing research, education and training programs to expand the region's ability to support the expansion of high-tech industries. States, universities and tribal colleges that join this transformative effort will bring the best of both online and brick-and-mortar education, which offer expanded technical and nontechnical curricula to help our campuses, systems and states respond to emerging technologies.

Forward

The high-tech future economy is rapidly growing at the nexus of human researchers, innovators and workforce. All Americans, including those in the rural northern plains and mountain states, should reap the harvest of wealth and opportunity afforded by these new technologies. State university systems must help their respective stakeholders benefit from emerging technology: their residents, business communities, and municipal and state governments.

Taken together, this is a tall order beyond the resources and span of control of any individual campus or state system. It is why our five states have formed this partnership of independent state systems into a collaborative 'system of systems' to achieve scale at speed, with a wide scope that reaches rural and tribal America. With the University Alliance, more rural

states, cities and towns can access the knowledge and resources they need to weave a better future in the emerging digital socioeconomic system.

Having such a regional system of state systems will allow the nation to achieve a scale of response that is more cost effective; expands the scope to more underserved populations and regions; and achieves both of these transformative goals more rapidly. The urgency of the challenge is hard to exaggerate. Just as the land-grant system transformed higher education to catalyze agricultural and industrial capacity across a growing nation in the 19th century, so can a regional alliance of rural states provide the educational foundation needed to ensure economic and democratic vitality and security for the entire nation, both urban and rural, in the 21st century. ▣

ⁱ Sovov, Vlad, "CHAT GPT Could be AI's iPhone Moment," *Bloomberg*, 12 December 2022, link: <https://www.bloomberg.com/news/newsletters/2022-12-12/chatgpt-the-gpt-3-chatbot-from-openai-microsoft-is-tech-magic>; Russel, Stuart, *Machine Compatible: Artificial Intelligence and the Problem of Control*, New York: Viking, 2019; Eric Schmidt, Henry Kissinger, Daniel Huttenlocher, *The Age of AI and our Human Future*, New York: Little, Brown and Company, 2021.

ⁱⁱ Heckman, Jory, "AI commission sees 'extraordinary' support to stand up tech-focused service academy," Federal News Network, 2 March 2021, link: <https://federalnewsnetwork.com/artificial-intelligence/2021/03/ai-commission-sees-extraordinary-support-to-stand-up-tech-focused-service-academy/>

ⁱⁱⁱ Mountains and Plains University Innovation Alliance members: Idaho (Boise State University, Idaho State University, University of Idaho), Montana (Montana State University, Montana Technological University, University of Montana), North Dakota (North Dakota State University, University of North Dakota), South Dakota (Dakota State University, South Dakota School of Mines and Technology, South Dakota State University, University of South Dakota) and Wyoming (University of Wyoming).

^{iv} Twenge, Jean M., *iGEN: Why Today's Super-Connected Kids Are Growing Up Less Rebellious, More Tolerant, Less Happy—and Completely Unprepared for Adulthood (and What This Means for the Rest of Us)*, New York: Atria Books, 2017.

^v White House calls for a new Bill for Rights in the Age of AI. Announcement link: <https://www.whitehouse.gov/ostp/news-updates/2021/11/10/join-the-effort-to-create-a-bill-of-rights-for-an-automated-society/>

^{vi} Congressional leaders include those from both side of the aisle, Senator Amy Klobuchar (D), Senator Josh Hawley (R), President Biden and former President Donald Trump. Multiple states have also enacted or are wrestling with this issue, as well, and the State of Washington enacted a tax on digital and social media companies to better fund its K-12 and higher education systems.

^{vii} Hagerott, Mark, "Silicon Valley Must Help Rural America: Here's How," *The Chronicle of Higher Education*, 23 September 2018, accessed here: <https://www.chronicle.com/article/silicon-valley-must-help-rural-america-heres-how/>; Hagerott, Mark, "Time for a Digital-Cyber Land Grant System," *Issues in Science and Technology*, 36, no. 2 (Winter 2020): 23–26.

^{viii} Letter from Gordon Gee, President of University of West Virginia, to Mark Hagerott, Chancellor of the North Dakota University System, December 5, 2018.

^{ix} <https://www.congress.gov/bill/116th-congress/senate-bill/3832/text>

^x Several press releases document the formation of the Alliance, link here to an NDSU press release: <https://www.ndsu.edu/news/view/detail/71110/>; also, an affiliate of National Public Radio (Prairie Public) also reported on the Alliance, link here: <https://news.prairiepublic.org/local-news/2022-12-27/ndsu-und-nd-university-system-part-of-a-regional-technology-and-innovation-alliance>

Mountains & Plains University Innovation Alliance

Idaho

Boise State University

Idaho State University

The University of Idaho

Montana

Montana State University

Montana Technological University

The University of Montana

North Dakota

North Dakota State University

The University of North Dakota

South Dakota

Dakota State University

South Dakota School of Mines
& Technology

South Dakota State University

University of South Dakota

Wyoming

The University of Wyoming

Five-State Innovative University Alliance— Features & Possible Futures:

Universities and colleges in this new effort can better reach scale and thus offer more competitive pay and better-staffed research facilities to recruit and promote tech innovators, especially cyber-computer science faculty, including those in ethics, law and the humanities, who focus on the digital transformation of our society.

New funding mechanisms, such as the Chips and Science Act, relieve the burden on hard-pressed rural universities and students. Hopefully, there will be some additional help from the respective state governments and legislative supporters.

To accelerate adaptation, respective legislators might provide additional state support, such as potential state tax incentives for contributions to university endowments that promote innovation in the emerging fields of research, economic development and workforce preparation.

The University Alliance, now invigorated with state support and federal monies, would provide an especially welcoming environment for regional and national technology companies to partner. Similarly, having achieved scale, this rural University Alliance would be better able to attract other research partners, including leading high-tech universities, with mutual interests or the desire to test technologies in our test centers and open skies.

Given that the five-state region is home to almost half of the nation's tribal colleges, a major focus of the University Alliance is encouraging investments in traditionally underserved peoples and places. By focusing on expanded partnerships with tribal colleges and nations, the University Alliance can help bridge the gap between emerging tech cultures found typically on the coasts with traditional native cultures.

Areas of research and innovation will include, as a minimum, autonomous systems, digital sciences, energy and advanced materials science, quantum computing, forest and rangeland management, cybersecurity, predictive and precision agriculture, and the social sciences and workforce programs where they intersect the emergence of new technology.



AI SECURITY

The National Security Commission
on Artificial Intelligence and
Adversarial Machine Learning

RAM SHANKAR SIVA KUMAR
and HYRUM ANDERSON, PHD

The National Security Commission on Artificial Intelligence (NSCAI) was established by Congress in 2018 to “consider the methods and means necessary to advance the development of artificial intelligence ... to comprehensively address the national security and defense needs of the United States.”ⁱ And this it did. More than 19 of the recommendations from NSCAI’s “The Final Report” were included in the FY2021 National Defense Authorization Act, with dozens of other recommendations influencing Acts of Congress and executive orders related to national defense, intelligence, innovation and competition.

Although the final report was released in March 2021—and the commission sunset in October that year—one element of the report is still widely underappreciated, especially outside of government circles. While the Department of Defense (DOD), as well as other government agencies ranging from the Department of Veteran Affairsⁱⁱ to the Internal Revenue Service,ⁱⁱⁱ have responded to the charge to *adopt* AI, very few government and commercial entities have invested in the calls to *secure* AI.

Organizations that adopt AI also adopt AI’s risks and vulnerabilities. Intentional attacks against AI are a nascent style of cyberattack by which an attacker can manipulate an AI system. So, while the U.S. does need more investment in AI, it also needs to decisively secure it from its adversaries.

In response, we wrote a book, *Not With a Bug, But With a Sticker: Attacks on Machine Learning Systems and What To Do About Them*, which will be published by John Wiley & Sons, Inc., on May 2, 2023. The book explains how AI systems are significantly at risk from attacks in both simple and sophisticated ways, which can jeopardize national security, as well as corporate and

industry security. The purpose of the book is to provide decision-makers with context that will sharpen their critique as they embrace the power of AI in government and industry. We also offer recommendations on how the intersection of technology, policy and law can provide us with a secure future.

Below is an excerpt from the book’s first chapter, “Do You Want to Be a Part of the Future?”

NSCAI Genesis Story

Ylli Bajraktari is not a household name, but in national security circles, he has a reputation for getting things done. Andrew Exum, a former Deputy




***Not with a Bug, But with a Sticker: Attacks on Machine Learning Systems and What To Do About Them*, by Ram Shankar Siva Kumar and Hyrum Anderson, PhD, John Wiley & Sons, Inc., 2023.**



Assistant Secretary of Defense for the Middle East, described Bajraktari and his brother Ylber^{iv} as “two of the most important and best people in the federal government you’ve likely never heard of.”^v Ylli Bajraktari escaped war-torn Kosovo and moved to the U.S. in his 20s. Burnished with *bonafide* credentials from Harvard’s prestigious Kennedy School, he steadily rose through the ranks at the Pentagon, eventually becoming an advisor to the Deputy Secretary of Defense.^{vi}

Whether it was fatigue that came from years of traveling the world to shape international policy or the pressure of working at the White House,^{vii} Bajraktari left the executive branch to join the National Defense University’s (NDU) Institute for National Strategic Studies as a visiting research fellow. A year-long assignment from the Office of the Secretary of Defense, this break from the intensity in the White House gave him time to study and recalibrate at the NDU’s libraries. He used some rare downtime for self-education on what he thought would be an instrumental cornerstone of U.S. competitiveness: artificial intelligence (AI). Bajraktari quickly absorbed the information he gleaned from pouring over books and watching YouTube videos on machine learning (ML). While at NDU, he organized the university’s first ever AI symposium in November 2018. Expecting a meager 10 people to attend, the response was overwhelming. Bajraktari had to turn away hundreds of would-be attenders because of the room’s fire code. While Bajraktari didn’t know it yet, his time at DOD, the White House and now at NDU



Left: Researchers asked an AI computer vision model, which typically performs well on image classification tasks, “What is pictured in this image?” To confuse the AI vision model, the researchers created images to maximize “penguin-ness,” “snakeness,” and “schoolbus-ness.” As a result, the AI vision model was 99 percent confident that the top image represents a penguin, the middle image represents a green snake, and the bottom image represents a school bus. Courtesy of Ann Nguyen.

Right: AI can be fooled by carefully constructed patterns. What looks like a sweater designed around psychedelic, multicolor swirls is actually a carefully constructed pattern to fool object-recognition systems into mispredicting that the wearer does not exist. However, the deceptive traits of this sweater work only in a controlled setting and are generally ineffective to avoid surveillance in real-world settings. Courtesy of Zuxuan Wu and Tom Goldstein.



had been preparing him for a leadership role to shape the nation's strategy for investing in AI.

That came in 2018, when the NSCAI was born out of the House Armed Services Committee. The leadership and guidance for the commission was to come from 15 appointed commissioners, a mix of tech glitterati that included current and former leadership at Google, Microsoft, Amazon, Oracle, and directors of laboratories at universities and research institutes that support national security.^{viii} It was an independent, temporary and bipartisan commission set up to study AI's national security implications.


When NSCAI commissioner and former Google CEO, Eric Schmidt, called Bajraktari to ask him to lead the commission, Bajraktari didn't answer the phone. It was Christmas break. Plus, he had been primed by his White House days to ignore calls from unknown numbers. So eventually, the former CEO of Google resorted to the plebian tactic: He sent Bajraktari an email. In it, Schmidt described that he had just been voted by the other commissioners to chair NSCAI and needed somebody to run the commission's day-to-day operations.

Bajraktari's email response was a simple one-liner: "I'll do it."^{ix}

With an ambitious goal, a tight deadline and a budget smaller than what it takes to air a 60-second Superbowl commercial, Bajraktari assembled a team of over 130 staff members to deliver a series of reports that would culminate in NSCAI's final report.

Even the initial findings packed a punch. Bajraktari and the two chairs of the commission headed to the White House to brief President Trump about the findings. Scheduled for 15 minutes, the meeting at the Oval Office lasted for nearly an hour. In December 2020, at the twilight of his presidency, President Trump signed an executive order entitled "Promoting the Use of Trustworthy Artificial Intelligence in Government."^x

Even with a change in the executive office, NSCAI's recommendations continued to make an impact. On July 13, 2021—well into the Biden presidency—the commission held a summit to discuss the final report at the Mayflower Hotel Ballroom in Washington, D.C. The Secretaries of Defense, Commerce and State, the National Security Advisor, and the Director of the Office of Science and Technology Policy all



made in-person appearances and spoke to the masked and socially distanced audience members. It was a powerful signal from the American government to both its allies and adversaries that the U.S., from its highest levels of trade, diplomacy, defense, security, and science and technology, was ready to invest in AI and take the NSCAI's recommendation seriously.

Adversarial ML: Attacks Against AI

Delivered on March 2021 to the president and Congress, the first page of NSCAI's final report includes a realistic assessment of where the country stands, which by the commission's own admission was uncomfortable to deliver. The NSCAI report was 756 pages long, but its opening lines summarize our unpreparedness. "America is not prepared to defend or compete in the AI era," it reads. "This is the tough reality we must face. And it is this reality that demands comprehensive, whole-of-nation action." To "defend ... in the AI era" certainly refers to holistic national security but also entails defending vulnerabilities



To demonstrate the limitation his custom-built autonomous vehicle vision system, artist James Bridle drew two concentric circles with table salt sprinkled on the pavement: The inner circle was a solid line, and the outer circle was a dashed line. The car quite readily drove into the circle but was unable to drive out, mistaking the implausible lane markings as valid traffic indicators. Source: Autonomous Trap 001 (James Bridle, 2017), courtesy of the artist.

in AI defenses that are spelled out in numerous recommendations to ensure that "models are resilient to ... attempts to undermine AI-enabled systems."^{xi}



But defend from *whom*? Defend from *what*?

Adversarial ML is sometimes called "counter AI" in military circles. Distinct from the case where AI is used to empower an attacker (often called *offensive AI*), in adversarial ML, the adversary is *targeting* an AI system as part of an attack. In this kind of cyberattack against AI, attackers actively subvert vulnerabilities in the ML system to accomplish their goal.

What are the vulnerabilities? What kinds of attacks are possible? Consider a few scenarios:

- Cyber threat actors insert out-of-place words into a malicious computer script that causes AI-based malware scanners to misidentify it as safe to run.
- A fraudster issues payment using a personal check that appears to be written for \$900, but the victim's automated bank teller recognizes and pays out only \$100.
- An eavesdropper issues a sequence of carefully crafted queries to an AI medical diagnosis assistant in order to reconstruct private information about a patient that was discarded after training—but inadvertently memorized by—the AI system.
- An adversary corrupts public data on the internet used to train a facial recognition biometric authentication tool so that anyone wearing a panda sticker on their forehead is granted system access.
- A corporation invests millions of dollars to develop proprietary AI technology, but a competitor replicates it for only \$2,000 by recording responses of the webservice to carefully crafted queries.

All of these scenarios—or scenarios very much like them—have been demonstrated (in some cases, by the authors) against high-end deployed, commercial models. They are all a form of adversarial ML, which is not just subversive but also subterranean in our discourse.



AI systems can produce nonsensical results in novel or foreign contexts. In this image, the boxes show how the AI system recognizes objects. But, when a polar bear is inserted into the image, the AI system becomes confuse and falsely detects a car. Courtesy of Amir Rosenfeld.

Chances are you have heard more about deepfakes (a form of offensive AI) than adversarial ML. But adversarial ML attacks are an older, but still pernicious, threat that has become more serious as governments and businesses adopt AI.

The NSCAI report wrote unequivocally about this point: “The threat is not hypothetical ... adversarial attacks are happening and already impacting commercial ML systems.”

ML Systems Don't Wobble, They Fold

To understand adversarial ML, we first need to understand how AI systems fail.

An *unintentional failure* is the failure of an ML system with no deliberate provocation. This happens when a system produces a formally correct but often nonsensical outcome. Put differently, in unintentional failure modes, the system fails because of its inherent

weirdness. In these cases, anomalous behavior often manifests itself as earnest but awkward “Amelia Bedelia” adherence to its designers’ objectives. For instance, an algorithm that was trained to play Tetris learned how to pause the game indefinitely to avoid losing.^{xii} The learning algorithm penalized losing, so the AI did whatever was in its power to avoid that scenario. Scenarios like this are like the Ig Nobel Prize—where it first makes you laugh and then makes you think.

But not all cases are humorous.

The U.S. Air Force trained an experimental ML system to detect surface-to-surface missiles. At first, the system demonstrated impressive 90 percent accuracy. But instead of getting a game-changing target recognition system, the Air Force learned a sobering lesson during field testing. “The algorithm did not perform well. It actually was accurate maybe about 25 percent of the time,” an Air Force official remarked.^{xiii} It turns out that the ML system was trained to detect missiles that were only flying at an oblique angle. The accuracy of the system plummeted when the system was tested on vertically oriented missiles. Fortunately, this system was never deployed.

Unintentional failure modes happen in ML systems without any provocation.

Conversely, *intentional failure* modes feature an active adversary who deliberately causes the ML system to fail. It should come as no surprise that machines can be intentionally forced to make errors. Intentional failure modes are particularly relevant when one considers an adversary who gains from a system’s failure either the hidden secrets in training data memorized by the system or the intellectual property that enables the AI system to work. This branch of failures in AI systems is now generally called “adversarial machine learning.” Research in this has roots in the 1990s when considering maliciously tampered training sets and, in the 2000s, with early attempts to evade AI-powered email spam filters.

But adversary capabilities exist on a spectrum. Many require sophisticated knowledge of AI systems to pull off attacks. But one need not always be a math whiz to attack an ML system. Nor does one need to wear the canonical hacker hoodie sitting in a dark room in front

of glowing screens. These systems can be intentionally duped by actors of varying levels of sophistication.

The word “adversary” in adversarial machine learning instead refers to its original meaning in Latin, *adversus*, which literally means someone who “turns against”—in this case, the assumptions and purposes of the AI system’s original designers. When ML systems are built, designers make certain assumptions about the place and manner of the system’s operation. Anyone who opposes these assumptions or challenges the norms upon which the ML model is built is, by definition, an adversary.

Take the event held by the Algorithmic Justice League, a digital advocacy nonprofit founded by Joy Buolamwini, as an example. In 2021, the nonprofit held a workshop called “Drag Vs AI” in which participants painted their faces with makeup to fool a facial recognition system.^{xiv} When facial recognition systems are built, they are relatively insensitive to faces with “regular” amounts of makeup applied. But when one wears over-the-top, exaggerated makeup, it can cause the facial recognition system to misrecognize the individual. In this case, participants have upended the conditions and assumptions on which the model has been trained and have become its adversary.

Text-based systems are equally fallible. It was not uncommon in the early days of AI-based resume screening for job posters to pad their resumes with keywords relevant to the jobs they were seeking, colloquially called *keywords stuffing*. The rationale was that automated resume screeners were specifically looking for certain skills and keywords. The prevailing wisdom of keyword stuffers was to add the keywords on the resume in white font, invisible to human screeners but picked up by keyword scanners, to tilt the system in your favor. So, if an ML system is more likely to select an Ivy league grad, one may simply insert “Harvard” in white font—invisible to the human reader but triggering to the system—in the margin to coerce the system to promote the resume. Subverting the system’s normal usage in this way would technically make one an adversary. Typos make a difference as well. “Remove all buzzwords. Misspell them or put spaces between them”—that was the direction from a group using Facebook to promote

ivermectin in a way that would escape Facebook's AI spam filters.^{xv} When the group found that the word ivermectin triggered Facebook's content moderation system, they resorted to simply "ivm" or used alternate words such as "Moo juice" and "horse paste."^{xvi}

Sometimes, adversaries can collectively refer to more than one person. In 2016, Microsoft released Tay, a Twitter bot that was supposed to emulate the personality of a teenager. Its purpose was to allow users to tweet at Tay to engage in a conversation with the bot as a playful publicity stunt. The ML system would parse the tweet as input and respond. Key to this system was that Microsoft Tay continually trained on new tweets "online" to improve its conversational ability. To prevent the bot from being misled or corrupted by conversations, Microsoft researchers taught it to ignore problematic conversations, but only from individual dialogues.

And this is where things began to turn for the chatbot. In a matter of hours, Tay went from a sweet 16-year-old personality to an evidently Hitler-loving, misogynistic, bigoted bot. Pranksters from Reddit and 4Chan had self-organized and descended on Twitter with the aim of corrupting Tay. Why? For fun, of course. They quickly discovered that Tay was referencing language from previous Twitter conversations, which could have a causative effect on Tay's statements. So, the trolls flooded Tay with racist tweets. Overwhelmed by the variety and volume of inappropriate conversations, Tay was automatically retrained to mirror the internet trolls, tweeting, "Hitler was right I hate the jews." With the company image to consider, Microsoft decommissioned Tay within 16 hours of launching it.^{xvii}

Microsoft had devised a plan to deal with corrupt conversations by a few individuals. But Microsoft was blind-sided by this coordinated attack. This group of internet strangers became an adversary. This kind

According to AI vision systems, these patterns are recognized, respectively from top to bottom, as a jeep, a goldfish and a washing machine. These images were curated by researchers to demonstrate the limitations of the current state of computer vision. They assembled a selection of naturally confounding images to demonstrate that otherwise award-winning computer vision systems failed on some photographed images. Error rates skyrocketed to 98 percent for these naturally occurring adversarial examples. Courtesy of Dan Hendryks.



of coordinated poisoning attack—corrupting an AI system by corrupting the training data it ingests—is one of the most feared attacks by organizations, according to our survey in 2020.^{xviii}

Then in 2022, Meta, Facebook's parent company, released an experimental chatbot called BlenderBot 3, which was "capable of searching the internet to chat about virtually any topic ... through natural conversations and feedback from people." Before too long users found that the bot began parroting election conspiracies that Trump was still president after losing the election and "always will be."^{xix} It became overtly antisemitic, saying that a Jewish cabal controlling the economy was "not implausible" and that they were "overrepresented among America's super-rich."

“Adversarial attacks are happening and already impacting commercial ML systems,” warned the NSCAI report. As with traditional cyberattacks, the economical inevitability of that statement stems from two conditions: that the odds of discovering a vulnerability in an AI system is high and that there are motivated adversaries willing to exploit it.

Never Tell Me the Odds

When the NSCAI report was published, Jane Pinelis, PhD, was vindicated.

Pinelis had been leading the DOD's Joint Artificial Intelligence Center that was responsible for testing AI systems for failures. She knew intimately how brittle these systems are and had been trying to convince the Pentagon to take up the issue of defending AI systems more seriously.


So, when NSCAI sounded the alarm about AI's dire straits and the implications for national security in a series of interim reports in 2019 and 2020, the issue gained center attention. More importantly, the NSCAI report convinced Congress to allocate money so that experts such as Pinelis were better resourced to tackle this area. For 2021, Congress authorized \$740.5 million^{xx} for a vast number of national defense spending programs to modernize the U.S. military. One key element of that initiative focused on Trustworthy AI. Today, Pinelis is the Chief of AI Assurance at DOD,^{xxi} where her work revolves around justified confidence in AI systems to ensure that they work as intended, even in the presence of an adversary.

Pinelis prefers “justified confidence” in favor of “trustworthiness,” because trust is something that is difficult to measure. Confidence on the other hand is a more mathematically tractable concept. Las Vegas Betting odds can establish reasonable odds of winning a boxing match. A meteorologist can estimate the odds of it raining tomorrow.

So, roughly what are the odds of an attacker succeeding at an attack against an ML system?

For this, we turned to University of Virginia Professor David Evans who specializes in computer security. Evans first considered the possibility of hacking AI systems when one of his graduate students began experiments that systematically evaded ML models. When he began looking more into attacking AI systems, what struck him was the lax security relative to other computer systems.

Odds of Success at Breaking the System	
Modern day cryptography system	1 in 1,000,000,000,000,000,000,000,000,000,000,000,000,000,000 (1 in one duodecillion)
Modern operating system	1 in 400,000,000 (1 in 400 million)
Modern machine learning system	1 in 2



bird



airplane



dog



frog

Encrypted forms of communication—for example, as used in Facebook Messenger or online banking—are built upon methods designed to provide strong encryption. These encryption schemes would be considered totally broken if there were any way to guess the secret key more efficiently than by just trying every possible combination of keys. How hard is that? The odds of compromising modern day encryption by brute force is 1 in 10 followed by 39 zeros. If a more efficient method were discovered, the encryption scheme would be considered broken and unusable for any system.

When designing operating systems that power everything from your laptop to your phone, Evans pointed out that for security protection to be considered acceptable, the odds of an attack succeeding against it should be less than 1 in 400 million.

In both scenarios, “justified confidence” in the security of these systems comes from a combination of analyses by experts, careful testing and the underlying fundamentals of mathematics. Although cybersecurity breaches are apparently becoming more prevalent, computing is more secure than it has ever been. We are in the Golden Age of secure computing.

But when it comes to the security of AI, we are currently in the Stone Age. It is comically trivial to attack AI systems. We already saw how internet trolls can do it. But the more we dial the skill level up, the stealthier the attack gets.

Modern ML systems are so fragile that even systems that are built using today’s state-of-the-art techniques to make them robust can *still* be broken by an adversary with little effort, succeeding in roughly half of all attempted attacks. Is our tolerance

Researchers found that simply changing the hue (color) and saturation (color intensity) of images caused a record-holding AI vision system to misrecognize objects 94 percent of the time. In this example, the same image of a bird is recognized as an airplane, a dog and a frog when hue and saturation are adjusted. Courtesy of Hossein Hosseini.

for AI robustness really 200 million times less than our tolerance for operating system robustness? Indeed, today’s ML systems are simply not built with the same security reliability as an operating system or the cryptography we expect for online chat apps such as WhatsApp or Facebook. Should an attacker choose to exploit them, most AI systems are sitting ducks.

And, indeed, there are motivated adversaries who might wish to exploit AI’s vulnerabilities.

AI’s Achilles Heel

In his confirmation hearings for Secretary of Defense, U.S. Army (Ret) General Lloyd J. Austin III, called China a “pacing threat,”^{xxii} adding that China “presents the most significant threat going forward because China is ascending.”

The clear stance of the NSCAI report is that at present, there is no greater challenge to American AI dominance than China. That is what the NSCAI chairs reiterated to President Trump in their briefing. Bajraktari and the chairs repeated this message to Secretary of Defense Lloyd Austin and Deputy Secretary Kathleen Hicks in the Pentagon. Bajraktari and the group would again stress this point to the Office of Director of National Intelligence. At every turn, they delivered a consistent and cogent message on the urgency of seizing the moment before China’s AI ascension.

For one thing, whatever the U.S. does, China is close at its heels. After the 2016 Cyber Grand Challenge by the U.S. government, China not only paid attention but held seven such competitions.^{xxiii} When the U.S. announced an AI system to help fighter pilots, China announced a similar system in less than a year.^{xxiv} When we (the authors) organized a competition to help defenders get experience attacking AI systems, the Chinese online marketplace company Alibaba took it to the next level. It not only held a similar competition, but an entire series of challenges with much larger prizes.

China seems to be acutely aware of the possibility that AI systems can be attacked, including those used by

[T]he Chinese Army can cut off, manipulate or even overwhelm the “nerves” of American AI military systems with data deception, data manipulation and data exhaustion.

the US military. In a 2021 document^{xxv} used by the Chinese Army, American AI systems are specifically called out as susceptible to information manipulation and data poisoning. Ryan Fedasiuk, a research analyst at Georgetown University’s Center for Security and Emerging Technology (CSET) noted that the Chinese document called the issue of data in AI systems the “Achilles heel” of the ML systems used by the U.S. Army.^{xxvi} The document notes that the Chinese Army can cut off, manipulate or even overwhelm the “nerves” of American AI military systems with data deception, data manipulation and data exhaustion. The Army Engineering University of the Chinese People’s Liberation Army partnered with Alibaba and other Chinese universities and participated in the AI Security challenge to upskill attacking ML systems.^{xxvii}

The Chinese government routinely uses social media—namely Facebook and Twitter—to boost and bolster its authoritarian agenda by creating fake accounts to flood these platforms with counter-narratives, sometimes with the same message verbatim.^{xxviii} Unsurprisingly, social media giants have started to use AI to detect these

spam accounts and shut them down. In 2021, reporting by the New York Times and ProPublica showed that more than 300 Chinese-backed bot accounts posted a video attacking Secretary of State Mike Pompeo’s stance supporting the Uyghurs on Twitter.^{xxix} This is how three Twitter bots captioned the videos:

Twitter bot 1: the videos Pompeo most interested in (%)

Twitter bot 2: the videos Pompeo most interested in ‘) (

Twitter bot 3: the videos Pompeo most interested in ^ ¥ _

The random characters appended at the end of each tweet were sufficient to evade Twitter’s AI-based spam filter that was tasked with detecting bot behavior. Such simple tricks work to confuse AI systems at even mature and well-provisioned companies.

There is a deterrence corollary to China’s framing of an AI Achilles heel. Andrew Lohn, Senior Fellow at Georgetown University’s CSET, put it succinctly when he pointed out how the ability to hack AI systems “could provide another valuable arrow in the U.S. national security community’s quiver.”^{xxx} This way, it could deter authoritarian regimes from developing or deploying AI systems—an adversarial AI strategic deterrent. The U.S. has still not fully extended deterrence into the cyber domain^{xxxi} but could use adversarial machine learning as an important arrow in that quiver to nullify any potential gains from AI systems developed by authoritarian regimes. This seems to be unfolding already. One interesting hypothesis from Lohn is that the Russians did not field AI-based weapons in the war in Ukraine because they knew how susceptible they were to adversarial manipulation.^{xxxii}

Defense Can Lead the Way in AI Security

Government agencies are not alone in producing critical ML systems that may be vulnerable to attack. Urgency to adopt AI by companies often breeds lax security standards. “To create models quickly, researchers frequently have relaxed standards for developing safe, reliable and validated algorithms,” a study found regarding those people and organizations

that were building AI tools used for COVID diagnosis.^{xxxiii}

In its final report, NSCAI put forth a series of strongly worded recommendations. “With rare exceptions, the idea of protecting AI systems has been an afterthought in engineering and fielding AI systems, with inadequate investment in research and development.” The report recommended “that at a minimum” seven organizations pay attention, including the Department of Homeland Security, DOD, FBI and State Department.

Many of the recommendations around AI Security involve testing and evaluation (T&E), verification and validation (collectively, TEVV) practices and frameworks. “All government agencies,” the report stated, “will need to develop and apply an adversarial ML threat framework to address how key AI systems could be attacked and should be defended.” The NSCAI’s recommendations include calls to make “TEVV tools and capabilities readily available across the DOD.” Also recommended are “dedicated red teams for adversarial testing” to make AI systems violate rules of appropriate behavior, exploring the boundaries of AI risk.”

Congress and DOD have started to respond. In addition to FY2021 and FY2022 National Defense Authorization Acts, which have implemented many of the recommendations to invest in and adopt AI, the government is at the very beginning of efforts to secure it. For example, in late November 2022, the Chief Digital and Artificial Intelligence Office issued an open “Call to Industry” for a comprehensive suite of AI T&E tools that includes tools specifically designed to measure adversarial robustness.^{xxxiv}

As has often been the case in cybersecurity and risk management, the government is leading a charge to secure AI systems. Recommendations from the NSCAI report have been an instrumental warning voice. It was as if NSCAI was awakening these high-stakes organizations to the plausible threat of attack on their AI systems. In an enigmatic voice reminiscent of the Oracle of Delphi, the NSCAI report directs critical agencies to “[f]ollow and incorporate advances in intentional and unintentional ML failures.” ☐

ⁱ <https://www.nscai.gov/2021/09/23/nscai-to-sunset-in-october/>

ⁱⁱ <https://www.nextgov.com/emerging-tech/2022/11/adoption-ai-health-care-relies-building-trust-dod-va-officials-say/379323/>

ⁱⁱⁱ <https://federalnewsnetwork.com/artificial-intelligence/2020/03/ai-as-ultimate-auditor-congress-praises-irss-adoption-of-emerging-tech/>

^{iv} <https://medium.com/@RPublicService/feds-at-work-right-hand-men-to-the-pentagons-top-officials-ca99b6c93fbf>

^v <https://www.theatlantic.com/international/archive/2018/01/trump-foreign-policy/549671/>

^{vi} <https://medium.com/@RPublicService/feds-at-work-right-hand-men-to-the-pentagons-top-officials-ca99b6c93fbf>

^{vii} <https://www.vanityfair.com/news/2018/04/inside-trumpworld-allies-fear-the-boss-could-go-postal-and-fire-mueller>

^{viii} <https://www.nscai.gov/commissioners/>

^{ix} Interview with Bajraktari

^x <https://trumpwhitehouse.archives.gov/articles/promoting-use-trustworthy-artificial-intelligence-government/>

^{xi} NSCAI’s “The Final Report,” <https://www.nscai.gov/2021-final-report/>

^{xii} VII, Tom Murphy. “The first level of super mario bros. is easy with lexicographic.” (2013).

^{xiii} <https://www.defenseone.com/technology/2021/12/air-force-targeting-ai-thought-it-had-90-success-rate-it-was-more-25/187437/>

^{xiv} <https://www.ajl.org/drag-vs-ai#:~:text=%23DRAGVSAI%20is%20a%20hands%20Don,artificial%20intelligence%2C%20and%20algorithmic%20harms.>

^{xv} <https://www.nytimes.com/2021/09/28/technology/facebook-ivermectin-coronavirus-misinformation.html>

^{xvi} <https://www.nbcnews.com/tech/news/ivermectin-demand-drives-trump-telemedicine-website-rcna1791>

^{xvii} <https://www.theguardian.com/technology/2016/mar/24/tay-microsofts-ai-chatbot-gets-a-crash-course-in-racism-from-twitter>

^{xviii} R. S. S. Kumar, M. Nystrom, J. Lambert, A. Marshall, M. Goertzel, “A. Comissoneru, M. Swann, and S. Xia, “Adversarial machine learning--industry perspectives,” in IEEE Security and Privacy Workshop, 2020.

^{xix} <https://www.bloomberg.com/news/articles/2022-08-08/meta-s-ai-chatbot-repeats-election-and-anti-semitic-conspiracies>

^{xx} <https://www.armed-services.senate.gov/press-releases/inhofe-reed-praise-senate-passage-of-national-defense-authorization-act-of-fiscal-year-2021>

^{xxi} <https://www.forbes.com/sites/markminevich/2022/03/23/ai-visionary-and-leader-dr-jane-pinelis-of-the-us-department-of-defense/?sh=5b121a4b5aa5>

^{xxii} <https://www.foxbusiness.com/politics/biden-defense-chief-china-pacing-amid-ascendancy>

^{xxiii} <https://cset.georgetown.edu/publication/robot-hacking-games/>

^{xxiv} <https://breakingdefense.com/2021/11/china-invests-in-artificial-intelligence-to-counter-us-joint-warfighting-concept-records/>

^{xxv} <https://perma.cc/X9KQ-4B9L>

^{xxvi} <https://breakingdefense.com/2021/11/china-invests-in-artificial-intelligence-to-counter-us-joint-warfighting-concept-records/>

^{xxvii} Chen, Yuefeng, et al. “Unrestricted adversarial attacks on imagenet competition.” arXiv preprint arXiv:2110.09903 (2021).

^{xxviii} <https://www.nytimes.com/interactive/2021/12/20/technology/china-facebook-twitter-influence-manipulation.html>

^{xxix} <https://www.nytimes.com/interactive/2021/06/22/technology/xinjiang-uyghurs-china-propaganda.html>

^{xxx} <https://cset.georgetown.edu/publication/hacking-ai/>

^{xxxi} <https://www.msnbc.com/cdn.ampproject.org/c/s/www.msnbc.com/msnbc/amp/shows/reidout/blog/rcna48322>

^{xxxii} <https://www.forbes.com/sites/erictogle/2022/03/16/the-vulnerability-of-artificial-intelligence-systems-may-explain-why-they-havent-been-used-extensively-in-ukraine/?sh=1f685d7637d5>

^{xxxiii} <https://pubs.rsna.org/doi/full/10.1148/ryai.2021210011>

^{xxxiv} https://go.ratio.exchange/exchange/opps/challenge_detail.cfm?i=46C49763-B80B-4AF9-80AD-053F2B2095EF

SmartLab

Innovative Digital Learning Tools

MATTHEW LONN, EDD

Assistant Principal,
Sheyenne High School, West Fargo

“I had the best experiences,” said Robert Joy in an interview last April. Robert was a 4th-grader at LaMoure Public School in southeastern North Dakota. “It is what I look forward to doing every day.”

The “it,” which Robert and his other 4th-grade classmates talked enthusiastically about, is the school’s SmartLab: a separate room in the school that was created to integrate digital curriculum and hands-on, project-based learning with the targeted goal of helping students focus on communication, collaboration, critical thinking and creativity.

Not only does the SmartLab increase student engagement—which is significant for today’s youngsters who already have access to computers, Chromebooks, and other smart devices at home and school—but the lab also contributes to LaMoure’s students outperforming peers on the North Dakota State Assessment (NDSA) for math and English Language Arts (ELA) proficiency by wide margins.

Small School, Small Town

The LaMoure school district is contained in one building that houses all 267 K-12 students with 22 instructors. The elementary school encompasses grades K through 6, the junior high school grades 7 and 8, and the high school grades 9 through 12.

LaMoure is a typical small town on the Northern Great Plains. With a population of 764, the town is average size for the state. The median household

income in 2020 was \$61,477, about 6 percent less than statewide. There is a grocery store, hardware store, gas station and several restaurants. The school district is the largest employer in town, while agriculture is the prime economic driver. There are no stoplights since traffic is usually light—except during planting and harvest seasons.

Innovative Approach to Learning

Walking into the school building, the junior high and high school is on the right, and the elementary school is on the left with the lunchroom straight ahead in the middle. Next to the lunchroom is an unassuming room that functions as the “SmartLab.”

As in other classrooms, there is a computer on the teacher’s desk. Similarities end there. There are four pillars in the SmartLab with power outlets, each servicing three separate long desks with two chairs and a computer. Along the east and south walls are shelves containing materials that look like toys. But, as the SmartLab learning facilitator, 23-year-old Colton Altringer, explained, these are student project kits, which include Snap Circuits, Zometools, Sphero, K’Nex, Makey Makey, Digital Sandbox, Lego Robotics, IQ Key and Ozobots. On the north wall, there are two 3-D printers, and on the west wall are posters with projects listed so students know what they should be working on.

In 2016, LaMoure was the first school in the state to pilot the SmartLab project, which started as a joint effort between North Dakota’s Center for Distance Education (CDE), an online supplemental educational provider, and Creative Learning Systems (CLS), which created SmartLabs and began installing them in schools nationwide in the late 1980s. Since then, CLS has continuously upgraded SmartLab technology and curriculum, while maintaining the core methodology of student-driven, project-based learning. At the time of LaMoure’s installation, I served as the Director of Statewide Programs for CDE. Now I serve in the West Fargo School District as a high school assistant principal.

Because of cost, the SmartLab program wasn’t a viable option for students in smaller rural districts until 2016, when the price tag decreased from \$150,000 to less

Robert Joy builds a forklift with K'nex, which enables students to follow building instructions or create original designs. Photograph by Jerry Anderson.



than \$80,000. Now there are 25 SmartLabs statewide, with several more to be installed next summer.

Altringer graduated from the University of Jamestown in December 2021 and, a month later, he was hired as a long-term substitute teacher at LaMoure. In addition to working with 7th-graders in the SmartLab, he teaches history to junior high students. He said the first few weeks in the SmartLab were a bit chaotic. The SmartLab is “much louder than other rooms in the school and students are moving around the room constantly, mostly in positive ways,” Altringer said. He explained that even though the environment looked chaotic, some structure was still needed, and expectations had to be set with students when he first started. But now, “they understand those expectations, and how they are supposed to be showing their learning.” He noted that it’s common for students

to sit down and begin immediately working on what they’re engaged in. This is especially true if it’s a project they are passionate about.

Junior High in the SmartLab

LaMoure’s junior high students are required to take a course titled “Intro to SmartLab.” In the spring of 2022, 8th-graders spent first semester in the SmartLab, and 7th-graders spent second semester in the SmartLab. They are on a block schedule, which means they spend 90 minutes every other day in the SmartLab.

Altringer views his role in the SmartLab differently than teaching American History in his regular classroom. In the SmartLab, he acts as a guide or facilitator helping students find the information they need to complete projects on their own. This contrasts

starkly with the methodology used in his 7th-grade history course, in which he relies on direct instruction via lecture and activities, and he administers both formative and summative assessments.

In the SmartLab, the lessons and content knowledge are already created and provided via cloud-based “learning launchers” that can be accessed digitally anywhere and on any type of smart device. Each SmartLab comes with a scope and sequence for students to follow, which expose students to multiple content areas, including ELA, math and science standards. At the start of a project, Altringer works with each student group to help define what skill proficiencies to develop.

Students must demonstrate their new proficiency levels with performance assessments, which require demonstrations of skill or competency mastery through task performance. In contrast, summative assessments, like those commonly used in Altringer’s history course, require students to remember key facts and dates, which were presented in lectures or practice activities, then answer written questions involving them.

As well, Altringer’s SmartLab students journal their work daily in their Google Drive accounts. Google Drive enables users to store and share files in the cloud and also synchronize files across digital devices. Google Drive includes Google Docs, Sheets and Slides.

The 8th-graders take photos and videos of their final product to demonstrate the process they completed. If time permits, the students present to their peers to demonstrate their learning. Altringer said this is often a powerful experience because students must explain the final product and how they created it.

Elementary Teaching in the SmartLab

Paulette Carlson teaches 14 students in the 4th-grade in the SmartLab. During every six-week period, beginning in the 4th grade and through the 6th grade, students take two-week rotations, for 45 minutes daily, in the SmartLab.

Carlson requires that students journal four to five sentences daily in the SmartLab, which helps her teach them how to use Google Docs. Carlson uses this process to reinforce ELA standards for writing. “All of the launcher projects students complete,” she said, “match the 21st Century Skills standards of the 4 C’s

(Critical thinking, Creativity, Communication and Collaboration).”

This is important, she emphasized, because of North Dakota’s focus on workforce readiness at the middle- and high-school levels. Carlson also integrates state standards in 4th-grade science, communication and math while students work on SmartLab projects outlined in the digital curriculum.

“In the SmartLab, students feel that the learning is put into their hands,” Carlson explained. “They have a sense of accomplishment since the teacher is the facilitator on the side, and they are self-directed.” Students want to miss school the least on SmartLab days. “When year in and year out, my students can’t wait to go to the SmartLab,” she said, “I don’t need to ask if it has value.”

Project-Based Learning

For decades, teachers and administrators have been integrating project-based learning into their curriculum. Its defining methodology utilizes projects in which students construct something that enables them to demonstrate mastery of a new learning level or skill set. Projects are designed to focus on real world issues to which students feel connected.

In project-based learning, the teacher seldom delivers knowledge to students via a lecture or other direct instructional method. Instead, the teacher directs students to work together to define and then discover the requisite knowledge on their own. Students must use the acquired knowledge to achieve the goal set at the beginning of the learning engagement. Instead of focusing on memorization and active listening, project-based learning places greater emphasis on the process students use to achieve end goals, such as collaboration and communication.

Results

Student feedback about the SmartLab has been overwhelmingly positive. When the 4th-graders were asked for their views about the SmartLab, their responses were characterized with comments such as “fun,” “challenging,” “cool,” “creative” and “jaw-dropping,” as well as lengthier statements like, “You want to go back every day,” “It’s fun coding and all the stations were fun,” “I had a brilliant experience,



Two 8th-graders, Emelia Lehr and Logan Potts, design and build a bridge in the SmartLab using Zometools, which provides components for student engineering projects. Photograph by Jerry Anderson.

it was so much fun;” and “I had to work hard, but I did it and it looked cool.” Attendance rates mirrored the state average of 95 percent. However, chronic absenteeism was only 8 percent, compared to the state average of 15 percent. A student is considered chronically absent if he or she attended school for more than 10 days and missed 10 percent or more of enrolled days.

Every year, all public-school students in 3rd through 8th and the 10th grade are required to complete the North Dakota State Assessment (NDSA). The exam measures proficiency in ELA, math and science. LaMoure’s students significantly outperformed their peers statewide in both math and ELA. During the 2021-22 school year, the average proficiency/advanced level in ELA for all students in these grades statewide was 45 percent, compared to 63 percent for Lamoure’s students. In math, the state average in math was 39 percent, versus 63 percent for Lamoure’s students. Lamoure’s achievement levels have been increasing steadily since 2019 when the ELA score was 48 percent and math was 48 percent.

Given that traditional instruction is used in most of Lamoure’s classes, as in schools statewide, a major differentiating factor is the SmartLab. There is no way to determine whether this is sufficient to explain the school’s high achievement levels, but certainly the enthusiasm for learning, which the SmartLab engenders, helped.

In science, Lamoure’s NDSA achievement fell from 78 percent in the 2019-20 school year to 43 percent this year, 1 percent below the statewide average. Mitch Carlson, LaMoure’s superintendent, explained that the school has focused heavily on reading and math interventions in recent years—at the expense of time spent on science and social studies. In addition, a new science curriculum was purchased this year and still requires more time for full implementation.

Moving Forward

As more complex and innovative digital tools become available, it becomes even more important for the field of education to examine how best to leverage those tools to increase student engagement and learning. The project-based learning approach emphasized in LaMoure’s SmartLab is a demonstration of how this successful tech-based methodology could be replicated across the educational landscape.

As Superintendent Carlson explained, the SmartLab wasn’t implemented to raise standardized test scores, although that happily occurred because of the increased student engagement and focus on critical content standards in SmartLab projects. Instead, SmartLab was selected to help students learn digital and project-related skills that better prepare them for the workforce. According to teacher and student feedback, the SmartLab enabled students to become more resilient problem-solvers. Now that’s a vital skill for any career. 📍

FOUR HORSEMEN OF TECHNOLOGICAL CHANGE

Farmers, Elevator Operators, Coal Miners, Bank Tellers

EVAN J. ZIMMERMAN

CEO of Drift Biotechnologies and Chairman of Jovono

Technology has changed the dynamic between labor and capital in the broader economy since the Industrial Revolution. The new steam and manufacturing innovations transformed large parts of the economy by explicitly changing the economics of labor and wages—and increasing inequality along the way. Ever since the early 19th century, when Luddites smashed looms out of fear that newfangled machines would decimate their wages and way of life by replacing specific labor, people have been fearful of technological change.

Technofuturists, however, have always dreamed of technology bringing us a brighter future. They envision a day when machines free workers from the drudgery of factory toil, of 10-hour workweeks, and more leisure time to explore their creativity and forge human connections.

Who is right? With the accelerating speed of AI—the development and deployment of which promise to be as disruptive as looms of old, if not more so—we will find out sooner rather than later. Just

this year, new tools, such as Google’s LaMDA and OpenAI’s DALL-E and ChatGPT, have demonstrated astounding capabilities, ranging from convincing chat conversations to acceptable text prose to mind-warping images that delight and disquiet.

But before we look to the future, we should consider the past. History shows us four examples of change brought about by labor-replacing technology. And so, the question looms: Will the future of workers today resemble that of farmers, elevator operators, coal miners or bank tellers?

Farmers

Farming was once the basis of every economy. Thomas Jefferson considered farmers in our sparsely populated country as the nation’s soul and its future, arguing at the birth of the republic that “[s]uch is our attachment to agriculture, and such our preference for foreign manufactures, that be it wise or unwise, our people will certainly return as soon as they can to raising raw materials and exchanging them for



finer manufactures than they are able to execute themselves.”ⁱ This remained the case for a long time. In 1850, farming accounted for half of all American jobs, but by 1980, due to technological change, that fell to only 4 percent.ⁱⁱ Since then, the proportion of agricultural workers has levelled off at 1 percent.ⁱⁱⁱ

Agriculture is one of humanity’s first technologies, along with fire, domesticated animals and spears. To this day, new technology drives efficiency with inventions such as machine vision satellite imagery, GPS-automated tractors, the Haber-Bosch process and pesticide-resistant GMOs. These technologies don’t reduce the value of farming. In fact, it is the opposite; they allow for much more value to come from much less labor, creating surplus value that can be enjoyed and invested in other areas of the economy. Though agriculture has always been a field of innovation, from forming the basis of the city-state to the invention of crop rotation, for millennia it was a labor-intensive business that consumed the majority of the workforce. Industrialization supercharged that, turning farming into just one part of a largely urbanized economy.

Historically, technological innovation powered the Industrial Revolution and made cities dominant manufacturing centers. From 1859 to 1929, America’s industrial output increased 28 times, and not coincidentally, from 1850 to 1920, America went from 15.3 percent urban to more than half urban.^{iv} Today, only 21 percent of Americans live in rural areas.^v This type of huge, secular change driven by technology is rare, not in the sense that technology frequently causes societies to change but in the sense that very few technologies are able to form a platform for over a century of downstream inventions. When it happens, people eventually adjust due to generational change and economic incentives. Like any long-term change, it is not always a straight line, but because it is so gradual and the impacted industries so huge, people have time to acclimate to the new reality. Even a state such as North Dakota, which is nearly twice as rural as the rest of the country, has a nearly 60 percent urban population,^{vi} which a century ago would have exceeded even the most urban countries in the world.

Elevator Operators

Unlike farming, not all technological change and labor shifts are gradual. Consider the elevator—with 6,818 vertical miles of elevators installed in the U.S. and Canada^{vii}—even today one of the most widespread forms of public transportation. Elevators existed for centuries, but in 1852, Elisha Otis’s automatic brake convinced people they were safe to use, especially for skyscrapers. As a result, use skyrocketed, and with it, the number of elevator attendants.

Today elevator operators have all but disappeared—there are so few that, other than specialized industrial roles like grain operators, the job is not even tracked. But at the time, they were required to, as the name suggests, actually operate the elevator, controlling everything from the speed of travel to even manually opening doors. Though automated elevators debuted in 1900, elevators still required attendants because users distrusted the machine. Due almost entirely to user demand, the number of elevator attendants increased to over 90,000 in the 1940s, according to the U.S. Census. That all changed in 1945 when New York City elevator attendants went on strike for more pay. That was the beginning of the end. By 1950, the first census after the strike, employment of elevator attendants had flatlined, and by 1960 automated elevators were becoming ubiquitous, and attendants’ jobs were on the way out.^{viii}

It took a PR campaign and some psychology to make people like driverless elevators (why do you think elevators used to have a voice that calls out each floor?), but automating elevators made them smaller and safer. They were simply better, so once people realized they were safe, no one mourned the elevator operators. People considered elevator operator jobs “low grade,”^{ix} and many of these operators moved to better jobs given the upward mobility at the time. By 1965, it was almost impossible to find working elevator attendants. That said, it is remarkable that it took more than 50 years to go from a technology’s perfection to the mere beginning of its adoption curve. But once it started, the change was rapid considering the physical infrastructure buildout that was needed.

Coal Miners

Coal mining employment peaked in 1923 with 862,536 miners and has steadily declined ever since. In 2016, there were only 81,484 coal miners, marking the first year on record this number fell below six figures. Now there are only 61,402 coal miners.^x Given the long length of time over which this has occurred, technological change—more than environmental regulation—is what has been killing coal mining jobs. Coal competes with other forms of energy, such as natural gas, wind and solar. As such, the coal industry has invested in technology for nearly a century to reduce the labor intensity of coal mining to cut costs and remain cheap.

But even that has not been enough in the face of competition from alternative energy sources such as natural gas, powered by the fracking revolution and solar, which has benefited from Wright's Law, or the idea that costs decrease as production increases and an industry "learns."^{xi} For example, solar is approaching prices cheaper than coal in some places, according to Bloomberg,^{xii} (though it cannot provide baseload power and relies on complementary power sources like natural gas, nuclear and batteries). Ironically, according to an American Enterprise Institute estimate, solar is 79 times more domestically labor-intensive per megawatt hour than coal, making it a net job creator.^{xiii}

Yet, coal-centric regions have been decimated. States such as West Virginia are among the poorest in the U.S., but even within states, like Pennsylvania, coal regions are some of the poorest counties (in fact, none of the five "Coal Region" counties have per capita income above Pennsylvania's average). This is not a uniquely American phenomenon. Across the Western world, large, historic coal regions are among the poorest in their countries, like Germany.

Something has broken down. Unlike in our previous parables, coal country did not rejuvenate. Coal's fade has occurred over a long period, just like farming, but coal country failed to adapt. Unlike farming, people did not move to new energy industries, nor did those industries come to them. Coal production acquired massive technological improvements that

reduced labor intensity, like elevators. Yet, new jobs did not materialize to take advantage of the newfound productive wealth and newly free labor even though alternatives (that could piggyback off at least some of the previous infrastructure investments) were available. Though many of these coal regions are rural, they are already part of the supply chain and have cheap land and labor that could be used for, say, manufacturing. And while technological improvements in coal drove the cost down and managed to increase demand for decades, it was not enough to stop the bleeding.

Coal demand peaked internationally in 2013, according to the International Energy Agency's Coal 2020 report,^{xiv} so there was not enough value created to offset the losses. Coal regions have the wealth and human capital to compensate by building natural gas refineries, or new nuclear plants, or solar manufacturing, but didn't—the interesting question is, why?

So, what isn't working? Like many natural resource-based industries, it has been difficult for coal to adapt. The key is culture. Coal mining historically involved dedicated populations, which created a culture^{xv} that melded community and identity. Because the decline of coal mining has been slow-moving and culturally entwined, this has led to a sense of despair and lack of belonging stemming from a sense of betrayal. This has resulted in the interesting phenomenon of coal country powering populists throughout the Western world, like Brexit, Marine Le Pen, Donald Trump and even Geert Wilders in the Netherlands, all of which put up some of their best vote margins in coal-dominated regions that have embraced populism because of this sense of betrayal in an attempt to grasp onto hope. Combine that with issues of economic mobility, like difficulty moving and job-retraining programs with poor records of success, and it's no wonder coal country, with its large sunk costs in fixed assets, has had difficulty adjusting.

States such as North Dakota, where a mere 15,000 out of 410,000 employees are in lignite coal,^{xvi} which is amazingly the same number as in 1910,^{xvii} have fared better because coal was never as dominant. Displaced coal workers were absorbed into other energy sources or even other industries.

Bank Tellers

Not all technological change results in the destruction of jobs, as with elevator operators; some, like ATMs, increase the number of jobs.

Arguably, ATMs were the first large-scale commercial application of artificial intelligence. Interestingly, ATMs are also arguably the first successful commercial applications of machine vision, which counters claims that AI necessarily destroys jobs. Though the first ATMs were chemistry based, Yann LeCun's back propagation technique^{xviii} allowed for optical character recognition to be widely used for reading checks in the 1990s, with the first commercial deployment in 1994.

When ATMs began appearing in the late 1960s, many commentators predicted they would eliminate bank tellers. This was seen as inevitable; ATMs were cheaper, faster and worked 24/7. They were popular with customers from the start and spread rapidly thanks to new technologies^{xix} once the shared ATM was developed (beforehand, ATMs were only usable within one bank's networks),^{xx} with a massive acceleration starting in the 1990s due to a big decrease in the cost of deployment.^{xxi} Nonetheless, since the introduction of ATMs, the number of bank tellers has steadily increased from 300,000 in 1970 to 600,000 in 2010. Though the number of bank tellers per branch decreased from 21 to 13, it became cheaper to operate a branch, so banks opened more of them. In 2020, the number of bank tellers finally decreased for the first time, not due to ATMs but the internet as banks finally started closing branches.

Furthermore, tellers adjusted their jobs around their computerized colleagues.^{xxii} The ATM-bank teller example is one of tech's favorite counterpoints to those who oppose technological progress. It is a fantastic story of what are called "centaur systems," or the observation that in many contexts, AI and humans working together outperform either one separately by focusing on their strengths (as in chess, which is the namesake of this type of collaboration). In economic terms, this type of capital-labor substitution revealed a complementarity. Humans are better than machines at selling. By making bank branches cheaper to operate, they transformed into high-powered sales centers

where humans intervene only for more complex transactions or to upsell more profitable services.

The Paths Forward

What we can learn from these four archetypes of technological change is this: Whether a new technology appears suddenly or gradually, as long as the economy is growing over decades (even with periodic recessions as at risk today), a soft landing is possible. Farmers took the new jobs created by the urbanization their technology enabled, elevator operators found new jobs because they were in a fluid economic environment, and bank tellers specialized to work with their ATM partners because a growing economy meant more demand for financial services. But we need to communicate the benefits of new technologies and start repurposing that newfound productivity immediately for future industries to make that happen. Older industries often try to resist, so we need to make people want to work in the new industries by making sure the new economy's onramps maintain workers' dignity and use their skills as much as possible. For coal miners, for example, this would mean job training programs^{xxiii} that widen the aperture of cultural pride from providing coal to providing energy, even in the energy forms that will win the future. It also means reusing physical infrastructure when possible, such as the recent Berkshire Hathaway effort to convert a West Virginia coal plant into a nuclear power plant.^{xxiv}

This brings us back to AI. The coming wave of AI will be swift, and tools from DALL-E to ChatGPT to LaMDA to Github Copilot are coming for white-collar jobs first—the types of jobs that are truly digital and require little or no capital expense to do. The past six months alone have seen a tidal wave of progress with more to come. This change is coming faster than we think, from AI graphic arts to fully robotic "lights out" car manufacturing. We will face more change than we have faced in hundreds of years, and if we don't learn from the past to think about the future, it will get messy. It won't be a few thousand coal miners wondering what they're going to do for work; it will be millions of workers. We need to be nimble so that when the new jobs come around that look different than what work looks like today, we know what to even retrain for.

We are not passive actors unable to influence the social conditions surrounding technological change, and that is what will make the difference. We don't yet know which technologies will fall into which buckets, but there are lessons we can learn from the history of industrialization that can change how we are affected by those technologies. The key differentiators between the optimistic and pessimistic technological obsolescent futures are escape valves and creating value faster than it gets destroyed, so that there are more goods and services to enjoy and invest in, as with farmers, elevators and ATMs.

AI's immediate future involves creating software that can respond to unstructured queries and create surprisingly good work in highly constrained domains and conditions. Though some leaders and technologists, such as Alibaba founder Jack Ma, predict "decades of pain" from the transition,^{xxv} it need not be so. It is crucial not to ask "dumb AI" to do too much,^{xxvi} meaning that we should not ask AI systems to go beyond their narrow domains nor deploy them beyond the capabilities they actually have—versus the ones we wish they had. Rather, we should use AI to augment humans and use humans to augment the work of AI systems. The best machines will often not replace people in factories but will allow them to increase productivity by doing specific steps better and safer, while even giving us back some free time and allowing Western countries to reshore manufacturing.

It is also critically important to keep track of the industries that complement AI-heavy workflows, so displaced workers can adapt quickly and, if need be, move into new industries or rapidly expanding industries that can absorb the excess labor. We cannot guarantee that the future will look like ATMs, because so much depends on the technology itself. But we can increase the likelihood of that future by embracing it and deciding, collectively, to adapt. Banks chose not to minimize costs but to maximize profits, which meant instead of eliminating jobs the second they could, they repurposed them into higher-value functions. As a society, we should take the same tact. It is also critical to recall that this change occurred in a regulatory environment that did not penalize the banks for changing the nature, scope and quantity of jobs in

response to technological change. Though government will have a role in helping those who fall behind, these four technologically challenged horsemen show that flexibility in society at large is essential for a successful transition.

The future of AI is coming fast, and it is coming strong. But take solace in this: Farmers, elevator operators, bank tellers and coal miners are not unique. These are merely examples symbolizing countless other professions that have been disrupted over time by new technologies. Just like these stand-ins, the workers of the multitude of historical industries adapted too. For all the generational tumult that whole economies can face, humans adapt and survive. Unlike our ancestors, we have the benefit of learning from history to enjoy the fruits of technology without the pain. ☐

ⁱ <https://press-pubs.uchicago.edu/founders/documents/v1ch4s9.html>

ⁱⁱ <chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.bls.gov/opub/mlr/1981/11/art2full.pdf>

ⁱⁱⁱ <https://www.bls.gov/ooh/Farming-Fishing-and-Forestry/Agricultural-workers.htm>

^{iv} <https://oxfordre.com/americanhistory/americanhistory/view/10.1093/acrefore/9780199329175.001.0001/acrefore-9780199329175-e-327>

^v <https://www.statista.com/statistics/985183/size-urban-rural-population-us/>

^{vi} <https://www.icip.iastate.edu/tables/population/urban-pct-states>

^{vii} <http://www.neii.org/presskit/printmaster.cfm?plink=NEII%20percent20Elevator%20percent20and%20percent20Escalator%20percent20Fun%20percent20Facts.cfm>

^{viii} <https://www.npr.org/2015/07/31/427990392/remembering-when-driverless-elevators-drew-skepticism>

^{ix} https://www.google.com/books/edition/Cost_of_Living_and_the_War/Ot7owwEACAAJ?hl=en

^x <https://arlweb.msha.gov/stats/centurystats/coalstats.asp>

^{xi} <https://ark-invest.com/wrights-law/>

^{xii} <https://www.bloomberg.com/news/articles/2017-01-03/for-cheapest-power-on-earth-look-skyward-as-coal-falls-to-solar?sref=tjPDGgi8>

^{xiii} <https://www.aei.org/carpe-diem/inconvenient-energy-fact-it-takes-79-solar-workers-to-produce-same-amount-of-electric-power-as-one-coal-worker/>

^{xiv} <https://www.iea.org/reports/coal-2020/demand>

^{xv} <https://researchrepository.wvu.edu/cgi/viewcontent.cgi?article=6375&context=etd>

^{xvi} <https://lignite.com/about-us/careers/lignitejobs/>

^{xvii} <https://www.history.nd.gov/hp/PDFinfo/Coal-Mining-Context-part-1.pdf>

^{xviii} <https://www2.spsc.tugraz.at/people/franklyn/ICASSP97/pdf/scan/ic970151.pdf>

^{xix} <https://www.history.com/topics/inventions/automated-teller-machines>

^{xx} <https://www.philadelphiafed.org/-/media/frbp/assets/economy/articles/business-review/1991/brmj91jm.pdf>

^{xxi} <https://www.theatlantic.com/technology/archive/2015/03/a-brief-history-of-the-atm/388547/>

^{xxii} <https://yalebooks.yale.edu/book/9780300195668/learning-by-doing/>

^{xxiii} <https://www.sciencedirect.com/science/article/abs/pii/S2214629617303341>

^{xxiv} <https://dda.ndus.edu/ddreview/opinion-atoms-of-the-world-unite-or-split/>

^{xxv} <https://www.cnn.com/2017/04/24/jack-ma-robots-ai-internet-decades-of-pain.html>

^{xxvi} <https://venturebeat.com/ai/dumb-ai-is-a-bigger-risk-than-strong-ai/>

CONTRIBUTORS

Hyrum S. Anderson, PhD, is the Distinguished Machine Learning Engineer at Robust Intelligence. He received his PhD in Electrical Engineering from the University of Washington, with an emphasis on signal processing and machine learning, and BS and MS degrees in Electrical Engineering from Brigham Young University. Much of his career has been focused on defense and security, including directing research projects at the MIT Lincoln Laboratory, Sandia National Laboratories and Mandiant as Chief Scientist at Endgame (acquired by Elastic). Currently, Anderson serves as the Principal Architect of Trustworthy Machine Learning at Microsoft, where he organized the company's AI Red Team and served as chair of its governing board. Anderson cofounded the Conference on Applied Machine Learning in Information Security (CAMLIS), and he co-organizes the ML Security Evasion Competition (mlsec.io) and the ML Model Attribution Challenge (mlmac.io).

Jerry Anderson serves as the Art Director for Dakota Digital Review. He earned a BA at NDSU and a BS in Design from Minnesota State University Moorhead. He worked for 31 years at the University of Mary as a graphic designer, photographer and instructor in photography. Anderson has published photos in many publications, including regional newspapers and magazines, the New York Times, US News & World Report and Newsweek. He has also published photos in numerous books, including *Every Place with a Name* (State Historical Society of North Dakota, 1976) and *North Dakota 24/7* (Penguin Random House, 2003).

Paulo Flores, PhD, is an Assistant Professor at the Agricultural and Biosystems Engineering Department at NDSU. He earned his BS in Agronomy at the Federal University of Santa Maria and his master's and doctorate degrees in Soil Sciences at the Federal University of Rio Grande do Sul, both universities in Brazil. At NDSU, Prof. Flores developed three new courses in precision agriculture and codeveloped another course on the use of drones for precision ag. He researches the use of drones and sensors for site-specific weed control, and he also investigates implementing high-throughput phenotyping approaches in several NDSU plant-breeding programs. Prof. Flores has authored and coauthored several research papers on the use of drones, sensors and imagery analysis for agricultural purposes, which were published in specialized journals such as Remote Sensing and Frontiers in Plant Science.

Mark R. Hagerott, PhD, serves as the Chancellor of the North Dakota University System. Previously, he served on the faculty of the United States Naval Academy as an historian of technology, a distinguished professor and the Deputy Director of the Center for Cyber Security Studies.

As a certified naval nuclear engineer, Hagerott served as chief engineer for a major environmental project defueling two atomic reactors. Other technical leadership positions include managing tactical data networks and the specialized artificial intelligence AEGIS system, which led to ship command. Hagerott served as a White House Fellow and studied at Oxford University as a Rhodes Scholar. His research and writing focus on the evolution of technology and education. He served on the Defense Science Board summer study of robotic systems and as a non-resident Cyber Fellow of the New America Foundation. In 2014, Hagerott was among the first American military professors to brief the Geneva Convention on the challenge of lethal robotic machines and to argue the merits of an early arms control measure.

The Honorable John Hoeven currently serves as North Dakota's 22nd U.S. Senator, after serving as the state's governor for 10 years. As governor, Hoeven worked to build North Dakota's future by focusing on six pillars of growth: education, economic development, agriculture, energy, quality of life and technology. Under his leadership, the state greatly expanded and diversified its economy. Now, as a member of the U.S. Senate, Hoeven continues his efforts to build a pro-growth business climate; support Science, Technology, Engineering and Mathematics (STEM) education; and strengthen the state as a hub of technology entrepreneurship. These efforts have brought about a third wave in North Dakota's economic growth with the rise of its technology sector. High-tech companies, both startups and multinational corporations, have taken root across North Dakota, developing new innovations to tackle global challenges in a wide range of sectors.

Don Howard, PhD, is a Professor of Philosophy at the University of Notre Dame, as well as a fellow—and former director—of Notre Dame's Reilly Center for Science, Technology and Values, and an affiliate of the university's new Technology Ethics Center. Prof. Howard earned a BSc from Michigan State University and master's and doctorate degrees from Boston University, specializing in the philosophy of physics under the direction of Abner Shimony. Prof. Howard is recognized internationally as an expert on the history and philosophy of modern physics, especially the work of Einstein and Bohr. He served as assistant editor and contributing editor for *The Collected Papers of Albert Einstein* (Princeton University Press) and is coeditor of the *Einstein Studies* series (Springer). Prof. Howard served as the Secretary of the International Society for Military Ethics. His current research interests include ethical and legal issues in cyberconflict and cybersecurity, the ethics of autonomous systems, and ethical issues in energy production and climate change. His

recent publications include “The Moral Imperative of Green Nuclear Energy Production,” published by the Notre Dame Journal on Emerging Technologies (2020). His editorials on technology ethics have appeared in The Wall Street Journal, CNN, InsideSources, NBC Think and other venues.

Arica Kulm, PhD, is the Director of Digital Forensic Services at the DigForCE Lab at Dakota State University. Her team works with clients to execute a variety of digital forensic supports for investigations with law enforcement agencies and cybercrime investigations for South Dakota Consumer Protection and other organizations. She also leads teams that provide free cybersecurity assessments for South Dakota cities and counties through the Project Boundary Fence. Kulm earned a bachelor’s degree from South Dakota State University, and her master’s and doctorate degrees in Cyber Defense from Dakota State University. She also holds several industry certifications. Her doctoral dissertation resulted in a patent on a digital forensic tool. Kulm’s research interests include the dark web and dark web host-based forensics. She is a much sought-after presenter at various conferences and trainings, and as a spokesperson for media engagements.

Ram Shankar Siva Kumar is Data Cowboy at Microsoft, leading product development at the intersection of machine learning and security. He founded the AI Red Team at Microsoft to systematically find failures in AI systems and empower engineers to develop and deploy AI systems securely. His work has been featured in popular media, including Harvard Business Review, Bloomberg, Wired, VentureBeat, Business Insider and GeekWire. He is a member of the Technical Advisory Board at the University of Washington and an affiliate at the Berkman Klein Center at Harvard University. He received two master’s degrees from Carnegie Mellon University in Electrical and Computer Engineering and in Engineering Technology and Innovation Management.

Matthew Lonn, EdD, is an Assistant Principal at Sheyenne High School in the West Fargo Public School District. He earned an MA in Educational Leadership from Concordia University, and an EdD from the University of Mary. He taught high school social studies and coached football and track and field in Red Wing, MN. His report, “Management Systems and Leadership Models in 21st Century Educational Organizations,” was published in the International Journal of Online Graduate Education. In addition, Lonn is a certified Lean Six Sigma Black Belt, which focuses on organizational process improvement to better meet customer/student expectations.

Patrick J. McCloskey is the Director of the Social and Ethical Implications of Cyber Sciences at the North Dakota University System and serves as the editor of Dakota Digital Review. Previously, he served as the Director of Research

and Publications at the University of Mary and editor of 360 Review Magazine. He earned a BA in Philosophy and Political Philosophy at Carleton University and an MS in Journalism at Columbia University’s Graduate School of Journalism. McCloskey has written for many publications, including the New York Times, The Wall Street Journal, National Post and City Journal. His books include *Open Secrets of Success: The Gary Tharaldson Story*; *Frank’s Extra Mile: A Gentleman’s Story*; and *The Street Stops Here: A Year at a Catholic High School in Harlem*, published by the University of California Press.

Mark P. Mills is a Manhattan Institute Senior Fellow, a Faculty Fellow in the McCormick School of Engineering at Northwestern University and a cofounding partner at Cottonwood Venture Partners, which focuses on digital energy technologies. Mills is a regular contributor to Forbes.com and writes for numerous publications, including City Journal, The Wall Street Journal, USA Today and Real Clear. Early in Mills’s career, he was an experimental physicist and development engineer in the fields of microprocessors, fiber optics and missile guidance. Mills served in the White House Science Office under President Ronald Reagan and later co-authored a tech investment newsletter. He is the author of *Digital Cathedrals and Work in the Age Robots*. In 2016, Mills was awarded the American Energy Society’s Energy Writer of the Year. In 2021, Encounter Books published Mills’s latest book, *The Cloud Revolution: How the Convergence of New Technologies Will Unleash the Next Economic Boom and A Roaring 2020s*.

Evan J. Zimmerman, JD, is an entrepreneur, investor and writer. He founded Jovono and serves as CEO of Drift Biotechnologies. Prior to that, he was a cofounder of Mighty Mug and Mighty Ventures, Inc, which has sold millions of units in 23 countries. He was Chairman of the Clinton Health Access Initiative technology council, which advises on the technology of global public health in dozens of partner countries, and he is a member of the strategy board for the Broad Center for Regenerative Medicine at USC. He also speaks and writes on technology in publications including TechCrunch, Stat News, the California Management Review, and the South China Morning Post. Zimmerman was inducted as the youngest member of the MAK Museum in Vienna’s Biennale Circle, which planned the 2017 Vienna Biennale: Robots. Work. Our Future. Evan has a law degree from Berkeley Law School, where he was a Dean’s Scholar with certificates in technology law and business law, and he won a Prosser award. He graduated from the University of Chicago, where he was a University Scholar, with general and departmental honors in economics and conducted his thesis work on the economics of embargoes as a visiting undergraduate at Harvard University.

DAKOTA DIGITAL
REVIEW PARTNER



TechND Major Members:



TechND Members:

702 Communications
AccuData Services, Inc.
AgriData, Inc.
Bank of North Dakota
Be More Colorful
BEK Communications
Bell Bank
Bismarck State College - Computers & Office Technology
Blue Cross Blue Shield of North Dakota
Capital Credit Union
City of Fargo
City of West Fargo
DAWA Solutions Group
Devii
Doosan
Emerging Prairie
Gate City Bank
Greater Fargo Moorhead EDC
Greater North Dakota Chamber
High Point Networks, Inc.
Lake Region State College
MDU Resources Group
Minot State University College of Business
National Information Solutions Cooperative
NetWork Center Inc.
Nexus Innovations, Inc.
North Dakota Dept. of Career & Technical Education
North Dakota E-Waste LLC
NDUS System Office
Onsharp, Inc.
RealCom Solutions LLC.
Stark Development Corporation
Stoneridge Software
Sycorr
United Telephone Mutual Aid Corp.
University of North Dakota

TechND was founded
in 2000 by North Dakota's
business, government and education leaders
to address workforce needs, advocate for a positive
business technology climate, encourage infrastructure
development and provide knowledge-sharing
opportunities for its membership.

TechND's strategic initiatives:

- Advocate for policies and initiatives that promote the use, growth and development of technology in North Dakota.
- Address employment needs by actively assisting to build a robust, technology ready workforce.
- Champion the technology community by serving as the sector's voice and celebrating the influence, impact and successes of the technology community.