leidos

Q&AI

AI holds the power to do incredible things. As you're about to read, we help our customers use AI to fight cancer, improve combat readiness, modernize the energy grid, and so much more.

But while AI has been portrayed as a magic wand, it can't fix everything. When you look beyond the hype, you'll find AI is effective when used in smart solutions to specific, data-intensive problems. Effective adoption of AI, like any other technology, requires narrowing in on practical applications that provide real value to you, our customers.

To illustrate this point, we've selected several use cases that demonstrate effective AI solutions and how they were designed. We also selected a number of other important topics surrounding the field of AI, from ethics to the evolving state of the workforce. We interviewed several of our best and brightest minds in data science to bring you their perspectives from the ground floor of the AI revolution.

AI is arguably the top emerging technology our customers are eager to understand and use. But while AI can do amazing things, it must be implemented with care. Our goal is to help you look beyond the hype and adopt AI in ways that add real value to your mission. We call this the "Leidos way," and we hope each of the following conversations provide valuable insights that help tell this story.

Enjoy!

Ron Keesing AI Director

Contents

Q&AI

How does AI reshape the energy grid?

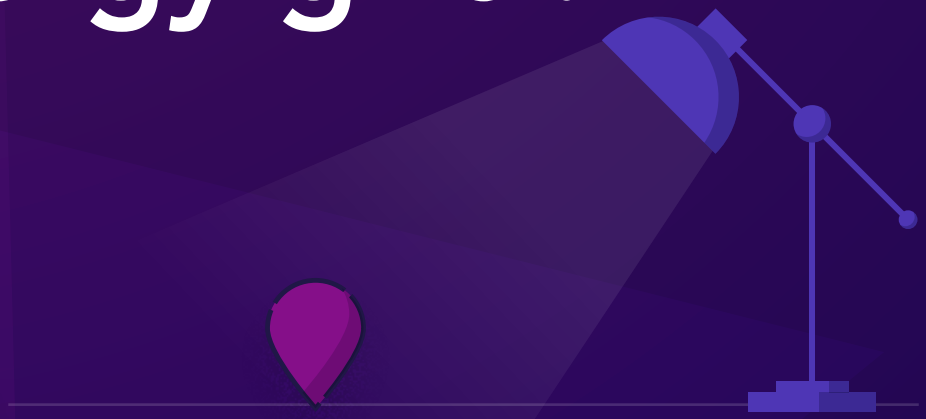

"We all want clean energy. We all want wind, solar, and distributed energy, and we all want cheap and green and reliable. But this is fundamentally impossible without understanding the current state of the grid."

Josh Wepman, CTO Commercial Energy

Powering the world requires an enormous amount of field infrastructure. In the U.S. alone, this includes roughly 180 million poles strung together by more than 5 million miles of conductor cable, a sweeping and interconnected grid that has expanded and evolved over time. Understanding the state of this grid is essential in order to modernize it effectively, but it's a remarkably difficult and expensive thing to do. Now AI is changing that. To learn how, we welcome Josh Wepman, Leidos Vice President and CTO for Commercial Energy Solutions. Wepman's team has created an AI-powered solution called Infrastructure Insight that gives utility companies a faster, cheaper, and better way to take inventory of the energy grid.

Q: First, why is it so important to modernize the energy grid?

Josh: To make it cleaner, more reliable, and more affordable. We need to build a 21st century energy delivery and management system, but the only way to do that is to get our arms around what we've already built over the last hundred years to figure out the right way to change it. We can use AI to do that. We can use the cloud to do that. We can use processing and machine learning to do that. But as long as the inputs are garbage, the outputs won't be useful, and it will reduce progress. What we need is a more effective way to understand the grid as it is today so we can tackle the most exciting challenges of the future: decarbonizing the planet, democratizing energy production, and letting everybody have the relationship with energy they'd like to have.

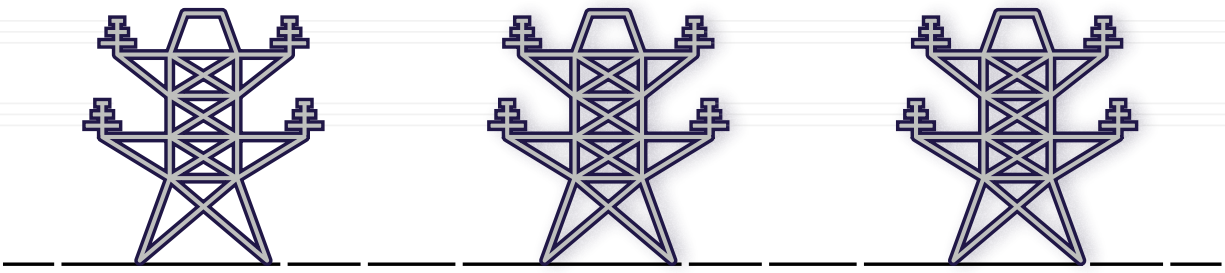

Q: What makes it so difficult for utility companies to account for their own field infrastructure, and why is it so important for this to change?

Josh: Power companies have massive spatial infrastructure out in the field. They've got lines, poles, and other equipment spanning huge geographies. As you think about the evolution of the energy marketplace, you can think of it in two dimensions. If you think of an x-axis as a shift from analog to digital, all of those assets across all of this infrastructure are moving from analog assets to digitally empowered assets. They have microprocessors. They've got communications assets to tell us about their state and performance. We're getting data rich, and we know much more now than we've ever known before.

If you think about the y-axis, we're moving from a procedural operating model to a more analytical and data driven model. We've always run the power grid based on a whole bunch of assumptions about what's happening out there. We run calculations and we try to estimate as best we can, but it's always been very procedural, not very data-driven. What to build, where to build it, how big should it be, what risk does it represent over the next 20 years; it all depended on a narrow and generally certain view of the future. Now that we've gone across that x-axis from low data to high data, we're beginning to make more progress migrating from procedural grid operations to data-driven analytical modes.

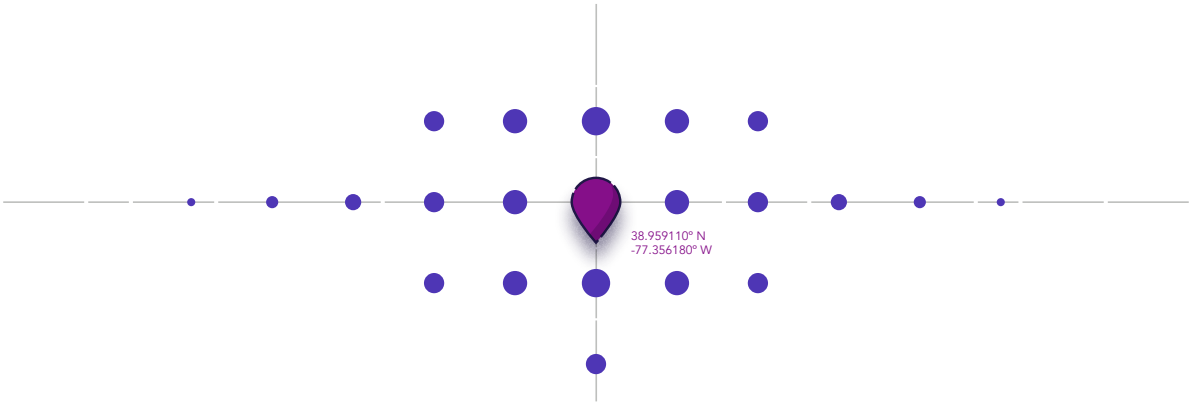

One of the things that's really important about making that leap from procedural to analytical is better input data. The key challenge that our customers face is that an 'approximate' understanding of their field infrastructure has always been good enough for their procedural analyses. They've always worked in rough approximations. But now that they've got all this measurement data, actually knowing what is out in the field, where it is, and how it relates to each other is more important than ever before.

Q: How does AI now help us do that?

Josh: This is what's called a conflation program. Conflation means I'm reconciling what's in my digital system with what's really out in the field. What we recognized was that if we flew LIDAR missions to find out where the macro infrastructure are, we can use machine learning and feature extraction to say 'these are poles and these are lines.' And we're talking about millions of poles. So for the first time ever we actually know where all these poles are in a consistent, latitude-longitude scheme. And for a great deal of these locations there are images from Google Street View, so now we can query Google Street View, and instead of trying to train machines to look absolutely everywhere for the small assets—the electrically connected devices that live on each individual pole and are important to actually delivering quality power to your home—we now know precisely where to look. That reduces the problem space dramatically.

We were able to take our data science and engineering prowess in this company, our electrical engineering design prowess, and our mission software prowess, and put all those things together. We took these images, used engineers to identify poles, cross arms, transformers, fuses, reclosers, and street lights. And we literally brought our engineers together and labeled thousands of example images to establish ground truth. And we had our data science and engineering team use TensorFlow to build training images to show what all of these assets are. We trained and we tested, and we trained and we tested to validate. And what we found was that we became very, very good at teaching machine vision and convolutional neural networks to take in millions of images of poles and produce inventories of those poles without ever sending a person out into the field.

Q: When it came to working with AI tools to develop this solution, what was your critical success factor?

Josh: Intellectual curiosity. We said 'philosophically, this ought to work. We ought to be able to train the images. We ought to be able to train a classifier.' Our lead developer spent the better part of a weekend just being curious. What he found was there are a thousand ways to do this, but boy do a lot of people rave about TensorFlow and how easy it is to get started. So he implemented it, wrote a little bit of Python glue code, and tried it out. He labeled a few images, processed them, and found that it worked really well. A shockingly low number of inputs produced surprisingly good outputs. If you're a developer, if you understand how to read code, and you're interested in trying things out, the friction of getting started and becoming dangerous is shockingly low. There's nothing about what we did that is fundamentally not solvable by intellectually curious people, which means we have to continue to push and evolve.

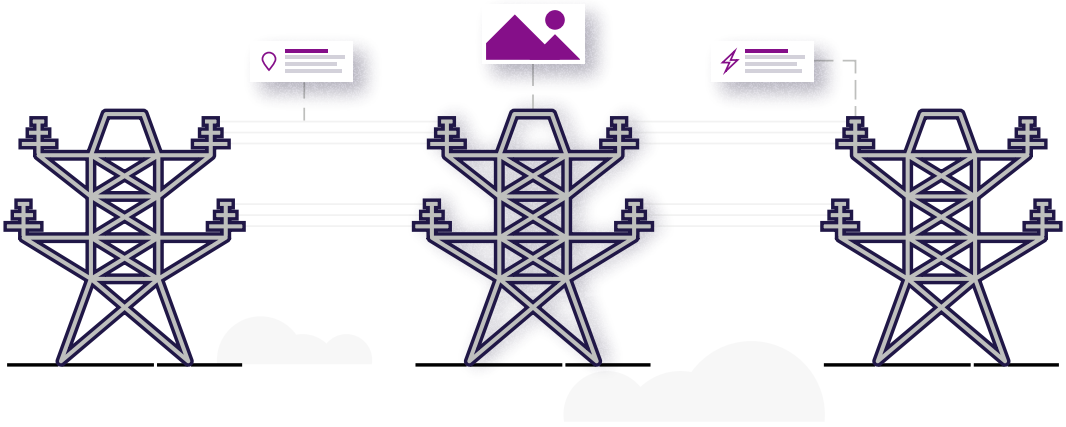

How does AI predict cyberattacks?

"It all comes down to the data — identifying data sources which could give you additional time to take preventative measures against an attack."

Graham Mueller Senior Research Scientist

Cyberattacks are designed to be hidden, but AI is shining new light on these threats earlier than ever before. Leidos is a prime contractor on an IARPA research program that uses AI-ML to observe the early stages of cyberattacks through unconventional signals, including tweets and other open source data. The research, which was featured in *WIRED*, demonstrates how AI-ML is accelerating threat intelligence and changing the cybersecurity landscape. To learn more, we welcome Dr. Graham Mueller, a senior research scientist on the project.

Q: Why has detecting cyberattacks in the early phases traditionally been so difficult, and why is this a problem uniquely suited for AI-ML?

Graham: First, cyberattacks are designed to be hidden. If a hacker is setting up infrastructure to deliver malicious email campaigns, for example, they don't want you to know about it. They don't want to be identified, and they will actively take measures to be hidden. The second problem is called a "web scale" information problem. Even if you set up the ability to monitor these things in real-time, the amount of information that is produced on a daily basis is huge. This is one of the main reasons ML solutions are so useful in this area. Traditional cyber defense systems are often signature based, which means they protect against things they've seen before based on certain rules. This approach is very brittle, because a hacker could just change the name of a malicious domain so that it would no longer be blacklisted. ML approaches can be so effective because they're probabilistic. They can take something you've never seen before and make some prediction about it based on how the system is trained, so you don't have to rely on a strict, rule-based system to block attacks.

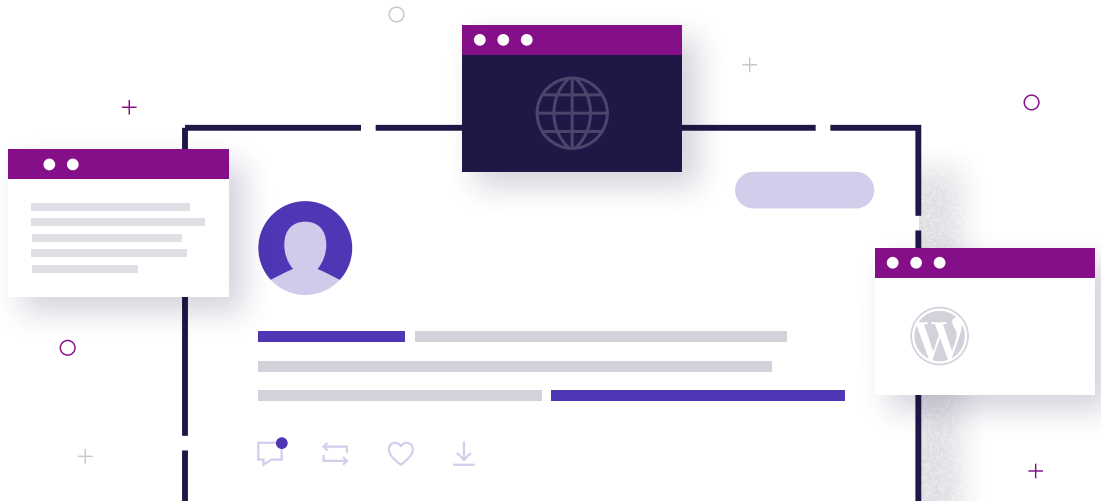

Q: The CAUSE program which you worked on revolves around AI-ML in cyber threat intelligence. What was the goal of the program?

Graham: Cyberattacks generally develop in the phases outlined in the cyber kill chain. In this chain of events, one of the early stages is the surveillance stage, where a hacker scans their target's infrastructure, or gathers other open source intelligence about the target's key personnel for targeted phishing emails. The attacks evolve from there. Our goal was to observe the early stages of an attack using unconventional data sources which serve as indicators or signals of the attack. Tweets are one of these unconventional signals, but there are many other data sources we looked at, including content on the dark web and open source software repositories. Much of our research was devoted to developing useful sensors from these unconventional data sources which were used as input into our ML driven prediction models.

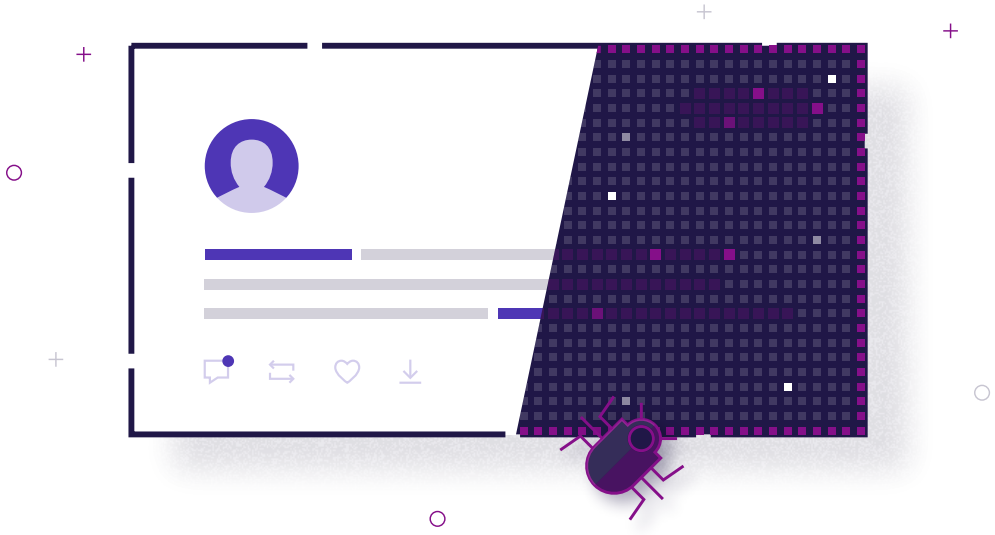

Q: What else do you see at the intersection of AI and cybersecurity? What other problems are you and your team of researchers looking into?

Graham: The big thing I think about is managing information. With the cybersecurity and threat intelligence landscape, part of the future with AI-ML is to bubble up the most pressing threats to organizations. There's so much information coming at you that it's very difficult to know which attacks are the most threatening. AI-ML tools have the potential to help us identify these threats both as they occur and at the scale that they occur and allow us to respond.

AI predicts cyberattacks

Q: In your experience, what are the critical success factors in developing AI-ML solutions that solve real problems?

Graham: It all comes down to the data — identifying data sources which could give you additional time to take preventative measures against an attack. There's always a need for providing actionable information to decision makers, rather than just a high-level overview of what's happening. You need the right data that tells you specific things relevant to your networks and systems in order to adjust your defenses accordingly.

One of the bottlenecks in developing machine learning applications is the need for large, labeled data sets which are used to train the ML models. There's a significant effort needed to first collect the data and manually annotate it. This is a hugely time-consuming thing to do. The approach we took is to use "weak supervision" to label data very quickly in order to leverage it. This allowed us to develop custom cyber-focused event extraction tools which we used as input into our ML models.

Q: What's next? What do you hope to achieve in future AI-ML projects?

Graham: Our overall goal is to provide real value to the cyber security world by providing actionable information and decreasing detection time. We're focused on providing indicators of malicious activity in real-time. Moving forward we want to quickly develop algorithms which extract important information from huge streams of data and summarize it. That's what is so exciting about our research.

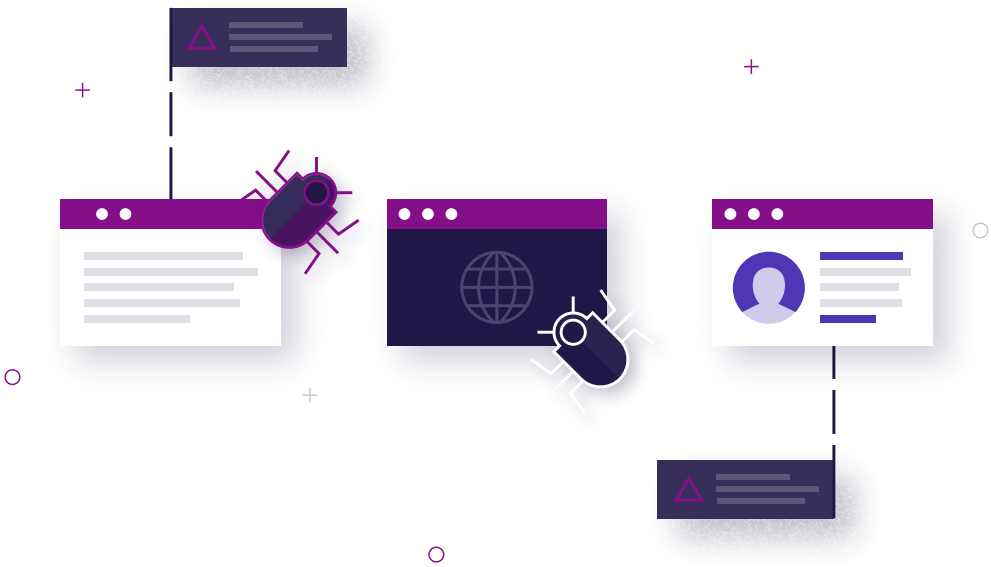

How does AI help us fight cancer?

"There are a lot of ways that we can find correlations in data sets, but actually understanding causation is a big challenge. That's where AI comes in."

Ron Keesing Director of AI and Machine Learning

Q&AI

We're diving into the computational war on cancer, where AI helps address the challenge of knowledge fragmentation. A Google Scholar search for cancer-related articles in 2018 alone returns 143,000 results, or roughly 2,750 research documents published every week. Now there's an AI system that can read and understand them. The solution, developed as part of a DARPA program called Big Mechanism, uses natural language processing (NLP) and big data analytics to extract knowledge from this enormous volume of scholarly papers. Through the program, a Leidos-led team of data scientists demonstrated that applying AI to this literature can help fill knowledge gaps, generate new hypotheses, understand important causes and effects, and develop more targeted treatments. To learn more, we welcome Ron Keesing, Director of AI and Machine Learning at Leidos, who explains how AI gives us greater potential to help people succeed in their battle against this horrible disease.

Q: What makes the cancer domain an ideal place to test new AI capabilities?

Ron: Obviously, cancer is a really important problem. It's also a very rich domain because there are literally thousands of cancer-related papers published every week. There's an enormous amount of human knowledge and understanding about the disease—far more than any individual human can read and understand. So it was an ideal place to test the hypothesis that maybe machines can understand large, complex causal mechanisms even better than humans by looking across all of the knowledge that's available.

Q: What was the most important goal of the program?

Ron: "Big Mech" was about understanding causal mechanisms that underlie really complex phenomena like cancer. There are a lot of ways that we can find correlations in data sets, but actually understanding causation is a big challenge. In this program we talk about mechanistic causation, or a physical mechanism that makes one thing drive another thing into a new state. Big Mech was all about combining human understanding of mechanistic causation in science with what we can draw from data to learn even more.

Q: Cancer-related academic papers written by pathologists were your primary data sources. What's in these papers that is so valuable?

Ron: They describe research findings in human terms. Most papers describe individual experiments and report specific phenomena: for example, one protein causing another protein to become altered. There are also papers that describe models of how more complex systems work. Typically, those types of papers are written in review journals by teams of scientists. They might describe a complete model of, let's say, a cancer pathway with multiple proteins that drive a signal through that pathway. These are really important because they typically represent a lot of humans coming together and thinking about how a whole system works. A big part of Big Mech was not only to read individual facts in the cancer literature, but also connect them up to these models that humans have built about cancer. Doing this allows those models to grow and expand as quickly as the rate of new scientific discovery.

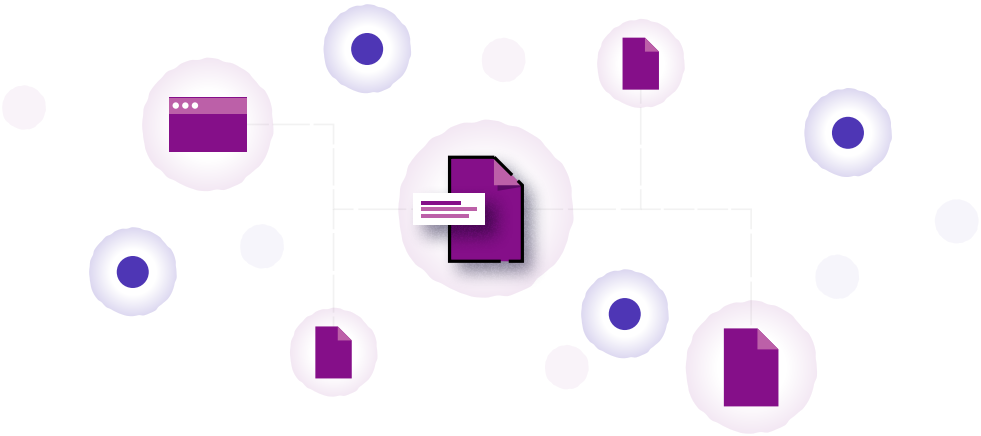

Q: What subsets of AI did you find most effective, and how did they work on the program?

Ron: We used NLP to find descriptions of things like causation. This worked by the AI actually reading human descriptions of new findings in the literature, and connecting them up to models that have already been built out and agreed upon by humans. Because these models are so large and complex, these subtle connections are easy for humans to miss. There might be an obscure journal where a researcher shows that a key protein in a pathway causes another protein to increase in concentration in a very specific situation. That's an important piece of knowledge, and we can use NLP to extract it from the literature.

Another important part of this is extracting information represented in diagrams and tables. Humans communicate not only in direct statements, but also in organized forms like diagrams and tables that are very rich in information. A lot of times what's described in the text is the most exciting positive information that was found, but often what's not described is the negative information—the things that didn't work. It turns out that the negative information is really important for machine understanding of the literature. Machines are great at finding things that might be true. Negative information helps prune down machine-generated hypotheses to the things that are far more likely to be true.

Q: What impact do you hope this program will have in the ongoing war on cancer?

Ron: In cancer biology, when there's a difficult individual case at a hospital, they often form something called a molecular tumor board to come up with a plan for how to treat that patient. The technology developed in Big Mech has already been used by these boards in some cases to help recommend individual treatments based on connections found by AI that no human had ever found before.

In addition to assisting individual treatments, this technology is also helping us discover new potential drug targets and drug designs. One of the big challenges across the cancer domain is finding the right places in the pathway that drugs can target. We understand some of those targets, which receive billions of dollars of research in the pharmaceutical industry. Big Mechanism research was able to identify promising new targets to go after so we can design new drugs for them as well. That's especially important in rarer forms of cancer that haven't received as much research attention.

Q&AI

Will AI steal our jobs?

"Even if we don't know exactly what the future will look like, the message is to be adaptable."

Dr. Alric Althoff Leidos Research Scientist

Should we fear the automation we also desire? It would be comforting to know a computer couldn't do our jobs, yet we're creating ones that can. With emerging AI we're discovering many cognitive functions we used to consider exclusively human—emotional intelligence, interpersonal skill, creativity, problem solving, critical thinking—are actually not. Unnerving to some, these advances will reshape the labor market and change the way we work. According to a report from the World Economic Forum, roughly 36 million Americans hold jobs already highly exposed to automation.

This isn't necessarily cause for alarm according to many experts including Dr. Alric Althoff, a senior research scientist at Leidos, who believes if adopted responsibly, AI-powered automation will create economic opportunity rather than diminish it. This might require significant reskilling, but like previous technological revolutions he believes humans will lean on our adaptability as a species. To learn more we welcome Alric.

Q: Let's start with the big question. Should we be concerned about AI taking our jobs?

Alric: We can't say for sure that the things we get paid to do now are safe from automation. But even though automation will have a major impact on the workforce, this evolution will be unexpected. People will still have jobs, but they will look different. We don't usually predict these things well, and it's easy to have a pessimistic outlook on change. But historical information gives us reason to be optimistic.

With every technological revolution, the planet's population has increased. The number of available jobs, and the number of available things to do given new technologies, has grown. Even if we don't know exactly what the future will look like, the message is to be adaptable. Lucky for us, adaptability is something human beings have in spades. There's really been no change that we haven't somehow adapted to.

Q: To what extent have we already given ourselves over?

Alric: We already rely on automation everywhere. Humans don't make light bulbs or ship packages without automated tools. These things used to seem like incredible innovations, but they're now ordinary and expected. Today, we're beginning to see AI algorithms do things we're used to thinking of as human activity. There's some anxiety around that. What if we discover the things we used to think of as human are really not unique? Does it mean that we continue to build tools to assist ourselves until we transform ourselves into something completely different? That's where the anxiety comes in for many.

Music, for example, is something we consider so human, but we've largely already given ourselves over to machines to produce a track. We used to think of image recognition as a human task, but the algorithms that do this are actually much simpler than we thought they needed to be. Another example is building computer chips, which is a process dominated by algorithms. There's a popular and justified belief that no one human being knows how a computer chip is put together from top to bottom. We might understand the algorithms behind it, but it would be a stretch to say there's one person who actually has a complete understanding of every step.

Q: What's the most impressive way you've seen machines taking on human tasks?

Alric: I've worked on a project with someone who is involved with program synthesis, which is basically when you define the goals of an algorithm, and it writes the code for you. It's not a really fast or well-developed technology. The reason why this is impressive to me is that we think of coding as a high-tech job that we won't be able to automate soon. But the other day I saw some code that I would think of as reasonably complicated written entirely by a machine using only a set of tests that the program had to pass.

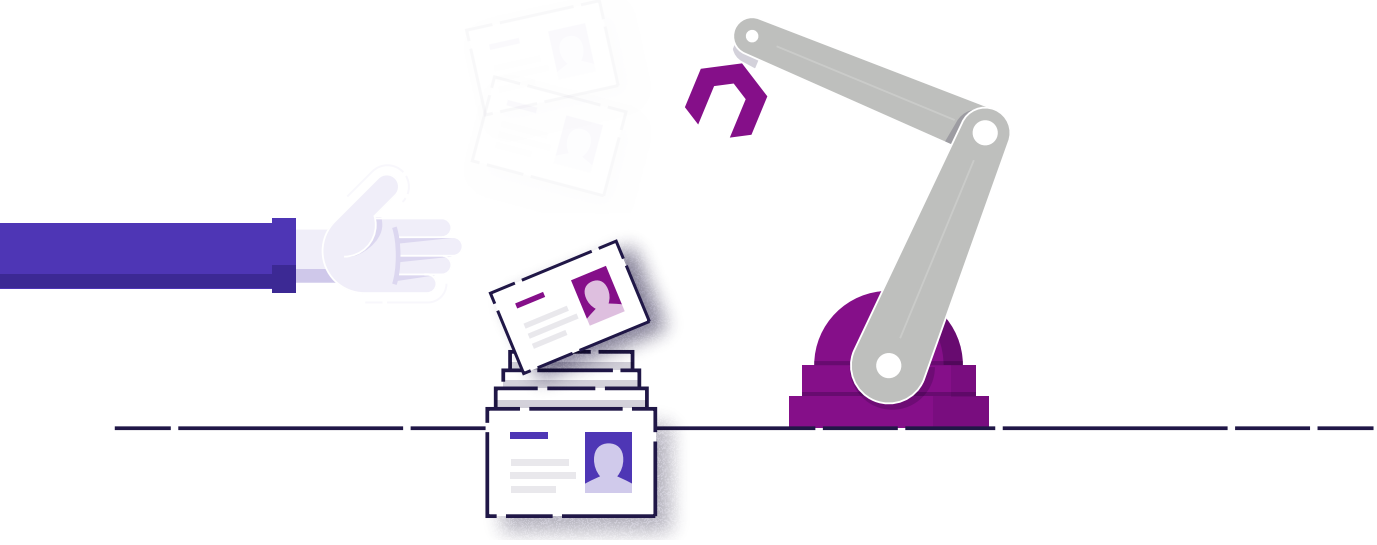

Q: What big conversations should we have as we hand over more and more responsibility to machines?

Alric: We're at a point, much like we were in the Industrial Revolution, where we can make decisions about the consequences of our actions. In the Industrial Revolution, we could have said, "Wow, there's all this smoke in the air. We're generating a lot of pollutants. This isn't great for the environment. We should probably change. We should probably not do this." That discussion could have started 100 years ago and prevented a lot of the global warming trends we're seeing now.

There's an impetus on us to ask these types of questions relative to the AI revolution. We have to decide as a society and globally what our responsibility is. We need to decide what we want to automate and what we don't. We have to carve out roles for humanity in the future we're building, because underneath all of this is economic pressure based on demand. It's human demand, so we have to decide where we're willing to draw the line between our desires and reality. This is the big challenge, but if we all start to really understand the connection between our decisions and their consequences, we can make better decisions.

We also need to ask ourselves important questions around security and trust. These are major concerns when we remove ourselves from critical decision-making processes such that we're trusting algorithms trained on data. The more trust we place in automated systems to do things for us, the more openings there are for nefarious actors to inject slight variations in those processes or try to influence them. Preventing this sort of exploitation is really the crux of where counter-adversarial AI is today. So we're working this angle, but we also need to take responsibility and modulate our demands based on understanding of the robustness of the state of the art.

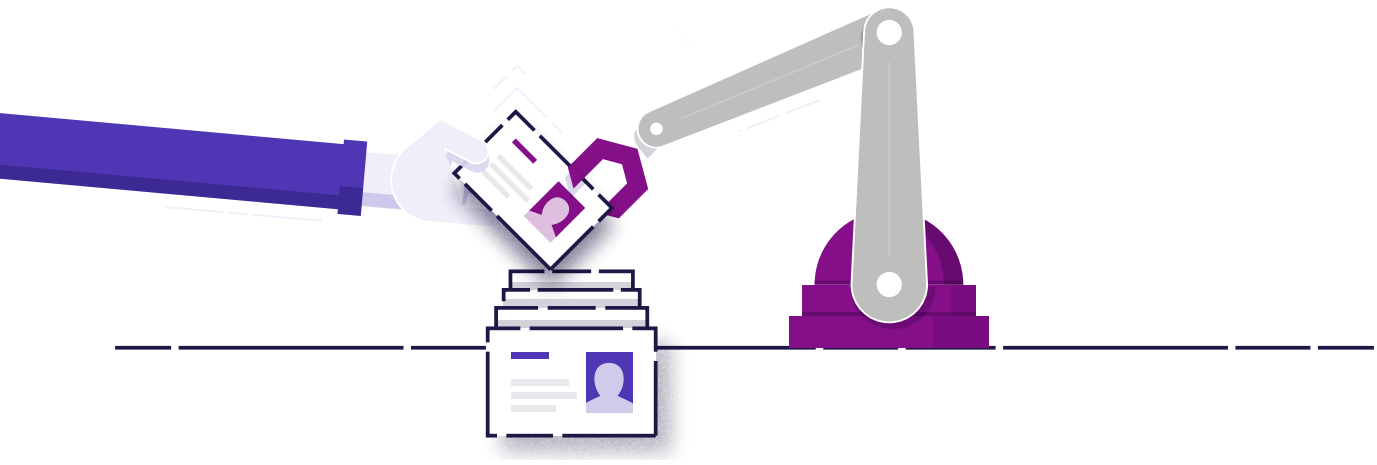

Q&AI

Confronting AI's inclusion problem

"We can't develop technology first and think about its consequences later. It's both together."

Dr. Shirley Cavin Data Science Manager, Leidos UK

The ethics of AI remain a hot topic as we begin the new decade, specifically surrounding algorithmic bias. Like a child to a parent, machine learning models mimic the behavior of their human creators who are prone to biased thinking. We've seen AI perpetuate and even amplify injustice in hospitals, courtrooms, and the workplace, and warning flags are up in higher education.

As this is happening, there's significant research pursuing breakthroughs that would allow our AI to meet anti-discrimination standards and policy. Leidos data scientists are confronting the problem by engineering models with greater transparency. One particular program in the intelligence community sheds light on data classification decisions made within the inscrutable black box of deep neural networks. At last summer's AI Palooza, Dr. Vicente Ordonez Roman presented his team's award-winning research, sponsored by Leidos, which studied the mechanics of why machine learning amplifies gender stereotypes in image recognition and natural language processing.

How can we trust AI when it's earned a healthy distrust in the past? To learn more, we welcome Dr. Shirley Cavin, a forceful advocate for greater diversity and inclusion in the field. Recently she sat on a panel, hosted by Leidos and Scotland Women in Technology, and discussed important topics regarding ethics in AI.

Q: According to the headlines, AI is not working for everyone. What will it take to change this?

Dr. Cavin: The bottom line is we need more accountability. Often we don't really know how our AI makes decisions. This might be okay with trivial matters, but definitely not when algorithmic bias causes discrimination or other harm to people. More broadly, we need to be aware of the damage our technology can cause. As technologists, researchers, and scientists, sometimes we get lost in the depth of the challenge itself. We become so focused on solving the problem that we don't fully think through how it will be used. AI is powerful because of its potential to

solve a lot of our complex problems, but when we include technology in human activities, we need to think about serious things like safety, security, reliability, and fairness. We can't develop technology first and think about its consequences later. It's both together. It's important to note that while we want to promote responsible technology, we also don't want these factors to kill innovation.

Q: Where does your passion about this issue come from, and why should others care as deeply as you do?

Dr. Cavin: Inclusion is a societal issue which makes us all vested stakeholders. We don't want to discard any member of our society. I guess my passion comes from my soft side. I love technology, but I also have a family. I'm a mother. So I try to bring all these things forward so that technology improves our lives. We can either do technology for the sake of technology, or we can do the much smarter thing and understand its consequences and impact on human life. I've listened to a lot of talks by researchers who are doing very, very interesting things with AI. But often when I raise questions around the consequences of that technology, these important questions go unanswered.

Q: How do machines become biased?

Dr. Cavin: AI is the type of technology that enables systems (software and hardware) to do activities that require "intelligence," which is the ability to use information, knowledge, and experience to make decisions and perform actions to achieve a set of goals. AI aims to replicate human behavior and the human way of learning, which will occur within the context of the communities in which it exists. If those communities have forms of social and cultural bias or discriminatory attitudes towards certain people and ideas, the AI can easily adopt those as well and even amplify them. This type of bias might also occur by design. As AI systems are designed, the designers could bring their own personal bias in the design of a system, and a set of discriminatory preferences could make their way into the system.

Third, bias could occur as the result of the data that is used to train and test the AI system. If this data is not a close and true representation of the ecosystem where it's intended to be used, the results may be biased toward a certain portion of the population. However, if this is overcome and a true data representation is used, this still may not guarantee the reduction or elimination of social and cultural bias. For example, in communities where minority sectors or views belong to small portions of the population, they will be uncommon in the data used for training and testing. All of these considerations, if not well-thought or understood, may lead to mistrust and concerns about using AI in our communities, which would be very unfortunate because of AI's great potential to benefit society.

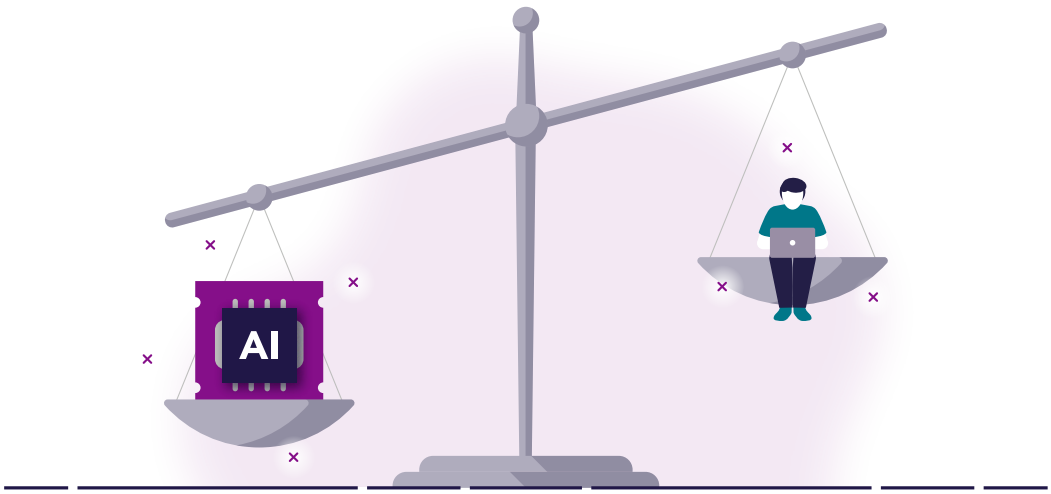

Q: Where have you seen the most progress?

Dr. Cavin: Frameworks that enable us to think farther are helpful when we develop programs, which is why I am quite pleased governments are taking the lead. For example, the European Commission is running an ethical framework pilot, a framework that we as part of technology driven organizations should follow. But leaving ethics issues only to ethics professionals isn't enough. For me, it needs to be a collaborative environment where everyone is participating in these discussions. Our common goal should be to make AI safe and inclusive, but also to promote the technology and continued innovation.

Q: What are the ingredients required to build this trust?

Dr. Cavin: Accountability, transparency, and fairness. When something goes wrong, there should be a human in the loop, a responsible party who can take ownership and fix the problem. We must also be transparent about the system's design, training, testing, and performance, which means we need to understand it ourselves. Finally we must promote fairness by enabling the use of data that's representative of the population with no bias toward or against particular sectors of people, and by making sure results are valid and outcomes respect the whole ecosystem's interests and the views of where it is intended to be used. Other key ingredients to build trust on AI systems are safety and sustainability. These are particularly important to consider before the deployment and public use of any AI solution, therefore efforts should be put in place to safeguard the well-being of the communities and their early AI adoption.

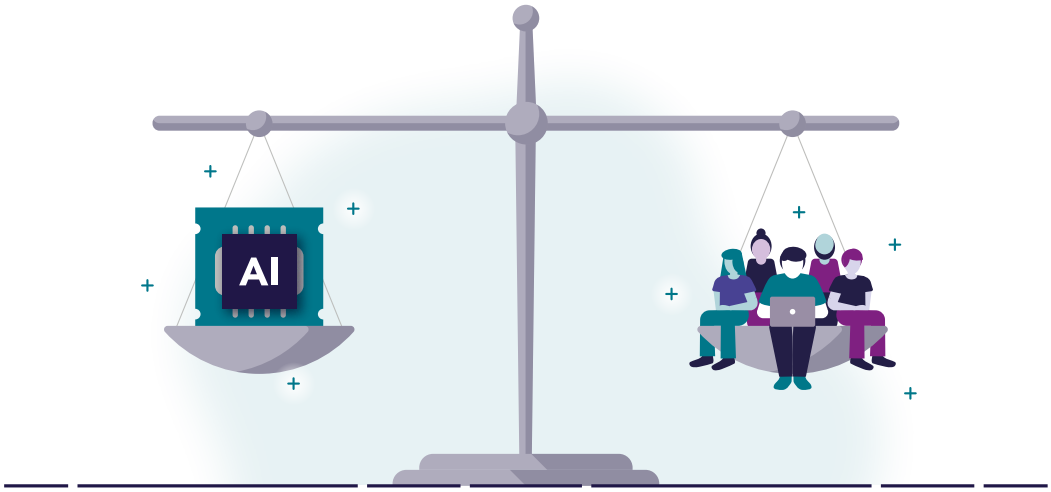

How does AI improve combat readiness?

"You don't want bad data, but bad data is a reality."

Dr. Donald Horner Senior Analyst, Defense

Soldiers prepared for combat are not only well trained, but also well supplied. Combat readiness depends heavily on an efficient chain of supplies, including mechanical widgets for all types of complex warfighting machinery. But often getting the right part to the right place at the right time is made difficult by requisition software fraught with bad data, which can prohibit the accurate prediction of supply needs.

A Leidos program with the U.S. Army used AI and machine learning (ML) to help solve this problem, and was recently featured on Open Data Science. To learn more we welcome Dr. Donald Horner, who was Principal Investigator on the program. Dr. Horner is a graduate of West Point (B.S.), MIT (M.S.) and Stanford (Ph.D.). He is a senior analyst at Leidos and a former engineer at Lockheed Martin's Skunkworks.

Q: To illustrate the importance of an efficient military supply chain, can you remember a time from your service when you weren't well supplied?

Dr. Horner: During the Bosnia conflict in 1995, I was the forward commander of a heavy transportation truck battalion. The battalion was particularly large given that we had several American companies augmented by Dutch and German companies. Our 'heavy trucks' moved everything from ammo to fuel to M-1 Abrams tanks to M-2 Bradley fighting vehicles. For whatever reason, our heavy trucks were burning up alternators at an unusually high rate. We had no idea why. And, because these alternators were typically not in the Army supply system in the quantities needed, we had several inoperable heavy trucks—which meant that cargo did not move.

Q: How did you solve the problem, and how might AI-ML have helped you back then?

Dr. Horner: My battalion supply technician found a truck dealer in Sarajevo that carried the alternators we needed. We probably bought 25 of these alternators, which were anything but cheap. And, I paid for all 25 with my U.S. Government American Express card. Big invoice. True story. AI-ML is particularly helpful in these non-standard requisitions because they can predict and prescribe items for resupply. Often times these predictions and prescriptions are items that humans might easily miss. AI-ML and deep learning in particular uncover esoteric but important clusters and relationships between variables in the data that humans don't see.

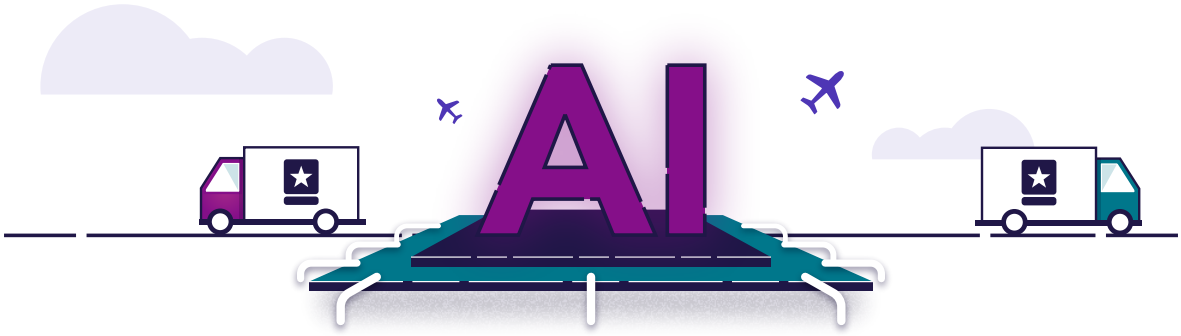

Q: The program you write about uses AI-ML to solve a major problem in military supplies requisition. How would you summarize the problem?

Dr. Horner: The Army has this wonderful logistics data set, but it had a lot of errors. An example of this can be something as simple as a person had entered a "one" instead of a lowercase "L" for a stock number, but you can imagine how those errors propagate through the Army system worldwide. These errors can lead to major delays in the supply chain. In the meantime you've got a piece of equipment, like an Army Black Hawk helicopter, that is not operational because it needs the part. It can take weeks. And all the while you've got an important weapon system that's not operational.

Low and behold, this is normal with data sets across academia and across industry. You don't want bad data, but data is a reality. More often than not, we start with bad data sets. So the first thing you have to do is cleanse the data. We started using a cleansing algorithm to see if we could teach it to identify the flaws. We found that we could use AI and ML protocols to accurately identify the flaws. And in the process, we found other flaws no one knew were there. Then we wrote some code that allowed us to fix the flaws almost instantly.

Q: What was the key to success for the AI-ML tools you used?

Dr. Horner: The key is you've got to teach the algorithms the data by repetition, repetition, and repetition. With gargantuan numbers of repetition, the algorithm detects very hard-to-find patterns. And by this pattern analysis, AI-ML can then cleanse the data instantly when it would have taken humans months. So the work of data cleansing and data correction is highly dependent on this deep learning.

Q: What are you most proud of about the program?

Dr. Horner: An efficient requisition process means combat readiness is improved, and combat readiness means fewer soldier will be put at risk. That's what this is all about. The whole military supply system is based on supplying equipment to enhance soldier survival. This AI-ML program enhances readiness and increases soldier survivability by getting our warfighters the right materiel at the right time.

Q: What do you envision next for AI-ML in the military supply chain?

Dr. Horner: We're going beyond predictability and now creating AI-ML systems that are very prescriptive. You get on Amazon today and order a pair of sneakers. You say yes I want that pair of sneakers. What does Amazon do? They have algorithms that immediately tell you that if you bought that pair of sneakers, you probably should have this pair of shoe laces. It's gone beyond predictability to prescribing for you, the consumer, what you should have. That's exactly the transformation I see happening in military supply logistics.

Another example is the Smart Cities initiative—well known in the private sector with success in cities such as Columbus, Ohio and San Francisco. 'Smart Hub' is the military analog to Smart Cities. Think of 'Smart Hub' as a hyper-IT interconnected base which yields more efficient operations while delivering more visible, highly rationalized, just-in-time, last-mile distribution and logistics services to troops in contact at reduced costs. Smart Hub hyper-IT connectivity refers to multiple connections with real-time information through an array of remote sensors and a robust suite of data analysis tools, including AI-ML. The predictive, prescriptive modeling strengths of AI-ML are a perfect fit for Smart Hub.

How can AI help anticipate human behavior?

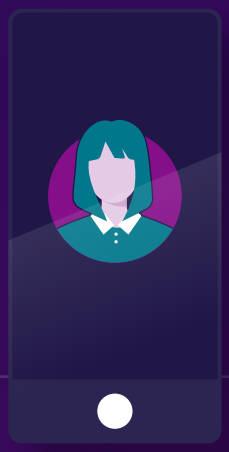

"AI can help us understand not just how people act, but also why."

Dr. Jonathan Pfautz Chief AI Scientist

We generate several quintillion bytes of data every day, and almost all of it describes our activity in the world, and how we interact with each other. This volume of data is continuing to expand with more and more ways to connect, work, and play – online. With so much of our daily lives and social interactions involving computers, you might think it would be easy for AI to predict human behavior. However, you might be surprised, according to Dr. Jonathan Pfautz, who says even our most advanced AI still struggles to connect causes and effects in human behavior. How does AI make these connections more visible? To learn more we welcome Dr. Pfautz, a chief AI scientist at Leidos and former program manager at DARPA. In a recent episode of the *Voices from DARPA* podcast, he shared a perspective on building computer models of humans to help understand what and why we humans do what we do. His thoughts led to the creation of a DARPA program called SocialSim, which has relied on the talents of scientists and engineers at Leidos.

Q: Why is it so hard to predict human behavior?

Dr. Pfautz: The science around human behavior is still new, compared to the study of the physical world, where there's a lot less controversy about how to "do good science." We use telescopes to study the sky, and microscopes to study cells, but we don't have the same kind of reliable, calibrated instruments for understanding how humans do what they do.

Q: Why has behavioral science lagged behind?

Dr. Pfautz: Many types of human behavior are called "wicked problems." It's complex and hard to do good science, and even with all this new data, we're dealing with uncertain and evolving situations, all the time. People change, technologies change, and the world throws us curveballs. A success in using AI to anticipate an event today might or might not work in two hours, two days, two weeks, two months, or two years. How definable is human behavior anyway? How stable is it? The complicated nature of these questions makes it extremely difficult to create a valid model that could help anticipate future human behavior.

Even with new, large data sets, social and behavioral research still fights to connect causes and effects. If you rewind all of behavioral science back fifteen years, you couldn't have predicted that everyone would be carrying a smartphone today, and how that's changing behavior. New technologies, like world events, dramatically change our behavior – whether it's how we drive, walk down the street, or interact socially.

Q: How do you see AI beginning to accelerate the science of human behavior?

Dr. Pfautz: AI techniques help in processing huge amounts of data (think of billions of tweets, blog posts, web pages) to form models of "what matters," or "features." Companies are using this data to do things like help anticipate and resolve traffic jams, share relevant local news, and target advertising. New AI techniques focused on "transfer learning" are one of many AI approaches that acknowledge that the problem of understanding and anticipating human behavior from online communications is going to continue to change (Can we learn from data on one problem to better understand another problem?).

Q: How do you see this progress improving society?

Dr. Pfautz: The opportunity for improving the human condition has never been more apparent, from encouraging global conversations on climate change, to rapid responses to humanitarian crises. One example is ability to look at data, and figure out how to "nudge" people towards the outcomes that benefit them. "Nudge science" essentially helps shape behavior in positive ways that are really simple. The classic example helps users on a government website answer the question, "How much money do I want to save for retirement?" If the default option is the one with the best results, people will pick the default more than anything else. We gravitate towards certain behaviors in certain situations, especially in these computer-mediated situations. AI helps us run those experiments on large numbers of people to nudge them toward doing what's best for them.

Take, for example, a program I ran at DARPA called "Warfighter Analytics using Smartphones for Health," or WASH. Your average soldier will see a doctor at most once a year. How do we help them when there isn't this notion to, "Hey, just go see a doctor." Well, maybe the general patterns of behavior from the smartphone you carry around all the time can help warn you – "get a check up."

Q: Conversely, how is it dangerous?

Dr. Pfautz: AI is a game-changer, for better or worse. It's a challenge when it starts to relate to our daily lives. Massive amounts of data are available on our activities, and this data is a commodity that is widely available. In the US, there is a wealth of private information out there. This means that if this data is useful, our adversaries can also learn from it. They can learn a lot about our behavior, creating opportunity for foreign influence. My existential fear is that with increased AI capabilities, foreign governments that don't share our values about freedom of speech and the right to privacy, will be able to understand human behavior better than anybody else. And, so it's no longer an arms race, it's a behavioral science race.

Q: What sort or work is Leidos doing to help our customers navigate these changes?

Dr. Pfautz: A few years ago, I started a program at DARPA, called "Computational Simulation of Online Social Behavior," or SocialSim. This program focuses on building computer models of how communications play out online. These models simulate what might happen when information or misinformation is shared with a complex network of people. Leidos has been a key contributor to this program, providing data sets across multiple social media sources, and supporting the testing and evaluation of simulations – to see if we can accurately anticipate how false information might spread online. Other teams on the program have focused on applying AI techniques to create simulations ranging from the spread of information on malware to the spread of misinformation about well-intentioned volunteer organizations in Syria.

How does AI optimize oil and gas production?

"The beauty of AI is that it can be applied to many other industries to solve their complex problems."

Chung-Yan Shih Senior Strategic Data Scientist

The United States is now the world's largest producer of petroleum and natural gas. Nearly a million wells across the country produce roughly 11 million barrels of crude oil and 4.3 million barrels of natural gas per day. While new technologies have improved oil and gas extraction, the big question for the industry is how to reliably predict well production. To help solve this puzzle, researchers like Dr. Chung-Yan Shih, a senior data scientist at Leidos, are turning to AI. His research for the DOE's National Energy Technology Laboratory (NETL) has helped reveal how machine learning can offer answers to the oil and gas industry as they strive to make their wells more efficient and keep up with demand.

Q: Why is it so important for oil and gas companies to forecast well production, and what challenges do they face?

Chung: Productivity prediction is critically vital for oil and gas companies because it drives business decisions and helps the industry plan for the future. But, since resources are stored underground, it's been hard to get accurate predictions, and even the best estimates come with significant unknowns. Predicting how much a well can produce throughout its life is referred to as the estimated ultimate recovery (EUR). Initial productions (e.g., first 12-month production) are usually used as proxies to indicate the overall performance of a well since a well often declines fast after its initial peak production. In the long run, many factors may affect the EUR's accuracy. Each well is unique, and the relationships of parameters affecting well production and capacity are complex.

For example, wells differ significantly based on their location, the time they have been in production, their design, and the resource extraction technology that is used, to name a few. However, it's even trickier than that. The reservoir that a well taps into may vary in geologic properties, such as thickness, thermal maturity, or gamma ray levels. Two wells with the same design at different locations are likely to perform differently. As a result, it has been a challenge to precisely predict EUR and to understand the relationships between factors and production.

Q: How does AI help us understand these relationships?

Chung: AI is well-suited to help make sense of complicated data sets. One of the key challenges for oil and gas companies is to identify the parameters that drive production. As you can tell from my previous responses, these can be fairly complex. Traditionally, production decline curves and type curves are used to estimate the EUR. Scientists and engineers history-matched curves to well production and averaged the selected curves as type curves to represent the shale play. The process relies heavily on the experiences of the scientists and engineers. However, AI can help to extract production patterns from production data to identify and understand parameter relationships. NETL has highly skilled geologists who estimate geologic properties, and sufficient data for production and well completion, so when we were first asked to help with this problem, we saw great potential in machine learning applications.

Our goal was to figure out patterns from that data and identify key drivers of production. We tested nine regression machine learning algorithms from simple linear regression to more complicated neural network algorithms. We used these algorithms to predict the first 12-month production of a well, which is a common proxy of EUR. These algorithms took variables from well design and geology, such as a well's lateral perforated length, the gamma ray value of the rock formation, and the amount of proppant, which is sand or a man-made material used in the hydraulic fracturing process. From there, we were able to determine different factors that drive production.

The first study provided great results, but we wanted to take it further because there will always be a gap between the production from the first 12 months of a well's life and the final EUR. And that is what we did in a second study. We used deep learning—a subset of machine learning based on neural networks—to understand a well's production changes over time. We used the same concept of adopting both well completion design and geology parameters to predict the entire production profile of a well, which we used to calculate EUR. Thanks to the capability of deep learning to extract knowledge from complex relationships, we can predict the performance of a well even before it is drilled.

Q: What was the most important benefit to your customer?

Chung: DOE recently launched an Office of Artificial Intelligence and Technology to "transform DOE into a world-leading AI enterprise." NETL is interested in determining if machine learning is useful and beneficial to subsurface analysis. The goal with the first study was to understand how the oil and gas industry could better use technology to efficiently extract subterranean resources. The results showed how machine learning could be applied to using existing data to determine production key drivers. AI now appears to provide a way to estimate production for future well development. Through the first and second studies, we've also helped NETL identify potential R&D projects to focus on to help improve resource recovery, such as well stage and spacing optimization, or the impact of different compositions of fracture fluid.

Q: How does this work impact the future of the oil and gas industry?

Chung: The industry wants to run wells and extract resources as efficiently and cost effectively as possible. We've shown that with AI, they can estimate well production performance based on the well completion design and geologic properties at the well location. Combining the experiences of geologists, engineers, and drillers, they can have a strategic plan for well deployment and optimization to maximize recovery and revenue. Machine learning has also helped us identify and extract complex interactions and relationships from numerous parameters, which the industry can use to design better wells.

Q: What other industries might benefit from this technology?

Chung: The beauty of AI is that it can be applied to many other industries to solve their complex problems. AI is a data-driven approach that is capable of learning complex patterns in the data. Often, the information provided by AI validates the assumptions and reasoning of subject matter experts (such as reservoir engineers and geologists). Sometimes, AI provides a fresh way of thinking and tackling a problem. While those results might be surprising at first, they usually offer valuable insights to help us understand more about the problem. For us, it is always good to work with the experts in the field to make sure AI is learning the tasks we would like it to learn. There have been cases where AI "cheated" in learning in an ingenious way that might be hard for a human to detect. Data scientists and domain experts working together can ask the right questions to learn about a client's goals and fill in the data gaps. These studies have broken new ground, and they show that the types of problems that AI and machine learning can solve are virtually limitless.

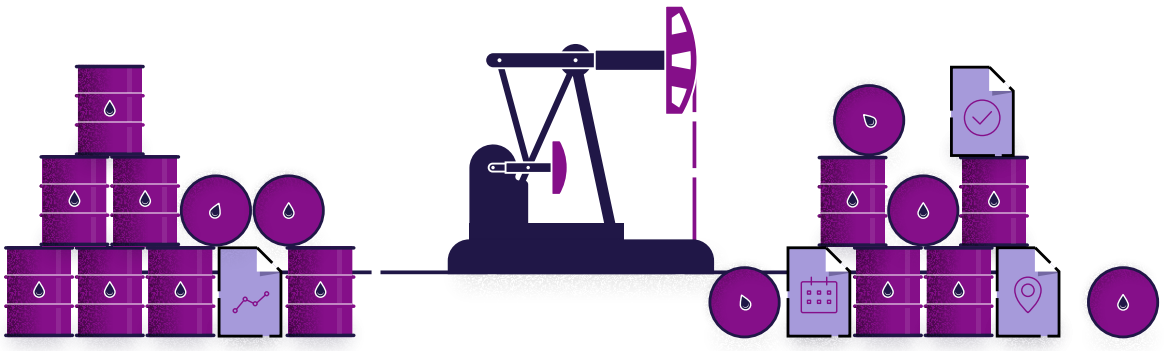

Can we prevent another AI winter?

"The conversation about AI should be a positive one, but also a clear one."

Ron Keesing AI Director

The conversation about AI mostly centers on its future, but it's also important to explore the lessons of its past. AI history teaches us sustained progress is no given. It's exciting today, but excitement wears off. At least twice before the field has experienced periods of decline so harsh they're commonly referred to as "AI winters." Why did this happen, and how can we prevent the next one? To learn more we welcome AI Director Ron Keesing. At CES Government 2020, Ron moderated a panel discussion among senior government and industry executives who explored lessons learned from previous AI winters and how to promote sustainable progress in the field.

Q: If AI progress occurs in seasons, where are we now?

Ron: It's certainly clear that we're in the middle of an AI spring, or perhaps even an AI summer. We're making a lot of discoveries and there's a lot of excitement in the field. However, I also see a growing irrational exuberance. For every successful AI project, many others fail. I think there's been too much hyperbole used and overoptimistic promises made. In fact there are a lot of things AI won't achieve until we solve some huge challenges related to artificial general intelligence, or AGI, which is a topic for another day. Let's just say our current systems have a lot of limitations and there are many aspects of them we don't even fully understand.

Q: Where does this irrational exuberance come from?

Ron: When it comes to AI we seem to have a hard time separating hype from reality. It captures our imaginations because it blurs the line between human and machine intelligence. But too often it's portrayed like a magic wand that can accomplish anything. The conversation about AI should be a positive one, but also a clear one. AI isn't easy, and getting value from it looks different in every situation. In fact research by Gartner shows the vast majority of AI projects will likely fail to deliver. When AI doesn't live up to its hype it can lead to widespread disillusionment. When that happens, funding dries up and R&D stalls. The previous two AI winters lasted for many years.

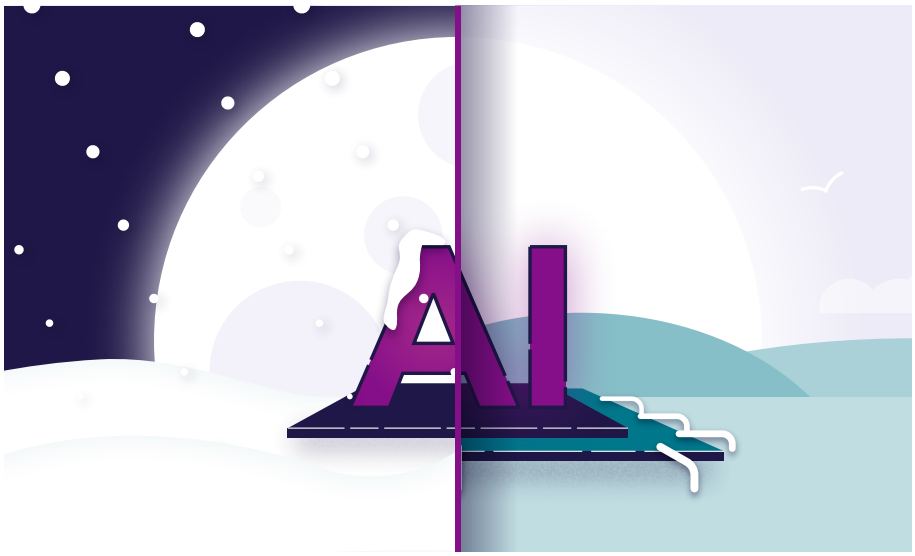

Q: Were they as harsh as the term "AI winter" implies?

Ron: It's mostly just a term used to make the point that we've been through these before, and they should serve as a cautionary tale. What's more important is that we're able to separate hype from reality, which is building AI that delivers sustained value is not easy. It's not a given that we'll make good on all of the promises out there. The message is to not get caught up in the hype, but to navigate this with sobriety.

Q: How can the government and tech community help prevent another AI winter?

Ron: By proactively addressing certain things that often cause AI to fail. I see at least three major challenges. First, we need a clear-eyed view of what's possible. I always make the point that AI can solve certain problems really, really well. But we're far from achieving AGI, or even building systems that can adapt to new domains and new situations. So it's really important to invest in AI projects that are actually supported by technology and not just hype.

Second, we need to put strong policy in place to govern the use of AI. We need to ensure it's fair, transparent, resilient, and secure. Right now our AI tools are a black box, which means they make decisions we can't understand, scrutinize, or protect. And once we have the right policy, we also need technology that bridges the gap between today's black box AI and our policy objectives. We and many others are working hard to help build solutions that bridge that gap, but in many cases we don't yet have solutions. There's a real risk we won't be able to build AI that's consistent with the policy we're creating, and this could lead to another AI winter scenario.

Third, building AI systems requires planning for a very different technology life cycle, and it's important to make smart investments that account for that life cycle. A lot of AI projects start with somebody doing a proof-of-concept, but there's a lot more to it beyond doing data science in a lab. There's actually a huge step to turn that initial demonstration into a working system that's connected up into real data. And once that AI system is deployed, it needs to be updated over its lifetime as the data and the environment changes. Models, data, and tools need to be updated continuously or the system builds up what's known as hidden technical debt. When that happens, your AI system stops delivering value and actually becomes an expensive liability.

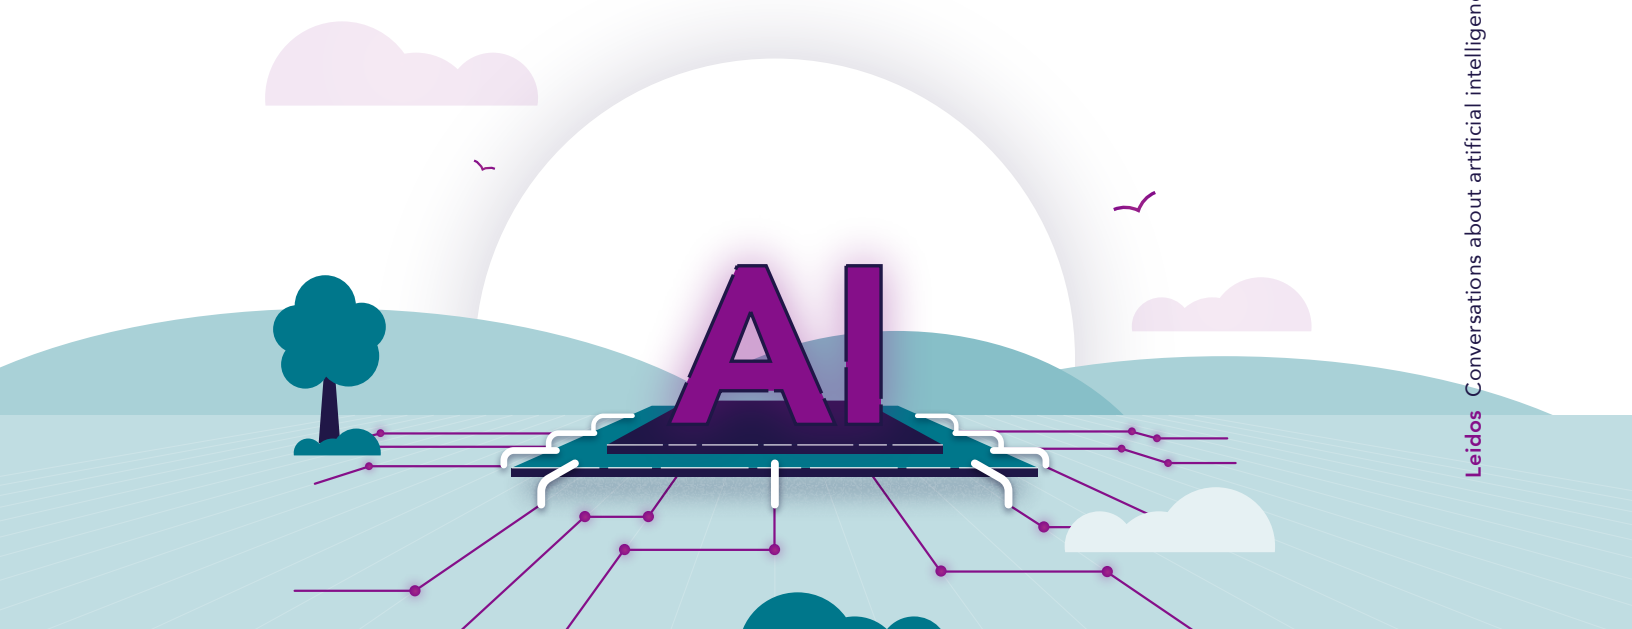

Q&AI

DEEP LEARNING

ALGORITHM

MACHINE LEARNING

How to speak AI

DATA SCIENTIST

NEURAL NETWORK

"It's important to speak the same language because as technology leaders, if we aren't clear, people will be misinformed."

Dr. Julie Rosen Leidos Chief Scientist

Q&AI

Artificial intelligence methods are based on mathematics, but the words of AI are also remarkably important. AI speak can sound like a foreign language. Like any language, speaking it fluently starts with command of basic AI words and phrases, jargon that is commonly used but often difficult to interpret consistently. For an introductory AI vocabulary lesson we welcome Dr. Julie Rosen, Leidos Vice President, Chief Scientist, and Technical Fellow Chair.

Q: It seems like the explosion of AI has also led to the explosion of AI buzzwords. Why is it so important to clarify the meanings of these words?

Dr. Rosen: When trends emerge or expand quickly, confusion often arises. The rapid maturation of AI has led to certain parts of the lexicon becoming unclear and confusing. We need strong working definitions to base not only technical conversations, but also broader ones involving funding sponsors and consumers of AI findings. It's important to speak the same language because as technology leaders, if we aren't clear, people will be misinformed. As far as business goes, if it appears that we don't understand the nuances of AI speak, our customers will think we're not expert in the field.

Q: AI and machine learning (ML) are often used together. Are the two synonymous?

Dr. Rosen: Not quite. ML is a subset of AI, but there are other flavors of AI as well, including rules-based expert systems and knowledge graphs, and it's important not to conflate the two. Methods of AI can be thought of like a Russian doll set. AI is on the outside, then ML, then deep learning. A true data scientist should know the full range of models and methods to consider for a given purpose.

Think of ML as the reasoning engine that consumes data, applies algorithmic logic to recognize patterns in the data, and sends the analytic results to a human or another computing component. General ML methods work with a layer of observed data (input), a layer of analytic output (e.g., forecasts, matches), and a hidden (or latent) layer where the attributes are connected to each other, and to the input and output layers. If the data are appropriately curated for like-comparisons, then ML techniques have the ability to learn new patterns by operating with a few initial rules, which mature algorithmically as data are processed over time. The beauty of ML is it allows us to characterize situational reasoning without the need for a full complement of programmed rules and instructions, which are dangerously brittle in real-time decision making with uncertain or missing data.

Q: What do you mean by deep learning?

Dr. Rosen: Deep learning (DL) is a machine learning technique, a set of algorithms that model phenomena measured through complex data. The etymology of "deep" in deep learning refers to the depth of the model. Consider huge, complex data sets like the human genome, social media data, or satellite imagery. When you're trying to forecast outcomes or recognize patterns within the data, each measurement associates with many attributes, also called features. In these cases we need more advanced models with more layers to learn features among the patterns in the data sets. When you're working with unstructured text and imagery, for example, these features become deeply interrelated, and the multiple layers of the DL network can start to resemble a big, hairy fur ball of connections. Now you've got to start digging deeper into the possible connections to determine what patterns you can infer from them. Such digging can get computationally intense, which why a new kind of chip, called GPUs, are in demand to train deep learning models.

Q: Speaking of networks—what's a neural network?

Dr. Rosen: Think of a neural network model as a computer that mimics how the human brain works electrically. The brain functions through a complex set of neurons, its basic working units, which modelers call nodes. Like our brains, neural net models are highly interconnected. They can be hundreds of layers deep and very wide. The brain has all these synapses, which modelers call edges or links, between nerve endings in the various cells in the brain. You want to be able to follow the connections among all those cells. In math-land we call this a multiply connected network. So if you can explore deep into the brain, maybe you can find a network path or multiple paths that help you connect one nerve ending to another. Or in the case of mathematical modeling done with ML algorithms, you're looking for a path of similarity or a path of feature closeness. But to make these discoveries, the network itself has to be highly multiply connected. It has to be very deep, hence DL modeling, to answer very large and complicated non-categorical kinds of questions.

Q: You mentioned data scientist, which is a major career path in AI. Walk us through the difference between the terms data scientist, data engineer, and data analyst.

Dr. Rosen: First off, let me say that these terms are greatly overlapping in academia and in the general marketplace. At Leidos, we use the term data scientist to refer to researchers who create data models and algorithmic methods to analyze big, highly complex data sets in order to find patterns. Often these patterns are not yet prescribed and too difficult to detect with traditional rules and static statistical methods. These professionals typically operate at the basic R&D end of the maturity spectrum, proving the accuracy and efficacy of the AI model or method.

At the opposite end of the technological maturity curve are Leidos data analysts, who employ these proven models, methods, and tools to investigate mission-specific data in support of decision making. Data analysts' mission is to impact the bottom line. As such, it's very important these professionals know their consumer's domain and what questions to ask the data.

On a Venn diagram, the Leidos data engineer overlaps with data scientist (to get the models and methods matured to scale and deployed in an end-to-end pipeline) and data analyst (to make the pipeline usable and maintainable). Data engineers have similar skills to generalized system engineers and solution architects. They start off with specs from a customer and work with third-party vendors to integrate commodity tools with custom-developed and optimized models and methods. Data engineers have the additional charge of understanding the impacts of design options to ingest and curate these big, dynamic, fast-arriving, diverse, unclean, mission-specific data sets. The data engineer is charged with making the data scientist's accuracy-proven models run faster, at the scale of the incoming data, and output the analytic findings through visualization, explanation, or communication element.

Why is AI so difficult to scale?

"If you want to apply AI models across the enterprise, you need a way to scale your access to data and computing resources without breaking the bank."

Tifani O`Brien Chief AI Engineer

Converting an exploratory model into an enterprise capability is called scaling, and it can be remarkably difficult with AI. The simplest approach might be spinning up copies of the model in the cloud, but this sort of system can be easy to break and expensive to run, especially when each copy requires significant computing resources. Solving these types of problems when scaling across the enterprise will be an important next step toward realizing the potential of AI-ML. It's an active area of research at Leidos, where data scientists like Tifani O'Brien are solving scalability problems every day. To learn more we welcome Tifani, who is expert in scaling AI for the intelligence community in particular.

Q: Take us inside the process of scaling AI. What's the big challenge our customers are facing?

Tifani: When a data scientist creates a model on their own computer, running on a set of data collected on a local hard drive, it works well when you need answers to a specific, time-bounded question. Model training happens once, and the results apply to the data already collected. However, if an enterprise wants to apply this model to the data they collect going forward, and improve the model as more data is collected, you need a way to scale your access to data and computing resources without breaking the bank.

Let's say you want to do a search across all your data including text, video, and audio. Machine learning models can extract audio speech from a video and convert it to text, and even do machine learning translation of all the content that is available in one language. Your data scientists have demonstrated they can do this on a month's worth of data already collected, and now you want to scale this up to do the same processing to all your historical data and all the new data streaming into your enterprise so your researchers can search across time and media type.

Just embedding the model in a browser application limits your users to only the resources available to each of them running it, so it would be very inefficient for them to have to extract and translate speech for each search every time. So we choose instead to pre-process all the audio and save the results for later searching. We deploy in the cloud with the ability to spin up as many resources as needed for the current task. This flexibility can come at a high cost though if not done efficiently.

Q: What particular methods make scalability more efficient?

Tifani: One of the most effective ways we've found to scale AI is to deploy microservices in containers, managed by a container orchestration system. Orchestration makes it easy to spin up additional containers whenever the demand is encountered, but this can quickly turn costly if not throttled correctly. For our purposes, we configured our usage to apply the more expensive processing resources for monthly model retraining, but once we deployed the production model we could dial it back.

Determining the appropriate API design and architecture is another scaling technique. Correctly delineating the boundaries of a service allows it to be flexibly inserted at different points in the processing pipeline, potentially avoiding bottlenecks. Also, locating the API layer inside the containerized microservice itself avoids losing valuable time to repeated container startup, hardware allocation, and session startup activities. Carefully managing the data as it transitions through a pipeline of AI models is another opportunity to reduce the traffic and cost of scaling up. We use a fast cache for input data that multiple models are going to access, and then move that data to cheaper, slower storage.

Q: What are some of the biggest challenges you normally face?

Tifani: One challenge is filtering out noise and selecting the right model to run on different data based on its type. If you run AI models on the data they are not trained to handle, you spend resources on irrelevant processing that also produces poor results. For example, using a model that was trained to detect firearms in photos on something different, like application icons, can result in many false positives and use expensive GPU computing time.

Another challenge is dealing with data proliferation. We often work with customers who have terabytes or even petabytes of data. Every time you process data using machine learning, you end up creating new artifacts based on that data. Take a situation in which we pull out speech from video. We have the original video file, then we have the audio file, and we also have text files based on the audio–one in a foreign language and one translated to English. And if you add to that extraction of all the still images pulled from the video, you'll add thousands of images that you must run through your processing pipeline again.

Q: What's your best advice for overcoming these challenges?

Tifani: To work well, scalable AI depends not only on strong AI capabilities, but also on an understanding of the computational infrastructure and how it handles resource contention, data transport and availability, and service orchestration. Many people say they can build microservices but then are challenged when they try to process truly large collections efficiently. We've been successful in scalable AI because we establish careful service boundaries, design analytics to best use the compute resources available, apply knowledge of the domain to customize the processing workflow, and apply the right models to the right data.

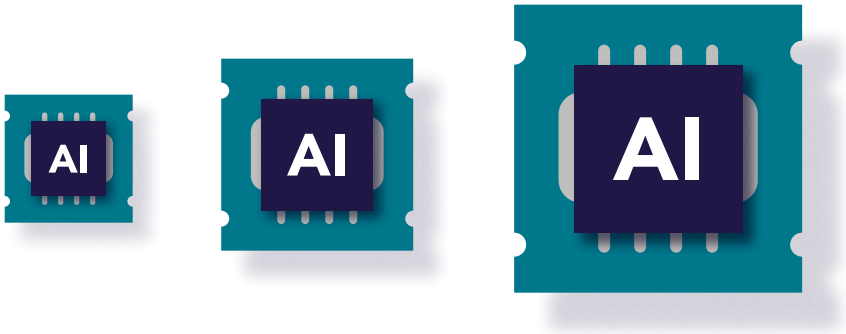

How can we trust decisions made by AI?

"The big secret about AI, machine learning, and deep learning is that 80 to 90 percent of the work is in the curation and preparation of the data that goes into a model."

Dana Moore Principal Engineer

AI has found itself at odds with the human need for reassurance. Trusting decisions made by AI and machine learning (ML) algorithms can be difficult when those decisions occur in a "black box," an inscrutable process that takes place without human supervision. As AI-ML becomes more widespread, pressure is growing to make sure these decisions are better explained and understood. But seeking this reassurance is about much more than warm feelings. Many who wish to adopt AI-ML often cannot afford to be wrong, notably when it comes to matters of national security. How can we trust decisions made by AI-ML? To learn more we welcome Dana Moore, author, lecturer, and Principal Engineer at Leidos.

Q: What's the big challenge when it comes to trusting AI decisions?

Dana: We're on the cusp of a new age of AI applications, from vehicular autonomy to guided surgeries. Machine learning and deep learning are the critical technologies in this explosion of AI-guided applications, but these models can be opaque at best—tough for people to understand. We want to understand the logic path these models follow to make decisions. We want to trust them and know why errors occur in order to understand how to correct them. We also want models that are provable. If you can demonstrate that a model is not wholly provable, then you can most likely demonstrate it is not valid, or that its security parameters have been breached.

Q: Why is it so important to trust decisions made by AI?

Dana: Trust is important to those who can't afford to be wrong. If you have a model that plays a board game against a human opponent, for example, you can afford to be wrong. But if you're working with self-driving vehicles, you simply can't. Trust and validity are important in many other domains as well, including medicine and national security. When the penalty for being wrong might be undue loss of human life, unlawful detention or prosecution, or in any case where poor recommendations might have dire consequences, it's very concerning. We get strange recommendations from Amazon all the time, even with years of purchasing data in those models. But these have no discernable effects on life or liberty. But apply the same faulty recommender engine to a "person of interest" scenario, for example, and lives and livelihoods are at stake.

Q: Trust and validity seem to be weak links right now in AI-ML. Why is this the case?

Dana: Provably correct code is not easily attained. Many AI-ML systems are simply too complex to demonstrate that a given outcome is provable. The best we can do in many cases is make a decent effort at demonstrating suitability to task and robustness against failure. But there are constraints even to this. It takes time and money to validate AI decisions, which means less-than-perfect (and generally unprovable) systems get deployed.

In traditional computer programming, you can trace outcomes and get a fairly good clue about the inner working of those programs. This came from a healthy skepticism in software engineering. Developers took cues from other engineering practices like structural engineering, and developed the ability to trace through code. Now we have the ubiquity of things like stack tracing, so that when you have a programming error, you can trace it back to its point of origin. Those kinds of things grew up in software engineering, but they're not yet part of AI-ML. A more realistic goal instead is the idea of "explainable AI," or the ability of AI-ML systems to explain their decisions to humans.

Q: What skills do you look for in AI-ML developers?

Dana: First, I understand why. Second, I understand why not. Third, I know when it will succeed. Fourth, I know when it will fail. Fifth, I know when to trust it. Finally, I know why it erred. These may turn out to be necessary but not sufficient constraints. That is, there may also be other elements that support explainability and model validation. Consider that even though these requirements may be met, there may still exist reasons why models or predictions are poor.

Q: If an AI decision can't be fully provable, what does "provable enough" look like?

Dana: Mathematic equations are provable, but machine learning and deep learning models aren't necessarily provable. There are certain output metrics that indicate precision and accuracy. What you really want are measures of performance and effectiveness from a model so that you can have confidence in its validity. If an AI system is well constituted and trained, has algorithms for prediction evaluation, and demonstrably produces reasonably high quality, true positive results, then that model may be suited for its purpose. But there's still a vital role for curation and feedback with regard to training data selection, feature selection, and reasoning algorithms.

Q: Going forward, what will it take to counter the growing discomfort about our ability to trust AI decisions?

Dana: The big secret about AI, machine learning, and deep learning is that 80 to 90 percent of the work is in the curation and preparation of the data that goes into a model. So the better job one does in that regard, the better models and predictions you come up with. As a consequence, systems begin to build a "track record" of trustworthiness, sensible decisions and apparent consistency. When we consider this question, the good news is that the subject of explainable AI is a hot topic of pursuit both in academia and industry. There's a near universal agreement that explainable AI and allaying mistrust are of paramount importance.

Q: What has led to this consensus?

Dana: Trusting AI has become a bit of a sensitive subject in an age when we hear so much about AI taking on humanlike decision-making roles in an ever-widening range of activities. In many cases, like operating a motor vehicle, we have always restricted these activities to the human domain. But the truth is that very few decisions are ever made by the seat of our pants in the modern era. For the most part, elaborate combinations of hardware and software systems support our every decision. Despite flaws and imperfections, we may have no alternative to relying on AI in the decision-making process. This, in turn, will lead to serious efforts to make sure we cover all the bases in creating trust of AI-ML decision making.

What are AI recruiters looking for?

"Perhaps the most valuable skill for an AI-ML candidate to have is adaptability in order to adjust to new and changing fields."

Mark Clark Senior Program Manager

Q&AI

Robotics. Thinking machines. Autonomous vehicles. Today, AI and machine learning (ML) are driving innovation and creating new careers in an ever broadening field. While landing a great career in AI-ML might seem improbable, new opportunities arise every day. Across industries, developers are creating machines that make decisions in much the same way as the human mind. AI-ML is not altogether new. However, these applications are increasingly running in the background, automating decisions, and appearing in places not seen before.

The skills and talents needed to create this AI-driven future are in high demand, and are critical to meeting important needs in the government sector from military readiness to cybersecurity. We spoke with Mark Clark, Leidos Principal Investigator and Senior Program Manager, about what it takes to build a career in the field of AI-ML.

Q: What are some of the major AI-ML trends that stand out to recruiters today?

Mark: First, people are excited yet fearful about what new technologies can do. Innovative software engineers and data scientists want to be a part of the group that is providing leadership in determining the capabilities and research directions of AI-ML. Recruiters are looking for candidates who aspire to shape the direction in which the field is moving. Second, universities are continuously challenged with preparing students to work in a field that is new and rapidly changing. To stay on the leading edge of the field, it's important for students to understand how educational resources can best be leveraged.

Q: How can AI-ML career candidates use their background to help them land a position?

Mark: Oftentimes a candidate's potential is hidden in their résumé, or not explicitly stated. Candidates looking to launch a career in AI-ML should always think about ways to go beyond the expected and list relevant participation in school or other independent activities on their resumes. A great example is designing an automated chess game or competing in robotics contests. Whether done in a classroom or not, these are experiences that stand out and are directly relevant in the field. Candidates can also use their resumes to point out nontechnical abilities, such as work ethic and a desire to grow. Conversely, recruiters should have the ability to look past traditional technical labels and field definitions and realize that mathematicians, statisticians, and engineers have valuable technical skills that can be adapted to data science and software design projects.

Q: What skills do you look for in AI-ML developers?

Mark: In a competitive recruiting environment, candidates with strong technical skills are in high demand. Employers look for coders and programmers who have put emphasis on developing strong mathematics skills, have the ability to leverage open-source technology, and are able to not only create applications from scratch, but also modify and understand existing applications. Throughout the education and training process, developers should have conversations about the specific programming languages needed to be successful. Of course, thinking critically and being analytical is always important—finding pitfalls in data collection, learning and applying new programming languages, and understanding user needs are all additional skills that potential job seekers need to possess.

Q: How are universities preparing students for AI-ML careers of the future?

Mark: Perhaps the most valuable skill for an AI-ML candidate to have is adaptability in order to adjust to new and changing fields. Oftentimes, students are best served by taking advantage of the new opportunities made available by their schools. George Washington University (GW) is a great example. For one, the school has embraced more interactive learning, which allows students to get hands-on experience with data science and experience the power of AI-ML. Schools like GW accomplish this by having students complete coursework via interactive notebooks. Completing projects and papers in this way allows students to integrate more cutting-edge solutions into the coding skills they are cultivating.

Q: What draws the best and brightest AI-ML talent to Leidos?

Mark: Leidos provides unique opportunities for candidates interested in AI-ML. We find that candidates are typically most interested in learning about the research we support, the relationships we have with research institutions, and how that relates to developing cutting-edge technology. So, what draws people to us when there are so many other companies right next door? The innovative talents, creativity, and diversity of our teams, project opportunities, and our internal referral network. These referrals are people that we would oftentimes be unable to reach with our traditional recruitment methods. We also reconsider candidates who come back to us with updated job skills and or professional experiences. These candidates readily exhibit their desire to grow and improve, and their interest in our work at Leidos.

Q: How do you see AI-ML developers affecting the technology landscape overall?

Mark: Even though AI-ML is an established field of research, it is still growing. Our process of developing technology and growing AI-ML capabilities provides important insights into how people are "re-thinking" the future. It is important to note that AI-ML is an enhancement to human intellect. As it stands, humans do not have the imagination to think of all the different ways machine learning can be applied, because we have inherent limitations. Leidos is working past the way AI-ML experts used to think, and we are investing in developers who are able to break down barriers and create new technologies that were previously unimaginable. Leidos researchers are working to get to a place in automation where manual tasks are more effectively completed by AI-ML empowered technologies that augment applications and make them more effective.

How does AI improve the way we develop computer systems?

"Machine intelligence has the advantage over human intelligence in that ML can view millions of results and extract information immediately in order to make the next decision."

Dr. Alric Althoff Leidos Research Scientist

The race to develop faster and more sophisticated computer applications is accelerating as AI and machine learning (ML) techniques continue to mature. While classical computers are capable of analyzing and processing certain types of data sets much faster than humans, the introduction of AI-ML extends research capabilities even further through simulated thinking. We talked with Dr. Alric Althoff, Leidos Research Scientist, who described how processes for developing computer systems are being expedited as smarter algorithms are combined with traditional research methods.

Q: How can AI-ML be applied when developing the next generation of computer systems?

Alric: Companies are developing new processors all the time, but integration with existing systems and components is tricky. A lot the integration time is spent ensuring that the applications a user cares most about will run fast enough and use a reasonable amount of power. In order to find a good final design, engineers test various design configurations and record what the effects are on the final results.

Now we're searching for good designs much more efficiently by applying AI algorithms to this testing process. Basically, a human can set up the algorithms to explore the possibilities, and the algorithm can run without human intervention to find good designs and hardware/software configurations. In practice, AI-ML can be used to identify good configurations faster than a human being, and allows non-experts to design better systems more easily.

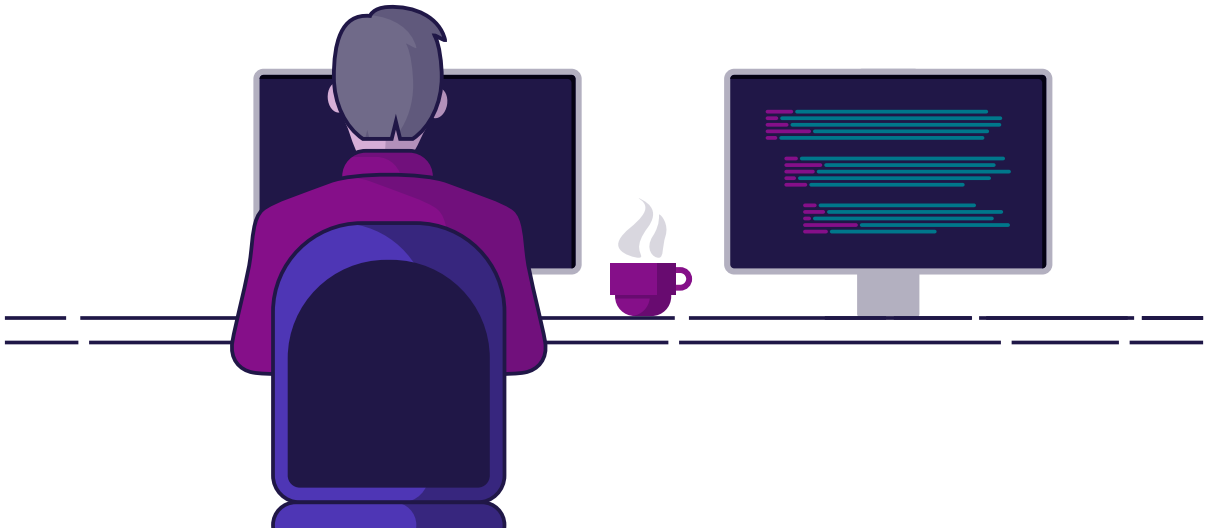

Q: What are the biggest challenges with using AI-ML in systems development today?

Alric: Many of today's challenges stem from the fact that AI-ML has received so much attention in certain domains while being relatively ignored in others. Due to the success of AI-ML in domains such as computer vision, there's a perception that AI-ML is ready to be integrated into areas where it hasn't gotten as much attention, regardless of the technical limitations that are still being addressed and overcome. So today, while ML methods are starting to become more mature, there's a lack of support in the workflow—particularly in the way that AI-ML algorithms interact with the tools that people already use. Often the algorithms don't behave well when you're trying to integrate ML into new processes or systems. So that is one of the biggest challenges in systems development—getting the workflow to play well with the ML algorithms.

Q: What are the benefits of machine intelligence in developing systems? In what ways is it superior to human intelligence?

Alric: For humans, there's a lot of complexity that needs to be managed in developing systems. A human might use a huge amount of time and energy on these things, but the machine has internal models that just keep track of that stuff. So where a human would lose sight of certain requirements, a machine can hold all of that material simultaneously in its view.

The advantage that a machine has over a human being in this situation is that ML can view millions of results and extract information from those results in order to make its next decision. In contrast, for humans, repeated experiments result in some fuzzy intuition about the nature of the problem. So a machine will be able to generate better priors over loss functions that it builds up about what we're trying to optimize over, even though we may not be able to state the nature of the problem from the outset.

These are specific methods in the field that allow us to leverage this machine equivalent of human intuition. This includes things like Bayesian optimization and reinforcement learning—that there's a feedback loop consisting of experimenting, observation, and the ML tool or algorithm adapting to the loss function values that it has learned in that feedback process at run-time. Glossing over the details, the takeaway is that Bayesian optimization achieves similar results in exponentially fewer samples than traditional supervised learning techniques. There are theorems that back that up, and this also aligns with our experience working with these algorithms. So applications where each evaluation of the loss function is very expensive will benefit from this approach.

Q: What domain interests you most?

Alric: Counteradversarial AI-ML is a really interesting domain. I have a background in statistical hardware security, and I often feel that the data science community treats threat detection, mitigation, and response more as if it were a data science problem than a security problem. By that I mean: collecting data, analyzing the data, designing algorithms that perform well given the data, and then publishing papers on the results. But on the security side, we know that you need to have a grasp of the underlying principles and design mitigations.

I think that we are going to see people who are security professionals focus on the hardness of the core problems of counteradversarial machine learning. They will be asking, in terms of safety and actual effectiveness, how strong should the emphasis be on using a stringent analytical approach? That is a big challenge, and I think that people are beginning to realize that. It's similar to the change that we're seeing across the machine learning research community as people tend to move toward 'de-biasing' algorithms, for example—moving toward robust algorithms in the context of how important it is that something be correct.

Q: What personally motivates your work in this field?

Alric: As human beings we observe and we interact with the world, and it informs us as to how we're doing. So we build up a worldview formed of a lot of little anecdotes and vignettes from our lives. Rarely is it a natural process to take a scientific outlook on the world, instead relying on a lot of assumptions. I think that's the major motivation for me. In the context of Bayesian optimization work, repeated trials contribute knowledge based on experiments rather than assuming a lot of prior information about the decision space. You can assume, but assumptions for human beings—and this is my personal feeling—don't really tend to serve us well at large scales. And so doing this sort of iterative fine tuning through automation is essentially applying the scientific method in data collection and experimental design. I think that is an important view to hang on to. It's also an important thing for the future of AI-ML, particularly as we move toward a world where we're relying more and more on algorithms to make decisions for us, to be aware of under-sampling, uncertainty, and applying models to regions where they do not generalize.

leidos

As a leading provider of data-rich solutions for the U.S. government, we understand our customers, their missions, and their data. Combining this with our expertise designing and building AI systems, we deliver solutions that provide immediate value for our customers. We take a disciplined, science-based approach to AI and ML that distinguishes our solutions.

For more information, visit **leidos.com/ai**