TAG

AI QUARTERLY

Q2 2024

MAPPING ARTIFICIAL INTELLIGENCE

ARTICLES / OPINIONS / INTERVIEWS

WELCOME TO THE INAUGURAL ISSUE OF TAG'S AI QUARTERLY

DAVID HECHLER, EDITOR

Before you get to the mapping portion of "Mapping AI," we have a feature package that's going to take you… well, all over the map. There's also a bit of time travel.

It starts with an old movie that's *not* "Back to the Future," but it has that feel. Because the 1991 film "Terminator 2: Judgment Day," which felt like action-packed sci-fi entertainment 33 years ago, feels uncomfortably prescient when you see it now. The battle between AI-powered bots and humans trying to prevent them from taking over no longer sounds like far-out fantasy.

Next we want you to meet our resident AI expert, who takes you on a journey into his past. Jay Wilpon is one of the earliest pioneers of speech recognition and natural language understanding technologies—a field in which he's spent his career—and he has the patents to prove it. And the knowledge. We think you'll find it fascinating to follow him on the arduous paths invention requires.

Fast forward as David Neuman imagines how AI will supercharge emerging technologies to power us into a better future (despite the warnings of danger ahead).

We end with two articles in the here and now. Shawn Hopkins lays out some of the challenges courts are grappling with as media companies sue AI shops for training large language models using copyrighted material without permission or compensation. And finally we asked Jay Wilpon to tell us when speech transcription will finally be good enough to allow us to dictate an introduction like this without having to check the transcript for errors.

After that we bring on the mapping. Our founder and CEO, Ed Amoroso, is the Map King. After introducing TAG's new AI advisory (following our cybersecurity and climate science practices), he presents a chart that describes the way companies are using and approaching artificial intelligence. It's followed by a detailed taxonomy that features 20 categories and 100 subcategories.

In the section below Ed's article, we provide an overview of the categories and short commentary on sample subcategories. And last, we interview the leaders of five startups—a selection of our Distinguished Vendors.

"ChatGPT says we should run Dive 24 Omaha."

CONTENTS

TAG AI QUARTERLY
Q2 2024

Lester Goodman, Director of Content

David Hechler, Editor

Analysts & Contibutors

Dr, Edward Amoroso
Joanna Burkey
Pete Dinsmore
David Hechler
Shawn Hopkins
Moriah Hara
David Neuman
Al Palimenio
John Rasmussen
Joe Sullivan
Jay Wilpon

Editorial & Creative

Lester Goodman
David Hechler
Miles McDonald
Rich Powell
Stephanie Amoroso

Research & Development

Matt Amoroso
Shawn Hopkins

Sales & Customer Relations

Rick Friedel
Michael McKenna
Laurie Mushinsky
Julia Almazova
Jane Mangiamele

Administration

Liam Baglivo

Dr. Edward Amoroso, Founder & CEO

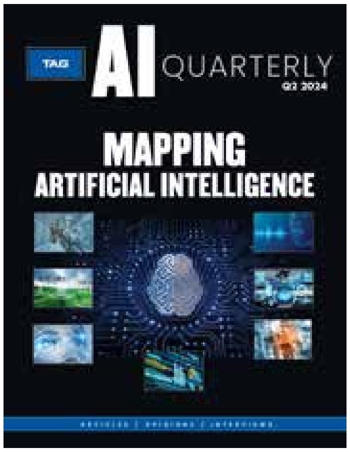

Volume 1, No. 1

Publisher: TAG, a division of TAG Infosphere, Inc.,
45 Broadway, Suite 1250, New York, NY 10006.

Experts and practitioners must recognize that best practices, technologies, and information about the artificial intelligence industry and its participants will always be changing. Such experts and practitioners must therefore rely on their experience, expertise, and knowledge with respect to interpretation and application of the opinions, information, advice, and recommendations contained and described herein.

Neither the authors of this document nor TAG Infosphere, Inc. assume any liability for any injury and/or damage to persons or organizations as a matter of products liability, negligence or otherwise, or from any use or operation of any products, vendors, methods, instructions, recommendations, or ideas contained in any aspect of the TAG AI Annual volumes.

The opinions, information, advice, and recommendations expressed in this publication are not representations of fact, and are subject to change without notice. TAG Infosphere, Inc. reserves the right to change its policies or explanations of its policies at any time without notice.

The opinions expressed in this document are those of the writers and contributors and in no way reflect those of TAG's Distinguished Vendors.

April 4, 2024

AI VIEWS

REVISITING TERMINATOR 2

IT HAS A LOT MORE TO SAY ABOUT AI, AND IT'S MORE OFTEN ON THE MONEY, THAN YOU MAY REMEMBER.

DAVID HECHLER

I was visiting my son in Michigan recently, and for some reason he suggested we watch "Terminator 2: Judgment Day." I'd only seen its predecessor, "The Terminator," and I saw that one during its theatrical release in 1984 (which was the only way we watched movies back then). I'd planned to see the sequel, which came out in 1991, but I never got around to it. Both films were written and directed by James Cameron and were well-received. The second, in particular, catapulted his career when it became the highest-grossing film of the year.

I enjoyed the movie, but this isn't a review delivered 33 years too late. What struck me all these years later is how prescient it was, how clearly it laid out the dangers of AI. And yet back then, I don't think it was perceived as frightening at all. There's nothing in the current **Wikipedia entry** that suggests anyone was expected to have an anxiety attack about our future. It was science fiction. It was entertainment. Yet, recent developments in AI have made it easy for us to understand the way humans were behaving in those movies. Because many of us have the same concerns right now. And the movies even predicted this would happen in this decade.

THE STORY

Let me briefly summarize the plots of the two movies. The **first film** lays out the backdrop. There will be a future war between a malevolent AI network run by artificial intelligence called Skynet, and humans are struggling to defeat this effort before Skynet gains sentience and wipes out humanity. In the future, the resistance will be led by a man named John Connor.

Skynet hatches a plot to gain the upper hand by sending an AI machine—the Terminator, played by Arnold Schwarzenegger—back from 2029 to 1984 in order to kill Sarah Connor before her son is born. The humans send back a soldier to protect her from this would-be assassin. At the end of the film, the Terminator is defeated at the cost of the soldier's life.

In the sequel, it's now 1995. John Connor is a young teen who doesn't know what to make of Sarah's stories about the future. He can't be blamed; his mother is locked up in a mental institution as a result of her "ravings." But the future isn't done with young Connor. Skynet decides to try again. This time it sends back a more advanced machine to kill John Connor. But it doesn't land alone. Again the resistance sends a protector—a reconditioned machine like the one that had failed to kill Sarah the first time. But this one is instructed to protect John Connor from this new threat.

Skynet's machine has all the advantages. The protector seems overmatched. But in the end, the resistance prevails. In the final scene, after the advanced machine is destroyed, the battered hulk that had helped the Connors survive tells them that he too must "die"—to ensure that humanity survives. John tries to argue with him, but Sarah understands that he's right. After a brief farewell, the robot leaps to his death.

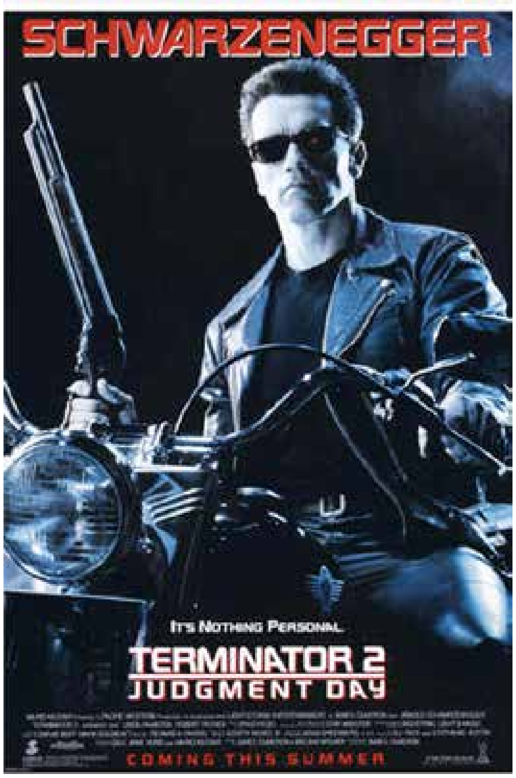

WHAT STANDS OUT TODAY

Arnold Schwarzenegger was famously the villain in the first movie. As the killer AI machine, he relentlessly hunted down Sarah Connor, killing every Sarah Connor he could find in the phone book in his quest to destroy his target. But in the second movie, he plays the AI defender—though neither the audience nor Sarah Connor understands this at first. When he arrives from the future, we all assume that the battle has once again been joined. But no, this time he's the "good guy"—or the robot equivalent.

BY TAKING THE VILLAIN AND REPROGRAMMING HIM THIS TIME INTO THE SAVIOR, THE FILMMAKERS ARE TELLING US SOMETHING IMPORTANT ABOUT TECHNOLOGY.

It was a brilliant stroke. By taking the villain and reprogramming him this time into the savior, the filmmakers are telling us something important about technology. The technology itself isn't evil. Nor is it good. It's neutral. It's the way humans *use* the technology that determines whether it advances constructive causes or destructive ones. At least until they achieve sentience and can act independently (assuming they ever do).

Another important theme was the way the humans insisted on anthropomorphizing Schwarzenegger's character. Some of the most engaging scenes were conversations between John, who never met his father (he was the soldier sent back to protect his mother and was killed by the Terminator), and the AI machine, who John seems to view at times as a father-figure. John tries to "humanize" him by suggesting phrases to sprinkle into conversation ("Hasta la vista, baby!") and teaching him how to smile.

But there are limits to how far this can go. John is baffled when the robot says he doesn't fear death. It's hard for the youngster to conceive of this. Speaking of which, the grand irony of the ending is that it was the *machine* that insisted he had to be destroyed in order to protect humanity from his ilk. The question we're left with as the screen turns black is whether, in his decision to self-destruct, the machine was merely following its orders to protect John Connor, or had actually empathized with his human companions and decided independently it was "the right thing to do."

THERE ARE EVEN DEEPFAKES

There's no mention of "deepfakes" in the film, but both of the AI machines are masters of the craft. They can flawlessly mimic the sound of anyone's voice, and they use the skill to their advantages. The killer bot can go one better, and he doesn't require a video camera. He is able to physically morph into a nearly perfect replica of Sarah Connor. Until his real mother appears nearby, John can't tell the difference, and it nearly costs him his life. It was another sign of the film's prescience: This was decades before **deepfakes** began to penetrate the public consciousness.

The foresight in this old film even extends to the dates. It's rather remarkable that we're now just five years away from 2029, the date that Skynet was on the verge of sentience and was battling to defeat humanity. Pretty good guess! Recent advances in AI have prompted serious conversations about whether something like this—sans the special effects—might really occur. I won't repeat the numbers from people in the field who now predict that AI will—or at least might—wipe out humanity. When I went back and found the sources of the information, I discovered that even skilled journalists had misrepresented the statistics.

They might just as well have used an advertisement from a scary old movie made two years after The Terminator: **Be afraid. Be VERY afraid!**

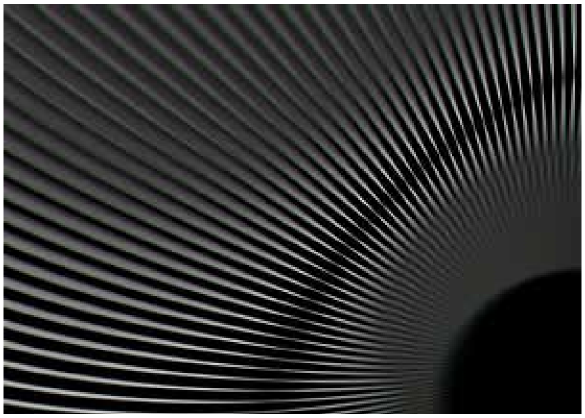

I'M TAG'S AI GUY

JAY WILPON

Now that we've finally launched this publication, I suppose it's only fitting that I introduce myself. In early 2023, I got a call from Ed Amoroso. He was looking to the future of TAG Infosphere, and he knew that AI was an essential component for enterprises and governments. They needed to prepare for the opportunities and risks that AI would invariably bring, especially in the security world. And he thought of me because we'd worked together for a long time. Ed had been a senior executive at AT&T responsible for the company's security, and I had been leading an AI team at AT&T Labs. In fact, I'd been one of AT&T's chief evangelists in AI for much of my 35+ years there. We discussed TAG's mission, and he asked if I'd like to join him on its quest.

Ed and I knew each other well. We'd discussed the future of the Labs, and we'd tag-teamed (no pun intended) presentations with AT&T's enterprise customers for decades. We both liked to talk (a lot) and help enterprises expand and protect their businesses. We also spoke our minds. I thought we were a perfect match at AT&T, and it didn't take me long to decide that the same would be true at TAG Infosphere. So, now I am a Senior Analyst leading our AI effort.

That's the straightforward part of my story. The more interesting development was how a kid who had no experience working with computers got hired to do this kind of work at AT&T in the first place.

HOW I GOT MY START

It was spring 1977, and I was interviewing for a job like any other soon-to-be college grad. I sat down in front of the screen at Bell Labs and Larry Rabiner said the word "seven" into a microphone. We went to lunch and when we came back there it was: the word "SEVEN" in big, bold capital letters on the Tektronix 4010 terminal. Geez, talking with machines. "This is f*@$ amazing!" I thought "It's '2001: A Space Odyssey.' I need to be part of this!"

A month earlier, I was sitting in my senior year statistics class at Lafayette College, where there were only five of us left in the math department. The professor said, “Are any of you interested in working at AT&T Bell Laboratories Research? They’re looking for someone with a math degree.” As a snot-nosed 22-year-old, I knew little about Bell Labs other than it was a major scientific laboratory. But it was close to my hometown, and spring break was coming up. And so I said, “Sure, I’ll go over there.”

OUR GOAL WAS TO TEACH A COMPUTER TO HEAR, UNDERSTAND, AND RESPOND TO HUMAN SPEECH. HUMAN-MACHINE COMMUNICATION IS WHAT WE CALLED IT.

I’d never heard of “artificial intelligence” or “AI.” Why would I? In the early ’70s, AI had entered a period of purgatory (a.k.a. the AI winter) when the realities and difficulties of creating and deploying useful AI could not compete with the promises and hype researchers and scientists had sold to the world. In 1969, the influential engineer and author John Pierce (also from AT&T) had published a scathing article on the prospects of speech recognition, calling the researchers “mad inventors and unreliable engineers.” He famously said that “the attraction [of speech recognition] is perhaps similar to the attraction of schemes for turning water into gasoline, extracting gold from the sea, curing cancer, or going to the moon.” These sorts of comments were repeated around the world. Funding for AI was cut everywhere.

So, there I was in 1977, sitting in a room at Bell Labs having no idea how my life was about to change, while AT&T, understanding the importance of speech recognition (not artificial intelligence—wink, wink), again decided to invent the future. I needed to be part of this—creating something from nothing!

After my interview, I went back to college expecting (or at least hoping) to hear that Bell Labs was inviting me to be part of the team, and contemplating whether I would give up the prospect of medical school to work there. And wondering how would I tell that to my parents. But nothing came. I called many times over the next few weeks and wrote letters (remember letters?). I even wrote a letter to John deButts, the CEO of AT&T. Yes, I was a real pain in the a**. And then, an offer arrived—as a summer intern. Hmm, not the big offer I was expecting. But I said yes immediately. I would prove my value and hope it was rewarded. And so my life in AI began on June 6, 1977, at AT&T Bell Laboratories Research in Murray Hill, New Jersey. Step one: learn to program a computer.

AT&T (2)

Above: Aaron Rosenberg, Stephen Levinson, Larry Rabiner, and Jay Wilpon in an AT&T speech computer room in 1978..

Right: Wilpon circa early 1990s at Bell Labs

Our goal was to teach a computer to hear, understand, and respond to human speech. Human-machine communication is what we called it. What were the words that a person said? What was the meaning of the words? How are words put together to form sentences (syntax)? What was the meaning and intent of the sentences that the person spoke (semantics)? And how should a computer respond to enable a conversation with a person? From a real-world, practical standpoint of a company like AT&T, with its hundreds of millions of customers, it was essential to be able to recognize any person, with any accent, age, or sex. We called that "speaker independence." Oh, and it would have to work if people were in their homes watching TV, in a quiet office, or in a crowded airport. There are a lot of moving parts and variability in human speech that we must consider, or model, for a machine to learn to "wreck a nice beach." I mean "recognize speech" *(damn speech recognizer!).*

AI AND MACHINE LEARNING (ML) ARE ABOUT PATTERNS—DISCOVERING THEM, MODELING THEM, FINDING THEM, PREDICTING THEM, GENERATING THEM, RESPONDING TO THEM.

FIRST WE LEARNED TO WALK

We began with just a few words: the digits zero through nine and the 26 letters of the alphabet. We would add more words later, assuming we could handle these 36. One of my first assignments was to record 100 people saying these words 10 times in a quiet, controlled environment. We were experimenting with machine learning algorithms to create a model from the voices of these 100 people that was representative of "everybody." We had no idea if this was even possible as no one had done it before. There were no open source databases, or any recorded speech that we could use. We had to create our own databases for years until the National Institute of Standards and Technology (NIST) entered the picture to help foster collaboration on data and testing across businesses, academia, and government.

This seemingly easy piece of the puzzle took us almost six months to complete We had to record audio on reel-to-reel analog tapes first, and then digitize them for storage on a computer. Oh, and by the way, the disks on our computers held a whopping 5 MB of data (as in megabytes). That is, we couldn't even store digitized audio files on our computers. (BTW, this spawned another line of research in our lab, leading to the inventions of MP3 and MP4 but those are stories for another time). So, we had to further analyze each audio file and only store a small set of mathematical features on the computer that represented the actual audio signal (even inventing new features, such as linear predictive coding or LPC). Invention is slow—and a lot of fun.

AI and machine learning (ML) are about patterns—discovering them, modeling them, predicting them, generating them, responding to them. It doesn't matter if it's speech, images, traffic patterns, or cybersecurity threats. In our case, we were experimenting with (and hopefully improving) various types of learning algorithms to determine if any were appropriate for recognizing patterns within a speech signal. And, since back in the good old 1970s most computers only had about 256 KB (as in kilobytes) of memory, any model we built, plus the speech recognizer, plus any application—all had to fit in 256 KB! You can imagine the limitations that put on researchers and engineers. Practicality gets in the way. Invention is slow, and hard, and still a lot of fun.

While this was going on, we were also exploring initial enterprise applications for AI. In 1979, we created and demonstrated the world's first voice dialing system. As in: "Siri, call home." Of course now everyone uses voice dialing. But in 1979, it was a new invention. This led to the first of my 50+ patents.

I loved my work. Discovering new things that can change the world is such a rush! It's hard to explain. For example, in the early days of speech recognition, all recognizers in the world could only recognize words that they were trained to find, like the ten digits. It would be impossible for them to recognize anything else. Knowing this, in the late 1980s we focused on simple use cases—those that would be crystal clear and super easy for people to use. We had a perfect first use case: automating the way people were asked to pay for operator-assisted calls. (People don't pay for phone calls anymore, but they did for most of the 20th century.) There were only five ways to pay for a call. So simple! That should make it easy for our speech recognizer, right?

A large team of researchers and developers spent several years building and integrating our speech recognizer into the telephone network to handle this application. However, in our early user tests, we realized that if you asked 200 million people to respond to a simple question asking how they wanted to pay for their call, such as: "At the tone, please say 'collect, calling-card, third-number, person-to-person, or operator,'" about 20% of them would say things like, "I'd like to make a collect call please." Yikes. The system doesn't know most of those words. And why are people being polite to a machine? Our AI would not work. That meant that 20% of 200 million people would try to use our system and would be completely frustrated. They had seen Star Wars IV, V, and VI by then, and of course machines could understand anything they said. Stupid system!

Everything stopped. BACK TO THE LABORATORY! Our speech recognizer needed to be able to "spot" the words the recognizer knew even if a person also spoke words that it didn't know. This led to probably the most frustrating 12 months of my professional life. New algorithm after new algorithm, new model after new model, experiment after experiment, failure after failure after failure. Absolutely nothing worked.

I was totally frustrated; ready to give up and move on to other AI challenges. I had one last idea. It was the winter holiday season of 1989. At the time, the idea seemed too simple to work, but what the hell. Try it anyway. I could set up the experiments and let the computer crunch away over a long holiday weekend. I would build my models and run the appropriate experiments, just as I had done many times over the past year.

I arrived on Monday morning expecting the same negative result. But this time something was different. The results were almost perfect. Almost too good to be true. I ran them again; same results. Holy S*&T! Enough with my test data. I needed to "talk" to the machine and see if this was real. Remember, this system only knew five phrases: collect, calling-card, third-number, person-to-person, and operator. That was it. I started easy: "Collect." The system responded "Collect" (good start). "Collect call please." The recognizer responded, "Collect." Two for two. Now an acid test. I said without pauses, "Mary had a little lamb collect its fleece was white as snow and everywhere that Mary went the lamb was sure to go." And there on the screen appeared the word "Collect." Oh, my g-d! We had done it. I ran down the hallway to my colleague's office and dragged him to my office. He had to see this now! He tried it and it worked.

SUCCESS AT LAST!

AT&T

Still from video of Wilpon at an AT&T speech lab circa 1984.

Over the next two years, working with a large team across all parts of AT&T, we productized our "wordspotting" algorithm and integrated it into the AT&T 4ESS and 5ESS switches (which ran the company's network). And in March 1992, AT&T deployed the first nationwide voice-enabled service, Voice Recognition Call Processing service, or *VRCP* for short. I could talk about those two years for hours—maybe later. VRCP handled 3 million phone calls per day (about 1.2 billion calls a year), making it the most widely accessed speech recognition AI system in the world at the time. It correctly recognized over 99.5% of the calls and was easily the most profitable speech recognition system in the world. It saved AT&T an estimated $3 billion during its first 10 years of operation. This service ushered in a new era of automation across enterprises—

AI automation. We convinced the world that speech recognition was real; that we had in fact extracted gold from the sea. Invention is slow, hard, and frustrating—and incredibly rewarding!

So, what's next? I was only 37 years old, and I wanted this feeling of discovery and accomplishment again. AI was alive and well. But after 15 years of work, we had only just scratched the surface. We were able to model and recognize a small set of specific words with a predefined syntax model of English that allowed the words to be combined in some rudimentary form to create sentences. In the lab, we even built a voice-controlled robot that we could talk to and have it move objects around on a table. However, it still could not understand what words and sentences meant, and we still weren't able to have a conversation with a computer.

THE SYSTEM WOULD TELL A PERSON TO SAY SOMETHING, AND IT WAS UP TO THE PERSON TO OBEY. WE WANTED TO FLIP THE EQUATION, PUT THE USER IN CONTROL.

NEXT STOP—UNDERSTANDING NATURAL LANGUAGE

To create an easy user experience, we had to go beyond words and learn (i.e., model) the semantics of language. How do we model the relationship and meaning among words? And how do we model and represent the intent/thought that a person was trying to convey by combining words together in a specific way? We call it natural language understanding.

If you think about it, when people called customer care 25 years ago, they were usually calling with symptoms of a problem or an action they wanted that they could explain in their own words. They would have a conversation with an agent, and hopefully resolve their issue. But could a machine take the place of an agent? For example, if someone said, "I installed your XYZ program on my computer and now my screen blinks blue," we could recognize the words, but then what? What did they want? Probably something like, "Hey, there is a bug in your software. I want to speak with someone who can fix it." Sadly, though, the first attempts at automating such interactions were annoying prompts that asked them to "press or say X" for this and "press or say Y" for that. Or, as we learned when we deployed VRCP, some people said things like, "I want to charge this call to my home." But what they "meant," in VRCP language, was "third-number." Customers were obviously frustrated, because they could not communicate with the machine in a natural way. We needed to advance AI.

Wilpon in a still from a video circa 1992

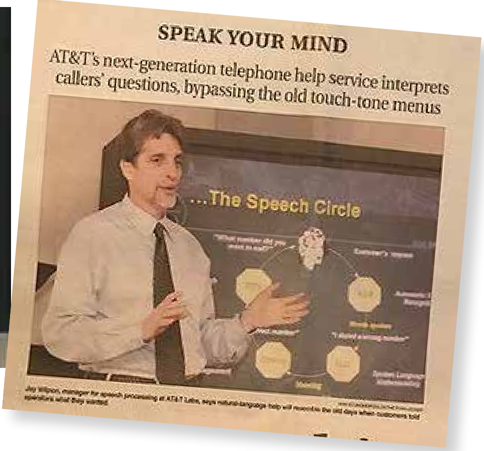

SPEAK YOUR MIND

AT&T's next-generation telephone help service interprets callers' questions, bypassing the old touch-tone menus

...The Speech Circle

Jay Wilpon, manager for speech processing at AT&T Labs, says natural-language help will [illegible] the old days when customers told operators what they wanted.

Article in The (N.J.) Star-Ledger

Fortunately, success had made believers out of people who had once thought AI was science fiction. After our success with VRCP, money started to flow, lots of money. And not just at AT&T. At established companies, startups, and government agencies across the country, more investment in AI was being made. We were able to hire more people. We were able to buy more computers and set up more labs. We could continue to expand AI.

To make AI easy to use, our group was going to change the user experience paradigm. With speech recognition systems, it was up to the user to know how to use the system. The system would tell a person to say something, and it was up to the person to obey. We wanted to flip the equation, put the user in control. Allow them to say anything they wanted and make it the AI's responsibility to figure out what they wanted. Instead of having people attempt to navigate complex menus and prompts, our goal was to give them a simple prompt—"How may I help you?"—and let them tell the machine what they wanted in their own words. At the time, no one thought carrying on a conversation with a machine was possible—except for a small group of people that I assembled into a new team.

And we did it again—something never before seen. In 2000, after eight years of dedicated work across many disciplines with many stakeholders, the first natural language understanding service anywhere was deployed for AT&T's customers. We obviously called it *How may I help you?* There were tremendous barriers that had to be overcome, including non-believers within AT&T and among the scientific community. But that's a story for another time. Three years later, we announced a new product line—AT&T VoiceTone—offering natural language understanding services to all our enterprise customers. Success!

After AT&T sold its speech and language AI and machine learning division a decade later, my team and I joined the privately held **Interactions Corporation**. There my career expanded to executive positions as SVP, Natural Language, and then as SVP, Corporate Development. I left four years later and began advising senior executives on AI.

Sometime along the way I was awarded the distinguished honors of **IEEE Fellow** and AT&T Fellow for my leadership in the creation of automatic speech recognition and natural language understanding algorithms and services. I appreciated the recognition. But for me the payoff has always been passion—and getting to work with some of the smartest people in the world.

That's why it was an easy decision when Ed Amoroso asked if I wanted to join his new adventure. I sense it's also passion that drives the people at TAG Infosphere.

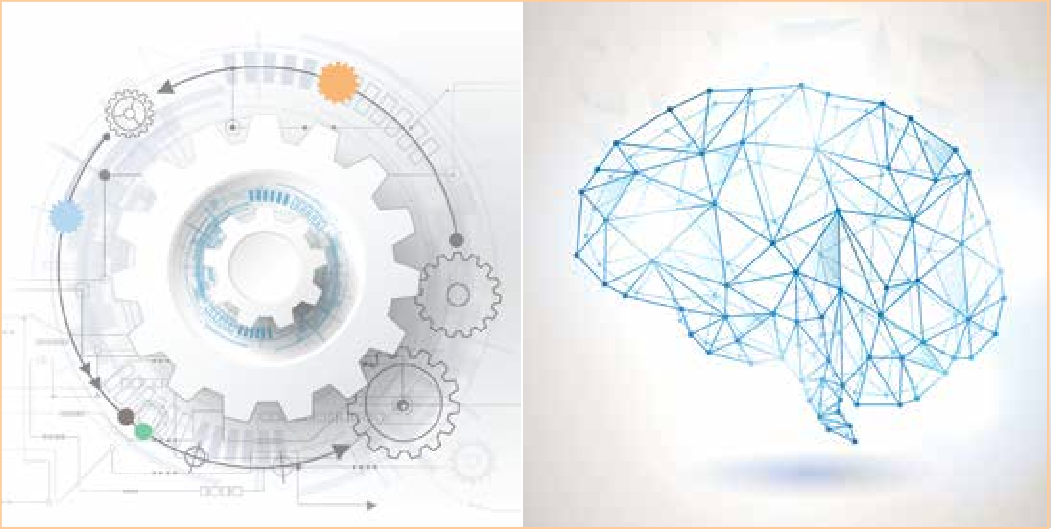

BEYOND AI: EXPLORING THE HORIZON OF TECHNOLOGICAL CONVERGENCE

DAVID NEUMAN

The recent acceleration in artificial intelligence is exhilarating, but it represents an incomplete picture. In addition to a powerful engine, AI needs a robust chassis to take us places. That chassis is the convergence of AI with other technological advancements, each playing a crucial role in shaping the future. Together they will catalyze unprecedented advancements, redefine paradigms, and unlock new opportunities for societal benefit—from autonomous healthcare robots to intelligent transportation systems and sustainable energy management.

The convergence will also address critical challenges, augment human capabilities, and shape a technologically advanced, sustainable, and human-centric future. This transformation's essence lies in leveraging AI's combined strengths with those found in other fields to foster interdisciplinary innovation and create solutions greater than the sum of their parts.

DIFFERENTIATING GENERATIVE AI FROM ARTIFICIAL GENERAL INTELLIGENCE

Just as a powerful engine drives the chassis of a car forward, AI serves as the driving force behind technological advancements, while other emerging technologies provide the chassis upon which this progress is built. Generative AI and Artificial General Intelligence are two key components in this vehicular analogy, each playing a distinct role in shaping the future landscape of AI. These two kinds of AI are at different stages of development and function differently. You need to understand them before you can fully understand how the engine and chassis function together.

GenAI excels in creating content, such as text and images, by recognizing patterns from large datasets. It operates within the scope of its programming and excels in tasks it has been trained for, yet it lacks a true comprehension of these tasks. On the other hand, AGI is the pursuit of technology's ability to understand, learn, and apply knowledge across a wide range of tasks, mirroring human cognitive abilities. Unlike GenAI, AGI is not confined to specific tasks but is designed to perform any intellectual task that a human can. AGI also encompasses the ability to apply common sense reasoning and to interact with the physical world in a way that GenAI does not. Therefore, while GenAI has made strides in specialized applications, the transition to AGI is a significant leap, requiring its capability to integrate and adapt knowledge broadly and profoundly, akin to a human's versatile intelligence.

While impressive at creating new outputs, GenAI may need help with tasks requiring genuine reasoning, critical thinking, or adapting to unfamiliar situations. Critics might argue this falls short of true intelligence, necessitating a more flexible problem-solving approach. Concerns exist that GenAI's capabilities are often overhyped, leading to misunderstandings about its nature and limitations. AI purists might argue that misrepresentation undermines the field's credibility and hinders progress toward genuine Artificial General Intelligence.

Here are some examples that illustrate the differences between these two AI domains.

GenAI:

- **Text Generation:** This encompasses tasks like writing poems, scripts, or even news articles. GenAI excels at replicating style and tone but struggles with more profound understanding and coherence.

- **Image Generation:** Creating and manipulating visual content like photographs or paintings falls under this domain. Again, GenAI excels at mimicking realistic visuals but may lack an understanding of composition or deeper meaning.

- **Specific Task Performance:** This broader category could include diverse areas like playing games, composing music, or even translating languages. While GenAI demonstrates impressive skills in these specific tasks, it struggles to adapt or apply its knowledge to new situations.

AGI:

- **Reasoning and Problem-Solving:** AGI is capable of going beyond just processing information, and instead can engage in logical deduction, inference, and problem-solving across various disciplines, such as medicine, robotics, transportation, and education. This implies understanding underlying principles and adapting to different scenarios.

- **Open-Ended Learning:** Unlike GenAI, which requires retraining for new tasks, AGI should be able to learn and adapt continuously, acquiring new knowledge and skills beyond its initial training data. This emphasizes a more flexible general intelligence.

- **Understanding and Common Sense:** True intelligence requires comprehending the world beyond data patterns. AGI can grasp concepts, relationships, and even common-sense reasoning.

- **Physical Embodiment:** While GenAI currently operates in the digital realm, AGI might require a physical body to interact with the world, navigate its complexities, and learn through experience. This includes visual and sensory perception as well as motor control.

- **Contextual Inference:** This would turn engines and chassis into intelligent vehicles that could not only sense a tire about to go flat but also infer what needed to happen next: find the nearest tire shop, ensure the right tire is in stock, make an appointment, and inform the driver, including providing directions. It also extends to advanced robotics and dangerous jobs currently performed by humans.

FORWARD-LOOKING USE CASES

So, what kind of chassis might AI be guiding in the near future? Imagine AI-driven robots intervening in emergencies, tackling the flames and chaos, which allows us to preserve human life by minimizing risk. And the closer they get to AGI, the more efficient they are likely to be, perhaps detecting the source of the fire more swiftly than their human counterparts can.

Technology like advanced robotics and AI will be used to shield humans from all kinds of dangers. For instance, from burn and inhalation injuries of firefighters, As a new grandfather, I marvel at the future possibilities for my granddaughter. I envision a day not far from now when she might be surprised by the dangerous jobs humans once did. “Grandpop, did humans really run into burning buildings to save people and fight fires?” she might ask. Hers will be a world where intelligent robots and human-operated instruments are the frontline warriors in crises, reflecting a blend of technology and valor that reshapes public safety.

The advancements won’t stop there. The confluence of genomics and artificial intelligence promises to usher in a new era of personalized medicine. Researchers will delve into our genetic makeup, unraveling the mysteries of diseases at a molecular level to provide tailored treatments that nip chronic conditions in the bud. Meanwhile, armed with Big Data analytics, public health officials will manage epidemics with unprecedented precision, leading to swift and effective health care delivery. As we have seen in the global COVID-19 pandemic, deadly viruses will morph many times. The reasoning of future technology powered by AI will predict the potential disease and vaccine efficacy changes to make proactive decisions to protect people and communities.

In a world rife with information, the challenge will be to sift through the noise. Artificial intelligence will be pivotal in combating misinformation, wielding natural language processing and deep learning to detect and dispel falsehoods. But technology isn’t a panacea; human discernment will be equally crucial, ensuring a balanced approach to maintaining the integrity of our public discourse. Citizens are bombarded with a wide variety of information mixing opinions and facts. But it is up to people to interpret this information. While technology could provide greater accuracy and analysis of the information we consume, I’m skeptical of how people will use it.

> **ARTIFICIAL INTELLIGENCE WILL BE PIVOTAL IN COMBATING MISINFORMATION, WIELDING NATURAL LANGUAGE PROCESSING AND DEEP LEARNING TO DETECT AND DISPEL FALSEHOODS.**

This balance between machine efficiency and human judgment will also extend to our energy systems. Fusing renewable energy with smart grid technology can transform our approach to electricity, driving efficiency and mitigating climate change. Similarly, smart cities represent the zenith of urban planning, optimizing everything from traffic to waste management, making city living more sustainable and enjoyable.

Artificial General Intelligence promises to revolutionize the judicial system by democratizing access to expert legal advice and representation. By leveraging AGI, we can envision a world where legal assistance is available not only to the affluent and those with specific connections, but accessible to all individuals, regardless of their socioeconomic status. This shift would have the potential to bring about greater equality within the judicial system, ensuring fair representation and justice for every member of society. Such advancements could lead to more equitable legal outcomes and contribute to dismantling systemic barriers, fostering a more just and inclusive legal landscape.

The way we learn and train for the future will also be redefined. Gone will be the days of static textbooks and passive lectures, replaced by interactive AI and holographic technologies that encourage critical thinking and lifelong learning. To be clear, as an educator myself I do not believe this will replace teachers. It will empower them. I recently had a young student tell me she believes if students "can't see it, they can't be it." This is where interactive AI and holographic technology could introduce a different way of learning. What would be the experience of a learner if, through advanced visualization, she could see the internal organs of the human body? How might technology influence other fields of science? In tandem, the automotive and transportation sectors will evolve with augmented reality and AI simulations, which will improve design, safety, and training.

Lastly, as our lives become ever more digitized, AI's role in cybersecurity will become central to our well-being. Sophisticated algorithms will guard against cyber threats, ensuring the digital world remains a safe space for innovation and connection.

In this tapestry of technological marvels, artificial intelligence stands not as a mere thread but as a vital weave that binds the fabric of future society. For my granddaughter, these advancements will form the natural backdrop of her life—a world where the once-fantastic is now routine and where each new day brings with it the promise of human ingenuity working in harmony with the machines we've created.

FINAL THOUGHTS

I recognize that my vision here will strike some as utopian. Other commentators have underscored the dangers AI may present—particularly AGI—in the coming years. Considering these in any technological advancement is realistic, but we must also challenge those perspectives and chase visions and dreams. I'm certainly happy Orville and Wilbur Wright weren't discouraged by the many obstacles they faced in achieving the first lighter-than-air powered flight.

Contemplating the horizon of other technological advances, we stand on the precipice of unprecedented opportunities. From the transformative potential of AI integration with diverse technological domains to the emergence of innovative solutions that address critical challenges, each use case conjures a future teeming with promise and possibility. Whether the convergence of genomics and biotechnology advances personalized medicine or the synergies between AI and cybersecurity fortify digital ecosystems, they underscore the transformative power of interdisciplinary innovation.

But it's up to us to seize the opportunities to chart a course toward a technologically advanced, sustainable, and human-centric future. Through collective imagination and concerted action, we can harness the full potential of technological convergence to shape a world defined by progress, prosperity, and inclusivity.

GenAI and the Copyright Dilemma

SHAWN HOPKINS

Over the course of history, advancements in technology have accelerated growth and unlocked unforeseen opportunities. Yet, along with these innovations come complex legal and ethical dilemmas, often outpacing existing regulatory frameworks. The swift emergence of affordable, accessible generative AI platforms capable of generating diverse content types such as text, audio, video, and images is a prime example. Governments around the world, including in the United States, are grappling with novel legal issues left unaddressed by current laws. Among the myriad concerns, copyright laws have become a focal point, illustrating the multifaceted challenges posed by evolving technologies.

The past year has seen various aspects of copyright law being tested in the courts, necessitating updates or entirely new legislative measures. One significant focus revolves around examining whether training large language models, vital for generative AI, constitutes copyright infringement when copyrighted material is employed without authorization or proper compensation. The recent **lawsuit** filed by The New York Times Company against OpenAI and Microsoft marks the latest development in a string of legal disputes involving artistic content creators and companies involved in generative AI (**GenAI**). Notable figures embroiled in similar disputes include actress Sarah Silverman, novelists Jonathan Franzen and John Grisham, as well as the photography syndicate Getty Images.

In order to provide the context for arguments surrounding the regulation of copyrighted training datasets, an overview of the creation process of GenAI models, along with a basic explanation of copyright law, is presented below.

Over the years, numerous rulings in U.S. courts have interpreted the rights of copyright holders, and these nuanced decisions shape the evolution of copyright law. Comprehending copyright principles, data utilization practices, and the legal implications of copyright infringement are among the issues explored by the U.S. Copyright Office USCO.

Established over 150 years ago as an integral part of the Library of Congress, the USCO oversees the implementation of copyright laws within the nation. Its Duties encompass registering copyright claims, documenting copyright ownership details, disseminating information to the public, and offering guidance to Congress and various governmental bodies on a diverse array of copyright matters. In addition, Congress has delegated the office authority to develop regulations concerning many areas of copyright law.

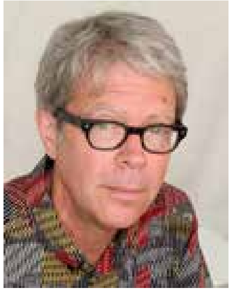

Actress Sarah Silverman and novelists Jonathan Franzen (top) and John Grisham are involved in legal disputes involving generative AI.

COPYRIGHT BASICS

Copyright law secures intellectual property by protecting original works of authorship as soon as an author fixes the work in a tangible form of expression. This protection can be extended to include many different types of works, including paintings, photographs, illustrations, musical compositions, sound recordings, computer programs, books, poems, blog posts, movies, architectural works, plays, and more.

Anyone has the right to claim copyright for their original work, which encompasses a bundle of associated rights. Pertinent to the current argument are rights such as the ability to (1) reproduce the work in copies; (2) create derivative works based on the original; and (3) authorize others to exercise these rights within certain statutory limits.

A derivative work refers to a piece based on an existing work but incorporating new, original, copyrightable material. Examples include motion pictures adapted from books, paintings inspired by photographs, and fresh interpretations of existing songs. Only the current copyright holder can authorize derivatives of this sort.

COPYRIGHT INFRINGEMENT

At the heart of the debate that recent developments in AI has provoked lies this question: Are AI-generated works reproductions, derivatives, or original creations? In other words, does AI output sufficiently distinguish itself to meet the standard of originality? Moreover, can a copyright be awarded to a machine?

The determination of whether a work qualifies as a derivative hinges on four basic rules: (1) The derivative must be original and copyrightable; (2) it should not merely replicate the original; (3) new material must incorporate preexisting copyrighted work; and (4) the creator must have contributed new elements or transformed the original. This raises the critical question of how “originality” is defined.

WHERE THE COURTS STAND

Daniel J. Gervais, PhD, dissects the nuances of law surrounding derivative work in "The Application to the Derivative Work Right to Literary and Artistic Productions of AI Machines." In his paper, he cites the view adopted by several U.S. Circuit Courts of Appeal in defining a derivate work. "A work will be considered a derivative work only if it would be considered an infringing work if the material which it has derived from a preexisting work had been taken without the consent of a copyright proprietor of such preexisting work." The courts in these infringement cases have focused primarily on the protection of the markets for the original work and potential derivatives made legally. These are important distinctions that are addressed by comments to the USCO.

Other Circuit Courts have taken different approaches over the years, focusing on the extent of originality and transformation. This approach measures the amount of transformation and variation from the original to the derivative. In one such case, the Seventh Circuit Court of Appeals found that a painting copied from still-life photographs of Dorothy from the movie "The Wizard of Oz" failed to qualify as an original work. The artist changed the background to a different scene from the movie, but this was not persuasive enough. However, an exact replica painting from a photograph of a scene in nature was deemed original.

THE DEBATE ABOUT WHETHER THE TRAINING OF LARGE LANGUAGE MODELS, CRUCIAL FOR GENAI, CONSTITUTES COPYRIGHT INFRINGEMENT WHEN COPYRIGHTED MATERIAL IS UTILIZED WITHOUT AUTHORIZATION OR PROPER COMPENSATION IS JUST ONE ASPECT OF THE MUDDLE SURROUNDING COPYRIGHT LAW.

AI MODEL TRAINING

In its simplest form, a GenAI model is a combination of computer algorithms working together to emulate human thought. These algorithms process large amounts of data and recognize patterns to help operationalize the decision-making process. Different types of structured and unstructured data are involved with "training" a model. These datasets come from (1) information that is publicly available on the internet; (2) nonpublic information obtained from third parties through commercial arrangements; and (3) information that users or human trainers create and provide.

Models are made up of large strings of numbers (called "weights" or "parameters"), which software code interprets and executes. The most powerful models consist of billions of weights. Each weight roughly reflects the statistical relationship between different datapoints in different scenarios. As models "learn" during the training process and become better at predicting the next thought or process, their weights update to reflect this improvement. When asked for a response, the model uses its weights to create a new response each time it is asked.

The fundamental operational logic of GenAI lies in generating outputs based on the analysis of billions of observed data points and their interrelationships in past instances. The likelihood of a sequence recurring is determined by how frequently one data point follows another in the observed datasets, assigning higher probabilities (weights) to more frequent sequences. For instance, when considering written text resembling a specific author's style, the program considers recurring phrasings, colloquialisms, and character perspectives commonly observed throughout their body of work. Such stylistic attributes serve to differentiate one writer from another. The GenAI model systematically characterizes every facet of the style and computes all conceivable combinations, with some combinations recurring more frequently than others. This systematic approach enables a GenAI model to predict text in the style of a specific author when provided with a scenario conducive to their writing style.

All that is background for last year's initiative undertaken by the USCO. On August 30, 2023, it published a **notice of inquiry** as part of its study regarding copyright issues raised by GenAI. Multiple respondents representing copyright holders and GenAI developers provided comments in support of their respective positions. There were two submissions of particular note asserting where the protections should lie.

WHAT COPYRIGHT HOLDERS SAY

The American Society of Composers, Authors and Publishers (**ASCAP**) presented **arguments** for the need to protect against infringement of existing copyrighted material used in the AI model. Titled **"A new federal right of publicity is necessary to adequately protect creators,"** it argues that "generative AI tools—unlike human creators—are not merely 'inspired,' 'guided,' or 'informed' by the protected works of others. Instead, these tools are based on the wholesale copying and ingestion of works by particular creators, whose content directly affects the behavior of the algorithmic model underlying the AI tool." The comment goes on to state, "These **AI tools** that are built using exact copies of sound recordings of music creators enable mimicry of those creators on a scope and at a quality never seen before. Without allowing the artists and creators to control their voice and likeness, this technology will create both consumer confusion and serious financial harm to the original music creators." The implications of harm are not exclusive to that of music but all forms of artistic creation.

ASCAP's position is that AI generated material, no matter how unique it may be in the arrangement of data, is clearly an unoriginal derivative since each bit and byte were potentially gleaned from some previous copyrighted material. Clearly, ASCAP wants the bar of transformation and originality raised for GenAI to also include style. This inference being that style, at least as interpreted by computers, is an amalgamation of pieces from preexisting work.

In the lawsuit **filed by The New York Times**, it has been alleged that, following the training of GenAI on various datasets, a process known as "fine-tuning" occurs. During this process, the AI is adjusted using specific types of works to replicate their content or style more accurately. Human feedback is provided to reinforce desired outputs, thereby enhancing the AI's ability to mimic the desired characteristics.

The New York Times sued OpenAI and Microsoft over the use of the newspaper's content to train AI.

WHAT GENAI DEVELOPERS SAY

The other respondent of note—this one opposing new copyright standards—was **OpenAI**. Its statement argues that "generative AI does not retain training data and verbatim repetition of such data is considered a bug by developers." The process makes every effort to prevent the repetition of training data, the statement went on. Nonetheless, if repetitive data exists in different datasets, the AI may naturally associate certain words or concepts, akin to human cognition. For instance, if the AI model is prompted with "how do I love thee?"—the opening line of Elizabeth Barrett Browning's Sonnet 43—it would likely generate the expected response, "let me count the ways." Given the widespread familiarity with this opening line across various contexts over the years, such an association is almost instinctive, OpenAI said.

Furthermore, the company argued, style is not simply a sum of parts plucked from artistic works. Rather, it’s a more sophisticated unlocking of artistic expression broken down into complex logical choices. Hence, the GenAI model is choosing the next phrase, musical note, color, etc., based on learned patterns of association. This is also what humans do when creating artistic works. For example, Michelangelo imitating Donatello or Greta Van Fleet sounding like a modern Led Zeppelin.

WHAT DOES THE FUTURE HOLD?

Given these arguments and the inconsistent judicial history, updating copyright law to accommodate the advancements of GenAI presents significant challenges. Similar struggles are evident among governing bodies worldwide grappling with these issues. For instance in February 2023, the U.K. abandoned a proposal that would have permitted extensive use of copyrighted music for training AI models without licensing or compensation. Additionally, the EU has introduced some safeguards through its Digital Single Markets Directive, allowing rightsholders to opt out of having their works utilized for AI training. However, a piecemeal approach across jurisdictions introduces complexities. These complexities were among the primary concerns of the Writers Guild of America regarding the use of AI in film and TV.

The decision to regulate or not presents a dilemma. One argument posits that allowing the utilization of preexisting works as training data jeopardizes entire industries, potentially leading to the displacement of numerous workers. And this would create great hardships for companies and people.

Others point out that technological advancements throughout history have given rise to new industries while rendering others obsolete. Why should AI be treated differently? Consider Johannes Gutenberg and his printing press, which replaced scribes and displaced some monks. Similarly, the introduction of word processing by Wang in the early 1970s and subsequent advancements significantly reduced the secretarial workforce, estimated to have declined by over 50% through 2015.

The debate about whether the training of large language models, crucial for GenAI, constitutes copyright infringement when copyrighted material is utilized without authorization or proper compensation is just one aspect of the muddle surrounding copyright law. We’ve only scratched the surface of this multifaceted argument, one that will undoubtedly challenge courts and legislative bodies worldwide for years to come. As evidenced by past court rulings on copyright infringement, there exists no clear-cut standard defining what is legal and what isn’t.

Interestingly, **OpenAI’s charter** emphasizes a commitment to leveraging any influence over the deployment of GenAI for the collective benefit, while avoiding the facilitation of AI applications that pose harm or overly concentrate power. Meanwhile, previous court decisions have primarily prioritized safeguarding the markets for original works. Similarly, ASCAP’s adoption of the “six principles for AI” advocate a “Humans First” approach, underscoring the importance of respecting human creators’ rights and work. And OpenAI seems to agree. Innovation, OpenAI proclaims, should not come at the expense of creators.

Is it possible the sides may yet find common ground?

FOR DOCTORS AND JOURNALISTS, IT'S BEEN AI'S HOLY GRAIL

TO BE ABLE TO RECORD AN INTERACTION AND LET A MACHINE TRANSCRIBE IT HAS LONG BEEN THE DREAM. IT'S FINALLY HERE—ALMOST.

DAVID HECHLER

I have spent most of my career as a journalist. The information a reporter gathers mostly boils down to documents and interviews. And many interviews are conducted on deadline, so you're writing or typing as fast as you can—hoping what you're getting is accurate. But sometimes for a feature article, or one that requires more depth, you want to record the conversation to capture it all. But that means you have to spend a lot of time transcribing. Or at least you did until recently.

For many journalists the Holy Grail has long been speech transcription software that could save us the time and drudgery of transcribing recorded interviews ourselves. And I, for one, was happy to celebrate when it appeared that advances in artificial intelligence had finally produced speech recognition software that worked. But like so many "saviors" that promise a whole new way of life, it has not delivered all it had promised. At least not yet.

I will show you just a few examples of what I mean in this article, which was transcribed by the program I use—and required hours to fix. We've [highlighted in brackets] *the transcript's errors to make them easy to see.*

But I'm not here to kvetch. I'm happy to have it. It's just that when an expert on this subject joined our company, I couldn't resist asking some questions. I wanted to understand the challenges software developers face in creating the kind of product I want. That's why I asked Jay Wilpon, one of the true pioneers of speech recognition technology and now TAG's AI expert (see the article by him in these pages), to sit for this interview. I was hoping he'd tell me that perfection is just around the corner.

David Hechler: Jay, have you ever used software to transcribe an interview or a speech?

Jay Wiipon: Yes. I use it all the time for first drafts of articles and to dictate text messages. The accuracy is imperfect. It's good enough to be able to use to get my thoughts down quickly. And I like talking better than typing.

Hechler: Have you done work professionally related to this field?

Wiipon: Yeah, I started working on speech recognition before there was speech recognition—back in the mid 1970s. In 1977, I joined AT&T Bell Labs and helped create a lot of the technology that led to the modern dictation software that you're using today.

Hechler: I noted that the company Dragon, which I'm going to refer to, was founded in 1975, although Dragon NaturallySpeaking didn't come out until the '80s. So we're talking about a similar time frame. When I became aware of this technology, I remember that doctors often recorded their patient notes on Dictaphones. And then they gave their tapes to an assistant to transcribe them. I'm sure those doctors were eager to be able to avoid paying assistants to type their notes. They were probably some of the early people eager for the technology we're talking about. And then, as a journalist, I was constantly looking for a product that would obviate my need to sit down after a recorded interview and type away. And then have to play back, and rewind, and listen to it over and over—to save me all that time and spare me that drudgery. And Dragon NaturallySpeaking was the first one that got good reviews, that sounded like maybe this is possible. But it was, in its earliest form, designed for one speaker, to dictate kind of like the doctors did. And that person dictating had to devote quite a bit of time to train the software to recognize his voice, his inflections, his accent. So the individual had to do a lot of work just to have it be somewhat accurate. It was clear to me that this wasn't what I was looking for. It wasn't for an interview, because I would have had to train it myself. And then it wouldn't have been trained for the person I was talking to. And since that person was changing all the time, it didn't seem like it was ever going to work for me. From what I read it did get better over time, but it wasn't close enough for me to decide to buy it. Because that's what I would had to do: spend several hundred dollars and then try it out. Does this match your recollection of the slow progress on the way to getting technology that could transcribe interviews?

Wiipon: You gave a good summary of the challenges that were in place almost 40 years ago. The first dictation product probably came out in the early 80s. Yes, Dragon also came out around that time. There were a number of companies that were offering dictation software. And one of the variables you actually missed. For the early dictations, you… had… to… speak… one… word… at… a… time. Not only did you have to train it to understand you, but you had to speak in that unnatural way.

Hechler: Yes, I remember that now. It seems that there was mutual training going on. There were a lot of ways in which the technology trained the human to behave a certain way. What you just said is going to be hard for me to transcribe because the software isn't going to know how to add those pauses you inserted a minute ago. I'll have to add dots between the words.

Wiipon: The point I was making was that being able to recognize what before was sort of isolated speech, one word at a time, to be able to move to continuous speech was a scientific achievement. And it took a a decade once we could recognize isolated speech reasonably well, to be able to recognize continuous speech very well. The second variable, you've got to look at at the computing capabilities. You had kilobytes of memory, or megabytes maybe back in the '80s. Even if you could train something, even if you had the data, the machines couldn't physically handle it in real time. So everything always started very small: 10 digits, alphabet, the dictation piece you talked about with doctors, especially radiologists. They were one of the first users of dictation engines. And it may sound like that was hard because they were saying a lot of polysyllabic words. But radiologists have a very short vocabulary. Most of it is boilerplate. They say one or two words and it generates a paragraph. So actually, it was very useful for them to increase their efficiency.

YOU DON'T WANT TO LEAVE A DEATH PENALTY CASE TRANSCRIPT IN THE "HANDS" (QUOTE, UNQUOTE) OF A MACHINE. AGREED?

Additionally, things like Dragon—you said something about hundreds of dollars. When Dragon came out in the '80s, it was around $10,000. Who else was going to be able afford that besides doctors? So it wasn't a cheap piece of software. A lot of research went into it. Research that the enterprise private sector did, a lot that the government sector did. Dragon was partially funded by DARPA. And DARPA allowed people to form their own companies. And the Bakers [Dragon's founders] did that with Dragon. What changed the equation on the price was IBM coming out with a dictation product in the '80s for a few bucks. Then in 1999, IBM released its ViaVoice program for $200. And so literally overnight Dragon had to drop their price to be competitive. Dragon wasn't long for this world. There were a bunch of mergers in the speech industry. The company eventually merged into Lernout & Hauspie and then ScanSoft, which eventually merged into Nuance.

Hechler: So it was only within the past 10 years that this has changed substantially, at least in my experience. You've now got Otter.ai, more recently Google Live Transcribe, Microsoft Word Dictate transcribes, and most recently OpenAI Whisper has joined the party. I haven't tried them all. But they seem to be much better. You agree?

Wiipon: They're infinitely better for a few reasons. One, 20 years ago the vocabulary sizes were still very small. Hundreds of words, maybe thousands of unique words. And you had to train it, and you had to know the accents. Now, speech engines can pretty much recognize any word that you could say, and it's easy to add new vocabulary. You can also speak normally. You can speak in a crowded environment and it will still pick it out based on newer algorithms and signal processing. It's like night and day. It's like the caveman inventing the wheel versus having a Tesla. It's completely usable, except for things that are critical to my life. I'm not going to use it to do something that could affect the safety of my family, for example.

Hechler: That's a perfect segue to my next question. To add context, I don't think we've ever seen machines relied on to transcribe testimony in court. Humans are required for that work—they're called court reporters—even where courtroom testimony is recorded electronically. I expect that will continue for a long time, because accuracy is so important. You don't want to leave a death penalty case transcript in the "hands" (quote, unquote) of a machine. Agreed?

Wiipon: I think it's going to be that way for a while. It may be the case that the machine can transcribe it better. But right now it's not allowed. The law is the human sitting there with their transcriber machines. I think that will change over time. It's an obvious evolution.

Hechler: Where accuracy is *not* so important, transcriptions have crept into the culture. When you get a voicemail on your iPhone, you don't have to listen to it. Usually, there's a transcript you can read instead. Zoom meetings like the one we're on now can be recorded and transcribed simultaneously. So what do you think about these developments? And are there other examples you want to throw in?

Wiipon: Those things will continue to happen. Another place, to continue the media theme, is television shows, movies. The accessibility acts [under the ADA] requires closed captions. The funny thing is that up until maybe 10 years ago, things like YouTube didn't require that because they weren't broadcast over the waves. But now you see that if you go on YouTube, you can get transcriptions in pretty much real time of anything you see there, because Google can put their speech engine on there, collect the data, train it, and build great models. So, from an accessibility point of view, for people that are disabled—have a problem hearing—it's a boon to be able to apply this technology to make their world much better. It's happened pretty much organically. Because once technology gets good enough, people will use it.

But there's another aspect of this technology we should discuss. Even in the early days of dictation, the companies you mentioned—Dragon, IBM—the thing that distinguished their products more often than not wasn't their accuracy. It was the user interface that allowed users to easily dictate and easily edit on the fly, go back, change a word. If they weren't sure of a word, the recognizer might highlight it with an underline that you could go back to. So a lot of what made these companies differentiate themselves was the user experience they provided to enable people to use a flawed technology.

Hechler: Let me turn to my own experience using this software. For my work, it is essential that the transcript I get of an interview like the one we're doing right now is accurate, because we're going to publish it. That means I have to rigorously review the text against the recording. I'm always disappointed by the results. My only consolation, and it's substantial, is that I remember how much work it was when I had to transcribe it all myself. There have been times I have spelled out words or acronyms in the recording, and the program still got them wrong. Here's an example. It's not so much spelling it out as using an acronym. When I refer to a company's information technology department as the IT department, IT is too often lowercased. It can be pretty confusing when someone says something like, "I called IT and asked what the problem was." And then it comes out, "I called it and asked what the problem was." **[In the example above, the only time "IT" was capitalized in the transcript was the first time I used it.]** Have you run into these kinds of problems over the years in your work?

Wiipon: Yeah, this gets to language. Speech recognizers are there to listen and transcribe. They don't know about punctuation. They don't know about accents. They don't know about pauses. They don't know about the things that you're talking about, like acronyms. They have to learn those things. Up until a few years ago, you didn't get punctuation when you dictated. You just got a long string of words. If you want to send a text message and use dictation—even now, I say something and I add "period" or "comma," because it doesn't know where to put any kind of punctuation. The difference between dictation, which is just recognizing the words without meaning, and more advanced forms, which I believe the modern dictation engines are moving towards, is to understand a little bit of the meaning.

Hechler: We'll come back to grammar. But first I want to say that I recognize jargon is likely to confuse the program. But I expected it to improve over time, and it hasn't. I interview people about chief information security officers regularly. Really **[the program rendered this word as "Billy"]** all the time. The acronym is [spelling it out] CISO. And I pronounce it *seaso*. But sometimes I will actually spell it out. It doesn't even get all those letters right when I spell it out. So it seems clear to me that there is no machine learning going on with the product I'm using. My subscription plan allows me to teach the program five words or phrases per conversation. That's it. I presume it won't remember them the next time. It's very frustrating. Isn't there a better way for the company to handle this? Can't they build in some form of machine learning without requiring payments and my having to train it? This takes me back to Dragon NaturallySpeaking.

Wiipon: Well, it's always the case that things advance forward slowly. Once things become good and ubiquitous, then you make it part of the product. If there's things that everybody's got, you make them available, and you differentiate yourself with other things. I don't want to second guess what people's business models are for generating revenue, but there's a reason that you have a bunch of software that's open source. And there's a reason that there are tools which allow companies to be able to customize AI, in this case speech recognition, so that there's a business. But over time, it does improve.

Hechler: Yes. That's the reason I'm using this software now. It is so far better than anything that used to be available. I want to be clear that once I saw what it could do, I didn't hesitate to start using it. That had never happened before. And I'm not unhappy that I did. It has and does save me an enormous amount of time. It's just that I really expected that there would be improvements—either from what the company was able to offer, or that somehow it would learn from me, from my work.

Wiipon: How do you know it's not? It definitely is. [This came out: "I know it's not. It definitely is."] Take Siri. Let's say you've got five Jays in your contact list. And the first time you say, "Call, Jay" it may say, "There's five Jays, which one do you want?" After you've done it for a while, you say "Call Jay" and it knows which one you want and it dials it. So clearly the machines and the technology and the services are learning. They may take a while before they home in on the exact thing you want. But they need to get enough examples of what you're trying to say and what you mean.

IT'S NOT BUILT TO WRITE. [TRANSCRIPT: "IT'S NOT BUILT TO RIGHT."] IT'S BUILT TO RECOGNIZE THE WORDS YOU SAID. THINGS LIKE PUNCTUATION ARE PART OF MEANING.

Hechler: But it doesn't absorb changes I make. It doesn't even know I'm making changes, as far as I can tell. If it actually absorbed what I'm doing to the transcript before I export it into Word, it would learn. But that hasn't happened. So there's no feedback loop at all. How do I know that? Because it makes the same mistakes over and over and over. That's why it's been such a disappointment to me: it hasn't learned, as far as I can tell, anything.

Wiipon: Well, in that case you try other software. You should experiment.

Hechler: I will. In the meantime, let's return to grammar. I'll just say It's atrocious. Sentences go on and on and on. Commas appear with alarming regularity, often where [the program changed "where" to "when"] they have no business being in the first place. What's the problem? Why doesn't the software know how to write?

Wiipon: That's pretty funny. It's not built to write. [Transcript: "It's not built to right."] It's built to recognize the words you said. Things like punctuation are part of meaning. They're not part of syntax of a language. They're part of the semantics of the language. You put punctuation in and by and large you intend a particular meaning or something to happen. And, yeah, it's something that they are learning. Like I said, a decade ago you wouldn't get punctuations. Now, at least Microsoft puts punctuation in, and most of the ones that I've tried will put punctuation in. If I'm dictating while I'm on the phone, I'm in the car responding to a text message, I'll still put it in myself [by saying "comma, period," etc.], just to make sure it's there. But that is coming along.

Hechler: The program often hears the wrong word. Sometimes it seems way off **[like rendering "really" as "Billy."]** Other times, it's plausible, but in context it's clearly wrong. But the program doesn't understand context.

Wiipon: Yeah. Context, again. When you think of syntax, syntax doesn't know, by and large, the context. Context is usually a long-distance relationship across sentences, across paragraphs, or across the meaning of the document. You can see in the generative AI space, with ChatGPT, if you ask it to produce things, it's making good context relationships throughout a document. So in the generative case, it's working really well. Would you expect that to happen in transcription? Probably. I'm sure there is some of that in there already. But speech adds another level of complication. Generating text is a lot easier than first recognizing something that someone says and then trying to make it fit into the context and meaning that somebody wants. It's a lot harder to be able to do that.

Hechler: What do you expect will be the future of speech transcription? And how long do you think it will take for it to become really reliable? By which I mean, getting the right words and getting some sort of grammatical structure that is closer to what I would do if I were doing the transcript myself? Maybe 10 years?

Wiipon: It will constantly improve, but I don't think anything will be perfect in 10 years. Too many things to deal with. I am not sure where the advancements would be. I mean people consider ChatGPT an advancement. Not dictation. I think that's showing you what large language models can do. So the language part the machines are going to handle really well. The speech part is still a difficult thing. Interestingly there's not as many researchers working on the speech side as there are on the language side.

Hechler: Does the language side include grammar?

Wiipon: Yeah. When I say speech I mean for signal processing. You know, sounds come out of your mouth, they go into your ear, they go into a microphone, which goes into a computer. How are those analyzed, and used in such a way that a machine can recognize or distinguish one sound from another? That's still a science. And it used to be very sexy. It hasn't been as sexy as language in the past 25 years.

Hechler: Well, we will see. When I get the transcript of this conversation, I'm going to look for weird anomalies, to show readers what I'm talking about. I'll show them some sentences that might be really hard to figure out. Or maybe it'll surprise me. Maybe it understands that it's under special scrutiny today. And it will do its best work ever. But I suspect there will be a number of passages that will have us saying, "Huh??"

TAG

MAPPING AI

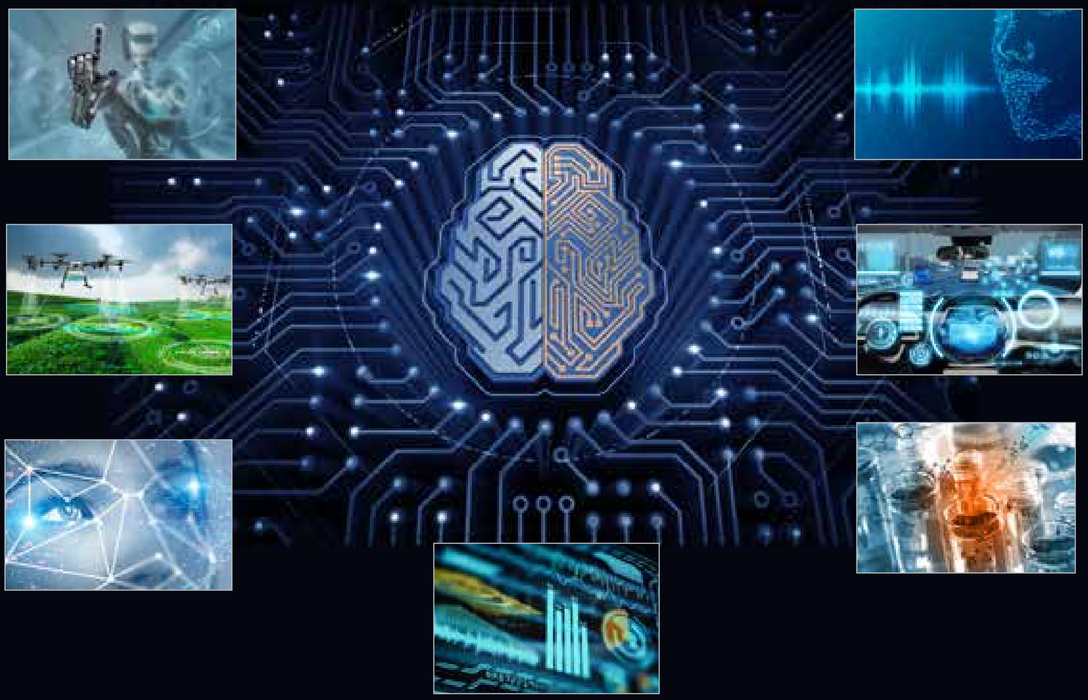

INSIGHTS, TRENDS, AND CATEGORIES

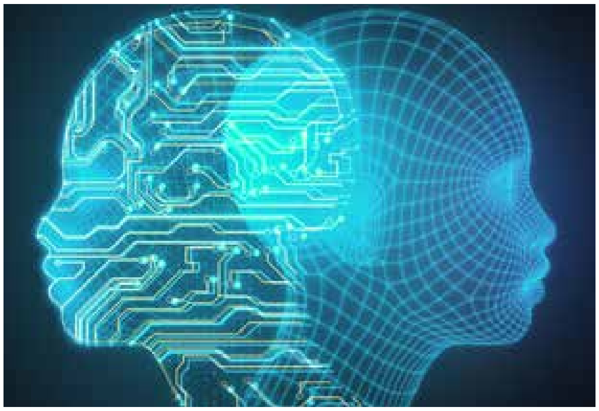

INSIGHTS, TRENDS, AND CATEGORIES IN THE TAG AI QUARTERLY EDITIONS

DR. EDWARD AMOROSO,
FOUNDER & CEO TAG INFOSPHERE

My roots run deep in computer science. My dad was an early computer scientist, receiving his PhD in the topic in 1966. I would earn the same degree in computer science 25 years later. Since then, I have been a professional computer scientist, starting at Bell Labs and continuing in many roles at AT&T. I also have spent the past 33 years teaching in the computer science departments at NYU, the Stevens Institute of Technology, and Monmouth University.

I suppose you could say that I'm in as good a position as anyone to provide commentary on trends in the discipline of computing. So, I hope you'll listen carefully and absorb what I am about to say, and I do not offer this point lightly. In 40 years of professional work in computer science, I have never seen a technology, and this includes browsers, the web, and the iPhone, with a greater potential to change society more than modern artificial intelligence (AI).

As the founder and CEO of the research and advisory company TAG Infosphere, headquartered in New York City and focused to date on cybersecurity and climate science, it was a natural progression for me to add artificial intelligence to our portfolio. Our customers, in both cyber and climate, were also beginning to demand guidance in this area, so without much hesitation we decided to launch our TAG AI practice in mid-2023.

OUR NEW TAG PRACTICE—AI

Having created advisories twice before (cybersecurity and climate science), I understood the time and work implications of opening a new practice. Our signature research approach is empirical, which implies that in a given area of interest, in this case artificial intelligence, what we do involves massive scans of the literature, publicly available information, research reports, and anything else we can get our hands on—and then devouring the information.

Our primary focus has always been on commercial innovations, and startup companies in particular. What this means is that since starting TAG AI, we're been in discovery, learn, and interview mode with startup companies trying to either sell their AI-based solutions directly, apply their AI focus to a specific domain, leverage AI for their main (non-AI) business activities, or determine the impact of AI on their organization or even country.

Naturally, such information from startup companies must be organized, and while it would seem poetic justice to allow AI to do this for us, we chose to develop our taxonomy based on human interpretation of what we see. I suspect that at some point this step might be unnecessary, but humans in 2024 are still very much part of the equation in terms of categorization, organization, and sense-making of commercial AI innovations.

UNDERSTANDING HOW ORGANIZATIONS USE AI

We have found it helpful to use a model of four concentric circles to explain how organizations leverage AI today for practical purposes (see Figure 1). As you'd expect, AI startups fall into this model at the core level but are a subset of the total universe of organizations. It seems true today that some companies will not need AI, but this will change. Our expectation is that soon every organization will know exactly how AI must be leveraged to survive.

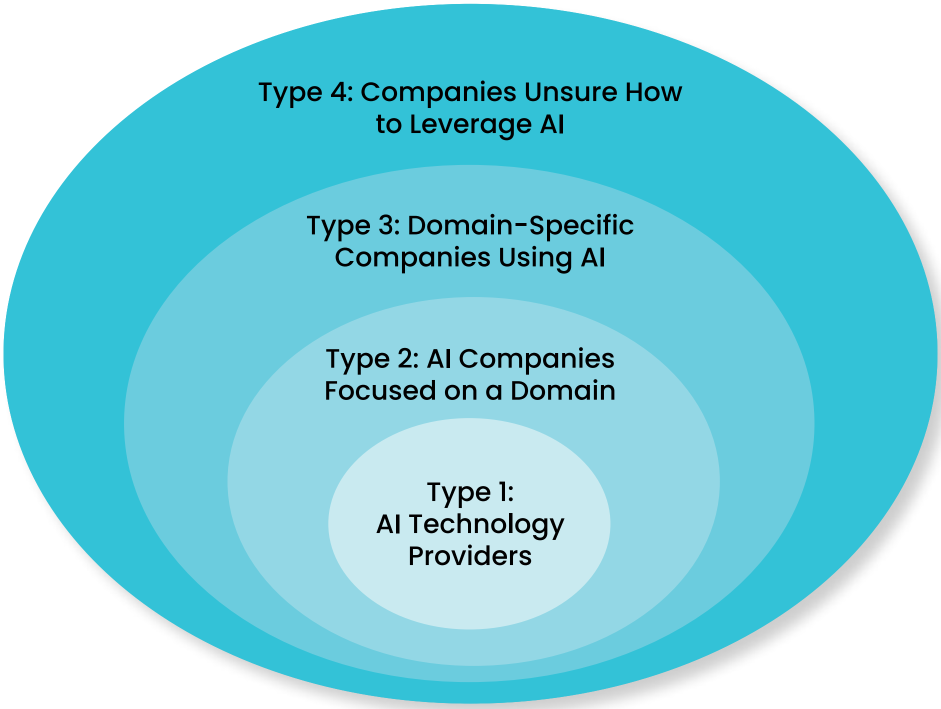
Type 4: Companies Unsure How to Leverage AI
Type 3: Domain-Specific Companies Using AI
Type 2: AI Companies Focused on a Domain
Type 1: AI Technology Providers

Figure 1. How Organizations Leverage AI

Many different companies can be mapped to the inner-most core layer of the model—referred to as Type 1 Core AI Technology Providers. For example, OpenAI, Google, and NVIDIA might be viewed as companies that see AI technology as their core business, even though Google has so many other businesses, as does NVIDIA. But from now on their core growth and value are likely to come from AI.

The next outer level of the model is referred to as the Type 2 AI Companies Focused on a Specific Domain. As examples, consider that Dark Trace, VineView, and Cerebras all aim their AI efforts on the specific areas of security, AgTech, and training. We would expect this grouping to grow dramatically as businesses and government agencies begin to recognize how fundamental AI is to a complete reinvention and improvement on how they achieve their core mission.

Type 3 Domain-Specific Companies using AI constitute a large category. In fact, it's becoming harder to find companies *not* in this grouping, because almost anything one can conceive these days in business or government can be enhanced or improved with AI. Obviously, companies that do not understand AI form the fourth group in the model. We recommend that you take a moment to reflect on where your organization falls into the model.

A NOTE ON GOVERNMENT TEAMS

It is also worth noting that we will often reference "companies" as synonymous with "organizations," and this can produce some degree of discomfort for groups such as federal government agencies or teams supporting state, local, regional, or municipal governments. If you fall into the category, rest assured that our references to companies will generally apply to you as well when we're talking about opportunities to leverage AI.

When we are talking about business opportunities to make use of machine or deep learning tools and technologies to create new products and services, then government teams will find this important for their own buying decisions, but will also find it useful to understand commercial trends in AI. This could lead to better regulations, laws, and policies from teams working AI issues in government.[1]

TAG AI TAXONOMY

The TAG AI Taxonomy is a breakdown of the various categories and subcategories within which we can identify startups and vendors that are developing, marketing, and in many cases currently providing commercial solutions. These solutions come from all over the world, even though many appear to be domestic US-based. We fully expect commercial attention in this area to be global, although access to commercial offerings in China and Russia are limited.[2]

Like prior similar efforts in cybersecurity and climate science at our parent company, TAG Infosphere, our TAG AI Taxonomy was developed iteratively based on open-source research, meetings with vendors, review of the literature, and discussions with enterprise teams, academic organizations, research teams, government agencies, and other AI stakeholders. We also try hard to integrate use of generative AI into our work as a hybrid contributor.

[1] *TAG Infosphere participates in many different forums related to government use of AI. For example, Dr. Amoroso participates in a Working Group on AI sponsored by the Governor's Office in the State of New Jersey. In addition, TAG now sells its Research as a Service (RaaS) through a contract with Carahsoft that can offer government buyers access to contracts registered with the US General Services Administration (GSA).*

[2] *We do not restrict our attention as industry research and advisory experts to regions outside China and Russia, but we do have some difficulty getting accurate information from companies in these regions. The war in Ukraine, still raging as we write these words, has also had an impact on getting useful commercial guidance from startups and vendors located in Ukraine.*

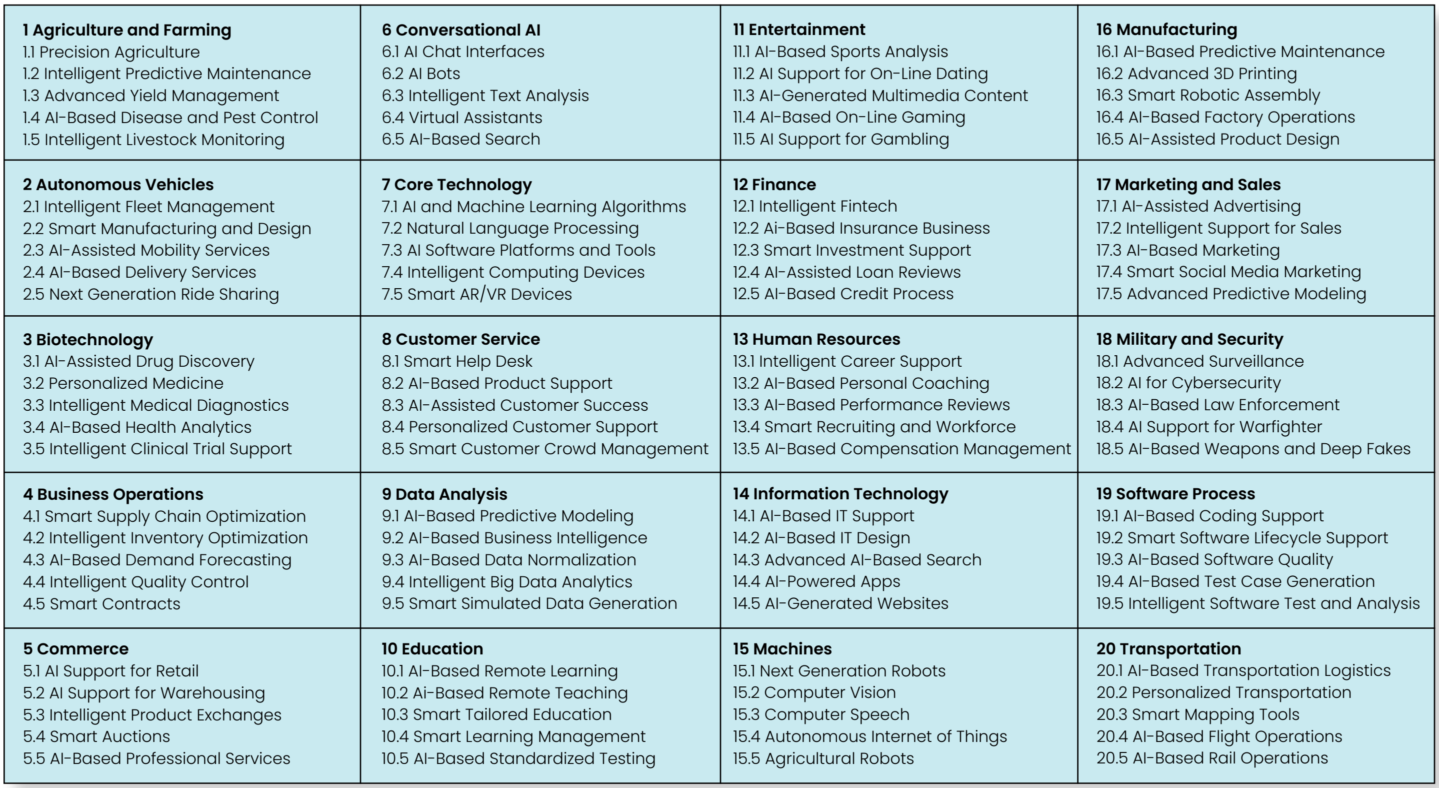

1 Agriculture and Farming 1.1 Precision Agriculture 1.2 Intelligent Predictive Maintenance 1.3 Advanced Yield Management 1.4 AI-Based Disease and Pest Control 1.5 Intelligent Livestock Monitoring	**6 Conversational AI** 6.1 AI Chat Interfaces 6.2 AI Bots 6.3 Intelligent Text Analysis 6.4 Virtual Assistants 6.5 AI-Based Search	**11 Entertainment** 11.1 AI-Based Sports Analysis 11.2 AI Support for On-Line Dating 11.3 AI-Generated Multimedia Content 11.4 AI-Based On-Line Gaming 11.5 AI Support for Gambling	**16 Manufacturing** 16.1 AI-Based Predictive Maintenance 16.2 Advanced 3D Printing 16.3 Smart Robotic Assembly 16.4 AI-Based Factory Operations 16.5 AI-Assisted Product Design
2 Autonomous Vehicles 2.1 Intelligent Fleet Management 2.2 Smart Manufacturing and Design 2.3 AI-Assisted Mobility Services 2.4 AI-Based Delivery Services 2.5 Next Generation Ride Sharing	**7 Core Technology** 7.1 AI and Machine Learning Algorithms 7.2 Natural Language Processing 7.3 AI Software Platforms and Tools 7.4 Intelligent Computing Devices 7.5 Smart AR/VR Devices	**12 Finance** 12.1 Intelligent Fintech 12.2 Ai-Based Insurance Business 12.3 Smart Investment Support 12.4 AI-Assisted Loan Reviews 12.5 AI-Based Credit Process	**17 Marketing and Sales** 17.1 AI-Assisted Advertising 17.2 Intelligent Support for Sales 17.3 AI-Based Marketing 17.4 Smart Social Media Marketing 17.5 Advanced Predictive Modeling
3 Biotechnology 3.1 AI-Assisted Drug Discovery 3.2 Personalized Medicine 3.3 Intelligent Medical Diagnostics 3.4 AI-Based Health Analytics 3.5 Intelligent Clinical Trial Support	**8 Customer Service** 8.1 Smart Help Desk 8.2 AI-Based Product Support 8.3 AI-Assisted Customer Success 8.4 Personalized Customer Support 8.5 Smart Customer Crowd Management	**13 Human Resources** 13.1 Intelligent Career Support 13.2 AI-Based Personal Coaching 13.3 AI-Based Performance Reviews 13.4 Smart Recruiting and Workforce 13.5 AI-Based Compensation Management	**18 Military and Security** 18.1 Advanced Surveillance 18.2 AI for Cybersecurity 18.3 AI-Based Law Enforcement 18.4 AI Support for Warfighter 18.5 AI-Based Weapons and Deep Fakes
4 Business Operations 4.1 Smart Supply Chain Optimization 4.2 Intelligent Inventory Optimization 4.3 AI-Based Demand Forecasting 4.4 Intelligent Quality Control 4.5 Smart Contracts	**9 Data Analysis** 9.1 AI-Based Predictive Modeling 9.2 AI-Based Business Intelligence 9.3 AI-Based Data Normalization 9.4 Intelligent Big Data Analytics 9.5 Smart Simulated Data Generation	**14 Information Technology** 14.1 AI-Based IT Support 14.2 AI-Based IT Design 14.3 Advanced AI-Based Search 14.4 AI-Powered Apps 14.5 AI-Generated Websites	**19 Software Process** 19.1 AI-Based Coding Support 19.2 Smart Software Lifecycle Support 19.3 AI-Based Software Quality 19.4 AI-Based Test Case Generation 19.5 Intelligent Software Test and Analysis
5 Commerce 5.1 AI Support for Retail 5.2 AI Support for Warehousing 5.3 Intelligent Product Exchanges 5.4 Smart Auctions 5.5 AI-Based Professional Services	**10 Education** 10.1 AI-Based Remote Learning 10.2 Ai-Based Remote Teaching 10.3 Smart Tailored Education 10.4 Smart Learning Management 10.5 AI-Based Standardized Testing	**15 Machines** 15.1 Next Generation Robots 15.2 Computer Vision 15.3 Computer Speech 15.4 Autonomous Internet of Things 15.5 Agricultural Robots	**20 Transportation** 20.1 AI-Based Transportation Logistics 20.2 Personalized Transportation 20.3 Smart Mapping Tools 20.4 AI-Based Flight Operations 20.5 AI-Based Rail Operations

Figure 2. TAG AI Taxonomy—How Organizations Leverage AI

This TAG AI taxonomy, as you can see, includes quite a few different categories (20) and associated subcategories (100). The implication is that AI is a rich area for startup activity, and any observer should expect to see massive commercial growth in each of these areas. This is good news for the technology sector, excellent news for technology investors, and even better news for entrepreneurs who seek to use AI as a means for solving problems.

In the sections below, we provide a general overview of each of the twenty categories with some brief commentary, assisted with generative AI, on the associated subcategories. We also provide a small representative sampling of interesting startups that map to each area. We hope this helps to spur additional research in AI innovation. We can assure you that it will serve as the basis for our own research and advisory work at TAG.[3]

In the following pages we present a selection from our volume of Insight Reports covering the categories in TAG's AI taxonomy. To see more visit **TAG's Publications Page.**

TAG

INSIGHT REPORT

OVERVIEW OF AI FOR AGRICULTURE AND FARMING

INTRODUCTION

This TAG Insights Report on Artificial Intelligence (AI) for Agriculture and Farming is intended to help companies, managers, practitioners, researchers, investors, and commercial vendors better understand current trends, issues, and market opportunities in this area. A list of representative commercial vendors working in various areas of AI for agriculture and farming is included. The five specific areas covered in this report include:

1. Precision Agriculture
2. Intelligent Predictive Maintenance
3. Advanced Yield Management
4. AI-Based Disease and Pest Control
5. Intelligent Livestock Monitoring

This report is intended for general and unrestricted use, but interested readers are encouraged to connect with the TAG research and advisory team for more information on the private TAG Research as a Service (RaaS) community that covers, discusses, and shares information on these topics in more depth and includes a wider range of startups, vendors, and companies.

OVERVIEW OF AI FOR AGRICULTURE AND FARMING

The following emerging global commercial opportunities involving AI for agriculture and farming solutions are covered in this report, including the listing of several viable commercial entities providing solutions on the market today:

- Precision farming benefits directly from AI-driven analytics and the ability to tailor the technology the farmer's specific needs. The automation of farm equipment, driven by AI, holds immense potential for labor-saving and productivity gains.
- Intelligent predictive maintenance benefits from AI by supporting the need to reduce maintenance costs proactively. Waiting for equipment to fail is a much more expensive option.
- Advanced yield management enables farmers to optimize resource allocation, enhancing crop yields and reducing wastage. Advancements in weather forecasting will aid in risk mitigation and yield optimization.
- AI-based disease and pest control are enabled by AI-powered analytics that focus on reducing the negative impacts of such problems.
- Intelligent livestock monitoring systems are improving animal welfare and productivity, while AI-based pest and disease detection systems minimize pesticide usage, fostering sustainable practices.

The combination of modern, advanced AI-based technology with the traditional practice of agriculture and farming provides excellent opportunities to improve business results and profitability while also creating new opportunities for more sustainable practices driving advances such as climate mitigation.

FOCUS AREA: PRECISION AGRICULTURE

Precision agriculture represents a paradigm shift in modern farming. AI applications in precision agriculture uses data-driven insights to optimize resource management, increase yields, and promote sustainable practices.

AI-driven analytics enable farmers to make informed decisions about irrigation, fertilization, and pesticide use, reducing waste and environmental impact. Remote sensing technologies, such as drones and satellites, equipped with AI algorithms, provide real-time data on crop health and soil conditions, aiding in early detection of diseases and pests.

AI also enhances livestock management by monitoring animal health and behavior, improving productivity and welfare. As the agriculture sector faces growing challenges due to climate change and population growth, AI for precision agriculture offers a path towards resilient and efficient food production. This technology continues to evolve, promising even greater precision and sustainability in farming practices for the future.

FOCUS AREA: INTELLIGENT PREDICTIVE MAINTENANCE

In agriculture, where machinery is the backbone of operations, the implementation of AI for intelligent predictive maintenance is driving increased efficiency and cost-effectiveness. AI-enabled systems use sensor data to continuously monitor the condition of farming equipment, such as tractors, harvesters, and irrigation systems.

Through predictive analytics, AI algorithms can detect subtle anomalies and wear patterns, thus allowing farmers to schedule maintenance before critical failures occur. This proactive approach reduces downtime, lowers repair costs, and ensures that machinery operates at peak performance, ultimately increasing overall farm productivity.

As AI technology advances, the predictive maintenance models will become more precise, optimizing resource allocation, and extending the lifespan of expensive agricultural machinery. With AI-driven intelligent predictive maintenance, farmers can embrace sustainable practices, streamline operations, and navigate the modern agricultural landscape with greater confidence and efficiency.

FOCUS AREA: ADVANCED YIELD MANAGEMENT

The use of AI for advanced yield management is a good example of how technology can improve the business conditions for farmers. AI's capabilities are now optimizing crop production, ensuring resource efficiency, and mitigating risks.

By leveraging data from sensors, satellites, and historical records, AI-powered systems provide real-time insights into soil conditions, weather patterns, and crop health. This data-driven approach enables farmers to make informed decisions about planting, irrigation, and fertilization. Machine learning algorithms analyze vast datasets to predict crop yields accurately, helping farmers optimize production and minimize waste.

As farmers and agriculture face increasing challenges due to climate change and a growing global population, AI's role in advanced yield management is pivotal in ensuring food security and sustainable farming practices. One can only hope that the technology truly advances the prospects for the farming and agricultural sector.

FOCUS AREA: AI-BASED DISEASE AND PEST CONTROL

AI is revolutionizing disease and pest control in agriculture, ushering in a new era of precision and sustainability. Farmers can use AI to combat threats more effectively to their crops while minimizing the use of harmful pesticides. AI-driven systems analyze vast amounts of data, including climate conditions, soil health, and pest populations, to predict disease outbreaks and pest infestations with unprecedented accuracy. This early warning system allows farmers to take targeted and timely actions, reducing crop damage and production losses.

Furthermore, AI enables the development of automated pest control solutions, such as robotic devices equipped with cameras and AI algorithms that can identify and selectively target pests, reducing the need for widespread chemical treatments. In an era where sustainable agriculture is important, AI-based disease and pest control offer not only economic benefits but also a path towards environmentally friendly and responsible farming practices.

FOCUS AREA: INTELLIGENT LIVESTOCK MONITORING

AI-based intelligent livestock monitoring is transforming the way farmers manage their animals, enhancing both productivity and animal welfare. Through a combination of sensors, cameras, and machine learning algorithms, this technology offers real-time insights into the health and behavior of livestock.

Monitoring systems can detect subtle changes in vital signs, allowing early identification of illnesses or distress, which leads to prompt intervention and reduced mortality rates. AI can also effectively and accurately predict breeding cycles and optimize feeding schedules, improving overall farm efficiency.

Beyond health and production benefits, intelligent livestock monitoring systems promote ethical farming practices. They ensure animals are treated humanely, with environmental conditions tailored to their needs. As agriculture evolves to meet global food demands while minimizing its environmental footprint, AI-based livestock monitoring stands as a crucial tool for modern, sustainable farming practices. It enhances animal well-being, boosts farm efficiency, and contributes to a more responsible and productive agricultural industry.

COMPANIES AND CONTRIBUTIONS

The companies listed below emerged as part of our research at TAG. Our goal in listing these fine firms is to provide a starting point for buyers, advocates, stakeholders, and researchers trying to make sense of the commercial landscape for artificial intelligence as a means for driving toward improved methods for agriculture.

PRECISION AGRICULTURE VENDORS

Agremo: Agremo uses AI and drone technology to analyze aerial imagery of fields and provide actionable insights for precision agriculture.

AgShift: AgShift uses computer vision and AI to automate quality inspection and grading of agricultural produce, as well as helping with disease or pest-related issues.

Blue River Technology (a John Deere Company)**:** Blue River Technology uses computer vision and AI to create smart farming equipment, such as See & Spray weed control.

Bushel Farm – Bushel Farm uses AI and data analytics to provide farmers with tools for managing and optimizing their operations.

Corteva – Corteva provides advanced solutions for agriscience with a portfolio of different solutions for farmers.

FarmWise: FarmWise develops autonomous weeding robots powered by AI to help reduce the need for herbicides and increase crop yields.

Gamaya: Gamaya uses hyperspectral imaging and AI to provide insights on crop health, allowing farmers to optimize their use of resources.

Prospera: Prospera offers AI-driven solutions for monitoring and optimizing crop health, providing real-time insights to farmers.

Resson: Resson offers a data-driven platform that utilizes AI to provide farmers with insights into crop health and pest management, as well as yield.

Taranis: Taranis provides an AI-powered precision agriculture platform using satellite imagery, weather data, and machine learning to monitor and optimize crop health.

INTELLIGENT PREDICTIVE MAINTENANCE VENDORS

Agrivi: Agrivi offers AI-driven farm management software that includes predictive maintenance features to help farmers anticipate and prevent equipment breakdowns.

Arable Labs: Arable Labs offers predictive analytics for agriculture, using AI to monitor and manage crop health and environmental conditions.

Ceres Imaging: Ceres Imaging uses AI to analyze aerial imagery and provide predictive insights for crop health and irrigation management.

Farmers Edge: Farmers Edge uses AI and machine learning to provide predictive analytics for agriculture, including equipment maintenance and weather forecasting.

Farmwave: Farmwave offers AI-powered software for agriculture, including predictive maintenance tools to keep farm equipment running smoothly.

Quantified Ag: Quantified Ag (owned by Merck) utilizes AI to monitor livestock health and behavior, helping farmers predict and prevent illness or issues.

Regrow: Regrow combines AI with remote sensing technology to offer predictive insights into crop health and potential issues, including equipment maintenance needs.

Tule Technologies: Tule Technologies utilizes AI and IoT sensors to provide predictive irrigation management for agriculture, helping farmers optimize water usage.

ADVANCED YIELD MANAGEMENT VENDORS

CropX: CropX uses AI-driven soil sensor technology to help farmers optimize irrigation and increase crop yield while conserving water.

Farmers Business Network (FBN): FBN uses data analytics and AI to provide farmers with insights on crop yield performance, input costs, and market intelligence.

Hiphen: provides AI-powered drone-based analytics for agriculture, helping farmers optimize crop yield through data insights.

Nutrien AG – offers AI-driven insights to help farmers optimize yield by providing weather, crop modeling, and field-specific recommendations.

Prospera: Prospera (Valmont Company) offers AI-driven solutions for monitoring and optimizing crop health and yield, providing real-time insights to farmers.

Sentera: Sentera offers AI-powered analytics for precision agriculture, including drone and satellite imagery analysis to optimize yield.

Trace Genomics: Trace Genomics uses AI to analyze soil microbiomes and provide recommendations for improving soil health and increasing crop yield.

AI-BASED DISEASE AND PEST CONTROL VENDORS

AgroScout: AgroScout uses AI and machine learning for early detection of diseases, pests, and nutrient deficiencies in crops through scouting and imaging.

Blue River Technology (a John Deere Company)**:** Blue River Technology's See & Spray system utilizes computer vision and AI to identify and target weeds in real-time, reducing the need for herbicides.

Burro: Burro develops autonomous robots for tasks like weeding and pest control, utilizing AI to target specific problem areas in the field.

FieldIn: FieldIn offers an AI-powered platform for pest and disease management, allowing growers to monitor and optimize pesticide use.

Proagrica – Proagrica solutions incorporate AI for crop disease and pest management by providing data-driven insights to farmers.

Spornado: Spornado uses AI and sensors to detect and monitor fungal diseases in crops and provides real-time alerts to farmers.

SwarmFarm Robotics: SwarmFarm Robotics uses AI-powered robots for precision spraying and targeted pest control in agriculture.

XAG – Chinese company, XAG, combines AI with drones to identify and address disease and pest issues in crops by delivering targeted treatments.

INTELLIGENT LIVESTOCK MONITORING VENDORS

Afimilk: Afimilk uses AI to monitor and manage dairy herd performance, including health, reproduction, and nutrition.

Connecterra: Connecterra offers an AI-driven platform called "Ida" that monitors cow behavior and health to optimize dairy farm operations.

CowManager - CowManager uses AI and IoT devices to monitor cow health and fertility, assisting farmers in making data-driven decisions.

EasyKeeper: EasyKeeper offers an AI-driven livestock management platform for small and medium-sized farms, covering health, reproduction, and records.

Ever.Ag – Ever.Ag employs computer vision and AI to monitor livestock behavior and health, providing insights to farmers for better management.

Livestock Water Recycling (LWR): LWR employs AI and sensor technology to monitor and optimize water usage and waste management in livestock operations.

Moocall: Moocall offers an AI-powered calving sensor that predicts when a cow is likely to give birth, allowing for timely assistance.

Sencrop: While initially focused on weather monitoring, Sencrop has expanded to provide livestock monitoring solutions that utilize AI for better herd management.

SomaDetect: SomaDetect uses AI and sensors to analyze milk quality and cow health, providing dairy farmers with real-time data for better herd management.

Vence: Vence develops AI-powered virtual fencing solutions for livestock, allowing farmers to control and monitor their animals' movements.

TAG

INSIGHT REPORT

OVERVIEW OF AI FOR AUTONOMOUS VEHICLES

INTRODUCTION

This TAG Insights Report on Artificial Intelligence (AI) for Autonomous Vehicles is intended to help companies, managers, practitioners, researchers, investors, and commercial vendors better understand current trends, issues, and market opportunities in this area. A list of representative commercial vendors working in various areas of AI for autonomous vehicles is included. The five specific areas covered in this report include:

1. Intelligent Fleet Management
2. Smart Manufacturing and Design
3. AI-Assisted Mobility Services
4. AI-Based Delivery Services
5. Next-Generation Ride Sharing

This report is intended for general and unrestricted use, but interested readers are encouraged to connect with the TAG research and advisory team for more information on the private **TAG Research as a Service (RaaS)** community that covers, discusses, and shares information on these topics in more depth and includes a wider range of startups, vendors, and companies.

OVERVIEW OF AI FOR AUTONOMOUS VEHICLES

The following emerging global commercial opportunities involving AI for autonomous vehicles are covered in this report, including the listing of several viable commercial entities providing solutions on the market today:

- AI optimizes fleet management by using real-time data analytics for route planning, predictive maintenance, and fuel consumption reduction, thereby increasing efficiency, and reducing operational costs.
- AI accelerates the design process of autonomous vehicles through generative design algorithms that can propose optimized structures and materials, and it ensures quality control on the manufacturing floor through machine learning-enabled robotics and inspection systems.
- AI enhances mobility services by personalizing user experiences through learning individual preferences, improving safety with dynamic response systems, and streamlining traffic flow with adaptive control systems.
- AI revolutionizes delivery services by employing autonomous drones and vehicles that can navigate to the doorstep, plan optimal delivery routes, and update schedules in real-time based on traffic and customer availability.
- AI transforms ridesharing by matching passengers with suitable rides in real-time, optimizing pickup and drop-off points, and integrating with urban transport systems to reduce wait times and improve convenience.

FOCUS AREA: INTELLIGENT FLEET MANAGEMENT

AI-powered fleet management opens numerous business prospects for both established vendors and startups across various domains. One avenue is optimizing operations by using AI for route planning, vehicle maintenance scheduling, and fuel consumption reduction. These improvements lead to cost savings and greater efficiency.

Predictive maintenance is another area of interest. AI systems that forecast maintenance needs help fleets minimize downtime and lower maintenance costs. Enhancing safety through AI, which monitors driver behavior and provides real-time feedback, reduces accidents and insurance expenses. Telematics and tracking, utilizing AI for real-time vehicle monitoring and driver performance assessment, contribute to improved security and efficiency. Sustainability concerns drive the demand for AI-optimized routing to reduce emissions, aligning with the green fleet management trend.

Data analytics is essential, with AI platforms converting fleet data into actionable insights for better decision-making. Integration services facilitate the seamless adoption of AI into existing fleet management systems. The rise of autonomous vehicles offers opportunities for managing and optimizing autonomous fleets. Customization caters to specific industry needs, such as logistics and transportation. AI transforms fleet management, offering practical solutions aligned with evolving fleet operator requirements.

FOCUS AREA: SMART MANUFACTURING AND DESIGN

Exploring AI for smart manufacturing and the design of autonomous vehicles reveals business opportunities for both established vendors and nimble startups. One avenue lies in optimizing manufacturing processes through AI-driven automation. This technology streamlines production, reduces waste, and enhances quality control, leading to cost savings and improved efficiency.

AI can also play a pivotal role in the design phase of autonomous vehicles, aiding in simulation, prototyping, and optimization. Vendors and startups can develop AI-powered tools that enable engineers to create more efficient and safer autonomous vehicles, potentially revolutionizing the automotive industry. Furthermore, AI can improve supply chain management by predicting demand, optimizing inventory, and reducing lead times.

Companies that offer AI-based solutions in this area can help manufacturers reduce costs and improve their overall competitiveness. Cybersecurity is another vital area, where AI can be used to protect autonomous vehicles from cyber threats, opening up opportunities for cybersecurity-focused startups. AI-driven smart manufacturing and autonomous vehicle design present a broad canvas for innovation, cost reduction, and improved performance, with vendors and startups at the forefront poised for significant growth.

FOCUS AREA: AI-ASSISTED MOBILITY SERVICES

The business of AI-assisted mobility services using autonomous vehicles represents a transformative shift in the transportation industry. Autonomous vehicles, powered by advanced artificial intelligence algorithms, are poised to revolutionize the way people move from one place to another. This burgeoning industry offers a plethora of opportunities and challenges for entrepreneurs, tech giants, and traditional automakers alike.

First and foremost, the adoption of autonomous vehicles promises to improve road safety by reducing human error, which is a major cause of accidents. This safety enhancement alone presents a compelling business case, as insurance costs could decrease significantly, and lives could be saved. Furthermore, AI-assisted mobility services offer the potential for increased efficiency and reduced congestion on our roadways. This means shorter commute times and less environmental impact, factors that are appealing to both individuals and governments focused on sustainability.

However, challenges like regulatory hurdles, cybersecurity concerns, and the high cost of developing autonomous technology must be addressed. Nevertheless, the business of AI-assisted mobility services using autonomous vehicles is on the cusp of transformation, and those who can navigate the complexities stand to reap substantial rewards in the form of safer, more efficient, and sustainable transportation solutions.

FOCUS AREA: AI-BASED DELIVERY SERVICES

AI-assisted delivery services using autonomous vehicles is rapidly gaining traction as technology continues to evolve. These services are poised to disrupt the traditional delivery industry, offering efficiency, cost savings, and convenience to businesses and consumers alike. As one would expect, drones will be a major future factor in the development of AI-based delivery services.

One of the most significant advantages is increased delivery speed and accuracy. Autonomous vehicles equipped with AI algorithms can navigate traffic, avoid obstacles, and optimize routes, ensuring timely and precise deliveries. This efficiency can translate into lower operational costs for businesses, reduced delivery fees for customers, and potentially even 24/7 delivery options. Moreover, AI-assisted delivery services can enhance safety by minimizing human error and decreasing the risk of accidents. This can lead to reduced insurance premiums and liability concerns for companies.

However, there are still challenges to overcome, including regulatory hurdles, public acceptance, and cybersecurity concerns. Companies in this space must also invest heavily in research, development, and infrastructure. AI-assisted delivery services using autonomous vehicles represent a promising frontier in the business world, offering the potential for improved efficiency, safety, and customer satisfaction. As technology continues to advance, this industry is poised for substantial growth and innovation.

FOCUS AREA: NEXT GENERATION RIDESHARING

AI-assisted next-generation ridesharing using autonomous vehicles is driving transportation innovation, poised to redefine the way people move within cities. This emerging industry combines artificial intelligence, autonomous vehicle technology, and ride-hailing services to create a more efficient, convenient, and sustainable urban transportation ecosystem.

AI-powered autonomous vehicles offer several advantages for ridesharing businesses. They can provide 24/7 availability, reduce operating costs by eliminating driver wages, and optimize routes to minimize travel times and congestion. This translates into competitive pricing for passengers and increased profitability for companies. Additionally, AI-assisted ridesharing promises enhanced safety through advanced sensors and real-time monitoring, reducing the risk of accidents and improving public perception of autonomous vehicles.

However, challenges such as regulatory compliance, public acceptance, and the high initial investment in autonomous technology must be addressed. Successful businesses in this field will need to collaborate with local authorities and continuously innovate to stay ahead in a rapidly evolving market. AI-assisted next-generation ridesharing using autonomous vehicles represents a promising business opportunity with the potential to revolutionize urban transportation, offering benefits for both companies and passengers while addressing environmental and congestion challenges in cities.

COMPANIES AND CONTRIBUTIONS

The companies listed below emerged as part of our research at TAG. Our goal in listing these fine firms is to provide a starting point for buyers, advocates, stakeholders, and researchers trying to make sense of the commercial landscape for artificial intelligence as a means for driving toward improved methods for autonomous vehicles.

INTELLIGENT FLEET MANAGEMENT VENDORS

AEye: AEye develops advanced perception systems for autonomous vehicles, which can be integrated into fleet management solutions for enhanced safety.

Aidrivers: Aidrivers provides autonomous vehicle solutions for industries such as mining and agriculture, including fleet management capabilities.

Aptiv: Aptiv provides advanced technology for autonomous driving, including software and hardware solutions for fleet management.

AssetWORKS: AssetWORKS FleetFocus can provide autonomous vehicle fleet managers with better control over diagnostics, fuel consumption, maintenance and more.

Aurora: Aurora is a startup working on autonomous vehicle technology and offers a platform for autonomous vehicle fleet management and ridesharing.

CYNGN: Cyngn Insight is a comprehensive platform to manage, monitor, and command your self driving vehicles.

DDS Wireless: DDS Wireless provides SaaS AI routing tools for paratransit, taxi, and school transport.

LogisFleet: Manage, track, and setup up compliance rules and checks with LogosFleet's vehicle management software.

Oxa: (AKA Oxbotica) A UK-based startup, Oxa, specializes in autonomous vehicle software and is involved in various autonomous fleet projects.

Robomart: Robomart is a startup that focuses on autonomous grocery delivery using self-driving vehicles with an integrated fleet management system.

SMART MANUFACTURING AND DESIGN VENDORS

Ansys: Ansys provides engineering simulation software that can be used for virtual testing and validation of autonomous vehicle systems.

Aras: Aras specializes in Product Lifecycle Management (PLM) software, which can be applied to manage the development and manufacturing processes of autonomous vehicles.

Autodesk: Autodesk offers software tools for computer-aided design (CAD) and engineering, which can be used in autonomous vehicle development.

Canoo: Canoo is known for its electric, autonomous vehicle platform, which can be customized for various transportation needs.

Clearpath Robotics: Clearpath Robotics offers autonomous mobile robots and simulation tools for research and development in autonomous systems, including manufacturing applications.

Cognata: Cognata provides a simulation platform for autonomous vehicle development, which can be used for design and testing.

CYNGN: Cyngn produces autonomous vehicles for managing warehouse operations.

Dassault Systèmes: Dassault Systèmes provides 3D design and engineering software, including solutions tailored for autonomous vehicle design and simulation.

Horizon Robotics: Horizon Robotics develops AI chips and software for autonomous vehicles and smart manufacturing applications.

Luminar Technologies: Luminar specializes in lidar sensors and perception software for autonomous vehicles, which can be applied to manufacturing quality control and design validation.

Magna: Magna is an automotive supplier that acquired autonomous vehicle technology from Optimus Ride to advance its R&D.

Motional: Motional formerly known as Aptiv Autonomous Mobility, combines expertise in autonomous vehicles with manufacturing and deployment capabilities.

NVIDIA: NVIDIA's AI and GPU technologies are used in various aspects of autonomous vehicle design, from perception and simulation to training neural networks.

Outrider: Outrider focuses on autonomous yard operations for logistics and manufacturing facilities, using AI to automate tasks such as moving trailers.

Qualcomm: Qualcomm's Snapdragon Ride platform is designed to enable rapid development of future-proof automated driving solutions.

ZF Group: ZF Group offers AI-powered manufacturing and testing solutions for automotive components, including those used in autonomous vehicles.

Zoox: In addition to autonomous ride-hailing, Zoox is involved in the design and manufacturing of autonomous electric vehicles.

AI-ASSISTED MOBILITY SERVICES VENDORS

Airobotics: While not exclusively focused on autonomous vehicles, Airobotics develops autonomous drones that could play a role in future mobility services.

Cognata: Cognata offers a simulation platform for testing autonomous vehicles, aiding in the development of AI-assisted mobility services.

Cortica: Self driving company that mimics the brain of a mammal, enabling vehicles to learn, identify scenarios before they happen, and react to dangerous situations.

Einride: Best known for their self-driving Einride Pod, Einride as a company is completely focused on freight hauling and trucking.

Euler Motors: Euler is an Indian electric vehicle technology startup that is focused on developing and manufacturing commercial electric vehicles.

Ghost Autonomy: Focuses on developing autonomous systems that can operate in harsh weather and low-light conditions.

Luminar: Luminar produces LiDAR technology used in autonomous vehicles to enhance perception and safety.

May Mobility: May Mobility offers autonomous shuttle services for urban mobility.

Mobileye: Mobileye, an Intel company, develops driver-assistance systems (ADAS) and autonomous driving technology, aiming to provide mobility solutions with AI.

Motional (formerly Hyundai-Aptiv): Motional, a joint venture between Hyundai and Aptiv, focuses on autonomous technology for mobility applications.

Nexar: Nexar's dash cams are part of their master plan to build the world's first "safe driving network".

Nissan: Nissan is actively working on autonomous driving technology as part of its Intelligent Mobility initiative.

Nuro: Nuro specializes in autonomous delivery vehicles and is exploring the use of AI for autonomous last-mile delivery services.

RoboSense: RoboSense develops LiDAR sensors and perception solutions that are crucial for autonomous vehicles' AI systems.

WeRide: WeRide is a Chinese startup focused on autonomous driving technology for mobility services and has conducted autonomous taxi trials.

ZF: ZF's platform is focused on autonomous driving technology and has collaborations with various partners for mobility services.

AI-BASED DELIVERY SERVICES VENDORS

Amazon Prime Air: Amazon has been developing Prime Air, an autonomous drone delivery service, aiming to deliver packages to customers' doorsteps.

BoxBot: BoxBot uses autonomous vehicles equipped with AI and computer vision for last-mile package delivery and offers an automated locker system.

Caterpillar: Caterpillar acquired Marble, which builds autonomous ground delivery robots and collaborates with companies for various autonomous delivery applications.

Fabric: Fabric (formerly CommonSense Robotics) focuses on micro-fulfillment centers and autonomous robots for automated order picking and delivery.

Flytrex: Flytrex offers autonomous drone delivery services for food and goods, partnering with restaurants and retailers for deliveries.

Kiwibot: Kiwibot operates a fleet of small delivery robots designed for last-mile delivery in urban environments. They use AI for navigation and obstacle avoidance.

Nuro: Nuro specializes in autonomous delivery vehicles and is known for its compact, purpose-built autonomous delivery vehicles designed for last-mile delivery of goods.

Postmates (Now part of Uber Eats): Postmates was one of the pioneers in on-demand food and goods delivery. They were exploring autonomous delivery options before being acquired by Uber Eats.

Refraction AI: Refraction AI's REV-1 is an autonomous delivery robot designed for both pedestrian walkways and roads, focusing on food delivery.

Starship Technologies: Starship Technologies focuses on small, sidewalk delivery robots equipped with AI and sensors for local deliveries of food and other items.

Swoop Aero: Swoop Aero provides autonomous drone delivery solutions for medical and healthcare supplies, particularly in remote and underserved areas.

Udelv: Udelv builds autonomous delivery vehicles for various goods, including groceries and packages, and uses AI for route planning and navigation.

Zipline: Zipline specializes in autonomous drone deliveries of medical supplies and operates in several countries, primarily in Africa and Asia.

NEXT GENERATION RIDESHARING VENDORS

Cruise: A subsidiary of General Motors, Cruise is developing autonomous vehicle technology for ride-sharing services.

easy mile: EZFleet is the electronic brain driving your fleet of autonomous vehicles.

Lyft: Lyft also had initiatives to develop autonomous vehicle technology for ride-sharing services. They partnered with various autonomous vehicle technology companies.

Navya: Navya is known for its autonomous shuttle solutions, including autonomous ridesharing services in specific locations.

Pony.ai: Pony.ai is a startup specializing in autonomous vehicle technology and has been testing self-driving vehicles for ride-sharing purposes.

Uber ATG: Uber's Advanced Technologies Group (ATG) is focused on developing autonomous vehicle technology for ridesharing and delivery services.

Waymo: Waymo, a subsidiary of Alphabet Inc. (Google's parent company), is one of the leaders in autonomous vehicle technology offering autonomous ride-hailing services.

Zoox: Acquired by Amazon, Zoox is working on autonomous ride-hailing services and autonomous vehicle technology.

TAG

INSIGHT REPORT

AN OVERVIEW OF AI FOR BIOTECHNOLOGY

INTRODUCTION

This TAG Insights Report on Artificial Intelligence (AI) for Biotechnology is intended to help companies, managers, practitioners, researchers, investors, and commercial vendors better understand current trends, issues, and market opportunities in this area. A list of representative commercial vendors working in various areas of AI for biotechnology is included. The five specific areas covered in this report include:

1. AI-Assisted Drug Discovery
2. Personalized Medicine
3. Intelligent Medical Diagnostics
4. AI-Based Health Analytics
5. Intelligent Clinical Trial Support

This report is intended for general and unrestricted use, but interested readers are encouraged to connect with the TAG research and advisory team for more information on the private **TAG Research as a Service (RaaS)** community that covers, discusses, and shares information on these topics in more depth and includes a wider range of startups, vendors, and companies.

OVERVIEW OF AI FOR BIOTECHNOLOGY

The following emerging global commercial opportunities involving AI for biotechnology are covered in this report, including the listing of several viable commercial entities providing solutions on the market today:

- AI-assisted drug discovery is revolutionizing the pharmaceutical industry by utilizing algorithms to sift through extensive biochemical data, enabling the rapid identification of promising drug candidates. These intelligent systems also simulate drug interactions virtually, minimizing reliance on physical trials in the early stages.

- Personalized medicine is becoming a reality with AI's capacity to analyze individual genetic and clinical data, crafting bespoke treatment plans that increase the likelihood of successful outcomes. AI's predictive power also plays a crucial role in pinpointing which patients will benefit from certain medications, thus personalizing healthcare like never before.

- In intelligent medical diagnostics, AI shines by providing unparalleled precision in interpreting medical imaging, leading to earlier and more accurate disease detection. Beyond images, AI's prowess in understanding and extracting pertinent information from written clinical notes is enhancing diagnostic processes.

- AI-based health analytics extend its benefits beyond individual care to public health, where it examines vast health record datasets to discern patterns, predict disease outbreaks, and inform proactive healthcare strategies. It also assists in refining healthcare operations by analyzing patient flows and resource allocations, ensuring better healthcare management.

- Intelligent clinical trial support through AI is transforming how clinical studies are conducted, from patient selection to real-time monitoring of trial progress. By predicting potential outcomes and swiftly recognizing both positive responses and adverse effects, AI is making clinical trials more efficient and adaptive.

FOCUS AREA: AI-ASSISTED DRUG DISCOVERY

AI-assisted drug discovery is a major advance in the pharmaceutical industry, streamlining a labor-intensive and time-consuming process. At its core, this innovation combines the power of artificial intelligence with vast datasets and advanced computational techniques to expedite the identification and development of new medications.

The process begins with data aggregation, where AI algorithms gather and synthesize a multitude of information sources, such as biological databases, chemical libraries, and medical literature. This data pool is then analyzed to identify potential drug candidates. Machine learning models, often powered by deep learning, play a pivotal role in predicting the molecular properties and biological activities of these candidates.

AI's predictive capabilities not only reduce the guesswork involved in drug discovery but also help researchers prioritize compounds for further investigation. Virtual screening and molecular docking simulations allow scientists to assess how these compounds interact with specific target proteins, predicting their potential efficacy and safety. As a result, the drug development timeline is shortened, saving both time and resources. Additionally, AI can uncover previously unnoticed connections and patterns within the data, leading to drug targets and therapeutic approaches.

Ultimately, AI-assisted drug discovery is important to patients in need of new treatments, as it accelerates the translation of scientific insights into life-saving medications, bringing us closer to conquering diseases that have eluded medical solutions. Such synergy between human expertise and artificial intelligence is reshaping the pharmaceutical industry and offering new avenues for innovation in healthcare.

FOCUS AREA: PERSONALIZED MEDICINE

AI-assisted personalized medicine is a major aspect of healthcare innovation, reshaping the way doctors diagnose and treat diseases. This transformative approach harnesses the power of artificial intelligence to tailor medical treatments and interventions to the unique genetic, physiological, and lifestyle characteristics of individual patients.

At the heart of AI-assisted personalized medicine is genomics. High-throughput DNA sequencing technologies generate vast amounts of genetic data, which AI algorithms analyze to identify genetic variations and mutations associated with specific diseases or drug responses. These insights enable healthcare providers to develop a personalized genetic profile for each patient.

With this genetic profile in hand, AI systems can then recommend the most appropriate treatments or therapies, optimizing their effectiveness while minimizing potential side effects. Machine learning models can predict how patients are likely to respond to different medications, allowing for the selection of the most suitable drug and dosage.

Beyond genomics, AI can incorporate data from wearable devices, electronic health records, and even lifestyle factors like diet and exercise. This approach ensures that healthcare recommendations are holistic and tailored to each patient's unique circumstances.

AI-assisted personalized medicine not only enhances treatment outcomes but also enables early disease detection and prevention. By continuously analyzing patient data, AI systems can identify subtle changes that may indicate the onset of a medical condition, allowing for timely intervention.

FOCUS AREA: INTELLIGENT MEDICAL DIAGNOSTICS

AI-assisted intelligent medical diagnostics also leverages the capabilities of artificial intelligence to enhance the accuracy and efficiency of disease detection and diagnosis. The process begins with the collection of patient data, which can include medical history, symptoms, imaging scans, and laboratory results. AI algorithms then analyze this dataset, using machine learning techniques to recognize patterns and anomalies that might not be apparent to human practitioners.

For medical imaging, AI can scrutinize X-rays, MRIs, CT scans, and pathology slides with remarkable precision. It can identify subtle abnormalities or lesions in images, helping radiologists and pathologists make more accurate and timely diagnoses. This is particularly beneficial in fields like radiology, where AI can assist in the early detection of conditions such as cancer, fractures, or neurological disorders.

For non-imaging diagnostics, AI can process a patient's symptoms and clinical data to generate differential diagnoses or recommend further tests. It can also assess the likelihood of specific diseases based on epidemiological data and medical literature, providing valuable decision support to healthcare professionals. Furthermore, AI systems continuously learn and adapt from new data, improving their diagnostic accuracy over time. They can also access a vast knowledge base of medical research and guidelines, ensuring that diagnoses are in line with the latest medical advancements.

Overall, AI-assisted intelligent medical diagnostics holds the potential to reduce diagnostic errors, expedite the diagnostic process, and improve patient outcomes. By combining the strengths of machine learning with the expertise of medical professionals, it represents a significant step forward in diagnosis and treatment.

FOCUS AREA: AI-BASED HEALTH ANALYTICS

AI-based health analytics is revolutionizing how healthcare organizations manage and utilize data to improve patient care, optimize operations, and make informed decisions. This technology leverages artificial intelligence to extract insights from healthcare data. The process begins with data collection from various sources within the healthcare ecosystem, including electronic health records (EHRs), medical devices, wearables, and administrative databases. AI algorithms then process and analyze this data, employing machine learning and deep learning techniques to identify patterns, trends, and correlations.

In clinical settings, AI-based health analytics can assist in patient risk assessment, predicting disease onset, and recommending personalized treatment plans. For instance, it can analyze patient histories to identify individuals at risk of chronic conditions or suggest tailored interventions based on genetic profiles and lifestyle factors. In hospital operations, AI can optimize resource allocation, bed management, and staff scheduling. It can predict patient admission rates, helping hospitals prepare for surges in demand, and reduce readmission rates by identifying patients at risk of complications.

Additionally, AI-powered natural language processing (NLP) can extract insights from unstructured data, such as clinical notes and medical literature, facilitating evidence-based decision-making by healthcare professionals. AI-based health analytics is also pivotal in public health, enabling early detection of disease outbreaks through data analysis and monitoring of social media trends. It aids in the development of targeted public health campaigns and resource allocation during crises.

FOCUS AREA: INTELLIGENT CLINICAL TRIAL SUPPORT

AI-based intelligent clinical trial support is improving medical research and drug development, making the process more efficient, cost-effective, and patient-centric. This approach leverages artificial intelligence to streamline various aspects of clinical trials. The process begins with the design phase, where AI assists in selecting trial endpoints, patient populations, and recruitment strategies. Machine learning algorithms analyze historical trial data and scientific literature to suggest optimal trial designs that are more likely to yield meaningful results.

Patient recruitment, a challenging aspect of clinical trials, benefits from AI. Intelligent algorithms can identify potential participants by scanning electronic health records and other data sources, helping trial organizers find suitable candidates faster and at a lower cost. During the trial itself, AI monitors and analyzes patient data in real-time, detecting adverse events and protocol deviations. This not only ensures participant safety but also allows for rapid adjustments to trial protocols, minimizing delays.

AI can also predict patient dropout rates and stratify participants based on their response to treatment, aiding in the design of adaptive trials that allocate resources more efficiently.

Furthermore, AI-enabled natural language processing can sift through volumes of scientific literature to keep researchers informed about the latest developments in their field, ensuring that trials are conducted with the most up-to-date knowledge.

AI-based intelligent clinical trial support is a game-changer in the pharmaceutical and medical research industries. It accelerates the drug development process, reduces costs, and enhances patient safety, ultimately bringing innovative treatments to market faster and improving the overall quality of healthcare.

COMPANIES AND CONTRIBUTIONS

The companies listed below emerged as part of our research at TAG. Our goal in listing these fine firms is to provide a starting point for buyers, advocates, stakeholders, and researchers trying to make sense of the commercial landscape for artificial intelligence as a means for driving toward improved methods for biotechnology.

AI-ASSISTED DRUG DISCOVERY VENDORS

AMPLY Discovery: AMPLY Discovery uses machine learning and synthetic biology to mine vast biological data to discover novel drug and nutraceutical candidates.

Atavistik: Atavistik is harnessing the power of protein-metabolite interactions to add a new lens to drug discovery.

Atomwise: Atomwise employs deep learning for structure-based drug design, enabling the rapid screening of potential drug compounds.

BenevolentAI: Using AI-driven knowledge graphs, BenevolentAI seeks to uncover novel drug candidates and repurpose existing drugs for new therapeutic uses.

Cloud Pharmaceuticals: Cloud Pharmaceuticals employs AI-driven molecular modeling to design novel drug candidates with a focus on efficiency and cost-effectiveness.

Data2Discovery: Data2Discovery's patent-pending knowledge graph and advanced data science technologies are being used to find breakthrough insights.

Deep Genomics: Deep Genomics integrates AI and genomics to uncover disease-causing genetic mutations and identify potential therapeutic interventions.

Exscientia: Exscientia is working to combine the power of AI and human creativity to make safer and more sophisticated drugs available to all.

Genesis Therapeutics: Genesis is working to develop new medicines using an advanced molecular AI platform.

Healx: Healx focuses on repurposing existing drugs for rare diseases using AI algorithms to match compounds with potential therapeutic benefits.

Iktos: Iktos combines AI and generative chemistry to design drug compounds with improved properties and higher success rates in preclinical testing.

Insilico Medicine: Leveraging AI for drug discovery, Insilico Medicine specializes in generative chemistry, target identification, and predictive biology to accelerate pharmaceutical research.

PathAI: PathAI utilizes deep learning to assist pathologists in diagnosing diseases more accurately through the analysis of pathology slides.

Pfizer: The large company is exploring means for using AI responsibly in the development of new drugs.

Pharnext: Pharnext employs AI to identify synergistic combinations of existing drugs, known as PLEODRUGs, for the treatment of various diseases.

Polaris Quantum Biotech: This company is offering the first quantum empowered SaaS for Drug Discovery.

Recursion Pharmaceuticals: Recursion combines AI and experimental biology to discover new treatments by analyzing cellular imagery and genetic data.

Relay Therapeutics: Relay Therapeutics puts protein motion at the heart of drug discovery to dramatically expand therapeutic possibilities.

Standigm: Standigm is a workflow AI-driven drug discovery company with offices in Cambridge, UK, Cambridge, MA, USA, and Seoul, South Korea.

Turbine.AI: Turbine.AI offers AI-powered solutions for drug discovery, including target identification, lead optimization, and patient stratification.

Unlearn: Unlearn develops generative machine learning methods to predict individual health outcomes.

Valo Health: Valo Health employs machine learning to optimize drug design, helping pharmaceutical companies identify promising compounds.

Xilis: Xilis is developing next-generation technologies to guide precision therapy for cancer patients and accelerate drug discovery and development.

PERSONALIZED MEDICINE VENDORS

Aitia: Aitia utilizes AI and causal machine learning to uncover insights from patient data, guiding personalized treatment decisions and drug development.

BluePrint Genetics: BluePrint Genetics employs AI for genetic diagnostics, supporting clinicians in identifying rare genetic diseases and tailoring patient care.

Congenica: Congenica's AI-driven platform aids in diagnosing rare genetic diseases, providing clinicians with actionable information for patient management.

Exact Sciences: Exact Sciences integrates AI into molecular diagnostics to assist oncologists in making treatment recommendations for cancer patients.

Fabric Genomics: Fabric Genomics combines AI with genomic data analysis to support clinical decision-making for inherited diseases and cancer.

Foundation Medicine: Foundation Medicine employs genomic profiling and AI to assist oncologists in tailoring cancer therapies based on the genetic makeup of individual patients.

Freenome: Freenome combines AI and liquid biopsy technology to detect cancer early and guide personalized treatment decisions.

Genome Medical: Genome Medical uses AI to provide virtual genetic counseling and clinical genomic support for patients and healthcare providers.

GenomOncology: GenomOncology's AI platform assists oncologists in matching cancer patients with optimal targeted therapies and clinical trials.

Genoox: Genoox uses AI and genomic data interpretation to help clinicians identify genetic variations and customize treatment plans for patients.

Kindbody: Kindbody utilizes AI-driven genomics to help clinicians identify genetic markers, assess disease risk, and tailor treatment strategies.

Personalis: Personalis employs AI to analyze cancer genomics data, aiding oncologists in optimizing immunotherapy and targeted therapy options.

PierianDx: PierianDx provides AI-powered genomic informatics solutions for healthcare institutions to interpret and apply genomic data in personalized medicine.

Prenetics: Prenetics offers AI-assisted genetic testing and personalized health solutions, helping individuals make informed health and lifestyle choices based on their genetic profiles.

SOPHiA Genetics: SOPHiA Genetics offers AI-driven clinical genomics solutions to aid in diagnosing and treating inherited disorders and cancer.

Tempus: Tempus harnesses AI to analyze clinical and molecular data, providing insights that inform personalized treatment decisions for cancer patients.

Variantyx: Variantyx provides AI-powered genetic testing and interpretation services for rare and inherited diseases, enabling personalized healthcare plans.

INTELLIGENT MEDICAL DIAGNOSTICS VENDORS

Aidoc: Aidoc specializes in AI-powered radiology solutions, rapidly identifying abnormalities in medical imaging scans to aid radiologists in diagnosis.

Aidpath: Aidpath leverages AI for histopathology analysis, supporting pathologists in diagnosing diseases more accurately and efficiently.

Butterfly Network: Butterfly Network utilizes AI in handheld ultrasound devices, enabling quick and accessible medical imaging for various clinical applications.

Deep 6 AI: Deep 6 AI offers AI-powered patient recruitment solutions for clinical trials, enabling healthcare organizations to identify suitable participants more efficiently.

DeepHealth: DeepHealth offers AI-powered diagnostic tools for radiologists, assisting in the detection and characterization of various conditions.

Digital Diagnostics: Digital Diagnostics offers AI-based diagnostic solutions for diabetic retinopathy, providing autonomous detection of this eye condition.

Embold Health: Embold Health utilizes AI and data analytics to assist healthcare providers in making more accurate and cost-effective diagnostic and treatment decisions.

Enlitic: Enlitic's AI-driven platform analyzes medical images to help radiologists and healthcare providers make more informed diagnostic decisions.

Google Health: Google Health incorporates AI into medical imaging and diagnostic tools, supporting healthcare professionals in diagnosing diseases.

MaxQ AI: MaxQ AI specializes in AI solutions for the interpretation of medical imaging, particularly for stroke and brain hemorrhage detection.

Niramai: Niramai employs AI for breast cancer screening, offering non-invasive and radiation-free detection solutions for early-stage cancer.

Paige: Paige utilizes AI in pathology to assist pathologists in diagnosing cancer and other diseases more efficiently.

PathVisio: PathVisio uses AI to analyze pathology slides and help pathologists detect and characterize diseases, including cancer.

Proscia: Proscia employs AI for digital pathology, streamlining the analysis of pathology slides and aiding in cancer diagnosis.

Quibim: Quibim specializes in AI-based medical image analysis for quantitative assessment of diseases, such as tumors and osteoporosis.

Subtle Medical: Subtle Medical uses AI to enhance medical imaging by reducing scan times and radiation exposure while maintaining image quality.

Ultromics: Ultromics employs AI to enhance echocardiography analysis, aiding in the early detection of heart diseases.

Zebra Medical Vision: Zebra Medical Vision uses AI algorithms to analyze medical imaging data, assisting in the early detection of various diseases.

AI-BASED HEALTH ANALYTICS VENDORS

Aetion: Aetion applies AI to real-world evidence to help healthcare organizations assess the safety and effectiveness of medical interventions.

Apixio: Apixio (merged with Claimlogiq) specializes in using AI to extract insights from unstructured healthcare data, supporting providers and payers in risk adjustment and quality measurement.

Availity: Availity uses AI to transform clinical data into actionable insights, enhancing quality measurement and care coordination for healthcare organizations.

Datavant: Datavant uses AI to connect and de-identify healthcare data, enabling secure data sharing for research and analytics.

Edifecs: Edifecs (acquired Health Fidelity) uses AI-driven natural language processing to improve coding accuracy and risk adjustment in healthcare coding and documentation.

Health Catalyst: Health Catalyst provides AI-driven analytics solutions to healthcare organizations, enabling data-driven insights and decision-making for better patient outcomes.

HealthVerity: HealthVerity employs AI for data governance and insights, helping healthcare organizations securely share and analyze patient data for research and decision-making.

Human API: Human API (part of LexisNexis) offers an AI-powered platform that enables healthcare organizations to access and integrate patient data from various sources.

Innovaccer: Innovaccer leverages AI for healthcare data integration and analytics, enabling care coordination, population health management, and cost optimization.

Komodo Health: Komodo Health uses AI to analyze real-world healthcare data, providing insights for life sciences companies and healthcare organizations.

OM1: OM1 uses AI to analyze real-world clinical data, supporting healthcare providers and life sciences companies in improving patient outcomes.

Redox: Redox offers an interoperability platform that uses AI to standardize and exchange healthcare data, facilitating seamless data sharing across health systems.

Saama Technologies: Saama Technologies provides advanced analytics and AI solutions for life sciences and healthcare to accelerate clinical trials and data-driven decision-making.

Tegria: Tegria (acquired KenSci) applies machine learning and AI to patient data for predictive analytics, supporting healthcare providers in improving patient outcomes.

Waystar: Waystar offers an AI-powered automation platform for healthcare operations, streamlining administrative tasks and reducing operational costs.

INTELLIGENT CLINICAL TRIAL SUPPORT VENDORS

Antidote: Antidote uses AI to match patients with clinical trials, simplifying the patient recruitment process for healthcare organizations.

Clinerion: Clinerion employs AI to accelerate patient identification and recruitment for clinical trials, improving trial efficiency and speed.

CliniOps: CliniOps provides AI-powered clinical trial management software to streamline data collection, monitoring, and reporting.

Medable: Medable uses AI and digital technologies to support decentralized and patient-centric clinical trials, improving trial access and engagement.

Medidata (a Dassault Systèmes company): Medidata offers AI-enhanced clinical trial solutions, streamlining trial management, patient recruitment, and data analytics.

Paradigm: Paradigm (acquired Deep Lens) employs AI to automate patient identification for clinical trials, improving patient access to cutting-edge treatments.

Phesi: Phesi utilizes AI and data analytics to optimize clinical trial design, patient recruitment, and trial execution for life sciences companies.

Science 37: Science 37 leverages AI and telemedicine to conduct decentralized clinical trials, making participation more accessible and patient centric.

THREAD Research: THREAD Research offers an AI-based platform for decentralized clinical trials, enabling remote data collection and patient engagement.

Trialfacts: Trialfacts employs AI to enhance patient recruitment for clinical trials, leveraging digital marketing strategies and analytics.

TriNetX: TriNetX offers an AI-powered global health research network to streamline clinical trial design, recruitment, and real-world evidence generation for healthcare organizations.

INTERVIEWS

AN INTERVIEW WITH WASIM KHALED, CEO AND CO-FOUNDER, BLACKBIRD.AI

UNVEILING THE POWER OF AI TO ADDRESS NARRATIVE ATTACKS CREATED BY MISINFORMATION AND DISINFORMATION

Modern businesses must now contend with a new form of threat known as a narrative attack. Rooted in disinformation and misinformation (see below for a discussion of the difference), the World Economic Forum states it is the #1 Global risk in 2024. There are now solutions that identify and reduce the risk of narrative attacks that are beginning to emerge, and organizations are now prioritizing this new threat vector that costs organizations more than $78 billion per year in financial and reputational harm.

In this interview, TAG's CEO and Founder, Ed Amoroso, discusses with Wasim Khaled, CEO and Co-founder of cybersecurity startup Blackbird.AI, the rise of narrative attacks, and how the company develops technology based on artificial intelligence to reduce the risks of such attacks on organizations. We expect platforms such as Blackbird.AI will soon become mandatory in most environments.

TAG: Help us understand the difference between disinformation and misinformation, and why such difference matters.

BLACKBIRD.AI: The terms "disinformation" and "misinformation" are often used interchangeably, but there is an important distinction between the two:

Disinformation is false, inaccurate, or misleading information that is spread with the intent to deceive or manipulate. Nation-states, cybercriminals, political groups, or other hyper-driven threat actors often create disinformation campaigns to influence public opinion, sow discord, or advance a particular agenda. Misinformation is false, inaccurate, or misleading information that is spread unintentionally, often by individuals who believe the information to be true. Misinformation can stem from honest mistakes, misunderstandings, or the need for fact-checking before sharing information.

Misinformation and disinformation can severely impact companies, governments, and society. The spread of false information erodes trust in institutions and leads to financial and reputational losses, political instability, public health risks, wasted resources, polarization, and decreased productivity. Addressing these challenges requires a purpose-built approach by leveraging AI to illuminate the narratives, the influence behind them, the networks they touch, the anomalous behavior that scales them, and the cohorts that connect them. Knowing this, organizations can make better strategic decisions during a crisis.

TAG: What is meant by a narrative attack and what are its implications for enterprise and government teams?

BLACKBIRD.AI: A narrative attack is a form of disinformation or misinformation that aims to shape public perception and cause financial, reputational, or societal harm to organizations such as corporations and governments. These attacks involve creating and spreading harmful narratives, defined as associations that shape perceptions about a person, place, or thing in the information ecosystem. The challenge to organizations is that narrative attacks can start with a single post, image, or video and quickly spiral out of control and become harmful, blindsiding companies that today are mostly unprepared for such an incident.

The implications of narrative attacks for enterprise and government teams are significant. Narrative attacks can lead to substantial financial and reputational losses, with publicly traded companies losing an estimated $78 billion annually due to disinformation-related incidents. Organizations targeted by these attacks will suffer from financial harm and negative public perception, damaging their brand and credibility. Narrative attacks also pose security risks and can increase the severity of security incidents, causing further financial and reputational harm. The narrative attacks also create the need to share information with the SEC, based on new regulations requiring companies to report “material” cybersecurity incidents within four business days, based on materiality determination, without “unreasonable delay,” as well as describe the incident’s material impact or reasonably likely material impact. Understanding the narratives around an incident or a fake incident within this short period to disclose is paramount information.

Examples of Use Cases that are impacted by narrative attacks created by misinformation and disinformation include:

Cybersecurity incidents
(real or fake breaches)

Geopolitical risk

Trust and Safety

Insider Threat

Supply Chain Risk

Critical Manufacturing attacks

Vulnerability Management

Crisis management

Stock Manipulation

Brand Reputation/Risk

Due Diligence/M&A investigation

Brand Reputation

Financial Markets Risks

Foreign Malign Influence attacks

Propaganda Campaigns

Health risks

Peacekeeping

Blackbird.AI offers an easy-to-understand Narrative Feed that provides a continuous and insightful summary of harmful narratives as they emerge and scale by providing a Narrative Risk Score (NRS) and highlighting critical elements.

TAG: How does the Blackbird.AI solution work and how would it be deployed and used in the typical organization?

BLACKBIRD.AI: We offer an AI-driven Narrative Intelligence Platform called Constellation that enables organizations to detect, measure, gain context, and prioritize narrative risk for critical decision-making. The platform analyzes text, images, and memes across the dark web, social media, and news in 25 native languages to surface emerging narratives, identify influential actors and networks, assess threats, and understand the impact of harmful online activity. By providing visibility into fast-moving narratives, identifying hidden threats and risks, and delivering actionable intelligence, Blackbird.AI empowers companies to respond quickly and strategically to mitigate the financial and reputational harm caused by narrative attacks.

In addition, Blackbird.AI offers an easy-to-understand Narrative Feed that provides a continuous and insightful summary of harmful narratives as they emerge and scale by providing a Narrative Risk Score (NRS) and highlighting critical elements, including cohort activities, bot network manipulation, and key influencers. It enables executive teams to understand the top risks impacting their organization.

We also offer a quick way to gain clarity around misinformation and disinformation through a new product called Compass by Blackbird.AI. It provides clarity to any online claim, article link, supported social media post, or video. When you ask Compass by Blackbird.AI a question or paste any link, it processes data in real-time from thousands of sources, checks claims, analyze results using Blackbird.AI's Narrative Intelligence Platform, and generates an accurate answer with footnotes and citation links.

Furthermore, Blackbird.AI's RAV3N Narrative Intelligence and Research Team comprises data science, national security, cybersecurity, and communications experts. Team RAV3N offers advanced benchmarking, threat development analysis, and learning programs to accelerate decision-making and empower leadership teams in understanding and responding to the evolving digital landscape of narrative attacks driven by misinformation and disinformation.

TAG: How do you envision the evolution of AI-driven narrative attacks in the coming years, and what role does Blackbird.AI play in this landscape?

BLACKBIRD.AI: Generative AI democratized the creation of misinformation and disinformation, making narrative attacks a more potent and widespread threat to every organization. This rapid escalation is partly due to the ease of creating warp reality content through user-friendly interfaces and advanced AI technologies.

The team at Blackbird.AI is more determined than ever to combat narrative attacks and contribute to a more informed and resilient global community. Disinformation is intertwined with societal polarization, human rights erosion, cybersecurity, and even the integrity of elections. We offer a platform all stakeholders can utilize, acknowledging that misinformation and narrative attacks are common threats. In doing so, we aim to contribute to a world where information empowers rather than manipulates, fostering a more informed, stable, and resilient global community.

"I like these new AI tools for farming. But how do I get the tractor through?"

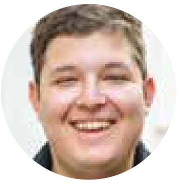

AN INTERVIEW WITH NEIL SEREBRYANY, FOUNDER AND CEO, CALYPSOAI

MANAGING SECURITY, ENABLEMENT, AND PERFORMANCE OF LLMS

The need to scan, flag, and alert on system behavior is central to cybersecurity and an important component of every enterprise protection architecture. It is thus natural that such capability be extended to the common use of large language models (LLMs), now deployed for a variety of tasks around the typical organization.

In this interview, we ask cybersecurity startup CalypsoAI to explain how their solution addresses this challenge and how their platform ensures that applications connected to LLMs are sufficiently protected from threats such as data leakages, information compromise, and other security attacks.

TAG: How does CalypsoAI detect and mitigate cyber threats associated with LLMs?

CALYPSOAI: Our platform is a first-of-its-kind solution for scanning, flagging, alerting, and protecting systems against generative AI vulnerabilities and internal and external risks in real time. Admins can evaluate model performance and behavior insights, such as decision-making processes, limitations, reliability, efficiency, and effectiveness. All LLMs, proprietary and third party, are enabled through the CalypsoAI platform, ensuring a company has 360-degree observability and visibility of generative AI usage. The platform identifies and blocks malware, stops intellectual property from leaving the business, and thwarts attempts to bypass system safeguards.

TAG: Can you explain whether security compliance issues have begun to arise with respect to LLM usage in the enterprise?

CALYPSOAI: Security compliance issues related to LLM use represent a growing area of concern. Organizations must consider data privacy, sensitive information handling, security of model outputs, as well as evolving regulatory requirements, to mitigate compliance risks and ensure the responsible use of these powerful AI technologies.

It's essential for the industry to understand the potential implications and considerations, for instance:

Data Privacy and Confidentiality: Enterprises utilizing LLMs must ensure the data they provide to models complies with data privacy regulations, such as the EU's GDPR or theCalifornia Consumer Privacy Act. Unauthorized or non-compliant data usage could result in legal and financial repercussions.

Sensitive Information Handling: LLMs generate text based on the input they receive. Enterprises must be cautious when using LLMs to handle sensitive information such as personally identifiable information , financial data, or confidential business strategies. Unauthorized disclosure could lead to compliance violations and reputational damage.

Model Outputs: Enterprises must consider the security of model outputs. This includes ensuring generated text does not disclose sensitive information to the wrong people, or introduce security vulnerabilities. Additionally, enterprises should have measures in place to detect and mitigate the risk of malicious actors exploiting LLMs, for instance by generating deceptive or manipulative information.

Regulatory Landscape: Regulatory bodies are increasingly paying attention to the ethical and societal implications of AI technologies, including LLMs. Enterprises must remain informed about emerging regulations and guidelines and proactively address compliance concerns.

Monitoring usage per model is a key capability of our platform, enabling creation of baselines to manage costs as usage scales.

TAG: Does CalypsoAI include the ability to monitor metrics like usage per model per LLM? What other metrics are relevant?

CALYPSOAI: Yes. Monitoring usage per model is a key capability of our platform, enabling creation of baselines to manage costs as usage scales. Cost-management solutions for LLM usage are essential for enterprises to successfully launch internal and external applications and stay on budget. The CalypsoAI platform also offers prompt tracking and auditability, which provides insights for user and model behavior. Understanding who is using the models, when, and why enables companies to create better training and processes, ensuring they remain compliant and maintain competitive advantage.

TAG: How does CalypsoAI ensure seamless integration of its platform with existing cybersecurity infrastructure within organizations?

CALYPSOAI: Because the CalypsoAI Platform is a SaaS solution, it provides a seamless API for quick and easy integration with cybersecurity infrastructure and all existing applications. Users can embed a few lines of code and be up and running in minutes.

TAG: In what ways do you foresee AI shaping the future landscape of cybersecurity, and how is CalypsoAI positioned to lead in this evolution?

CALYPSOAI: We expect AI to maintain a central role in shaping the future landscape of cybersecurity by enabling more proactive, adaptive, and effective defenses against increasingly sophisticated cyber threats. However, AI itself presents new challenges and risks that must be addressed to ensure the security and resilience of digital infrastructure and systems.

Its influence is expected to continue shaping the field in several ways, such as:

Threat Detection and Prevention: AI-powered threat detection systems can analyze vast amounts of data in real time to identify patterns indicative of cyber threats, including malware, phishing attempts, and other malicious activities. Machine-learning algorithms adapt and evolve based on new threat information, enhancing the effectiveness of cybersecurity defenses.

Automated Response and Remediation: AI-driven automation can enable faster and more efficient responses to cyber threats by automatically triggering actions, such as isolating infected systems, blocking suspicious network traffic, or applying patches to vulnerable software, much as our platform blocks private content from leaving an organization and prevents suspicious content from entering. Any automated solution that reduces the cognitive burden on cybersecurity teams allows them to focus on threat mitigation strategies.

Behavioral Analysis and Anomaly Detection: AI algorithms can analyze user behavior and traffic patterns to detect deviations from normal activity, which may indicate unauthorized access attempts, insider threats, or unusual system behavior. For instance, our platform offers customizable rate limits and access controls as preemptive security measures. Behavioral analytics, such as our auditing and tracking feature, helps detect and respond to threats that evade traditional signature-based detection methods.

Predictive Analytics and Risk Assessment: AI algorithms leverage predictive analytics to assess potential cybersecurity risks and vulnerabilities by analyzing historical data, system configurations, and emerging threats. We’ve built in full visibility through interaction monitoring and auditing capabilities. Predictive risk assessment can help organizations prioritize security efforts and allocate resources more effectively to address the most critical threats and vulnerabilities before they are exploited.

Enhanced Authentication and Access Control: AI technologies such as biometric authentication and behavioral biometrics can provide more secure and user-friendly authentication methods, reducing the reliance on traditional passwords and mitigating the risk of credential-based attacks.

Adversarial AI and Cybersecurity Arms Race: As AI technologies advance, concerns are growing about malicious actors potentially leveraging AI-powered attacks involving malware, adversarial machine learning, and social engineering. The cybersecurity community must continue to innovate and develop defensive strategies to counter emerging AI-driven threats and win the arms race between attackers and defenders.

Our generative AI security and enablement platform is leading an evolution in cybersecurity, uniquely positioned to define the future landscape of digital defense. By harnessing the power of advanced AI algorithms, our solution offers unparalleled capabilities in threat detection, prevention, and response, empowering organizations to stay ahead of emerging cyber threats confidently. Our tool goes beyond static signature-based detection methods, leveraging dynamic generative models to anticipate and adapt to evolving attack vectors in real time and providing holistic security coverage across the LLM ecosystem. Our commitment to continuous innovation ensures our platform remains at the leading edge of cybersecurity technology, ready to address the challenges of tomorrow's cyber landscape with unrivaled precision and efficacy.

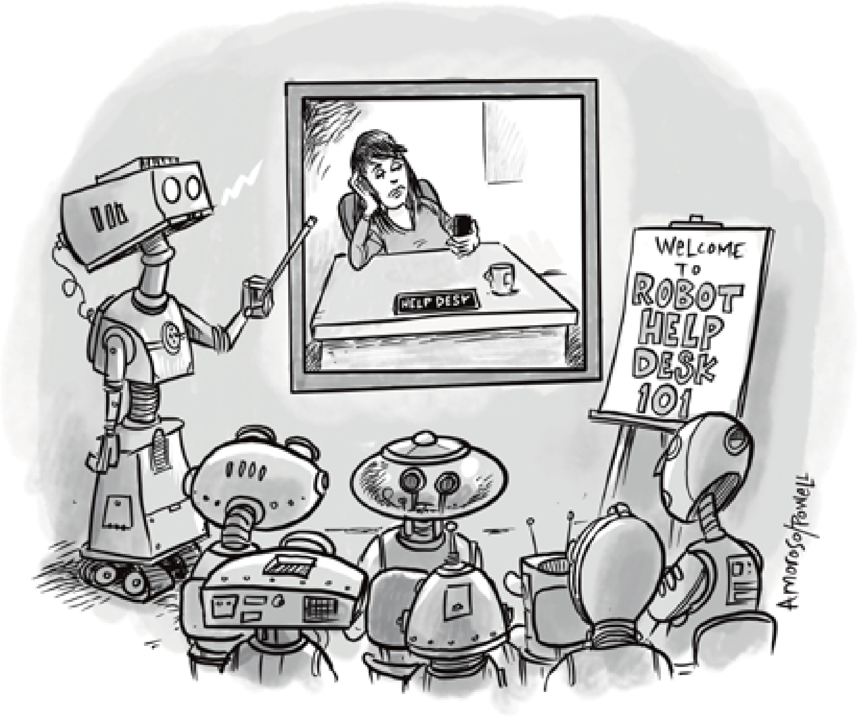

"You'll need to learn to mimic the lethargic disinterest of the human help desk operator."

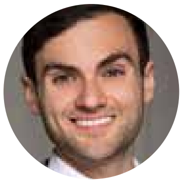

DANIEL CHRISTMAN,
DIRECTOR OF AI PROGRAMS,
CO-FOUNDER, CRANIUM

ESTABLISHING SECURITY AND COMPLIANCE FOR ARTIFICIAL INTELLIGENCE

The revolution in use of artificial intelligence (AI) is well-underway and enterprise teams must now grapple with the challenge to ensure security against threats and to establish compliance with applicable regulations and standards.

We recently spent time with startup company Cranium to better understand their approach to AI security and how they envision such work to help build trust in the use of Ai across the entire business and government landscape.

TAG: Can you share examples of how you see organizations using AI to transform their business today?

CRANIUM: Large organizations across various industries are leveraging the power of AI to drive innovation, improve operational efficiency, and deliver enhanced products and services to their customers. Some examples include the automotive manufacturers are using AI for autonomous vehicles, diagnostics, and predictive maintenance. Healthcare leverages AI for medical imaging, patient care, analyze images, identify diseases, and personalize treatment plans. Large manufacturing companies integrate AI into production processes for predictive maintenance, quality control, and supply chain optimization.

TAG: Do most customers of yours have a good idea of how and where they are using AI, or is it necessary to help them perform discovery?

CRANIUM: The level of understanding varies by company. Some customers may have a good grasp of AI technologies and their potential applications, while others may require assistance in performing discovery to uncover opportunities for leveraging AI in their business operations. It is always important to gain AI visibility and assess AI risk.

TAG: How does your platform work? What would customer do to deploy and use your technology?

CRANIUM: Our platform works by leveraging different sensors that we deploy in our customers environment to capture relevant information about the internal and external AI being leveraged. Our lightweight code sensor can be pointed at public and private code repositories to automatically capture an AI bill of materials for a particular AI system. We additionally offer a deeper connector that deploys in the customer's cloud environment to scan the AI development environments, such as Azure OpenAI Service, Azure ML, and AWS Sagemaker. Customers can then use all this information for continuous monitoring and compliance against internal governance frameworks, as well as broader regulatory requirements, using the AI Card. Customers can also leverage the AI Card to capture information about the AI being used in their supply chains.

TAG: Do you see an acceleration of external regulatory or compliance requirements and demands in this area?

CRANIUM: Absolutely. Just this week, the EU enacted the leading regulation for AI with the AI Act. In the US, most of the regulations exist at the state level and are still pending, but there are many in development. As of February, there are ~400 local AI related bills proposing regulations. Some of the notable regulations in development include California's proposed SB 294, 'The Safety in Artificial Intelligence Act'. One active law in New York, Local Law 144, began enforcement in July of 2023 and focuses on regulating automated employment decision tools, particularly those leveraging AI. Given the EU's AI Act, the development of local regulatory requirements is likely to accelerate, as state legislatures will have something to baseline their requirements against.

TAG: In your view, what is the future of AI in business and government? Will we see a hybrid arrangement of human and AI activity for most tasks?

CRANIUM: In the future, we see AI being a part of every key process for both enterprises and federal entities. Given the efficiency gains already being showcased for organizations leveraging AI, there looks to be no sign of slowing for AI use-cases being moved into production environments. There will likely be a hybrid arrangement between human and AI for all tasks, just like how the computer is symbiotic with every current workflow today.

There will likely be a hybrid arrangement between human and AI for all tasks, just like how the computer is symbiotic with every current workflow today.

AN INTERVIEW WITH MARK WOJTASIAK, VP OF PRODUCT, VECTRA AI

REDEFINING CYBER THREAT DETECTION WITH AI

Few cybersecurity companies have the depth of experience in applying artificial intelligence (AI) to cybersecurity as Vectra AI. They have developed a strong capability in using AI to accurately detect cyber-attacks, investigate and initiate response – and this clearly matches the needs of most modern organizations seeking to improve their defenses and thus resilience to attacks.

We recently spent time with the Vectra AI team to better understand how they are using AI to speed up cybersecurity detection and mitigation as well as how their customers are coming to depend increasingly on AI as a critically important component of their defense against automated attacks.

TAG: How does Vectra AI's AI-driven platform enhance organizations' ability to detect and respond to cyber threats in real-time?

VECTRA AI: We think about this in the most basic terms. At the core, an organization's ability to detect and respond to cyber threats in real-time comes down to three questions: can we see it, can we stop it, how fast can we see it and stop it? With attack surfaces constantly expanding, attacker methods evolving, new threats emerging, and a barrage of alerts, traditional methods of threat detection and response are overly manual, complex, latency ridden. On top of that, the shortfall of SOC resources and skills has only made matters worse and when it comes to the question of how fast can we see it and stop it? The answer is not nearly fast enough.

We have a premise that serves as the basis for our AI platform –modern enterprises are hybrid, thus all attacks are hybrid attacks. We argue in the modern hybrid enterprise, hybrid attacks are rendering traditional approaches to threat detection and response inefficient and ineffective. For SOC teams, detecting a hybrid attack is like finding the needle in a stack of needles. The only way to find the needle is to think like a hybrid attacker. Today, we think we have individual attack surfaces to manage – endpoints, networks, identities, clouds, email applications, etcetera, but hybrid attackers see one giant integrated attack surface. Integrated being the key word and our platform is designed to give defenders a real-time integrated view of attacks across the entire hybrid attack surface. This removes complexity and latency in detection, investigation and response processes and dramatically reduces SOC analyst workload.

TAG: Can you elaborate on the specific AI techniques and methodologies employed by Vectra AI to analyze network behaviors and identify malicious activities?

VECTRA AI: The Vectra AI approach to threat detection blends human expertise with a broad set of data science and advanced machine learning techniques. This model delivers a continuous cycle of attack intelligence based on security research, global and local learning models, deep learning, and neural networks. Using behavioral detection algorithms to analyze metadata from captured packets, our cybersecurity AI detects hidden and unknown attacks in real time, whether traffic is encrypted or not. Our AI only analyzes metadata captured from packets, rather than performing deep-packet inspection, to protect user privacy without prying into sensitive payloads.

Global learning begins with the Vectra AI Threat Labs, a full-time group of cybersecurity experts and threat researchers who continually analyze malware, attack tools, techniques, and procedures to identify new and shifting trends in the threat landscape. Their work informs the data science models used by our attack signal intelligence, including supervised machine learning. It is used to analyze very large volumes of attack traffic and distill it down to the key characteristics that make malicious traffic unique.

Local learning identifies what's normal and abnormal in an enterprise's network to reveal attack patterns. The key techniques used are unsupervised machine learning and anomaly detection. Vectra AI uses unsupervised machine learning models to learn about a specific customer environment, with no direct oversight by a data scientist. Instead of concentrating on finding and reporting anomalies, Vectra AI looks for indicators of important phases of an attack or attack techniques, including signs that an attacker is exploring the network, evaluating hosts for attack, and using stolen credentials.

Our AI prioritization engine combines thousands of events and network traits into a single detection. Using techniques such as event correlation and host scoring, our AI correlates all detection events to specific hosts that show signs of threat behaviors. We then automatically score every detection and host in terms of the threat severity and certainty using our own threat certainty index. Finally, we track each event over time and through every phrase of the cyberattack lifecycle putting special focus on entities that are strategic value to an attacker.

TAG: What sets Vectra AI apart in AI-driven cybersecurity, particularly in terms of scalability and adaptability to evolving threats?

VECTRA AI: I would say it's our approach. A decade ago, we created a methodology built on five core principles of applied AI for cyber security. 1. Start with the right problem statement. 2. Get

“The Vectra AI Platform helps our engineers and analysts take ambiguity out of their day and focus on what matters.”

the right data. 3. Build an ML engineering competency. 4. Unlock ML innovation with platform. 5. Continually validate and improve. Our methodology is rooted in integrating security research, ML engineering – the combination of Data Science and Engineering - and UX focused on one mission: use AI to find attack signal inside data at speed and scale. We have over 150 models spanning neural networks, supervised ML, unsupervised ML, and novelty detection, 12 references for MITRE D3FEND – more than any other vendor – and a network effect made up of over 1500 customers continuously validating and improving our AI detections for both existing attacker techniques and new ones we discover. There are cases where we’ve identified new attacker techniques and developed detections before they are published in MITRE ATT&CK which means our customers get continuous coverage for new attack techniques without any detection engineering work.

TAG: How does Vectra AI ensure minimal false positives and false negatives in its AI-powered threat detection algorithms?

VECTRA AI: We believe detecting and responding to what we call hybrid attacks in real-time can only be done with AI. AI is the only way to deliver SOC teams what they need – integrated accurate hybrid attack signal at speed and scale. We call it our attack signal intelligence and it uses AI to analyze, triage and correlate thousands of detection events a day spanning networks, identities, clouds, and SaaS applications per day. Instead of delivering thousands of alerts on individual threat events, our AI platform delivers single digit alerts per day on prioritized entities – both hosts and accounts – under attack.

In the most basic terms, our AI answers the three questions SOC analysts need answered every day they sit in front of their monitors: is this threat real, do I care, and how urgent is it? In other words, is it worth my time and talent. One of our customers put it best “the Vectra AI Platform helps our engineers and analysts take ambiguity out of their day and focus on what matters.” How we do it is simple. We leverage our pre-built, behavior-based, domain specific AI detections to make unknown attacks known. We us AI to integrate and automate threat event correlation to remove detection engineering latency. And most importantly, we use AI to shift the analyst experience from event-centric threat detection to entity-centric signal prioritization, thus reducing noise and workload thus maximizing the value of existing SOC talent.

TAG: Looking ahead, what role do you envision AI playing in the future of cybersecurity defense, and how is Vectra AI positioned to lead in this ongoing evolution?

VECTRA AI: Like I said before, today, detecting and responding to hybrid attacks in real-time can only be done with AI. We see the future as a fully AI-driven SOC. The first phase of evolution is all about the use applied AI for proactive defense - from identifying emerging attacker behaviors to detecting and prioritizing entities early in an attack campaign. We believe our AI ML approach; our attack signal intelligence and our entity-centric prioritization engine is at the forefront of this movement.

I see phase two of the AI-driven SOC comes in the use of generative AI for prescriptive defense related to threat investigations and response. We see this happening already with vendor adoption and use of LLMs to help SOC analysts reduce investigation workload and speed investigation times. Potentially, AI could be taken a step further and prescribe, even take the appropriate response action to contain or isolate the attack. Vectra AI has chosen to focus first and foremost on delivering the most integrated and accurate attack signal. We contend the more accurate the attack signal, the more compelling the application of LLMs for effective investigation and response.

I see phase three of the evolution of the SOC is AI for predictive defense. Given our understanding of attacker behaviors – our attack signal intelligence - combined with our approach and the network effect we enjoy from our 1500 plus enterprise customers, Vectra AI is well positioned to lead and innovate in predictive AI-driven defense.

"As a matter of fact, I think this class could be replaced with an app."

AN INTERVIEW WITH RICK CACCIA, CEO AND CO-FOUNDER, WITNESSAI

ENABLING SAFE AND SECURE AI

The need for security and governance controls to ensure safe and effective artificial intelligence (AI) usage across the enterprise has now emerged as one of the most important objectives for leadership teams. The goal is to protect the use of AI, including the data, but without compromising the benefits of AI for the organization.

We spent time with startup company WitnessAI to better understand how their platform provides the benefits. Our emphasis in the discussion is on how the platform ingests data, performs its process, and provides actionable insights and support for the modern enterprise security team.

TAG: What are the main functions supported in the WitnessAI platform?

WITNESSAI: The WitnessAI platform applies runtime guardrails to enterprise AI use. The product has three major functions. First, it provides visibility, via our Witness/OBSERVE module. You turn it on and it tells you which of the many AI projects/chatbots/LLMs on the internet your employees are accessing, and specifically what prompts and responses they are generating in those systems. It's got a database of thousands of AI domains on the internet and we use it to give you useful insights: which systems are your employees using, where do those systems host their data, do those systems use risky components, etc. It does the same for internal access to private LLMs, so you have a unified catalog of user AI activity.

Second, once you can see what your people are doing, you'll want to apply some controls, so our Witness/CONTROL module provides that via AI-oriented policy. By AI-oriented policy, I mean statements such as "data from client A can't be used by our internal LLM to generate documents for client B". Once you can see and control it, you want to protect it, so the third function is security. Our Witness/PROTECT module can redact or block sensitive data on the fly, detect jailbreaks and prompt injection, detect hallucinations based on training data, as well as other LLM-specific security actions.

The point is to enable safe and effective adoption of enterprise AI. Balancing safety (stop unwanted things from happening) and effectiveness (employees are more productive) is the key; it's not all about scary attacks.

The point is to enable safe and effective adoption of enterprise AI. Balancing safety (stop unwanted things from happening) and effectiveness (employees are more productive) is the key; it's not all about scary attacks.

TAG: Are you connecting with app and chatbots – and how is this accomplished?

WITNESSAI: We connect with public and private LLM-based apps and chatbots in three ways. We have a simple API, like a "Twilio for AI" that developers can use to apply our guardrails to prompts and responses as they embed chat into their apps. Next, we have a browser-based universal chat app that enterprises can restrict to a specific set of chatbots. This is for organizations that want all employee use to be through a controlled client and only with certain chatbots. Finally, we have a network connector that works with native apps such as VS Code, MS Word, etc to apply guardrails to co-pilot use. All of these work without any endpoint agent or plugin.

For all of this, we create a separate, encrypted instance of our platform and guardrails for each customer, in the cloud of their choice. We don't think enterprises are going to be comfortable shipping all of their AI activity data to a multi-tenant SaaS service. We can't see the customers' data, and we think that's the right approach.

TAG: What are the different security methods required for public versus private LLMs?

WITNESSAI: For public LLMs, visibility and data security are very important. Companies don't want copyright, confidentiality, or other issues when their employees are using random public chatbots. For private LLMs, visibility is easier, but data and topic access control become much bigger concerns. How do you ensure regulatory separation across customer data sets? How do you limit which topics can be asked of a private LLM? We address all of these. Of course, for both scenarios, prompt injection, jailbreak, hallucination, etc are also concerns and we add guardrails for those, too.

TAG: Do you expect your platform to be used increasingly for AI audit and compliance?

WITNESSAI: Actually, we think it's the other way around – audit and compliance are the current major use cases, and we think security will grow later. Today, the most common statement we hear from CISOs is "I don't know which AI apps my users are actually using. Can you show me where they are going and what they are doing?" We don't hear CISOs worrying about abstract advanced new attacks just yet. Our view is that today, spending on AI guardrails is compliance-driven, and as attacks become public in the future, spending will become more reactive and security-driven. Today, the most immediate use case for enterprises is some version of "my business users want to use chatGPT or some other chat system, and I want to let them do that, but in a way that is safe and that we can see."

TAG: What sort of future advances should we expect to see from commercial vendors supporting AI security in the future?

WITNESSAI: There are two different types of commercial vendors here. The first is legacy security vendors. When major shifts happen, such as mobile or cloud, and now AI, there is a view that the legacy guys will solve the new problems. In practice, it takes a very long time and occurs only after a lot of M&A. After they reach a certain size, legacy security vendors don't really build anything anymore. New problems are always solved first by new vendors. Some get acquired and a few rise to become the new platform players. The second type is the LLM vendors themselves. They will definitely address parts of LLM security, most likely around model security. For example, protecting against poisoning, drift, etc., areas that they are well-suited to solve for their own products. Enterprises may find that they need something that is consistent across LLM suppliers, in the same way that the cloud hyperscalers are very good at protecting their own clouds, but vendors that provide posture management across all of them are also doing well.

At WitnessAI, we are focusing on ensuring safe use of enterprise AI, rather than ensuring the security of the model itself. We don't think that is an area where the LLM providers will build solutions, why would they? Customers need cross-LLM solutions for observing activity, applying consistent data and topic policies, and securing information. We think that is a market that is going to grow very rapidly, and it's where we are focused.

"I wonder if we're using the same AI-matching software?"

TAG

DISTINGUISHED VENDORS

Q1 2024

Working with AI vendors is our passion. It's what we do every day. Following is a list of the Distinguished Vendors we've worked with this past three months. They are the cream of the crop in their area—and we can vouch for their expertise. We are eager to celebrate good technology and solutions when we find them. And the vendors below certainly have met that criteria.

BLACKBIRD.AI

Blackbird.AI protects organizations from narrative attacks created by misinformation and disinformation. Our AI-based Narrative Intelligence Platform enables you to know the harmful narratives that impact your organization/industry, the influence behind them, the networks they touch, the anomalous behavior that scales them, and the cohorts that connect them to help you make better strategic decisions when a crisis hits.

CALYPSOAI

CalypsoAI is the leader in AI Security and Enablement. As a trusted partner and global leader in the AI Security domain, CalypsoAI empowers enterprises and governments to leverage the immense potential of generative AI solutions and large language models (LLMs) responsibly and securely. CalypsoAI strives to shape a future in which technology and security coalesce to transform how businesses operate and contribute to a better world.

CRANIUM.

Cranium helps cybersecurity and data science teams understand everywhere that AI is impacting their systems, data or services. We secure your organization's AI and machine learning systems to ensure they are compliant and trustworthy, without interrupting your workflow.

VECTRA®

Vectra AI is the leader in AI-driven hybrid attack detection, investigation, and response. The Vectra AI Platform with patented Attack Signal Intelligence modernizes the SOC to rapidly detect, prioritize, and stop the most advanced cyber-attacks. Organizations worldwide rely on Vectra AI to find the cyber-attacks other tech can't.

WITNESS AI

The WitnessAI Platform enables safe and effective adoption of enterprise AI, through security and governance guardrails for public and private LLMs. WitnessAI provides visibility of employee AI use, control of that use via AI-oriented policy, and protection of that use via data and topic security.