COGNITIVE TIMES

STEVE NORDLUND
Vice President
Boeing Horizon X

DR. GREG HYSLOP
CTO and Senior VP of Boeing
Engineering, Test & Technology
Boeing

AI IN THE SKY

HOW ARTIFICIAL INTELLIGENCE IS ENABLING THE NEXT GOLDEN AGE IN AVIATION

DECEMBER // 2017

FEATURES

CONTENTS

Amir Husain
Editor in Chief

John King
Executive Editor
Production Director

Jon Coyle
Art Director
Designer

Victoria Salas
Producer

Milton Lopez
Producer

Erin Russell
Producer
Staff Writer

Marla Rosner
Staff Writer

Contributing Writers
Keith Estes
Jim Fitzgerald
Keith Moore
Gen. John R. Allen, USMC (Ret.)
Amy Baek



FROM THE EDITOR

by **Amir Husain**

AI policy & technology go hand in hand.

Artificial intelligence is too profound a development and too gigantic a shift to merely be a story about technology alone. AI certainly represents the epitome of our technical prowess, but its implications are far-reaching and go significantly beyond what we would expect solely from a technological development.

In our paper, originally published by the US Naval Institute's Proceedings magazine, Gen. John Allen and I made the case that the application of AI in the battlefield will give rise to what we call, "hyperwar." In that same paper, we stressed that AI would not just drive changes in the way war is waged, but would have profound implications for our national economy, industry, and even social stability.

Just two weeks ago (November 2017), the United Nations Convention on Conventional Weapons (CCWUN) met in Geneva. Ambassadors and representatives from over 100 countries debated—not for the first time—what controls should be put in place to ward off the threat of autonomous weapons. Unfortunately, the attendees could not even agree to a definition of what constitutes an autonomous weapon. With no consensus at hand, further discussions were pushed out to next year.

Meanwhile, developments in Russia and in China indicate that artificial intelligence will make its way into weapon systems very quickly. At the rate things are going, operational AI capabilities will likely be deployed in autonomous weapons in these countries prior to UN ambassadors even agreeing on a definition of what an autonomous weapon really is!

De-facto is outpacing de-jure.

The fact that bans provide little comfort is precisely what we at SparkCognition have been arguing now for the last couple of years. In my extensive national and international travels, I have personally met with four-star generals from NATO member states, leaders of allied nations, ambassadors, and policy makers. Many members of SparkCognition's senior leadership team have collectively briefed senior officers in the Pentagon and high-level representatives from Allied states.

Why are we spending our time doing all this?

Because we understand what AI means to the world. We know that if we are to truly lead the charge in building the AI-powered world of the future, we cannot bury our heads in the sand and avoid the political, policy, economic, and ethical aspects of AI. We choose to not shy away from things that seem difficult. It is only in this way that we can make a meaningful difference.

The time for policy makers, political leaders, and military commanders to understand artificial intelligence is now. If we want to secure future elections, prevent AI-powered mind hacking campaigns that create social instability, and ward off cyber-physical attacks against our infrastructure, then we must act now. If we want to develop the policies that will keep us at the forefront of AI research and allow us to apply the benefits of this technology to improve efficiency, safety and grow our economy while protecting those impacted by automation, the time to act is now.

We cannot hope that bans will protect us in the future. We cannot hope that even if a ban were in place, all signatories would respect it. Hope is not a strategy. We need hard work and we need realism.

This issue of Cognitive Times profiles our amazing partners, Boeing, Honeywell, and other leading companies who are committed to building the AI-powered world of the future with us. It is a fantastic journey—and we hope you join us.

Demystifying THE BLACK BOX

by **Keith Moore**

Consider the following scenario: A grocery store employee is responsible for ordering inventory and must place an order soon. The employee sees that a snowstorm is forecasted for the area. Based on experience, the employee knows that people act irrationally during storms. Over the past few years, every time a storm has been forecasted, the store runs out of perishable goods. Given this evidence, the employee feels justified in ordering additional inventory of milk and eggs.

If questioned by a manager, there is ample evidence to defend the employee's ordering decision—which is generally required in order to stay employed. There is no guarantee that the actions were the absolute best for the given situation, but a clear, discernable methodology was followed in order to achieve what the employee believed would be the highest utility. This is an example of heuristic decision-making.

This method of defensible decision-making is prevalent because it leveraged rational systemic analysis as well as intuitive experience. This flow could then easily be interpreted by another person, making it transitive in nature. Unfortunately for human intuition, the world is changing.

Now, consider an alternate example. A grocery store employee is responsible for ordering inventory and must place an order soon. The employee uses an algorithm that factors in a forecasted snowstorm. The algorithm also notes that the Yen is extremely strong versus the British Pound. On top of that, chicken fertility rates in Georgia have dropped below their three-year average and the hashtag #snowpocalypse is trending on Twitter. Given this evidence, the algorithm instructs the employee to order additional inventory of milk and eggs.

To most people, this decision-making process seems completely insane. The employee couldn't possibly use it as justification to order more supplies. However, puzzlingly, it produces a more accurate forecast for ordering supplies.

This is the newest problem in the algorithmic world: explanation and defensibility of a model. We know that deep learning algorithms work, but why they work breaks human intuition. This is problematic. Because of this, massive amounts of research have been poured into "explainable AI" to help demystify what is inside of the deep learning "black box." While this research is ongoing, there are a few highly promising technologies available today that we believe are going to help pioneer humanly intuitive, explainable AI.

Model-Agnostic Methods: Where the End Justifies the Means

Instead of spending efforts trying to decrypt what is happening within a neural network, it's possible to make black box assumptions about what the network is doing given the inputs and outputs. These solutions, deemed as model-agnostic, hope to create some approximation of what is happening within any model and explain the decisions made.

The technology leader in this space is an open-source tool called LIME (Local Interpretable Model-Agnostic Explanations) from the University of Washington, which takes a model-agnostic approach to determining the "why" behind any decision. LIME operates on the assumption that it is difficult to understand how the model behaves in all cases, but it is possible to understand how it behaves in particular situations.

LIME was designed to create interpretable linear predictions, which it does by agitating input data slightly and identifying how the output changes. LIME uses this to determine an approximate linear function capable of replicating similar local outputs to the model it is interpreting. This linear model developed by LIME can then be used to determine feature importance for an outcome.

While LIME clearly adds value to any machine learning model, it still only provides approximate insights into the driving forces behind why a model behaves the way it does. For example, in the above scenario, it might indicate that the largest influence on the decision to order more milk and eggs was the weather forecast, and the second most important was the Twitter hashtag #snowpocalypse. It could easily prioritize all of the inputs to try and justify the use of them to a person, but it does not actually explain the reasoning behind the decision.

Because of this shortcoming, a new technology is being developed, known as aLIME (Anchor Local Interpretable Model-Agnostic Explanations). aLIME is the If-This-Then-That extension to LIME. For example, in a dataset where the goal is to predict if a person's salary is more or less than $50,000, a series of features are given (Figure 1a). Lime is capable of highlighting the major features associated with the model's prediction (Figure 1b). aLime goes a step further by compiling a heuristic that can be used to intuitively explain why the prediction was made (Figure 1c).

Tools like aLime are a step in the right direction for explainable AI, but there are two major shortcomings:

1. ***These tools do not explain what is actually happening under the hood of a model***
2. ***They provide no information about how confident a model might be in its predictions***

Luckily, there are a few other methodologies that help to address these problems.

One promising topic worth mentioning is neural network composition analysis methods. Techniques such as Garson's algorithm, Lek's algorithm, or randomization show a lot of promise, but more research in them is necessary to justify their application to deep learning.

FIGURE 1: *An evaluation of LIME and aLIME on a sample data set*

(a) Instance

FEATURE	VALUE
Age	37< Age ≤ 48
Workclass	Private
Education	≤ High School
Marital Status	≤ Married
Occupation	Craft-repair
Relationship	Husband
Race	Black
Sex	Male
Capital Gain	0
Capital Loss	0
Hours per week	≤ 40
Country	United States

(b) Linear LIME explanation

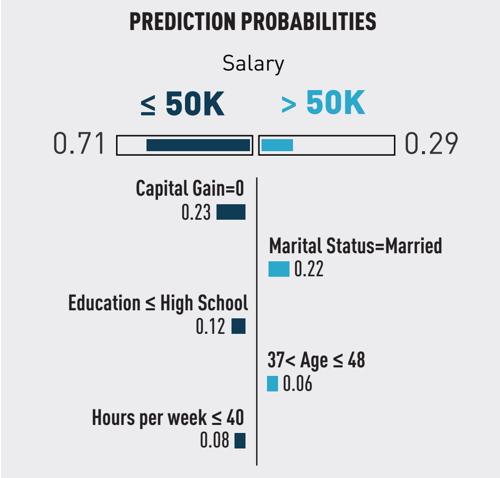

(c) aLIME explanation (anchor)

IF Education ≤ High School
THEN PREDICT Salary ≤ 50K

FIGURE 2: *Differentiation between a Bayesian network and a traditional neural network*

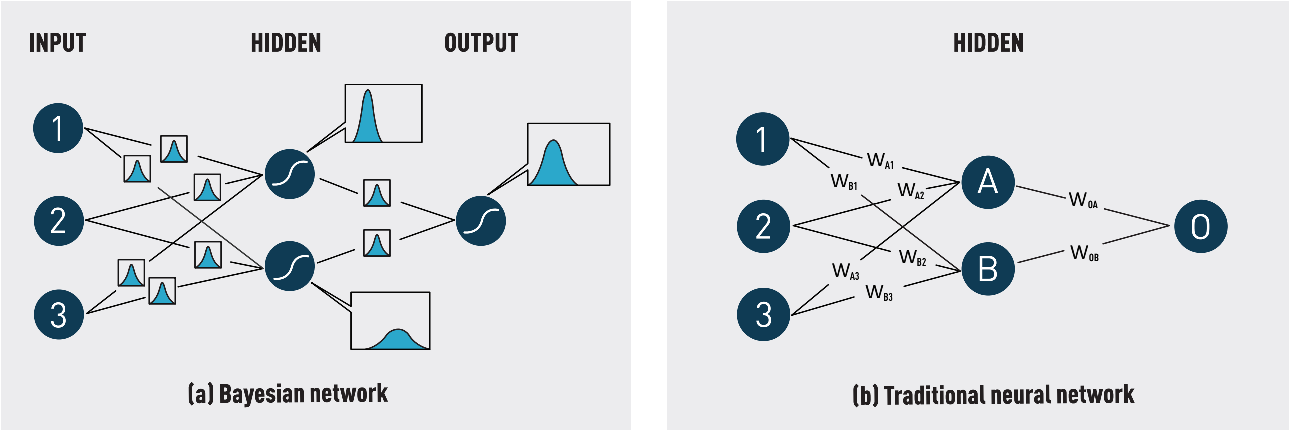

Bayesian Inference–Bayesian Deep Learning

A final promising approach to unraveling the secrets of the black box comes in the form of Bayesian inference. Whereas model-agnostic tools focus on small adjustments of data for analysis and compositional analysis tools analyze network weights and variable behaviors, Bayesian deep learning focuses on making neural networks follow certain statistical principles common in the modeling world—notably Bayesian inference. When combined with network composition or model-agnostic approaches to understanding a neural network, Bayesian deep learning opens up the possibility of a neural network capable of explaining its actions as well as how confident it was in the decision it made.

Unfortunately, the application of Bayesian inference to deep learning is far from straightforward. Bayesian inference is designed for primarily linear models. Because of the highly nonlinear nature of neural networks, the application of Bayesian inference is often intractable or computationally inefficient.

Over the past two years, many of these deficiencies have been overcome and key innovations have spurred drastic advancements in Bayesian inference for deep learning. The rise of probabilistic programming languages like Edward, developed out of Columbia University, have enabled the use of inference in deep learning frameworks like Tensorflow. These inference-based neural networks have already been proven to be more computationally efficient than comparable benchmarked machine learning systems. Instead of attaching a scalar value to every parameter in a network (i.e. connection weight), Bayesian deep learning instead fits a distribution to each parameter using some predefined number of sample points. An example of how these Bayesian networks appear different than traditional scalar networks can be visualized in Figure 2.

Based on Figure 2, it is apparent that a Bayesian neural network requires significantly more parameters than a traditional scalar one. However, that does not necessarily mean more computational resources are required to optimize the network. The approach taken by probabilistic programming tools allow for rapid computation despite the increased number of data elements. This probabilistic model understands the relationship between the likelihood of an expected output given a distribution of weights, parameters and known inputs.

Bayesian deep learning is still in its infancy, but it better explains the uncertainty within a model, and has been proven effective in mathematically proving why certain techniques in deep learning work. The future ahead for this technology is bright, but continued innovation will be necessary in order to gain parity with other heavily marketed technologies within the deep learning community.

While deep learning has raced ahead as the hero of the 21st century, it's important to consider all of the technology that must keep up in order to facilitate the decision-making process we have used for centuries. Luckily, promising technologies like aLime, Neural Network composition analysis, and Bayesian deep learning are providing useful insight into the inner workings of this increasingly complex technology. As long as continued innovation is maintained in developing explainable AI and ethical practices are employed in the development of deep learning systems, the potential for world-changing innovation in the 21st century is abundant. AI does not have to be perfect, but the decisions it makes must be defensible.

FP Foreign Policy

*Originally published on foreignpolicy.com | 11.03.17

THE NEXT SPACE RACE IS ARTIFICIAL INTELLIGENCE

by **John R. Allen** and **Amir Husain**

Nearly 60 years ago, then-Senate Majority Leader Lyndon B. Johnson seized his colleagues with a stark Cold War warning: Whoever wins the space race, he predicted, would gain "control, total control, over the Earth for purposes of tyranny or for the service of freedom."

The United States won that race not only by reaching the moon but by inspiring the next generation of scientists, technologists, and optimists.

Recently, Russian President Vladimir Putin echoed Johnson's forecast in light of the next great technological race: artificial intelligence, or AI. "Whoever becomes the leader in this sphere will become the ruler of the world," Putin said.

Johnson, leaning into the consequences of the Soviet threat, can be accused of hyperbole. Putin can be accused of the same and, perhaps, worse. But there is truth in their common understanding of technology's power, one that transcends generations and geopolitics. Right now, we fear, the United States is at risk of losing this critical race.

One of us commanded 150,000 troops from 50 nations; the other invents AI technologies that are used in—among other applications—energy, finance, and systems for our national defense. We have seen firsthand how AI will transform warfare, from autonomous flight control systems that can revolutionize air combat to algorithms that can give commanders an unprecedented and precise view of the battlefield. Soon, AI also will become the most potent enabler of competitive advantage throughout most areas of our society and economy, both in work and leisure—with consequences far beyond the usual debate about automation supplanting manufacturing jobs.

While the United States is the birthplace of AI and has historically been the home of the most important innovations and research institutions in this space, our global competitors are on our heels. China recently announced a multibillion-dollar AI development plan to lead the world in the technology by 2030. Russia is developing the next-generation MiG-41 fighter with AI that could control the aircraft at hypersonic speeds as fast as Mach 6. If we don't approach this contest with the same fierce focus we found during the Cold War, we risk losing a lot more than pride.

We don't need an advanced machine to calculate where our competition is outpacing us. We can see our weaknesses with the naked eye: The exclusion of AI from our high-level national agenda, a reduction in science and technology funding, and immigration curbs all harm our competitiveness. The question is whether we'll correct course before it's too late. We have lots of thinking to do about the problems we face and what action to take.

First, the same openness that encourages American researchers to lead the world in innovation also pushes their discoveries quickly into public view before they are even granted protection. Coupled with accessible online lectures from top universities, competitors copy our research easily. While we value America's culture of academic openness, U.S. companies need a faster patent process and government support that can give them some teeth in intellectual property disputes with foreign infringers.

Second, our regulatory regime makes it more difficult to build things in the United States and sell them to other countries, creating a market for foreign competitors who would otherwise not stand a chance. For years, the United States curbed exports of encryption technology and basic processors. This only led international competitors to fulfill demand, creating a market for themselves. When U.S. allies like Saudi Arabia, Pakistan, the United Arab Emirates, and Turkey needed access to unmanned aerial systems to prosecute the war on terror, these requests were delayed or denied. We have since lost almost all of these markets to Chinese exports and indigenous development. As the clamor for curbs on AI grows, we need to keep our competitiveness in mind when we put rules in place.

Third, China is publishing a larger number of papers than the United States about deep learning and beating us in fielding supercomputers. We need more public investment in AI research, not cutbacks in agencies that have been longtime backers of this field, as the current administration's budget proposes. We also need more science and technology funding to compete with the billions of dollars China is investing in its 2030 vision.

The History of
NATURAL LANGUAGE PROCESSING

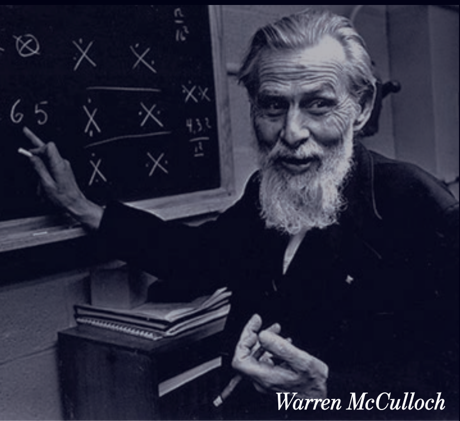
Warren McCulloch

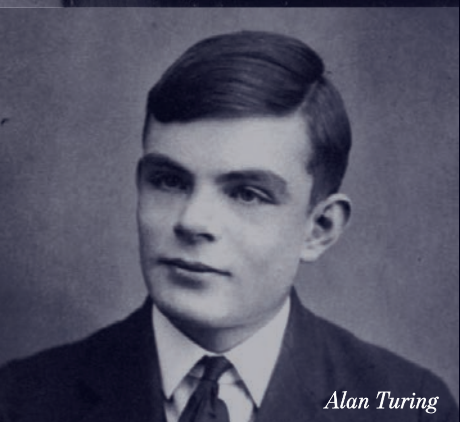
Alan Turing

Frank Rosenblatt


ELIZA

1943
McCulloch and Pitts mathematical model of neural networks created (beginning of deep learning developments)

1950
In 1950, Alan Turing publishes "Computing Machinery and Intelligence," a paper from which the Turing Test emerged. This is when NLP research truly began.

1957
Frank Rosenblatt presented the perceptron, forming the foundation for deep neural networks

1957
Noam Chomsky's "Syntactic Structures" introduced linguistics to the idea of "Universal Grammar," a rule-based system of syntactic structures

1964
ELIZA, a simulated psychotherapist, worked off of a few general grammar rules, was developed

1964
STUDENT, a program to solve high school algebra word problems, was written

60s-80s
Rule-based systems were being widely used for machine translation

1969
Roger Schank introduced the concep-tual dependency theory for natural language understanding

1970
SHRDLU, a natural language system working in a restricted domain with restricted vocabularies, was written

1970s
Programmers began to write "conceptual ontologies," which structured real-world information into comput-er-understandable data

1978
LIFER/LADDER, a natural language interface to a database of information about US Navy ships, was created

1986
Deep learning concept was first introduced

Late 80s
First statistical machine translation systems were developed (introduction of machine learning algorithms for language processing)

Late 90s-00s
AeroText used for information extraction & relational link analysis

1995
A.L.I.C.E. (Artificial Linguistic Internet Computer Entity) was created

1999
Deep learning theories sparked conversation again

2002
Most common topic-based model was developed

2010
First industrial applications of deep learning to speech recognition began

2011
IBM's Watson beats two human competitors on Jeopardy!

2017
Google Brain image enhancement uses deep learning to create detailed images by using tiny source images

2017
The innovative ability to understand tabular data was developed for an NLP solution, DeepNLP®

Present
Deep learning models are excelling at analyzing large texts including structured and unstructured data

IBM Watson on Jeopardy!

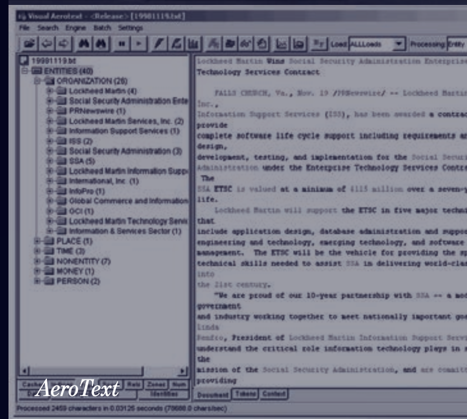
AeroText

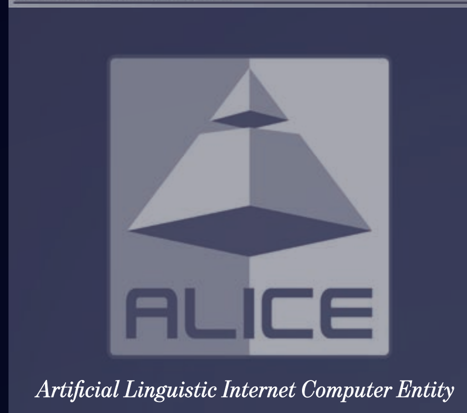

Artificial Linguistic Internet Computer Entity

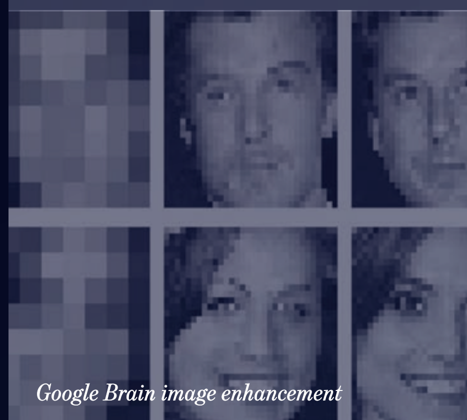
Google Brain image enhancement

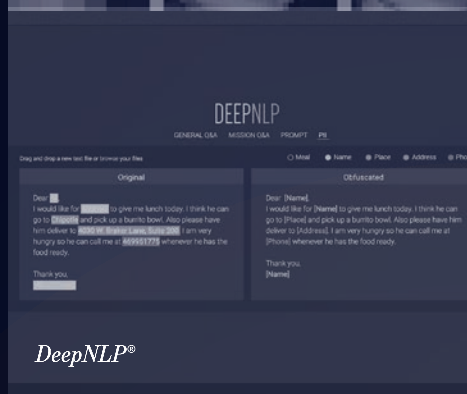

DeepNLP®

Bringing AI to Major League Baseball:

ENDING THE MARINERS' PLAYOFF DROUGHT

by **Keith Estes**

Ingrained in every player from Little League to the Major Leagues is the idea that baseball is a mental game. As Yogi Berra famously proclaimed, "Baseball is 90 percent mental and the other half is physical."

Of course, games are won on the backs of the incredibly talented players residing in Major League Baseball (MLB). But series, divisions, and ultimately World Series championships are won in the minds of the players, managers, and general managers of baseball organizations.

In a 162-game season, managers and general managers make countless decisions factoring in the opposing teams, upcoming schedule, and numerous other data points that affect the game of baseball. The decision-making ability of the front office is the main factor of success, and strings of bad decisions often result in a termination (e.g., the Seattle Mariners' disastrous front office from 2004–2015).

The MLB is at the forefront of the data analytics revolution in professional sports by housing entire analytics departments within front offices, hiring Ivy League data scientists, and replacing traditional, gut-feeling GMs with their analytical counterparts (like the Mariners replacing Jack Zduriencik with Jerry Dipoto after the 2015 season).

Collecting everything from Aaron Judge's league-leading maximum batted ball exit velocity of 121.1 MPH to Albert Pujols cellar-dwelling sprint speed of 23.0 ft/sec, the MLB has hordes of information analyzed by their team of data analysts, statisticians, and mathematicians to dive into the insights and patterns in the data. The next step in the analytics progression of Major League Baseball is surely AI-powered technology to support the front office in reaching better data-driven decisions.

Artificial intelligence is being used in similar instances across a wide variety of industries like oil and gas and utilities so the time is as opportune as ever for Major League Baseball. To put it frankly, utilizing these technologies is all about the data. Whether it's temperature, pump pressure, and vibration monitoring in an oil rig or batted ball exit velocity and player sprint speed, data is data to a machine powered by AI technology.

Artificial intelligence excels at finding subtle patterns and hidden insights in data sets of all shapes and sizes, particularly under complex or changing conditions. Specifically, machine learning, deep learning and automated model building are all AI-powered technologies able to influence baseball and the decisions made in America's pastime. These changes would assist in delivering a model that could learn complex patterns, incorporate human feedback, derive rules, and display reasoning for proposals given by the cognitive system.

MY OH MY, HELP THE M's

A successful proof of concept with an MLB team would be needed before breaking into the entire MLB. Like the Athletics in the mid-to-late '90s, the Seattle Mariners could use a breakthrough from a string of disappointing seasons. While Moneyball is no longer a novel principle, introducing cognitive analytics and AI technology into their front office to break free of their playoff drought could be the perfect fit.

A team plagued by historically bad luck in everything from injuries to blockbuster trades, the Seattle Mariners have enjoyed few successes since their record-tying 2001 season. From the good times of the early 2000s to the bad times of practically every season since, the Mariners have gone through four GMs, ten managers, and nine losing seasons.

Although they were slated to break their playoff drought in three of their last four seasons, the 2017 Mariners once again fell below expectations. With fears of stretching the drought to 17 seasons, the M's are the perfect example of a team that could benefit from the incorporation of AI technologies in their front office's decision-making. Specifically, three areas that stand to be improved by AI include scouting potential players, preventing player injuries/optimizing player productivity, and player management.

The next step in the analytics progression of Major League Baseball is surely AI-powered technology to support the front office in reaching better data-driven decisions.

MONEYBALL 2.0

Scouting has a lot of moving parts and certainly is not an exact science. For the Mariners, current GM Jerry Dipoto integrates a heavy dose of objective statistical analysis and sabermetrics (Moneyball 1.0) with subjective evaluations from scouts to constitute their scouting strategy. When building an AI model for the M's, Dipoto and the front office could develop ideas and parameters of what their ideal players look like for each position, a baseline that the model will try to attain.

For all MLB teams, scouts are held accountable for delivering reports on prospects' current and potential mental, physical, and medical factors and advise if a player is worth the team's resources. Although a subjective analysis from the scouts, natural language processing, an AI-powered technology, could pull insights from their information and enable a model to provide in-context answers regarding a specific player. This combination would present a smooth integration between the subjective and objective aspects of the Mariners' strategy.

Utilizing the wide range of databases that aggregate MLB data today such as Baseball Reference, Baseball Prospectus, or FanGraphs, an AI-powered model could also pull insights from databases that include player statistics, sabermetrics, medical information, injury reports, MLB payroll information, and more to give the Mariners players that fit their parameters. As teams chase the best players for the best prices to work under their salary cap, this AI model could garner insights and patterns across data flows rather than using undervalued statistics, essentially creating Moneyball on steroids.

> *As teams chase the best players for the best prices to work under their salary cap, this AI model could garner insights and patterns across data flows rather than using undervalued statistics, essentially creating Moneyball on steroids.*

For example, after digesting historical data on past and present MLB players and comparing it to the data gathered on a prospect, an AI model could find that prospect A shares a similar throwing motion, stature, and Fielding Independent Pitching rating as player C. These similarities could also be recognized by the model as indicators of similar future performance. With this in mind, player C excelled his first three years in the MLB but ultimately struggled with arm injuries and ineffectiveness until retiring two years later.

The Mariners would then have to decide if the string of similarities between the two is enough to warrant them away from offering a contract to the prospect. While analysis of these similarities are currently conducted by humans, AI-powered techniques can observe the data from every angle and provide insights at a much higher level.

A FAREWELL TO TOMMY JOHN?

Although scouting for quality players is important for building a good team, mitigating the effects of the biggest inefficiency in baseball—player health—deserves more attention from teams. Teams such as the Mariners currently rely on medical personnel to predict and prevent the risk factors associated with major injuries. However, through the utilization of Statcast data and sabermetrics, AI could provide better prediction and prevention of injuries as well as performance optimization on an MLB team's greatest assets: the players.

AI technologies are already being applied to the most important assets in industrial settings, so why should it not be applied to the multi-million dollar assets that MLB players are to their respective teams? The average salary for a Mariners' player is over $2 million per year, with Robinson Cano exceeding $24 million a year as the team's top earner. With this type of payout, the M's should be exhausting every tool and technology to ensure their players are at optimum performance and are not at risk for injury.

For AI-powered technologies to be used in injury prevention and performance optimization, access to Statcast data would be needed. Statcast, MLB's relatively new analytics tool, tracks players on the field through a combination of radar technology and cameras. The system generates

roughly seven terabytes of data per game (the size of approximately 3,000 Netflix movies), and much of this data is accessible only to teams.

The sheer amount of data generated means it would be impossible for a single analyst or even a team of analysts to sift through the entirety of player-movement data available over the course of a season. A Little Leaguer wouldn't be expected to hit Clayton Kershaw's curveball; relying on humans to analyze one season's worth (approximately 17,010 TBs for those interested) of player-movement data is equally ridiculous.

Uncovering hidden insights in player-movement data would give teams information on possible connections between the anomalies in movements that lead a player to be more at risk of an injury. Similarly, data on players' movements would allow an AI model to understand how to maximize the effectiveness of players and give teams a prescription on where the greatest inefficiencies lie.

For example, a model would potentially uncover that Jarrod Dyson, the Mariners' current center fielder, could improve his batting average and on-base percentage by taking advantage of the opposing team's defensive alignment with a drag bunt to a specific place on the infield. While Dyson might not make the perfect bunt each time and teams may catch on to the strategy, the model will continue to learn and adjust based on the data.

CONVENTIONAL THINKING IS GOING OUT THE WINDOW

On top of making decisions that affect the team long term like scouting and injury prevention, an MLB manager must make numerous day-to-day decisions. Anything from the upcoming schedule to the historical matchups of a hitter versus the other team's pitcher needs to be analyzed before deciding who to play and where they provide the biggest value add for the team on that given day.

Many factors can affect the outcome of a baseball game—rarely is there just one decision that can chalk up a win or be blamed for a loss. An AI model in player management could realistically factor in all analytical aspects of the game to provide a manager with an enhanced scouting report.

To level out the unpredictability in baseball, accounting for any and all analytics allows insight into the predictable factors that can impact an MLB game. An AI-powered model could conduct matchup analysis and prepare for different scenarios, allowing a manager to make better data-driven decisions such as delivering the optimum lineup against that day's opponent. Other game-day decisions that could benefit from a data-focused approach could be ensuring a rested lineup, efficiently planning pitching assignments with insight on the upcoming schedule of games, and player positioning within the game.

For this aspect, an AI-powered model could benefit the 2017 Mariners by allowing team manager Scott Servais to see that Felix Hernandez is slated to pitch against the Los Angeles Angels, a team that he has a history of bad performance against. While the Mariners are short on quality starting pitchers due to injuries, the model could prescribe an unconventional strategy to start a relief pitcher against the Angels and switch relief pitchers every two or so innings to keep the Angels off-balance. The model could provide insight into the predicted success of multiple pitchers versus King Felix against the Angels, and the manager could then make a better data-driven decision to give the M's the best probability of winning.

Where winning is in the front of everyone's minds, MLB front offices are pulling out all the stops to help their teams to victory. And if the Seattle Mariners are looking to finally end their 16-year playoff drought, reaching better data-driven decisions by bringing AI to the MLB on scouting, performance optimization, and player management will help them come out on top.

NOVEMBER

Nov
4

CBS News Uses UNIVAC Computer to Predict Election

Presper Eckert and John Mauchly developed UNIVAC for the Bureau of the Census. It was the first American computer designed for business and administrative use.

Nov
5

Internet Boasts 100 Million Sites

On this day in 2006, the number of sites on the internet tipped over 100 million.

Nov
7

Marie Curie's Birthday

Marie Curie was a French-Polish physicist that discovered the elements polonium and radium alongside her husband. She is best known for her research on radioactivity, and was the first woman to win a Nobel Prize and also the only woman who has won multiple Nobel Prizes.

Nov
8

National STEM/STEAM Day

This day is meant to inspire kids to explore and pursue their interests in science, technology, engineering, art, and math.

Nov
10

Microsoft Introduces Windows

In 1983, Microsoft Windows was initially released. Due to a couple delays, the operating system was not introduced to the public until 1985.

Nov
11

Veteran's Day

Originally called Armistice day to commemorate the end of World War II, the holiday was expanded in 1947 celebrates the service of all U.S. military veterans.

DECEMBER

Dec
3

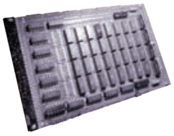

CDC Announces 7600 Supercomputer

Developed in 1968, this model is considered by many to be the first supercomputer, calculating at a speed of 40 megaflops.

Dec
9

Grace Hopper's Birthday

Grace Brewster Murray Hopper, was an American computer scientist and one of the first programmers of the Harvard Mark I computer in 1944 and invented the first compiler for a computer programming language.

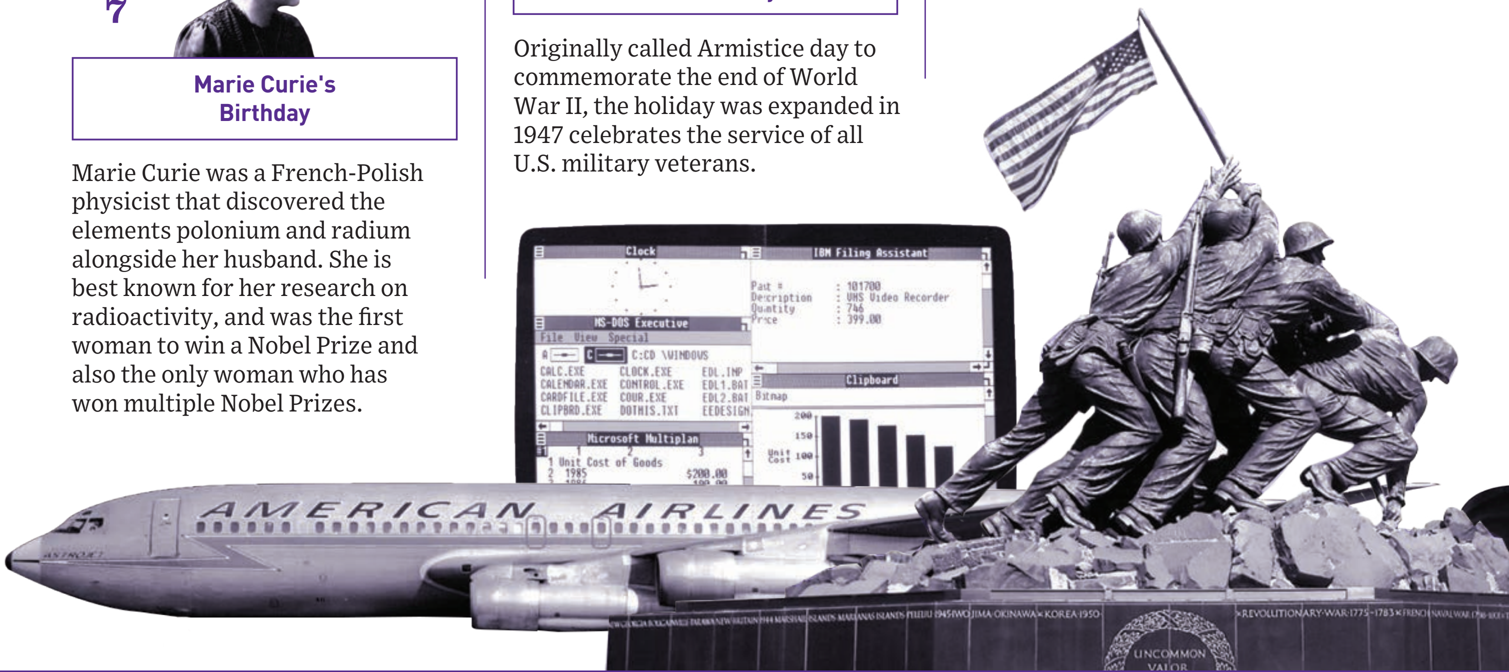

IMPORTANT DATES

JANUARY

TIME MACHINE 2017

Dec 13-14

Time Machine

SparkCognition's first annual summit, taking place in Austin, TX.

Dec 14

International Shareware Day

International Shareware Day was created to remind shareware users about the value they've gained through use of these programs.

Dec 25

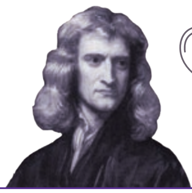

Isaac Newton's Birthday

Best known for his work on gravity and the three laws of motion (and a certain apple fable), Newton's work was instrumental in the scientific revolution of the 17th century.

Jan 8

Stephen Hawking's Birthday

World-renowned physicist Stephen Hawking was born on this day 1942, in Oxford, England.

Jan 11

Amelia Earhart's Solo Flight

In 1935, Earhart became the first woman pilot to fly solo between Hawaii and the United States. She landed in 18 hours and 15 minutes.

Jan 15

Wikipedia Day

Wikipedia was formally launched on January 15th, 2001, by Jimmy Wales and Larry Sanger. It has a total worldwide monthly readership of approximately 495 million.

Jan 21

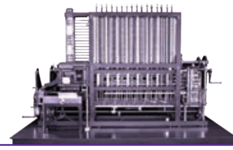

Invention of the First Mechanical Computer

In 1822, Charles Babbage conceptualized and began developing the Difference Engine, considered to be the first automatic computing machine. The Difference Engine was capable of computing several sets of numbers and making hard copies of the results.

Jan 29

Apple Macintosh is Released

Apple Computer, Inc. released its Macintosh computer with an unprecedented media campaign, including a groundbreaking TV commercial shown during the 1984 Super Bowl. Selling for about $2,500, the Mac used a Motorola 68000 microprocessor and had 128k of RAM (memory). The Macintosh's graphical user interface (GUI) was revolutionary and led to its rapid adoption in education, desktop publishing and graphic design.

Jan 29

BOEING

First Jet Passenger Service

In 1959, American Airlines used Boeing 707 jet airliners as the first jet passenger service across the U.S.

FROM HISTORY & TODAY

CHINA'S GROWING AI ADVANTAGE

by **Erin Russell**

Maintaining a leadership position in any endeavor requires significant investment and creativity. This is especially true for an industry changing as quickly as tech, when the race to market can make or break an entire ecosystem. The United States received a harsh reminder of this in the 1980s when facing fierce technological competition from Japan.

Japan was overtaking U.S. market share in the semiconductor industry, becoming the biggest global supplier in the mid-eighties. In response, the United States created a consortium of government and private companies called Sematech that banded together to share the costs and risks of developing computer chip-making skills. A significant investment in Intel by IBM enabled the company to create the microprocessor, and when Japan couldn't reproduce the invention, the country was relegated to the background of the field.

Since then, the US has led the world in developing dominant technology, especially in tech hubs like Silicon Valley and Austin. However, recent budget cuts to funding AI research indicate a troubling shift in priorities from the administration. In contrast, China has committed tens of billions of dollars in the coming years towards efforts to make the country an AI superpower by 2030.

"Whether you look at metrics like the number of publications and the number of patents; or qualitative metrics, such as participation in major AI conferences or success in different competitions, it's clear that China's becoming a major player in AI," explains Elsa Kania, adjunct fellow at the Center for a New American Security (CNAS).

A Harvard graduate, Kania's interest in Chinese policy was piqued by discussions around the Third Offset Strategy—seeking to outmaneuver advantages made by top adversaries primarily through technology. She began digging around in Chinese language materials available at the open-source level to get a clearer sense of the landscape, before landing her current position at CNAS, where she focuses primarily on Chinese defense innovation in artificial intelligence.

Kania specifies that while China's defense strategy has historically been reactive (for instance, developing counter-space capabilities and missiles to reinforce counter-intervention posture), their military is no longer playing catch-up. "The People's Liberation Army (PLA) is trying to actively change paradigms in warfare through pursuing military innovation in disruptive technologies, including artificial intelligence, where it sees that it has the potential to surpass the U.S. technologically and to apply these technologies militarily."

The New Global Arms Race

China is increasingly beating U.S. researchers to major technological milestones. When Microsoft triumphantly announced in 2016 the development of language comprehension software that could match the human ability to recognize and understand speech, the excitement was dampened by a tweet from Andrew Ng, Baidu's former chief scientist: "We had surpassed human-level Chinese recognition in 2015; happy to see Microsoft also get there for English less than a year later."

China's New Generation AI Development Plan, released in July, details its AI ambitions, a field leaders are hoping to grow to a $150B industry. While Kania mentions the plan is "not dissimilar" from U.S. plans in terms of policy focus, China has committed funding while the U.S. has cut it.

"China might take a more integrated and holistic approach to looking at the ramifications of AI in warfare and being willing to experiment and explore," she explains, "rather than taking more of a focus on AI for specific applications or specific problems, as the U.S. military has in certain respects to date."

A People Problem

China has a larger workforce than the U.S., though is for the most part less educated—which is central to their AI strategy. While the U.S. military views their people as their greatest asset, China and other populous authoritarian regimes are perhaps less willing to trust people, given concerns of control. In particular, the PLA continues to struggle with human capital and training, seeking to leverage AI to compensate for the shortcomings.

It remains to be seen if taking humans out of the loop in warfare will be an actual advantage. However, Kania points to former Deputy Secretary of Defense Bob Work anticipating that, unlike the U.S., authoritarian regimes like China and Russia might be more inclined to make the removal, seeing humans as more of a weakness.

"The Chinese military is pursuing research and development for a broad range of applications involving AI: everything from intelligent and autonomous unmanned systems; use in intelligence, surveillance, and reconnaissance; to support war games, simulations, and training; and also to provide intelligence support to command decision-making," Kania explains.

In fact, the state-run newspaper China Daily announced in August that the government had begun developing a new cruise missile system with a "high level" of AI.

> *The People's Liberation Army is trying to actively change paradigms in warfare through pursuing military innovation in disruptive technologies, including artificial intelligence, where it sees that it has the potential to surpass the US technologically and to apply these technologies militarily.*

Not Quite Battle-Ready

Kania thinks China still faces challenges, saying, "It's unlikely that the Chinese military would progress towards full autonomy across the board when artificial intelligence is at the current stage in its development. It still makes mistakes no human would make." (She does note that opinions on this can certainly differ within military ranks.)

Kania stresses that even in non-lethal military systems, errors like a mistake in a translation could dangerously inform decision-making. She references Stanislav Petrov, who is credited with preventing World War III in 1983 by questioning the results from an early warning system that indicated missiles had been launched from the U.S. She advises, "You need to have somebody who actually understands the system and how it works, understands what the intended output is likely to be, and is willing to question that in the case of an error."

While the U.S. still has an advantage in today's information-centric ways of warfare, a more intelligent battlefield is coming. China has made its position clear—AI is the way of the future—and now it's up to the U.S. to respond and learn from history to avoid being relegated to it.

AMERICAN OFFSET STRATEGIES
What They Tell Us About the Future Ahead

by **Ken Wisian** and **Marla Rosner**

One of the most famous concepts in evolutionary biology is the Red Queen hypothesis. It takes its name from a passage in "Through the Looking Glass," in which the eponymous queen states that "it takes all the running you can do, to keep in the same place." Similarly, the hypothesis states that organisms must continually evolve and adapt not only to gain an advantage, but simply to survive in a constantly changing environment. In this way, evolution can be characterized as an arms race between predator and prey, with survival going to those who adapt the fastest.

This same principle holds true for human beings. Our biology may be relatively static—at least in terms of our lifespans—but our technology is evolving at an ever-increasing rate. Thus, failing to innovate does not mean staying in place, but falling behind.

The United States military is at a critical juncture wherein this principle is more relevant than ever. The advantage the U.S. has enjoyed from its lead in precision weaponry is eroding, as other countries have reached and even surpassed this level of technology. Furthermore, much of the previous strategies and technologies of the U.S. are not viable in the increasingly common tactics of guerilla warfare. Essentially, the U.S. no longer has a clear advantage over many potential opponents. This means that not only is American victory far less assured, but also conflict is far more likely to break out, as other powers are no longer deterred by the military prowess of the U.S.

This is not the first time the United States has found itself in this position. During the First and Second Offset Strategies, the U.S. recreated its military policies around new technologies to maintain superiority in combat. Many experts in and outside the Pentagon have already recognized that the time has come again to create a new strategy. In an address on the topic in 2016, former Deputy Secretary of Defense Robert O. Work stated that he has been contemplating a new defense strategy in earnest since 2012. Longtime rivals such as China and Russia have been rapidly developing their military technology, and the U.S. needs to do the same to stay ahead. It needs a new way forward—a new offset strategy to preserve its advantage—and artificial intelligence (AI) can provide that path.


Titan Nuclear Missile

F-22 Raptor

The First and Second Offset Strategies

An offset strategy attempts to maintain a military advantage over potential enemies in order to deter conflict from occurring and assure victory if it does. These strategies are typically achieved through an asymmetry in technological or operational prowess, and have been a cornerstone of Western military policy since the end of the Korean War. For example, the First Offset Strategy, developed in the early 1950s, saw America focusing its military spending on stockpiling nuclear weapons to maintain its asymmetrical standing relative to the U.S.S.R.

However, once the Soviet Union caught up to the U.S. in nuclear arms, this strategy became ineffective. Exhaustive research programs were then implemented around reconnaissance, surveillance, and target acquisition. Technology such as surveillance and targeting sub-systems armed U.S. forces with superior intelligence and situational awareness, allowing the U.S. and its NATO allies to wield a mastery of precision weaponry, stealth, and battlefield intelligence of a type the world had never seen. The operation and use of this combination of capabilities came to be known as the Second Offset Strategy.

The Need for a Third Offset Strategy

Unfortunately, nothing lasts forever, particularly when it comes to technological leads. In the 40 years since the Second Offset Strategy was first implemented, the rest of the world hasn't simply been sitting idle. Many other countries—most notably longtime rivals China and Russia, as well as smaller regional players and exporters of weapons and equipment—have developed comparable precision weaponry and intelligence capabilities. Former U.S. Defense Secretary Chuck Hagel explained in 2014 at the Reagan National Defense Forum that "while we spent over a decade focused on grinding stability operations, countries like Russia and China have been heavily investing in military modernization programs to blunt our military's technological edge."

In the 40 years since the Second Offset Strategy was first implemented, the rest of the world hasn't simply been sitting idle.

The West now faces a significantly reduced strategic advantage compared to what it once enjoyed from its lead in precision weapons. In fact, Western militaries are struggling to effectively combat the increasing prominence of guerilla and insurgent warfare.

The strategic position NATO and other Western powers are in now is eerily reminiscent of Germany's losing strategy from World War II: forces and units of ever-increasing technological sophistication fielded in ever-shrinking numbers. Every time a new generation of aircraft is built, the quality of the aircraft jumps, but the price tag jumps with it, meaning that fewer aircraft can be bought. Over 1,000 of the previous top-of-the-line aircraft, the F-15 Eagle, have been built at a cost of roughly $27.9 million per unit.

Conversely, the more recent F-22 Raptor is one of the world's most expensive warplanes, estimated to cost anywhere from $137 million to $678 million per unit, even before accounting for flying costs[1]. Between costs of building, maintenance, and upgrades, the United States produced only 187 operational units of its top air superiority fighter, and largely leaves the planes grounded. Rather than revers-

[1] *Axe, David. "Buyer's Remorse: How Much Has the F-22 Really Cost?" Wired. Conde Nast, 14 Dec. 2011. Web.*

F-35 Lightning

ing the trend, the latest new warplane—the F-35 Lightning II—is already costing even more than its predecessor.

This is just one prominent example of a troubling pattern in U.S. and U.K. military acquisition. With the constant reduction in the size of military forces—regardless of how technologically advanced those forces may be—it is becoming harder for NATO forces to muster the bare minimum numbers to cover the areas necessary to win an encounter. Even the number of personnel has shrunk drastically, a reduction that became evident during the prolonged wars in Iraq and Afghanistan, both of which strained the American forces' ability to rotate personnel in and out of combat.

The pilot shortage is expected to worsen due to massive increases projected in airline pilot hiring. The U.K. has arguably fared even worse, with their fast jet fighter fleet reduced by almost 80% since the end of the first Gulf War. General Raymond Odierno, the former U.S. Chief of Army Staff, recently expressed concerns over similarly significant reductions in the British Army, wondering if the U.K. could continue to fulfill its international obligations with such reductions in force.

This strategy—focusing on costly platforms at the expense of numbers—didn't work for Germany in WWII, one of many cautionary tales from history for the U.S. and U.K. Relying on numbers alone is a high-cost strategy, and not one the U.S. and U.K. should or would implement. But there is a threshold where reduced numbers translate into a severe strategic and tactical disadvantage, regardless of the technological sophistication of individual assets.

The good and the bad news is that the U.S. and its allied Western nations have been down this road before on at least two occasions in history—before the implementation of each previous offset strategy. The similarities in each case are striking: Before the First and Second Offset Strategies, as well as in the present, the U.S. faced disadvantages in sheer numbers, while longtime foreign rivals had caught up technologically.

What is needed now is exactly what was needed in the 1950s and 1970s: new technological innovations that will re-establish Western military superiority by giving forces an asymmetrical advantage over potential antagonists.

Artificial intelligence is almost certainly the answer to this quandary. In its current state, AI could transform the range of military operations in drastic favor of its earliest adopters. This transformation includes advanced Intelligence, Surveillance, and Reconnaissance through computer vision-aided machine learning algorithms and deep learning; AI-powered autonomous decision-making support; advanced networked sensors for integrated fire control; miniaturized high-powered computing capacity deployed at the "edge;" high-speed networks; offensive and defensive cyber capabilities; autonomous swarming; and cognitive analysis of sensor data.

Defense means always remaining on the leading edge of technology. Standing still in military affairs is nonexistent—there is only innovating or falling behind. Artificial intelligence holds the capability to completely change the landscape of warfare, and that transformation will happen with or without the United States. Machine learning and AI are the future. If the U.S. wants to revitalize its military strategy with a new offset strategy and renewed technological dominance, it needs to join in on the technological revolution.

THE COGNITEAM

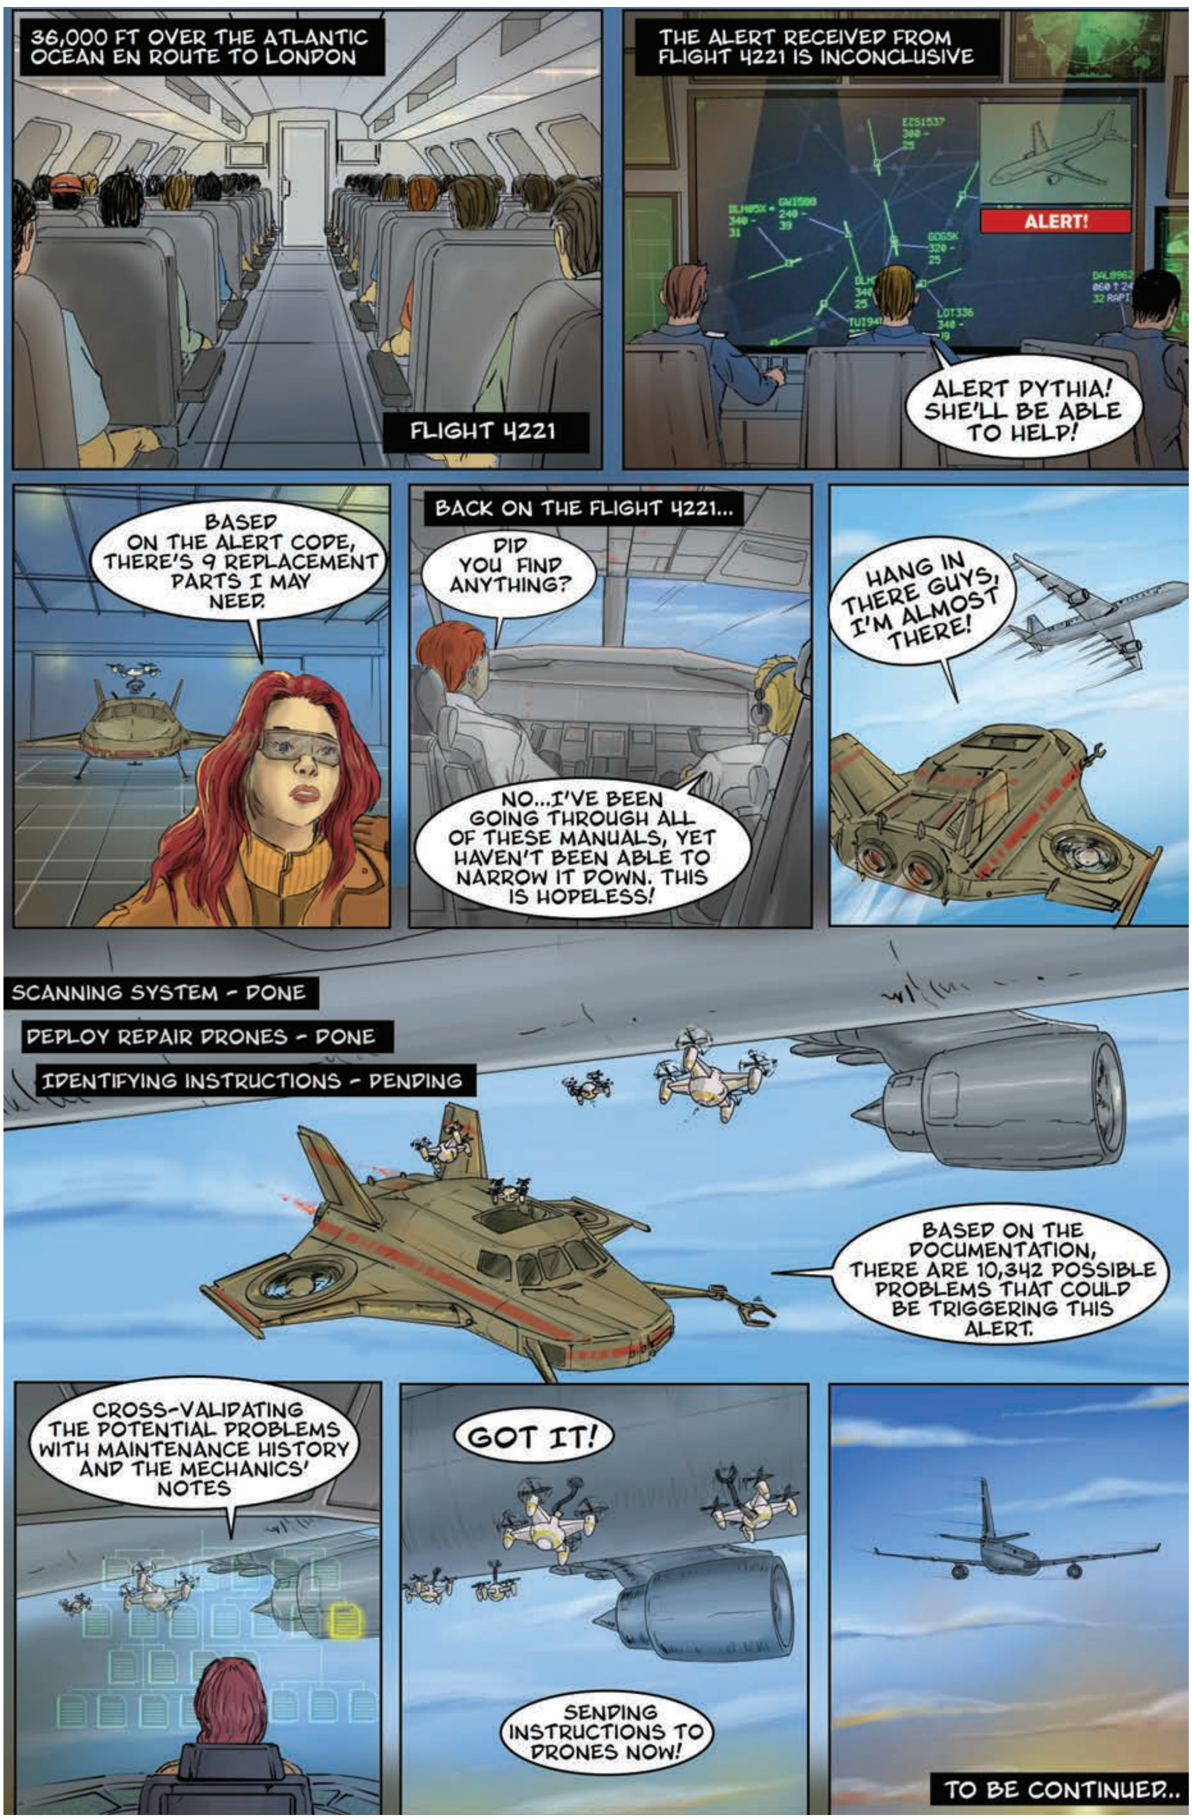

AI AND THE ART OF AIRCRAFT MAINTENANCE

by **Erin Russell**

It's 6:20 p.m. and you finally got to the front of the line. You've been at the airport for over four hours waiting on a plane to take you home, and it's your turn to ask a series of questions to a kind, attentive customer service agent who has been on the receiving end of an hour of verbal abuse. Unfortunately, you'll likely gain no new insights into your situation. It's the tail end of a long trip through multiple airports, and half a dozen different flights. You have an early presentation in the morning to brief the executive team on the meetings of the last week. You're tired, hungry, and ready to get home.

The plane that was supposed to take you home never made it. You overheard the airline crew talking about the minor maintenance issue which kept the plane grounded—a fault code which is a quick fix, but requires a very specific replacement part. Why is such a minor issue causing such major problems?

Flight delays are the bane of every traveler. However, as planes have essentially become flying computers and maintenance manuals have grown to the size of baby elephants, mechanics struggle to service aircraft within the 15- to 20-minute window allotted. There's a lot at stake in that time period: A small team of mechanics must diagnose a complex system (an Airbus 380 has approximately 4 million parts produced by over 1,500 companies) to determine if passengers must be re-routed or a part must be sent for in-depth analysis. The pressure of servicing aircraft within a tight schedule is a major challenge faced by carriers who are focused on safety first while maintaining customer expectations of speed and efficiency.

CHAD KARTCHNER
Director of Connected Aircraft Strategy
Honeywell

Chad Kartchner, Director of Connected Aircraft Strategy at Honeywell, is looking for solutions to the maintenance challenge. The multi-billion dollar aerospace division of the Fortune 100 conglomerate has a long history in the industry, employing around 40,000 people and keeping a strong focus on maintenance software. "Who isn't a Honeywell customer?" he explained. "Until I joined the company, I wasn't aware that every time I flew on a plane, something on that plane was probably manufactured and supported by Honeywell Aerospace."

Until I joined the company, I wasn't aware that every time I flew on a plane, something on that plane was probably manufactured and supported by Honeywell Aerospace."

Kartchner's dedication to innovation is what attracted him to Honeywell. As a teen, he spent a few formative years in Taiwan, where he was exposed to engineering as a profession. He completed his undergraduate degree in Electrical Engineering at BYU, then moved on to graduate studies, earning an MBA from Kellogg School of Management. "I made a conscious decision to try to focus on where I felt like I could add the most value, which were the things that I was passionate about," said Kartchner. "And for me, it was more about the process of integrating disparate pieces into an entire system."

This led him to the growing services division of the Honeywell. "We invest a lot of our time now looking at what our customers' experiences are and how we can make them better," Kartchner explained. "We can look at the end-to-end process, and use that as a driver for innovation and growth."

The Maintenance Dilemma

Aircraft maintenance proved to be an area particularly ripe for improvement. In one of his early exposures to aircraft maintenance, Kartchner learned about an approach sometimes called "shotgun maintenance." An experienced field engineer detailed this process to Kartchner: A mechanic has 30 minutes to fix an issue. After running diagnostic tests, a mechanic may make a judgment call and swap out the two or three parts most likely to have caused the problem, and hope for the best. This means a lot of wasted time and labor replacing parts that are perfectly fine. While it's possible to send the parts back to Honeywell to test for faults, the parts often stay at the hangar while a director follows the plane's status to see if the problem was fixed.

Kartchner spent a lot of time at hangars visiting with mechanics and directors of maintenance to understand this process. One suggested approach was to provide more training for the mechanics to diagnose the problem with higher accuracy. However, the complicated reality of the systems makes it difficult for mechanics to sufficiently analyze every possible issue.

Instead, inspired by IBM Watson's ability to help a doctor make a diagnosis, he began to explore the benefits of augmented systems. Rather than relying upon a mechanic's memory, a system with perfect recall would be able to answer specific questions to diagnose a problem and provide just-in-time training if necessary.

The AI Toolbox

Natural language processing (NLP) is the key to Honeywell's success. Honeywell has partnered with SparkCognition to use the company's DeepNLP® solution for centralizing maintenance information. Somewhat like a computer's search function on steroids, DeepNLP can ingest massive quantities of data and consolidate it into a database that is searchable using natural language—which means the user can have a more-or-less complete dialogue with the program. DeepNLP provides answers, not just links to keywords, and gives confidence scores as to the accuracy of those answers.

"SparkCognition has been a great partner," added Kartchner. "I've been able to frame up some use cases and some problems, and asked them to go tackle those technology challenges."

One significant advantage of NLP technology is the ability to preserve expertise in different forms—maintenance logs, handwritten notes, and data in tables. This aids maintenance across an enterprise, because it passes on "tribal knowledge" from shop books or other nontraditional sources. Operators can learn from others who have faced a similar issue.

"Once we create a system that is user-friendly enough that people want to use it, it's generally better," Kartchner adds. "I believe it will create this virtuous cycle, and be able to collect additional data that we can leverage to create even more value for the end user." He specifically points to repair reports as a potentially rich source of information for planning purposes.

As the industry continues to innovate, smarter planes are also helping schedule maintenance. Onboard systems log fault codes that were captured during the flight, and transmit that information to a Honeywell Global Data Center, which packages this as an email sent to a director of maintenance. This enables mechanics to be better prepared when the plane lands, both in terms of knowing what to expect (and doing additional research if needed) and planning the maintenance schedule.

In the future, more analysis can be done by an augmented system before an email even gets to the mechanic, who will know exactly what's wrong with the plane and know how to fix it. This means less time grounded by maintenance issues.

The Business Advantage

"AI is more expensive than paper manuals," Kartchner pointed out. "This technology does come at a price. But we are in the right place for trying this out is because of what's at stake." He references a recent EU law that requires passengers to be reimbursed for delays, which cost one airline $25M in six years.

Kartchner believes AI blends case-based and model-based problem-solving approaches to create the most effective maintenance system. "We are able to take the official records and create models from things like maintenance manuals and have a starting point that is really good," he explains. "Then we can augment those model-based approaches with knowledge that's gained over time. Using AI, we can make the system question and answer and back and forth, so you end up with not only the blending of the two approaches but the ability to go off-script and ask for additional information and get that information quickly."

"For this to work, it's all about scale," Kartchner elaborates. "We have to be able to automate the process of ingesting all of that information and learn quickly. The fact that we are doing that with great accuracy and not a lot of manual training is a tribute to the technology really getting to the place where what we are trying to do is feasible."

The Future of Flying

Kartchner doesn't see the future of maintenance in fancy gadgets—though he does think there's potential in hands-free augmented reality. He's more excited about enhancing user experience, like enabling dialogue with a computer to get assistance: "Being able to have a case and talk to a computer almost like a coach, and that coach has the ability to recall everything about that aircraft, and not only recall it but make smart recommendations based on all of that information."

The ultimate goal, for Kartchner? "It would be really cool to be flying somewhere and seeing a mechanic in a vest on a plane with the technology that we have created, and not only just using it but enjoying it."

As Kartchner and Honeywell continue to innovate and push the boundaries of possibility by applying AI technology to aviation maintenance issues, they're working toward making the grounded plane a thing of the past, and getting you home on schedule.

ARTIFICIAL INTELLIGENCE
Goes to Flight School

by **Jim Fitzgerald** and **Erin Russell**

A shortage of pilots in the Air Force has reached a margin of over 2,000 pilots in the coming years, a problem echoed by other branches of the military. Low recruiting numbers are compounded by declining retention of trained pilots, as civilian airline carriers (who are facing shortages of their own) have increased compensation to become more attractive options. In order to meet the country's defense needs, there's significant pressure to train pilots more efficiently and effectively.

On a fundamental level, aviation training in the Air Force, Navy, and Army hasn't changed much since the 1920s. While the equipment has become more sophisticated, the education methods remain reminiscent of machine assembly: Move each student along a conveyor belt of training using a building block approach, and discharge those who don't meet the criteria.

While this model ensures high-quality pilots, it also rejects candidates that may only need a little extra time in training beyond the standard "three strikes" policy. The government has already invested substantial time and money into these student pilots that attrite, so there are significant incentives to increase retention by both adding needed personnel and improving ROI.

As training methods are adjusted to align with the modern classroom, artificial intelligence (AI), virtual reality, and augmented reality (AR) have the potential to aid student pilots, instructors, and pilots fulfilling continuing education.

An Iterative Syllabus

AI can be harnessed to find insights into very large data sets on a periodic basis at affordable costs. A systematic, periodic flight syllabus review may reveal powerful suggestions for change. Inputs could include not only the scores of specific maneuvers on event grade sheets but also other data such as undergraduate major, weather on the training sortie, and days from last flight.

Given the bandwidth of the senior flight instructors, this syllabus review process is currently incomplete and infrequent. The cultural change of being willing to differentiate the syllabus within reasonable boundaries versus the cookie-cutter approach would be far more economical and produce better-trained pilots.

As training methods are adjusted to align with the modern classroom, artificial intelligence (AI), virtual reality, and augmented reality (AR) have the potential to aid student pilots, instructors, and pilots fulfilling continuing education.

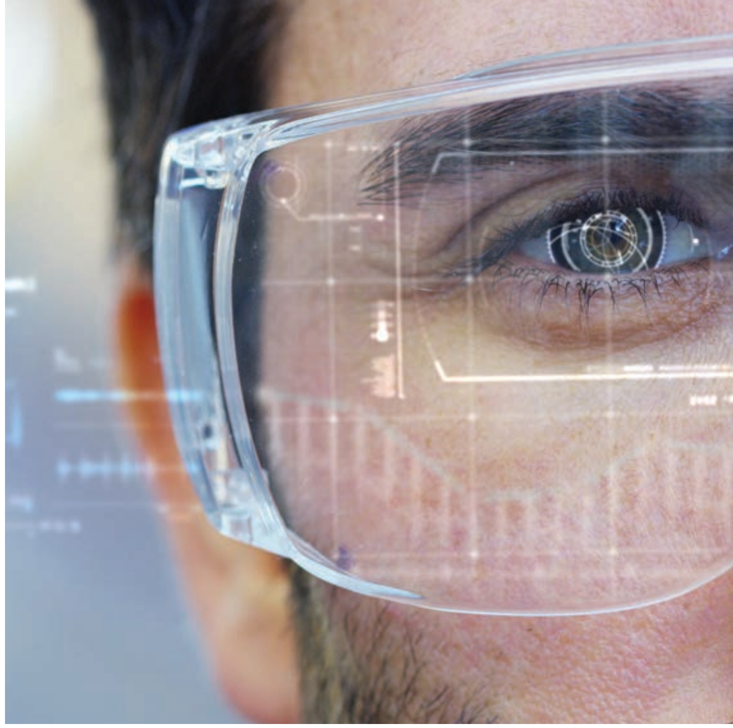

Personalized Training

Applying AI to the analysis of flight simulator data and individual flight training event data could yield even more awareness for the student in the training process. Currently, simulators rely primarily on human supervision, and analysis of simulator data is limited to post-event debriefs of the recorded mission. Thus, training could potentially suffer from instructor bias or not provide enough information for the student to improve upon. For example, a student could continually struggle to make timely decisions but not know which specific actions were being executed slower than normal.

Collection and analysis of biometric information from eye trackers and heart monitors could give a more complete picture of a trainee's performance. Because AI excels at collating information from different sources, it could show that, for example, the student spends too much time on an instrument panel scan or focuses on the incorrect sensor. With the available data, the instructor could spend more time helping the student prioritize their scan and ensuring they feel fully prepared before incorporating additional tasks.

There is also a large amount of individual flight training data that could provide key insights and thus faster learning to the flight students. In Navy pilot training, for example, each student completes hundreds of practice aircraft carrier landings at a runway prior to embarked aircraft qualifications. Using AI to quickly and routinely analyze throttle position and angle of attack during these practice sessions would add immediate value. This analysis would also invariably reduce the disqualification rate of these expensive carrier landing qualification training evolutions and produce a better trained naval aviator.

Augmented Reality

Augmented reality powered by AI could also be used to substitute or enhance ground-based simulator training. Specifically, AR mission rehearsals and "warm-up" flights could be used as refresher training on aircraft carriers where physical space is limited. As cockpit technology progresses toward sensor fusion, this will inevitably come to fruition and increase mission readiness with limited additional cost.

Looking Ahead: Continuing Education

After completing their training and earning their wings, all pilots must log hours on various topics to stay up to date. If AI were applied to create more personalized training, records could also be tracked over time, placing more emphasis on areas in need of improvement. A more robust tracking system could also provide clues on overall areas of training weakness and how to overcome such obstacles. This would be far more cost-effective than the current approaches to training and readiness.

What could this look like? Picture a chip that would stay with a pilot throughout their career (or lifetime), maintaining and analyzing training and proficiency records and eventually directing training. This would enable organizations to break the time-scheduled proficiency model and replace it with a truly as-needed paradigm.

AI will make pilot training more efficient and effective in the coming years. The rate of adoption and depth of application will directly correlate to cost savings and increased output in DoD pilot training. As a country, we must embrace AI to retain our global lead.

The OODA Loop: WHY TIMING IS EVERYTHING

by **Wendy R. Anderson, Amir Husain,** and **Marla Rosner**

Warfare is undergoing a shift of a historic magnitude. As countries invest more heavily in technology for the battlefield, countries without AI in their arsenal will find it increasingly difficult to compete.

It may seem like a stretch to claim that AI will soon become a prerequisite for military success, but the nature of how AI will reshape the battlefield makes this prediction all too real. While humans will continue to provide high-level input, machines will take over the decisions involved in planning and execution. With this new development brings a new type of conflict in which actions and reactions occur at near-instantaneous speeds, and the time required for the OODA loop—the decision cycle of orient, observe, decide, and act—will be reduced to nearly zero.

A fast, efficient decision-making process has long been recognized as crucial to the success of military operations. Accordingly, the concept of the OODA loop has been a key part of U.S. military strategy for roughly 70 years. Coined by United States Air Force Colonel John Boyd during the Korean War, this idea states that decision-making occurs in a four-step loop.

The first step, observation, involves an entity observing their surroundings and any occurring events or changes. For example, a pilot may see or hear an approaching plane.

In the second step, orientation, the entity processes the observed information, feeding it through their own previous knowledge and experience. In the orientation phase, the pilot will ask themselves such questions as "Is the approaching plane an enemy plane or a friendly non-combatant?" and "Do I have a clear shot?"

The third step, deciding, is when the entity settles on what they believe to be the best course of action, given what they have observed and analyzed, such as the pilot choosing to open fire on the other plane. The final stage of the loop, act, is taking the action that has been decided upon by the entity.

Boyd had observed that American F-86 aircraft were winning a majority of their dogfights against the Russian-made MiG-15s, despite the latter being stronger and more maneuverable aircraft. Despite their disadvantages, the F-86s provided a wider field of vision than their Russian counterparts, along with hydraulic controls that were easier and quicker to use. Boyd concluded that these advantages were key to the American victories, as they allowed the American pilots to shorten their OODA loop in the observing and acting phases. This strategic opportunity not only allowed the American pilots to act/react more quickly with better information, but also to outpace and thereby disrupt the OODA loops of their adversaries.

As Harry Hillaker, chief designer of the F-16, wrote:

The key is to obscure your intentions and make them unpredictable to your opponent while you simultaneously clarify his intentions. That is, operate at a faster tempo to generate rapidly changing conditions that inhibit your opponent from adapting or reacting to those changes and that suppress or destroy his awareness. Thus, a "hodge-podge" of confusion and disorder occur to cause him to over- or under-react to conditions or activities that appear to be uncertain, ambiguous, or incomprehensible.

In essence, the faster the OODA loop, the greater the advantage a military force can have over its opponents.

Currently, however, there are limitations on the OODA loop. Humans cannot go through these steps instantly. Observing one's surroundings, processing information, making decisions, and taking action require precious time, even under ideal conditions, so there is a limit on how short the OODA loop can become. Human processing and reaction time can vary based on outside factors such as fatigue, blood glucose levels, emotional state, and more. These factors can add significantly to the time it takes to complete an OODA loop and thus lead to poorer outcomes in decision-making. For example, one widely publicized study found that judges' likelihood of granting parole correlated most strongly with how long it had been since they had last eaten[1]. Furthermore, a human can typically hold, at most, only a few variables at one time in conscious thought. Choosing focus can often affect survivability: How many pilots have been so focused on getting their own "kill" that they were shot down by an unseen attacker? Despite how essential human cognition, reaction time, and decision-making are to combat, these processes are all imperfect and easily compromised.

Machines suffer from none of these flaws. They can process information and react to it nearly instantaneously. Machines are also less easily impacted by outside factors. AI can take in and process wider ranges of information than any human being ever could—holding thousands of variables in "conscious" thought at once—and then use that information to make to make logical and tactical decisions that they can execute almost instantly. Furthermore, machine intelligence can take the form of easily replicated software and run on inexpensive hardware, allowing deployment at scales sufficient to essentially enable an infinite supply of tactical, operational, and strategic decision-making.

Imagine an AI-powered autonomous drone rather than a human pilot in the dogfight described above. The observation step is about the same as a human; the autonomous drone uses different senses to "observe" the approaching plane, but the basic idea is the same. The orientation and decision phases, however, look very different. Rather than the limited information that a human may know and recall at any given time, the drone, in its orientation phase, has access to encyclopedic quantities of knowledge that it can analyze almost instantly. Using this analysis, it can then rapidly but accurately select the best course of action. In this way, the drone can move from noticing a plane to concluding it is an enemy that poses a threat to deciding it needs to be shot down, all with a speed beyond human capabilities. Its focus can be as divided as necessary. However, if another enemy combatant approached from behind while the drone was attempting to fire on the first plane, it could undergo another OODA loop to take evasive maneuvers without hesitation or distraction.

Many are still uneasy at the prospect of allowing AI to make the decision to fire. It is true that AI can speed up the OODA loop, even if introduced only partially into the decision-making process. It can augment intelligence by analyzing a vast quantity of data, boiling it down, and presenting it to humans in a simplified format. In this way, AI can give humans all the information they need to make optimal decisions while still leaving ultimate agency in their hands. Due to the speed of machines and the volume of information they can handle, the fastest—and therefore necessary—option will eventually be to remove humans from the loop entirely.

AI is not just a tactical advantage; sooner or later, it will become a necessity. Nations such as China, Russia, and many others are already investing heavily in AI. These countries will be making use of machine learning to tighten their OODA loops, regardless of whether their rivals do the same. Greater reliance on AI software is the only way to survive this new paradigm of warfare.

[1]Danziger, Shai, Jonathan Levav, and Liora Avnaim-Pesso. "Extraneous Factors in Judicial Decisions." Proceedings of the National Academy of Sciences of the United States of America 108.17 (2011).

BOEING LOOKS BEYOND THE HORIZON IN ITS SECOND CENTURY

by **Erin Russell**

For the first time in U.S. industrial and financial markets, the top five companies by market cap are all technology companies: Apple, Alphabet, Microsoft, Amazon, and Facebook. As internet pioneer Marc Andreessen remarked over six years ago, "software is eating the world." And as SparkCognition's President and CEO, Amir Husain, likes to add, "AI is eating software."

In an age when all relevant software companies are buying, integrating, and/or developing artificial intelligence (AI) technologies to enhance their solutions, Husain is right to point out that our future belongs to those who lead in AI innovation. We are living in the age of AI, and it will define the technological future across all industries.

It is a tough feat for traditional industry players to think and act like a nimble startup. While there is no guaranteed strategy for how to ultimately be successful in the age of AI (see Husain's book, The Sentient Machine, for the closest thing to a blueprint), there are examples from companies that have succeeded most throughout the transition. And there is at least one household name that's doing it well.

STEVE NORDLUND
Vice President
Boeing Horizon X

> *When we turned 100 years old, our chairman and CEO, Dennis Muilenburg, talked about the startup of the second century. We are all entrepreneurs of the second century of Boeing and therefore the second century of aviation and aerospace.*

Flying High

Shares of Boeing are up more than 71% so far in 2017, putting the company far ahead of second-ranked Apple, with a 46% gain. If software is eating the world, and if AI is eating software, how does an aerospace and defense company beat out the most financially liquid organization in America? By becoming an AI company.

Boeing's success in 2017 can be attributed to many factors—defense stocks shining from Trump's victory, North Korea's nuclear threats, recent mergers, Boeing's new 737 MAX and 787 Dreamliner commercial jets, demand for its Apache helicopters and missile defense systems—but the company's emphasis on AI is a much less reported story that has some insiders excited about the future.

In April of this year, InverstorPlace published an article about the top 10 AI stocks to watch. Amid the list's usual suspects (IBM, Microsoft, Amazon, GE, Google, Facebook, Apple), one company stood out. The article named Boeing as a top AI stock to watch, specifically calling out the venture capital arm, Boeing HorizonX, and its investment in AI leader SparkCognition.

While the combination of Boeing's commercial and military leadership has some analysts expecting profits to rise nearly 40% this year and average an annual growth rate of almost 20% for the next few years, it's the company's vision of an AI-enabled future that has company veterans and industry insiders most excited about Boeing.

The Second Century of Innovation

Since its inception, Boeing has shaped aviation and aerospace history. In an era of buzzy, agile startups, Boeing, which recently turned 101 years old, continues to lead the field in innovation by developing and strategically investing in promising technology and attracting top talent to the team.

Dr. Greg Hyslop plays a large role in turning innovations into reality. As Chief Technology Officer of The Boeing Company and Senior Vice President of Boeing Engineering, Test & Technology, Hyslop oversees 45,000 engineers worldwide.

So what do the next 100 years hold for Boeing? Hyslop sees artificial intelligence playing a major role in Boeing's future, generalizing, "The first century of aerospace was defined by the aerodynamicists, and the second century is going to be defined by chemists and mathematicians. It's going to be about materials, and it's going to be about algorithms."

Hyslop's colleague responsible for ensuring Boeing is at the forefront of materials and algorithms is Steve Nordlund, Vice President of HorizonX, Boeing's venture arm.

"I was one of those kids that kind of hung on the fence watching the planes take off and land," recalls Nordlund. "But I realized there was a lot more to the aviation industry than just the transportation end of it."

At first Nordlund followed his childhood dreams, attending Embry-Riddle Aeronautical University in pur-

DR. GREG HYSLOP
Chief Technology Officer and
Senior Vice President of Boeing
Engineering, Test & Technology
Boeing

> *The first century of aerospace was defined by the aerodynamicists, and the second century is going to be defined by chemists and mathematicians. It's going to be about materials, and it's going to be about algorithms.*

suit of a career as a pilot, but his priorities shifted after starting his studies. He moved into the tech industry and earned his MBA from the University of Florida. His talents aligned when he helped kickstart a startup with a colleague to build aviation drones, soon becoming one of the fastest-growing companies in the U.S. Boeing acquired the startup in 2008 and Nordlund came along for the ride.

"Working at Boeing is like being a kid in a candy store," Nordlund explains. "Every day, we are connecting people via airplanes and satellites, protecting people with our defense products and services, exploring the mysteries of the universe, and looking to inspire future generations of innovators."

Boeing HorizonX keeps its sights set beyond the horizon and its focus on three areas: technology enablers, new market opportunities, and disruption in aerospace. Nordlund notes that Boeing HorizonX is not driven as much by the financial return of its portfolio but rather the ability to bring innovation into the company.

"Boeing's always been really good at internal innovation. We literally have rocket scientists," he elaborates (in fact, Hyslop is a member of the Aeronautics Committee of the NASA Advisory Council). "By engaging with and investing in startups, we are building on the depth and breadth of our knowledge and expertise within Boeing. We want to tap into the innovation that exists in these companies to help shape the future of Boeing and the future of aerospace."

Boeing HorizonX leads with customer focus by looking at the business problem to solve, the market opportunity, and the strategy to get there, ensuring that the right technology solutions follow. Nordlund believes that if you're not solving a problem, if you're not making a difference, then it's really not going to penetrate into the marketplace."

"A lot of people ask me what an entrepreneur like me is doing at a 100-year-old company like Boeing," he adds. "And the most important data point in that question is the 100 years. When we turned 100 years old, our President and CEO, Dennis Muilenburg, talked about the startup of the second century. We are all entrepreneurs when it comes to shaping Boeing's second century—and with it, the second century of aviation and aerospace."

Autonomy is of particular interest to Boeing and the aerospace industry. Boeing HorizonX is exploring emerging commercial autonomy markets, including urban mobility. Nordlund sees promise in AI when it comes to autonomy: "The power of AI, tied into the autonomy and the autonomous operations and the machine learning, starts to take it to the next level. I think that's going to break open some new market opportunities for us and solve a lot of customer problems."

Hyslop agrees there is significant value in automation, primarily for safety reasons. Pointing to the example of a robot inserting fasteners into the fuselage, he clarifies, "The jobs we want to automate are the jobs that are dull, dirty, and dangerous." With the growing market for planes, factories must respond to increasing demands while maintaining the same level of quality. "That is really the driving force behind the advent of more automation and hence more an application for artificial intelligence in our factories," he adds.

AI in the Sky and the AI Race

Prior to leading the Boeing HorizonX team, Nordlund oversaw strategy for Boeing's Defense, Space & Security business unit. So he took note when Air Force Lieutenant General Jack Shanahan recently declared, "The Department of Defense should never buy another weapons system for the rest of its natural life without artificial intelligence baked into it."

Nordlund acknowledges that warfare is shifting, and AI is at the core of technologies for wearables, maintenance, and training that fundamentally impact the field. "It's important to focus and invest in these emerging technologies like AI to ensure that we're building up the strongest capabilities and infrastructure to enhance the defense of

our country and allies," he says.

Hyslop echoes these sentiments, adding, "There is so much investment pouring into this technology around the world, it's inevitable that it's going to be used in other defense forces. So it's important for us as a nation, I think, to understand how we're going to use it in our defense."

Hyslop references the most promising areas for AI as precision-guided munitions ("there's a real possibility that AI, just as precision-guided munitions, could significantly reduce the amount of collateral damage") and overall better decision-making ("We hear people talk about the fog of war...with AI-based systems, you've got access to sensor data, more sensor data than the human being has access to").

So how can the U.S. keep a competitive advantage with so much investment happening internationally? "There is clearly a role government can play in this, but we can't look to the government to outspend the rest of the world," says Hyslop. He adds that the government can make regulations to enhance the deployment of autonomous systems while also ensuring their safety.

Nordlund sees a few opportunities for the U.S. to keep a competitive advantage. The first is to keep investing in the right technologies, politics aside. Secondly, he hopes to make the acquisition system less bureaucratic, allowing competition to make the field stronger. (He notes that other countries are using e-commerce, allowing for quick access to new technologies). Finally, similar to Boeing HorizonX, he advises looking to the commercial space for innovative ideas: "I think the collaboration that can happen between government and industry could prove valuable as we prepare to address future threats."

Nordlund understands the integration of AI into military applications will require a thoughtful approach, but he doesn't necessarily see this as a unique challenge. "With any technology, you have to realize that both good and bad can come out of it. The focus should be on how we further the benefits," he says. "That takes leadership, which is where Boeing plays a role."

Bringing It Home

Nordlund's favorite part about working in the industry is getting off an airplane and seeing excited families waiting for arrivals. This passion also shows when asked about the most compelling technology he's worked with: the ScanEagle unmanned system.

"If we hadn't innovated around the business model, that technology may never have seen the battlefield," he says. "We were able to deploy an unmanned aerial vehicle that allowed for intelligence, surveillance, and reconnaissance work that led to support for our Marines and soldiers on the ground and our sailors on the waters. And I know that we ended up saving lives."

Both Hyslop and Nordlund emphasized excitement that the industry is at an inflection point—and disruption is coming. Nordlund sees it impacting all aspects of Boeing: "How we connect economies and people, how we protect the economies and homelands, how we explore space and beyond, and how we inspire young people to be the future innovators for the world. We are at the cusp of a revolution that's going to happen over the next couple of decades, and we're right here at the core of it."

Personally, Hyslop is looking forward to a world where terrestrial transportation is no longer the norm. "Aerospace changed the world when it started as an industry 100 years ago," he says, "and I think it's primed to enter in another golden age. Aerospace has always been very competitive, but it has changed the world."

It's clear Boeing isn't resting on its laurels as they enter their second century. Both Nordlund and Hyslop represent the culture at Boeing: lead and innovate. As for the next hundred years, Boeing intends is to be at the forefront of another golden age of aerospace. Hyslop is excited and optimistic: "Aerospace, at its root, has always been very innovative. It has changed the world. And we're in that phase again."

It is more than a revolution in military affairs;
it is a revolution in human affairs.

On Hyperwar

by **General John R. Allen, USMC (Ret.)** and **Amir Husain**

2 January 2018: Abandon Ship

The battle damage was devastating, and constituted the beginning of what the United States soon would discover was a widespread, strategic attack. The guided-missile destroyer had not recognized that its systems were under cyber attack before the situation turned kinetic.

The speed of the attack quickly overwhelmed the ship's combat systems. New developments were occurring in seconds or less. Before anyone could even react, the battle was over.

The captain had survived, but he was severely wounded, as were many crew members. Fires were burning out of control, and the ship was listing badly from flooding. Evidently, the autonomous platforms knew exactly where to strike the ship to both maximize damage and reduce the chances of survivability. With his capacity to command the ship now seriously compromised and the flooding out of control, the captain did what no U.S. skipper had done for generations—he issued the order to abandon ship.

The Coming Transformation

On only a few occasions has history witnessed fundamentally transformative changes in the way war is waged, enabled by technological developments. The employment of cavalry and the advent of the rifled musket are a few examples. Another such shift is coming to the field of battle, and those who are unprepared will fare poorly.

This time, the shift will be an innovation that warfare has never seen before: the minimization of human decision-making in most processes traditionally required to wage war. In this coming age of hyperwar, we will see humans largely providing broad, high-level inputs. Machines will do the planning, executing, and adapting to the reality of the mission, and take on the burden of thousands of individual decisions with no additional input.

Explaining Hyperwar

What makes this new form of warfare unique is the unparalleled speed enabled by automating decision-making and leveraging artificial intelligence. The implications of these developments are numerous and game-changing.

First of all, while human decision-making is potent, it also has limitations in speed, attention, and diligence. Machines do not suffer from these limitations, and machine intelligence can be deployed at scales sufficient to enable an infinite supply of tactical, operational, and strategic decision-making.

Secondly, AI allows for far better coordination and concurrency of action. With machine-based decision-making, sensors, soldiers, and weapons can be coordinated instantaneously, enabling the rapid massing of forces and the execution of kinetic action and subsequent dispersal. This coordination will far outpace what can be done under human direction.

Thirdly, AI can simplify logistics by enabling robotic soldiers and autonomous drones equipped with synthetic intelligence. The needs of robotic soldiers will not be as varied as those of a human soldier, nor will these machines be as indispensable as human lives.

Finally, AI technology will enable groundbreaking changes in training. Natural language processing systems can ingest hundreds of thousands of manuals and guides to augment human operators. AI systems can be trained via simulators rather than real-world experience. The system that evolves the best-performing neural network can then instantly transfer its training and knowledge to as many other systems as needed.

28 May 2027: An Autonomous Defense Rises

The artificially intelligent cyber defense system in the guided-missile destroyer was the first to detect what appeared to be an attempt at a major cyber intrusion. The initial attack and successful defense occurred within microseconds. The ship then detected a massive incoming swarm attack and forwarded threat information to the rest of the fleet, enabling other units to prepare for an impending attack.

The captain moved quickly, donning the augmented reality headgear and gauntlets to assimilate and react to the totality and complexity of the battle he was about to lead. With a sweep of his hand in virtual reality, he initiated the anti-swarm batteries. In that instant, naval warfare changed forever.

Hours later, after checking diagnostics that showed the health of his ship and crew, the captain reflected on the engagement. The attack had come seemingly from nowhere. The cyber defense system had detected the initial intrusion, and not only had it protected the ship, but it also had reasoned the attack was a precursor to something larger. This hypothesis had been formed, researched, and validated in less than a second. Within 10 seconds, the ship initiated battle stations on its own, and the captain donned his augmented reality ensemble. The entire battle had unfolded and ended in minutes.

AI systems had foiled a coordinated, complex cyber and autonomous swarm attack. The captain was struck by the realization that they had risked the ship at nearly every point where human actions and decisions were required. Though he was a master of the combat systems of the USS Infinity (DDG-500), he had just experienced the near mind-numbing speeds of AI-driven warfare. He had become the first U.S. commander to fight in the environment of hyperwar.

Is This a Revolution in Military Affairs?

The scenarios and discussion here provide a window into only a few ways in which AI will fuel the next great shift in warfare. The hyperwar these technologies will enable is a new paradigm we need to plan for. Near-peer opponents are already investing heavily in these technologies and have obtained operational AI-powered weapon systems.

If we are poised at the edge of hyperwar, we must explore the changes necessary to adapt. This situation will require understanding the moral dimensions of these advances, educating a new generation of leaders, and developing the AI-powered analytical systems and autonomous weapons platforms. The mental, moral, and physical challenges of hyperwar demand analysis and a searching conversation. Other nations are moving forward aggressively in this area. We must make the strategic investments both to be prepared to wage hyperwar and to prevent us from being surprised by it.

THE NEW REALITY OF CYBER THREATS

by **Erin Russell**

We live in a world dependent on critical infrastructure—power plants, electric lines, water treatment facilities, and fiber optic networks. According to CyberX, at least one third of industrial sites overseeing the centralized operations of this infrastructure are connected to the internet.

This interconnectedness has enabled next-generation technologies in remote control, automation, data collection, and predictive and prescriptive analytics, to name just a few. However, it also opens up vulnerabilities in security for those who know how to exploit gaps in cyber protection, and today, much of our critical infrastructure is under attack.

"There's a rapid uptick in people using various types of exploits to try to gain access to the private sector systems," explains Philip Chertoff, research fellow in the cyber program at the GLOBSEC Policy. "These attackers are motivated by a whole number of factors, whether it's espionage, which is frequently in the news; criminal reasons, as we saw with the ransomware attack at the NHS in the UK; to even what is frequently called script kiddies, which is basically doing it for some kind of ego or personal aggrandizement."

GLOBSEC is an international think tank focused on shaping global debates through research activities and connecting key experts on foreign and security policy. Their Cyber Resilience Program is developing novel areas of research on cybersecurity, ranging from policy recommendations for NATO and EU agencies on critical infrastructure protection and information exchange to security implications of emerging cyber threats like data manipulation. Philip Chertoff is one of the experts shaping those policy recommendations and the global cybersecurity discussion.

"There's been a democratization of cyber attacks so that more and more people have the ability to find tools online to attack systems," Chertoff continues. "And I think companies are especially now harder-pressed to manage an increasingly large attack surface."

Systemic Vulnerabilities

Our critical infrastructure was built well before the internet existed, before there was any concept of cybersecurity. Even the internet itself was developed without security in mind.

"The internet is built on an earlier world where people trusted each other," says Sean Smith, professor of computer science at Dartmouth College. "It was not built with security in mind."

In addition to his role as professor of computer science, Smith is Director of Dartmouth's Institute for Security, Technology and Society. He knows cybersecurity and his opinion of cyber disaster is a bit more nuanced than others. "A lot of my colleagues keep saying, 'We're going to have a cyber Pearl Harbor,' where there's going to be some sudden, massive attack from an enemy. Should we be worried about that or should we worry about a cyber

PHILIP CHERTOFF
Cybersecurity Research Fellow, GLOBESEC
University of Chicago

There's been a democratization of cyber attacks so that more and more people have the ability to find tools online to attack systems.

[1]*https://www.wired.com/story/russian-hackers-attack-ukraine/*

Love Canal?" [Referencing the infamous New York town slowly poisoned by toxic waste seepage.]

Along those lines, Smith sees three major risks with IoT: holes in an interface, dependence on legacy systems, and the connection to the physical world.

"As we scale from computers that look like computers to ones that are embedded in everything, and as we conduct that ramp-up without thinking about how we build these interfaces (which always seem to have holes), then we, as the society, are setting ourselves up for a really big problem," Smith warns. While patching vulnerabilities in a computer or on a phone is relatively automatic, Smith expresses concern that people would remember to patch devices connected to a thermostat, or even a train.

Smith isn't speaking in cybersecurity platitudes. Our infrastructure is vulnerable, starting with electric meters. For the 70 million (and counting) Americans with smart meters, security analysts have pointed to the radios embedded in the software of the machines which could potentially be reprogrammed to transmit in other frequencies, and result in a downed cell phone network.

Legacy systems are also something to consider. A company's reliance on an old system can obviously be problematic, as proven when planes were grounded in November of 2015 at Orly, a French airport running on Windows 3.1. If current systems are not built with security in mind, our future infrastructure will continue to be unstable.

"We're building tomorrow's legacy systems today," Smith explains. "The companies that created the stuff might not be around tomorrow."

The connection of cyber and physical worlds means significant danger from the anonymous internet. Smith specifically references Shodan, a search engine to look for embedded systems. "Every time somebody looks at it, they find interfaces that should never have been exposed, yet are now exposed on the open net," Smith explains. "Dan Tentler found a steel plant with a vat of something that was in the order of 1,000°C. It doesn't really matter what it is. If somebody on the open Internet without authentication can start controlling that vat, you've got a problem."

"Cybersecurity is a Process Issue"

Chertoff is particularly concerned cybersecurity has been treated as an IT issue, pointing out that the main attack vector for malicious actors tends to be users. "Cybersecurity is a process issue, not a technical issue," he clarifies. "For a lot of companies, this means that they're going to have to go through HR changes and reorganize how they actually operate their business in order to be more secure." This means going beyond simply adding antivirus software and rethinking processes that put sensitive data at risk.

Chertoff's main focus is around cybersecurity for critical infrastructure, which can include everything from a water cleaning plant to the electric grid. "There's a huge

> *What I hope for from AI is to start to minimize manual processes that analysts have to go through, a lot of this pattern recognition work. Then AI is handling the bulk of the data analysis when it comes to security incidents and events, and the analyst can really think primarily about what threats the organization is taking holistically.*

part to the national supply chain that people don't see, but which attackers will hone in on as a vulnerable point," he explains. "These types of attacks are meant to undermine the legitimacy of the government institutions, because one of the fundamental parts and responsibilities of government is to deliver critical services to its citizens."

Wired's June cover story, shedding light on Russia using the Ukraine as a training ground for larger attacks, has been particularly worrisome for those in the cybersecurity community. "The fact that it's only gone so far in previous attacks is not a question of capability," Chertoff warns. "It's a question of will."

The Secure Future

Though systemic vulnerabilities are frightening, there isn't an easy way forward. "AI can enhance human creativity," adds Michael Horowitz, professor of political science at the University of Pennsylvania. "And that's a double-edged sword." Unfortunately, many organizations appear to be turning away from digitization because they can't secure what they're digitizing. "It's a real shame because obviously digitization offers a lot of significant economic benefits to corporations," Chertoff adds. Chertoff sees promise in AI's ability to aid situational awareness, providing "a much broader perspective on possible vulnerabilities."

What does AI have to do with cybersecurity? At this point in the game, they're inseparable. It's estimated there are over 100,000 zero-day cyber attacks daily. The sheer volume of the cyberthreat landscape means it's impossible for human security analysts, regardless of capacity and experience, to keep up with the sophistication and adaptability of those cyberthreats.

Amongst a limitless potential of applications, artificial intelligence technologies are showing efficacy in the detection of, and protection against, a new generation of cyberthreats.

"What I hope for from AI is to start to minimize manual processes that analysts have to go through, a lot of this pattern recognition work," adds Chertoff. "Then AI is handling the bulk of the data analysis when it comes to security incidents and events, and the analyst can really think primarily about what threats the organization is taking holistically."

Chertoff also recognizes the benefits of having a plan of defense in case of a critical attack. "The National Energy Reliability Commission has been planning scenarios like this for a while, especially regarding natural disasters but increasingly about cyber attacks," he notes. "I say that actually the U.S. energy grid has gotten especially good at planning for this type of thing." Decentralization of a system is another way to increase resiliency, though Chertoff notes it could become a weakness if the individual nodes are not secured appropriately.

These threats emphasize the importance of developing an investment strategy and policies related to AI, as it undoubtedly becomes part of our future. Forming a path forward that minimizes vulnerabilities, provides a plan for critical infrastructure, and encourages development of AI will ensure that the technology is not a source of fear. As Horowitz notes, "AI is not a weapon. It will have an impact much broader than that."