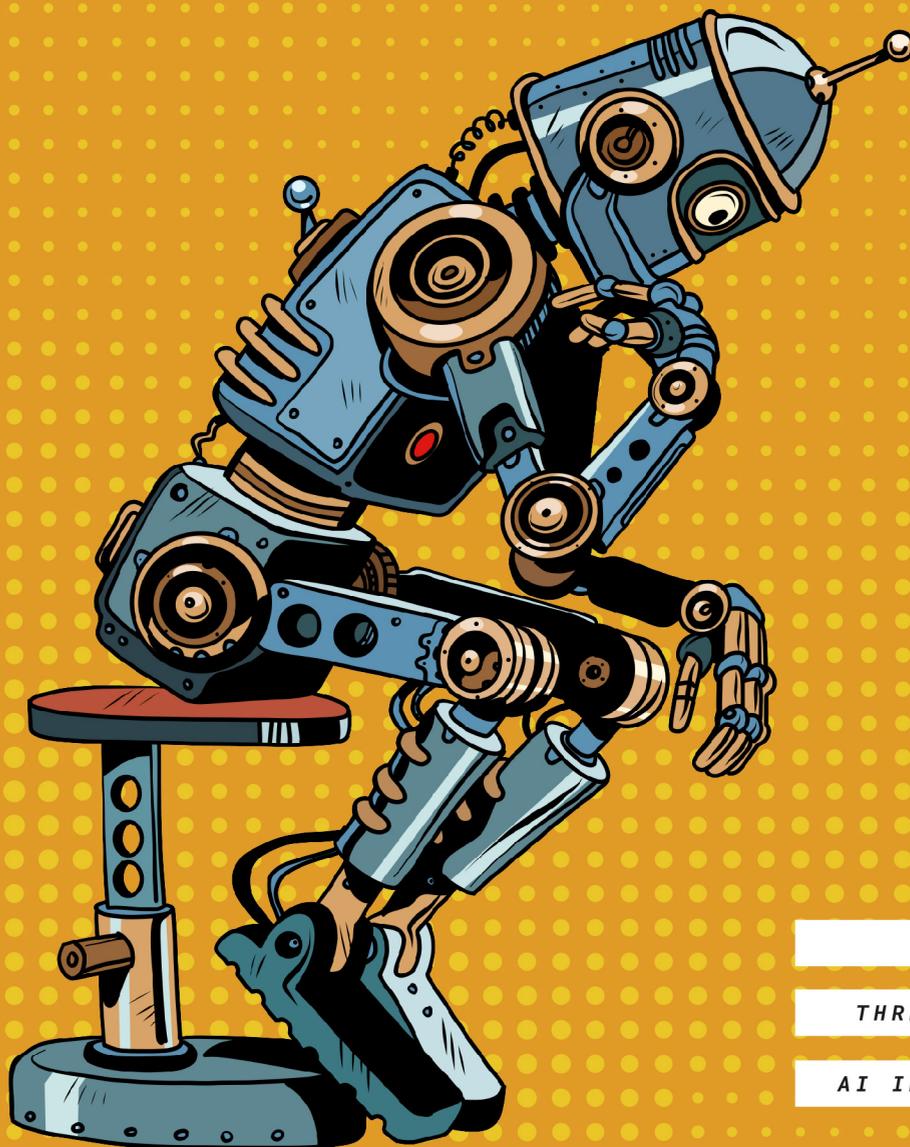


COGNITIVE TIMES

CAN AI TRULY *UNDERSTAND* LANGUAGE?

VOL. 5 NO. 3 // 2020



INSIDE

WHAT IS GPT-3?

THREE COMING SHIFTS IN AI

AI IN THE FINANCIAL SECTOR

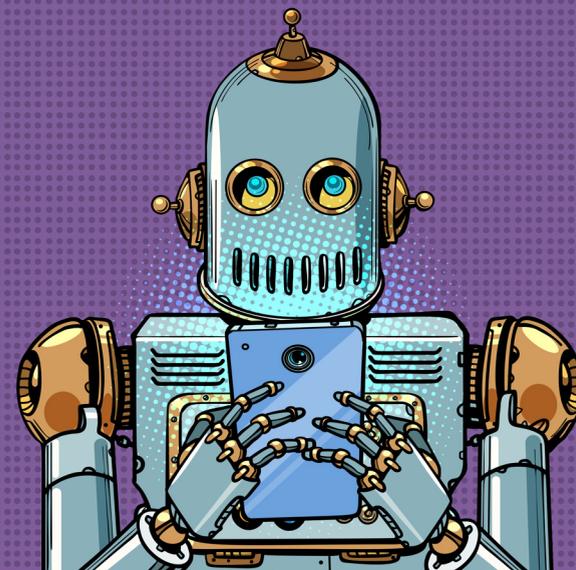
INNOVATION THAT'S OUT OF THIS WORLD



Innovators see beyond physical limitations to create amazing possibilities. Boeing is proud to work with partners who are committed to creating a better world.



FEATURES



32

WHAT IS GPT-3, AND WHAT DOES IT MEAN FOR NATURAL LANGUAGE PROCESSING?

Prof. Bruce Porter digs in and addresses many questions about the development of GPT-3 and the future of NLP.

BY BRUCE PORTER



38

THREE COMING SHIFTS IN AI

Amir Husain explores the future of AI and shares three concrete developments that he predicts are in the not-so-distant future.

BY AMIR HUSAIN



44

AI IN THE FINANCIAL SECTOR

Jezri Mohideen, Global Chief Digital Officer for Wholesale at Nomura, shares his journey from science to finance.

BY AMANDA CENTER

CONTENTS



09

In Conversation

Sameer Qureshi from Lyft on autonomous vehicles.

BY COGNITIVE TIMES STAFF

12

Elevator Pitch

Learn about Open Ocean Robotics' unique solution to a global problem.

BY COGNITIVE TIMES STAFF

14

Data Points

A roundup of recent tech news you need to know about.

BY PRISCILLA LIU



16

Tech Explainer

Unstructured data poses many challenges. What are the solutions?

BY COGNITIVE TIMES STAFF

18

Meet The Mind

Gil Gur Arie shares his perspective on the future of AI at Ford.

BY COGNITIVE TIMES STAFF

20

Market Leader

SkyGrid is pushing the boundaries of autonomous drone flights.

BY ZEHRA AKBAR

22

Action Plan

Your tech-savvy to-do list of stuff to read, see, and virtually attend.

BY COGNITIVE TIMES STAFF

26



24

The Jetsons & NLP

The future of natural language processing.

BY MAHE ZEHRA

26

Flyover Country

Fiction from the future.

BY AMIR HUSAIN AND STEVE NORDLUND



30

Interact with GPT-3

See GPT-3 in action.

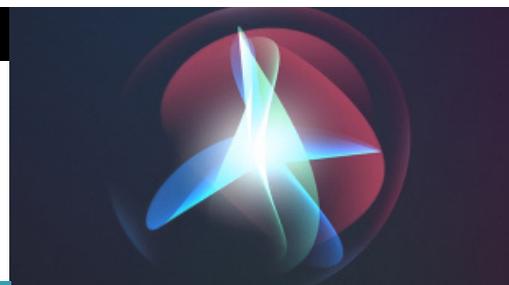
BY MAHE ZEHRA

50

Turing Talk

An interview with Siri.

BY PRICILLA LIU



An all-in-one drone app to automate every phase of flight.

No matter your mission, whether to inspect a pipeline, respond to an emergency, or secure a perimeter, we make it easy to autonomously execute your flight and detect objects in real-time with AI computer vision.



Did we mention it was free?

Visit skygrid.com/flight-control

EDITOR IN CHIEF

Amir Husain

MANAGING EDITOR

Mahe Zehra

EDITOR

Amanda Center

CREATIVE DIRECTOR

Milton Lopez

DESIGNER

Ahsan Naveed

STAFF WRITERS

Zehra Akbar

Pricilla Liu

CONTRIBUTING WRITERS

Prof. Bruce Porter is a two-time Chair of the University of Texas Computer Science Department and serves as SparkCognition's Chief Science Officer, where he leads the company's many R&D initiatives. Prof. Porter's research focuses on machine reading, a technology that holds tremendous potential for capturing knowledge for automated inference, question answering, explanation generation, and other AI capabilities. He has won the Best Paper Award at the National Conference on Artificial Intelligence, the College of Natural Sciences Teaching Excellence Award, the National Science Foundation's Presidential Young Investigator Award, and the President's Associates Teaching Excellence Award.



Steve Nordlund is vice president and general manager of Boeing NeXt, a business division dedicated to building a safe, sustainable and seamless future mobility ecosystem. Boeing NeXt is driving the company's approach to next-generation airspace management and global airspace integration, working with regulatory agencies and industry partners to ensure a future where autonomous and piloted air vehicles safely coexist. Nordlund oversees the development of Boeing's next-generation mobility platforms - a portfolio that includes cargo and passenger air vehicles. Since joining Boeing in 2009, Nordlund has held business development, business line and strategy leadership positions across the enterprise, including his role as vice president of Autonomous Systems for Boeing Defense, Space & Security. Most recently, he was vice president of Boeing HorizonX, a strategy organization that experiments with business models for emerging aerospace applications, leveraging nontraditional partnerships and a venture fund that invests in early-stage startups.



Image: freepik.com/peenat

To all doctors, nurses, and healthcare staff on the front lines of the COVID-19 pandemic: We are grateful for all that you do. Thank you for being lifesavers, and please stay safe!



application of that knowledge through the development of technology, that we will conquer this and future pandemics. While innovation in the form of a vaccine has been top of mind for many (including myself), there are several other technological milestones that were reached this year. The march forward of technology is incessant. As I have often said, technology is inevitable. The human effort that propels technology forward requires perseverance and a desire to pursue the art of the possible with tools of rationality and science. Ultimately, it is these pursuits that guarantee not only the health of our loved ones, but also the health of the world around us.

In that spirit of the perseverance of innovation, this edition of Cognitive Times takes you through many recent, exciting developments across a range of exponential technologies and breaks down how they are being applied to the industries that shape our world.

For example, Ford Motor Company's Chief Data and Analytics Officer talks to us about how the organization is placing a heavy emphasis on "modernization" and reaching what they call "data superiority". While Sameer Qureshi of Lyft provides us with a glimpse into his thoughts on autonomous, self-driving vehicles. We also cover transportation with Open Ocean Robotics founder Julie Angus, explaining the importance of the ocean data collected and analyzed by their vessels, and with a spotlight on SkyGrid's innovations in drone technology and unmanned traffic management software.

In the pages that follow, you will also find a fascinating discussion with Nomura's Wholesale Chief Digital Officer, Jezri Mohideen, on how technologies such as artificial intelligence and cryptocurrency continue to bring breakthroughs to finance and trading. And with deep learning techniques making strides in language with solutions like GPT-3, we take a deeper dive into what this technology means, and how people can interact with it. You may even see the masterworks of a couple of human-AI poet collaborators!

So, as we enter into a new year - one that may still be filled with challenges and change, but also, most definitely, of hope - let us remain as positive as we can, take care of those around us, and not lose sight of our commitment to constant innovation. Here's to 2021, and to the incredible discoveries, technological progress, and human happiness I hope it brings to us all.

Amir Husain

Founder & CEO of SparkCognition

"The Dreams We Set Out To Conquer": SC AI Human- Collaboration
"Don't Overthink": Line Gauthier

NOTES

FROM THE EDITOR

As 2020 comes to an end, we look back at a very... strange year. We have watched the entire world significantly adapt in a pandemic that continues to surge across the globe. Working from home took on an urgency and intensity that few amongst us would have imagined just months ago. Medical researchers persevered and delivered a major vaccine faster than ever before. But through this year of endless change, our fundamental need and commitment to prioritize science and technology has not wavered.

Throughout history, we have seen the many challenges technology has helped us conquer and the complex environments it has helped us navigate. More than anything, and more than ever before, COVID-19 provides further evidence that ultimately, it is our technological prowess that will be humanity's great saviour. It is the knowledge gained through the scientific method, and the



TEXAS

The University of Texas at Austin

MACHINE LEARNING LABORATORY

THE NEXT SCIENTIFIC FRONTIER

UT Austin is the proud home to the new Machine Learning Lab and NSF AI Institute for Foundations of Machine Learning. We create the technology that drives AI systems, enabling them to acquire knowledge, learn from real-world experience, and make predictions in complex environments.

<https://ml.utexas.edu/>

Protect critical OT control systems with AI



Increased digitization and the growing need for connectivity exceptions in OT environments are challenging traditional perimeter-based security models. Using AI, the DeepArmor Industrial solution prevents zero-day attacks even on isolated systems.

Start protecting your systems today.

LEARN MORE:
go.sparkcognition.com/deeparmor-industrial



BY COGNITIVE TIMES STAFF WRITER

SAMEER QURESHI

We ask Sameer Qureshi, Director of Product (Software), Level 5, Self-Driving Division at Lyft, about the future of the automobile industry.

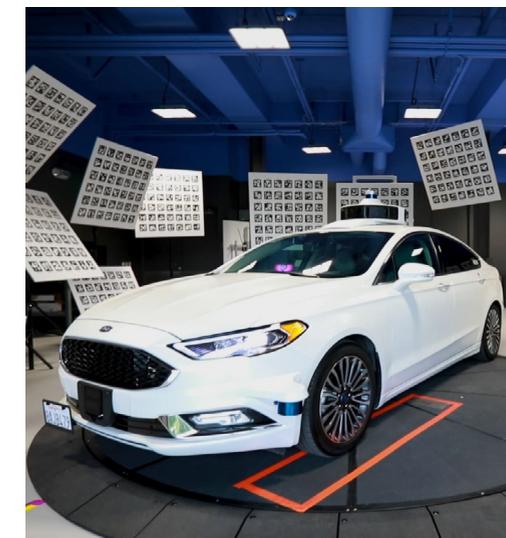
COGNITIVE TIMES: What has your career journey looked like so far? Could you tell us a little bit about how you got to Lyft?

SAMEER QURESHI: I've always been a software person. I went to Carnegie Mellon for a Bachelors in Computer Science, and to Stanford for a Masters in Compu-

ter Science. And I've always wanted to work on software that is very close to hardware (think operating systems, device drivers, robots, etc.). So I spent two years at IBM working in the AIX Core Kernel Team (AIX is IBM's variant of Unix). I also spent 8+ years at Intel working on device drivers for Intel's GPUs, specifically focusing on video - and along the way, I worked with Apple to help them enable what became FaceTime on the Mac, leveraging Intel's GPU capabilities.

Having spent many years working with hardware, I wanted a new challenge that wasn't directly working with silicon, but still had some aspect of hardware. So I decided to join Tesla to work on the Model X, with its iconic Falcon Doors. And until I joined Tesla, I didn't realize how much software goes into one of their cars. After taking Model X through launch, production ramp, and developing its various variants, I was asked to help the Autopilot Team with the Autopilot 2.0 launch - which at the time was a gargantuan task, given multiple changes in Autopilot leadership, and the immense time pressure we were in to bring Autopilot 2.0 up to parity with Autopilot 1.0.

Finally, after having spent 3.5 years at Tesla, I decided to look for a different opportunity, but stay within the Autonomous Vehicle industry. My belief is that the very first deployment of AVs will be for rideshare. Given that, Lyft was an obvious choice. They already had a large ridesharing network, were actively trying to bootstrap Lyft Level 5 (the AV division), while also building a solid team.



CT: There are so many competing approaches to self-driving cars; the Waymo approach, the Tesla approach, the GM approach, etc. What are the key differences between these approaches and which do you think has the best chance of working?

SQ: The good thing about all these approaches is that they all lead to self-driving cars. And then all these self-driving cars get deployed on rideshare networks! The best approach is one that enables AVs to be deployed at scale on rideshare networks, which is why Lyft has a two-pronged approach to self-driving. We have our Level 5 division which is focused on building the tech for AVs in-house, and an Open Platform team that partners with companies such as Waymo and Motional to get their AVs on our network.

Tesla and Waymo have significantly different approaches towards solving the problem, but that is driven by the fact that Tesla already has more than a million cars on roads worldwide, that need to be updated via over-the-air software updates for improved functionality and increased capabilities. Waymo, on the other hand, is building an AV from the ground up, and outfitting their cars with a lot more sensors and compute - but designing the car from a driverless perspective up front.

Both approaches involve significant amounts of machine learning techniques, that are designed to enable various features of the AV stack. Both approaches also have different pros or cons and different challenges. But the jury is still out on which approach has the best chance of working. If the industry knew what the correct answer was, they would all be following the same approach!

CT: In your view, when will more than 50% of the cars in the U.S. be self-driving capable?

SQ: This is still going to be many years out. Humans are inherently very unpredictable. I had co-authored a blog post that explained some of the challenges, and why AVs are hard. If all cars could become self-driving cars overnight with a snap of my fingers, the problem would be a lot easier. But a world where both human-driven cars and AVs need to co-exist makes things harder.

Having said that, we're already seeing Waymo offer some driverless rides in Arizona, and there are numerous YouTube videos of awestruck passengers (and bystanders). And Cruise has stated that they will be starting driverless rides in San Francisco by the end of the year. If this momentum continues, hopefully we will see AVs being deployed in more cities in the coming years.

CT: How do self-driving vehicles stand up to electric vehicles? Which is a bigger deal?

SQ: They are both trying to solve different problems, but are somewhat related. EVs are trying to solve sustainable transportation and push us towards a world where there are no carbon emissions generated by hundreds of millions of combustion engine vehicles worldwide. The vision is for all vehicles to be electric powered, fueled by clean energy sources, such as wind and solar. And we're

already seeing all major automotive companies worldwide selling EVs, Plug-in Hybrids, etc. Here at Lyft we just announced - in collaboration with the Environmental Defense Fund - our commitment to reach 100% electric vehicles on the Lyft platform by 2030. By working with drivers to transition to electric vehicles, we have the potential to avoid tens of millions of metric tons of GHG emissions to the atmosphere and to reduce gasoline consumption by more than a billion gallons over the next decade. This includes cars in the Express Drive rental car partner program for rideshare drivers, our consumer rental car program for riders, our autonomous vehicle program, and drivers' personal cars used on the Lyft platform.

AVs are trying to solve multiple problems. Safety is a huge one. Nearly 40,000 people die in traffic accidents in the US alone. Convenience is another big problem. Imagine you could read, sleep, eat, or work while your car drives you from work to home, or safely takes your children to their little league baseball game, or even provides mobility to blind and handicapped people. And lastly, the ridesharing industry would radically change. Combining ridesharing with AVs would put a significant dent in car ownership. If you could have access to a car nearly instantly that was cheap, safe, convenient, and only you and your family were the passengers, why would you ever want to buy your own car?

The two are, however, related. AVs have a prerequisite of being able to tap into a big power source. The standard 12-volt battery in a combustion engine car is not big enough to power all the hardware than an AV requires. Hence, all current AV platforms are based on hybrids or EVs - that allow the self-driving system (SDS) to draw power from the large electric battery.

CT: What type of machine learning is most effectively employed for autonomous vehicle applications?

SQ: There are surprisingly multiple components within the SDS that leverage machine learning techniques. Perception is the part of the SDS that is responsible for determining what the AV "sees" around itself. The Perception system is designed to recognize most objects around the AV (vehicles, bicyclists, pedestrians, traffic lights, etc.). Perception is historically where most of the ML is used to train and teach the AV to recognize things.

Then there are the Prediction and Planning subsystems of the AV stack. These were not as ML-heavy in the past, but the industry is shifting to leveraging more ML for these subsystems too. Both are described briefly in the blog post I also mentioned above -- and in fact, I also talk about how we use a combination of rule-based systems, and learned systems as we build the Prediction and Planning subsystems. We have also recently talked about how we're using machine learning to try and solve Motion Planning, which will pave the way for us to leverage data-driven techniques to solve autonomy.

At a high level, we end up using various ML techniques - some

that involve large scale labeling and training, some that rely on the output of another ML subsystem, some that are purely trained in an unsupervised manner using large amounts of data, and some that are powered by the data we collect from our ridesharing network. As we train our ML algorithms, we ensure that we're hitting the metrics and results we want (and are improving on them), while making sure that the datasets we use for our training are constantly updated and refreshed to cover any failure conditions we might have seen.

CT: What is your advice to young engineers looking to learn more about vehicular autonomy. Where do they start?

SQ: Vehicular autonomy will remain an exciting industry for many years to come - and there are many challenges and problems that will need to be solved, and hence there will be many opportunities for innovation, research, and breakthroughs along the way. Some have even called AVs one of the most difficult (and unsolved) problems in Computer Science.

My advice to young engineers would be to get involved in robotics, at school or at home. Having an understanding of how robots work, along with the understanding of how machine learning works - would be the ideal combination for someone wanting to learn more about AVs. In fact, I took a few robotics classes in my undergrad - which were my favorite classes at CMU.

Although a bit expensive, LEGO has an excellent product line called LEGO MINDSTORMS that teaches young engineers how to build and program simple robots. They even sell various different sensors that can be used with their kits. Additionally, Coursera, Udemy, and

Udacity have online classes for machine learning.

CT: So when do I get my completely safe, autonomous self-driving car?

SQ: It's not a matter of when, but where you'll be able to get into a self-driving car. Currently, Lyft has partnerships with Waymo (in Phoenix, AZ) and Motional (in Las Vegas, NV) where Lyft riders can already request self-driving (with safety drivers) rides. For complete self-driving cars, what we can tell you is that this will take a long time as we incrementally get riders used to self-driving cars. So, when you think about how AV's are going to be deployed on the Lyft network, it will be first based on what rides an AV can or can't accomplish. At first, AV's will only be able to perform a small percentage of all rides. AV's won't be able to drive in inclement weather, so we're going to need drivers to be able to service the large percentage of trips that AV's aren't able to perform in the beginning.

Lyft has an important role to play to get people to try out self-driving cars. We believe the key to transitioning self-driving cars from the development stage to public offering is a rideshare network. By bringing AVs to market through rideshare, and complemented by human drivers, we ensure riders have reliable transportation options as they begin to try self-driving rides. Introducing self-driving vehicles into the mainstream via a rideshare network is a seamless way to encourage broad adoption. It will be some time before AVs are deployed commercially at scale and it's important to note that rideshare only accounts for less than 1% of all Vehicle Miles Traveled (VMT) today, so we'll still need drivers even as that number grows and even when AV's are in play.



Images courtesy of Lyft



THE PROBLEM

Getting ocean data is tough – it’s a dangerous, challenging and costly environment to operate in, yet a \$2.5 trillion industry depends on it.

THE SOLUTION

Open Ocean Robotics uses solar powered drone boats to collect, analyze and send back data in real time.

SPACE

Blue economy, the economic uses of ocean and coastal resources in a sustainable way.

THE UPSIDE

Sending a ship to sea is costly. It requires experienced crew and isn’t always possible. This really hinders the ability of ocean industries to operate, governments to protect their coastlines and researchers to understand our oceans. Our drones can go out and collect that essential data safely, affordably and in many cases better than a crewed vessel.

THE RISK

Like other new and disruptive technologies, we are changing the way things are done and regulations governing autonomous boats have not kept up. This means that there is uncertainty as to how we obtain permissions to operate in various countries and in some cases, we need to work with regulators to help inform policy that will allow this



ELEVATOR PITCH

We ask an up-and-coming venture to sell us on its vision.

► Company	Open Ocean Robotics
Home Base	BRITISH COLUMBIA, CANADA
Founded	2018
Founders	JULIE ANGUS COLIN ANGUS

industry to flourish while ensuring not only safe, but safer operations at sea.

THE PITCH

Open Ocean Robotics makes it cheaper, easier and safer to understand our oceans using our proprietary solar powered self-driving boats that travel our oceans for months at a time collecting ocean data. We help ocean researchers, offshore industries, and governments operate more effectively and safely, providing solutions for hydrographic surveying, defence & security and environmental monitoring.

THE USE CASE

Currently we are working on a project to detect and stop illegal fishing in marine protected areas. Although critical waters are increasingly being protected with this designation, there is often little or no enforcement due to the

challenges of patrolling large and often remote areas. Our drones, equipped with a suite of sensors including radar, thermal camera, high definition daytime cameras, AIS receiver and a hydrophone, will be able patrol to these areas for months at a time seeking out illegal activities and reporting them as well as collecting evidence that can be used for prosecution.

THE FUNDING

Open Ocean Robotics raised its first round in 2020, and funded much of its research and development by leveraging grants and competition awards.

THE FUN FACT

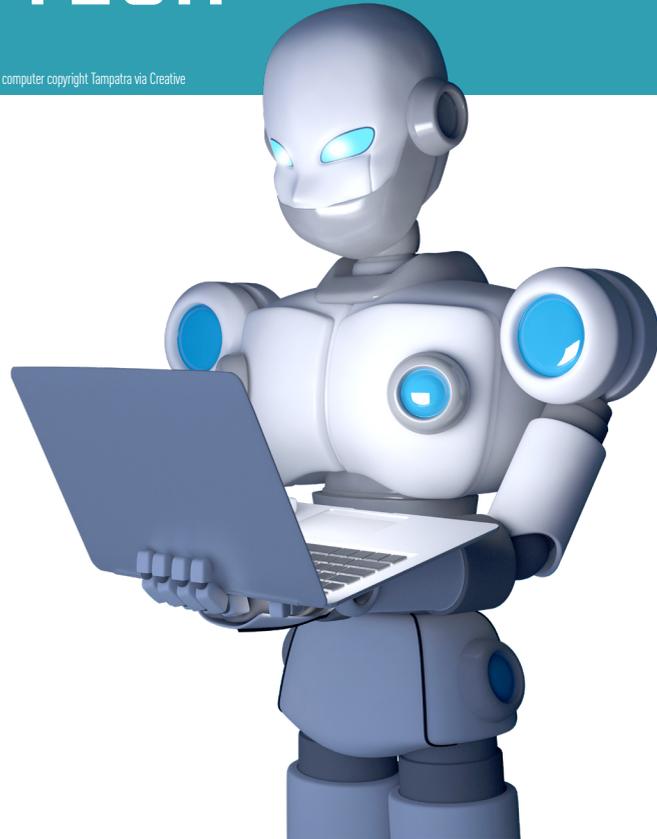
The company was founded by two people who love exploring the ocean in small boats and spent five months rowing across the Atlantic Ocean together.



Images courtesy of Open Ocean Robotics

WHAT'S HAPPENING IN TECH

Robot using a laptop, computer copyright. Tampatra via Creative



MICROSOFT AT THE FOREFRONT

> Microsoft has built an image-captioning algorithm that surpasses human accuracy in specific tests. The company utilized the AI system to update its assistant application for the visually impaired with plans to integrate it into Outlook, Word, Powerpoint, and other Microsoft products. According to researchers, the model uses VIVO that utilizes paired image-tag data to learn visual vocabulary. To benchmark performance, Microsoft entered the new AI system into the 'nocaps' challenge and currently ranks first on its leaderboard. (*"Microsoft bringing breakthrough AI image captioning to Word, PowerPoint, Outlook" - AI Daily*)

BY PRICILLA LIU

GPT-3 FOR THE MODERN BLOG

The AI-powered GPT-3 technology is engineered for content marketers looking to score well on Google with high-ranking blog posts. Users only need to enter keywords into the search bar and the system will evaluate top-performing content associated with the subject on Google. It then generates a unique idea by running the data through the GPT-3 text generator. GPT-3 can work on anything that integrates a language structure. This includes writing essays, answering queries, summarizing lengthy texts, translating languages, and even creating computer code. (*"OpenAI GPT-3: Everything You Need to Know" - Forbes*)

ANOTHER CURIOUS IBM INVENTION

Researchers from IBM and Pfizer have come up with an innovative AI model that predicts Alzheimer's at an early stage in healthy individuals better than basic clinical tests. It leverages a short language sample from a verbal cognitive test that a patient takes. The AI model can predict the onset of Alzheimer's with 71% percent accuracy, compared to standard clinical tests that are correct roughly 59% of the time. This study can result in the creation of an easily accessible metric that health-care professionals can use to evaluate the risk of Alzheimer's, making timely intervention possible. (*"5 Things to Know About IBM's AI Predictor for Alzheimer's Disease" - IBM News Room*)

NVIDIA'S BREAKTHROUGH RECORD

NVIDIA has another contribution to its name as its AI inference in MLPerf with A100 Tensor Core GPUs set records on data center, cloud, and edge benchmarks. The latest benchmarks presented four new tests that gauge performance in NLP, recommendation systems, speech recognition, medical imaging, and computer vision AI use cases. (*"NVIDIA Breaks 16 AI Performance Records in Latest MLPerf Benchmarks" - NVIDIA*)

SCORE ONE FOR EGGPLANT

Intelligent automation specialist Eggplant announced the launch of its new AI-powered software testing platform. It is designed to enhance the quality of digital experiences and the speed of deployment on any platform. Enterprises benefit from improved levels of automation and faster time to market that quickens the pace of the continuous delivery approach. Organizations can seamlessly deploy this solution on Eggplant's DAI platform that features cloud-based automation, monitoring insights, and fully-automated, self-healing test assets. (*"Eggplant Launches Cloud Platform For AI-Driven Intelligent Test Automation" - AI News*)



Reimagine industrial maintenance. Transform your operations.



Predictive maintenance programs must evolve to deliver tangible business value. By harnessing the power of AI, the SparkPredict solution generates actionable insights to protect your critical assets and improve your competitive advantage.

Discover the industry's most-deployed AI analytics solution for predictive maintenance.

LEARN MORE:
go.sparkcognition.com/sparkpredict

BY MAHE ZEHRA

THE CHALLENGES OF UNSTRUCTURED DATA

DeepNLP™ unlocks the information hidden in unstructured data.

80% of all data produced is unstructured. This includes PDFs, books, journals, audio, video, images, notes and many other media. The amount of data this translates to is beyond human comprehension. So much information that is waiting to be deciphered and put to use. So much data allowed to lie fallow.

Humans can understand this information of course, but they are not capable of handling data on the order of zettabytes. Analysis driven entirely by humans doesn't scale, adds human error and is a waste of time and resources.

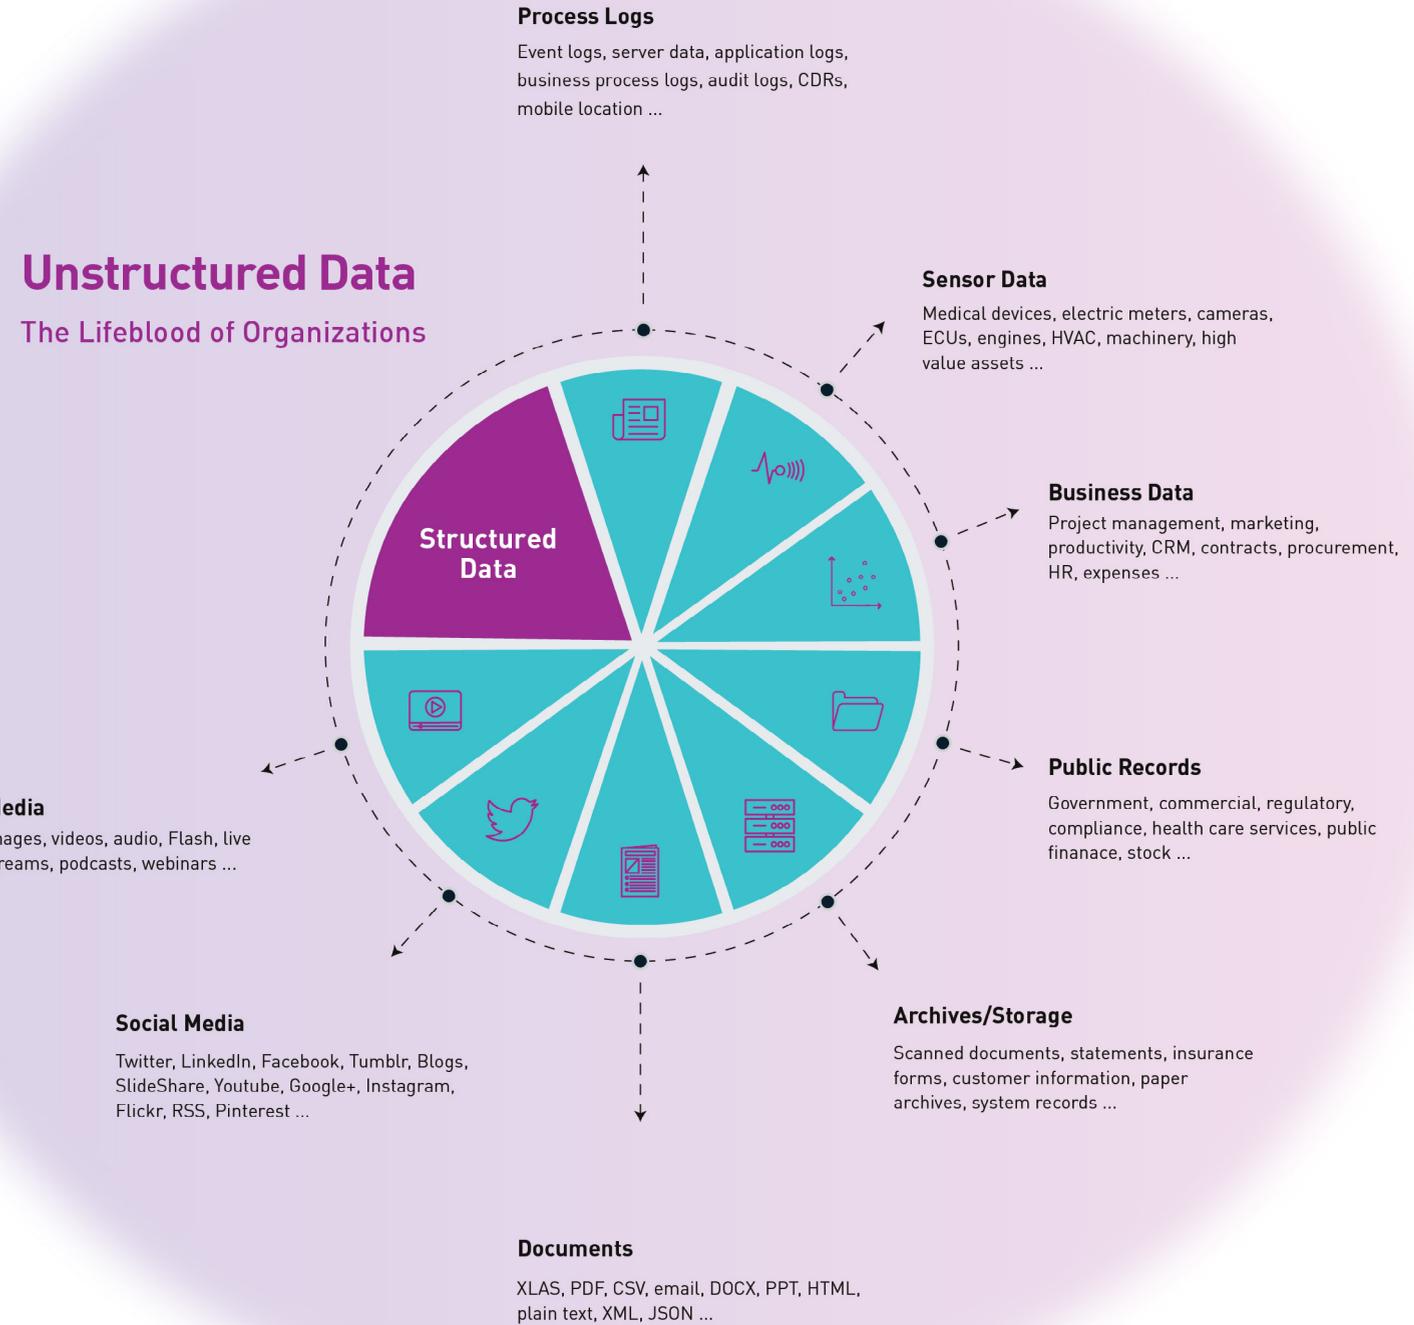
While machines aren't naturally inclined to process unstructured data, they can be taught. Natural language processing or NLP is a field dedicated to teaching machines to use and understand language in a human-like fashion.

The recent machine learning boom has revolutionized the subject, allowing it to flourish in new ways. This is because machine learning revolves around creating algorithms that can learn beyond their initial instructions. This allows programmers to feed text into a machine learning program and let it glean the rules itself, often using probabilistic and statistical models to figure out usage in a more natural way. This makes improving the model easier. Feed it more text and let it learn in a manner similar to the human brain.

SparkCognition's DeepNLP software digs deep into unstructured data using proprietary machine learning algorithms and offers insights and solutions.

Using deep learning, SparkCognition developed an advisory application for aircraft front-line maintenance staff. The application allowed technicians to conduct machine-to-human dialogue to troubleshoot asset failures and mechanical issues, with high accuracy. This solution lowered the cost of maintenance and improved asset availability by 10%.

This is just an example of what is possible when we learn to analyze unstructured data using NLP. To learn more about NLP and how it unlocks unstructured data for organizations visit: www.sparkcognition.com/product/deepnlp





Images courtesy of Ford Motor Company

MEET THE MIND BY COGNITIVE TIMES STAFF

GIL GUR ARIE

Join us as we talk to Gil Gur Arie, Chief Data and Analytics Officer (GDI&A), Ford Motor Company

CT: Gil, congratulations on joining Ford as the company's Chief Data & Analytics Officer. Please tell us a little bit about your background and the journey that got you to your present position?

GGA: Before I came to Ford, I was part of the Israeli defense force's intelligence corps where I used big data and AI for the purposes of counter terrorism and homeland security. Though to the casual obser-



/Run Program/

ver this might seem like a very different environment to Ford, I actually see a lot of resemblance from a data standpoint. In both cases the objective is the same: using data to drive value and as a building block of strategy creation. Making sense of the vast volumes and variety to accomplish known goals and reveal hidden opportunities.

CT: Given your experience, you are an expert at “sense making” from large quantities of data. How do you think this skill will be valuable when applied to the data Ford has access to? Will analytics be used to improve the customer experience? Improve manufacturing efficiency?

GGA: Every industry is now part of what we might call “the age of big data”. No matter what you do, you’re dealing with huge volumes of data, and a large variety of data signals and types. At Ford we’re looking to leverage it to achieve what we call “data superiority”. That’s the point at which we can use data to drive value not only in manufacturing and the customer experience, but also to improve dealership effectiveness, sales and marketing strategies, product development and Ford Credit. We aspire to use data to enhance the entire customer journey; from a customer interested in purchasing a vehicle, to entering the dealership, applying for credit, and improving the ongoing relationship with the dealer afterwards.

CT: Ford is one of the most iconic American brands of the past century. But in any organization with established traditions, change can be difficult to engineer. What challenges have you encountered as you seek to implement a data-first methodology?

GGA: Every other organization that I’m aware of that has been around for so long naturally finds itself with a tension between the existing ways of doing things and the new, the rapid, and the agile. Ford is now progressing with a new plan with our new CEO, Jim Farley, emphasizing the need for modernization. As a company, we’re making a non-linear leap to “modernize everywhere”. We’re pushing the boundaries of the enterprise when it comes to using data and AI, while at the same time decommissioning some of the heavy burdens we’ve carried forward until this point. For example, we can train a machine with artificial intelligence to spot the tiniest flaws on the assembly line. We also use an audio device and analytics to assure you get the right ‘click’ that means you’ve correctly inserted an important piece of cable into a vehicle connector. Setting these new capabilities on top of a legacy sys-

tem foundation can be challenging, but we have a solid approach that allows us to “productionalize” the new while rebuilding or replacing the dated components.

CT: Even after they decide to implement AI, large companies are often split on the question of whether a single, centralized analytics organization should own data efforts, or whether this function should be distributed. In the one instance you might have conformity, while in the other you get independence and speed. What approach are you taking at Ford, and why?

GGA: There is a very fine balance here, which is why Ford has decided to adopt a hybrid approach. While we have my team driving a centralized strategy, we also practice and implement a democratized platform for employees where they can leverage data themselves within their skills teams. For example, we have built what we call the Mach-1ML platform that can be used by teams to build their own machine learning solutions. In many cases, this

has already created efficiencies and brought times down from a few months to a few weeks, and in some cases a few days. We already have hundreds of different active users outside of my data and analytics centralized team.

CT: Please paint for us a picture of the data-enabled automobile of the future. How will data make the automotive experience richer and more enjoyable?

GGA: Firstly, it’s important to say that we’re no longer talking about the “automobile of the future.” The new F-150, the Mustang Mach-E, and the new Ford Bronco are really the most important vehicle launches in Ford’s recent history as they represent the next level of vehicle connectivity. Each of these vehicles will be equipped with our in-vehicle SYNC 4 technology that enables customers to make use of AI-driven technologies like smart navigation and voice recognition, while the vehicles themselves will keep customers out in front with regular bumper-to-bumper over-the-air updates. They are the cars of the future but they’re already here.

As a company, Ford has decided to place big bets on autonomous, electrification, and connected vehicle technology. For all of these, data is really the fuel to make the end goal and that means we need to have the right capabilities to understand how our vehicles are being used so that we can make the customer experience better. We also continue to expand the roster of tech companies and startups that we partner with to enhance our internal processes and connected vehicle technologies.

“As a company, we’re making a non-linear leap to ‘modernize everywhere.’ We’re pushing the boundaries of the enterprise when it comes to using data and AI...”

BY ZEHRA AKBAR

SKYGRID FLIGHT CONTROL™

Industry

Unmanned Aviation, Blockchain, AI

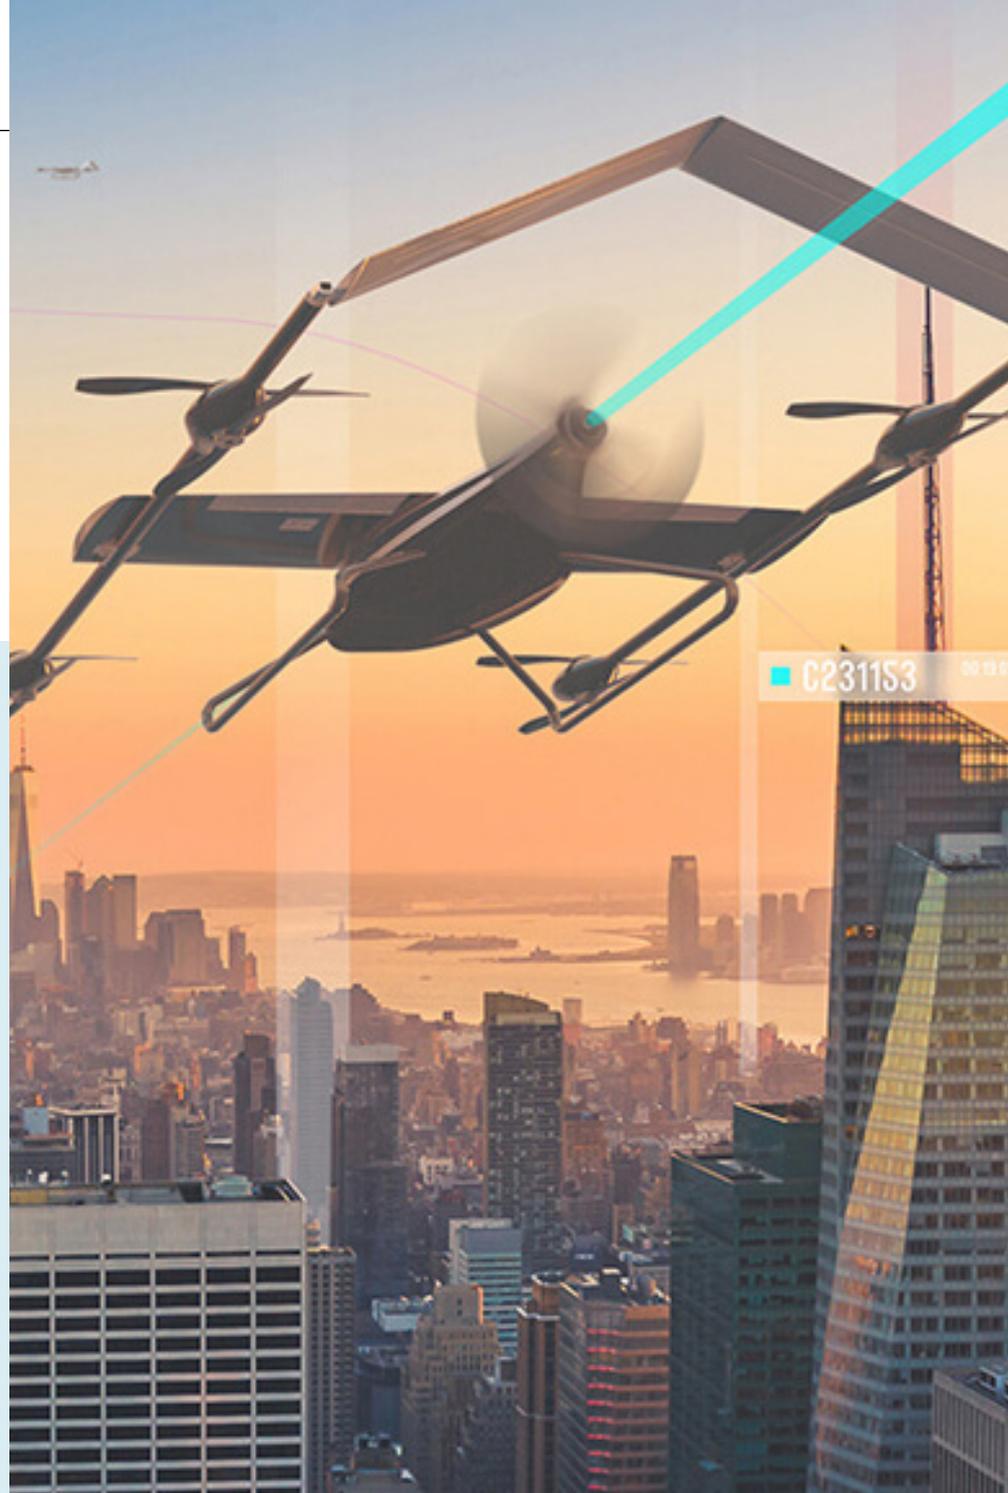
Location

Austin, Texas

As drone technology advances, its applications are evolving rapidly across the globe. Drones are delivering test kits and disinfecting outdoor surfaces to help fight COVID-19. They're assessing hurricane damage and delivering aid to affected areas. And they're surveying oil pipelines to inspect damage and prevent leaks.

From retail and logistics to healthcare and energy, drones are disrupting a wide variety of industries and upending outdated business models. Even with all this progress, drone operations remain complex and challenging. The burden falls on drone operators to manually plan, execute, and adapt their flights in the rapidly changing airspace. They're expected to monitor weather changes, avoid buildings and construction cranes, navigate around roadway traffic, and comply with shifting regulatory dynamics.

But this approach isn't scalable. As an organization's fleet grows, it's not feasible to manually plan and execute every flight that's delivering a package, conducting



Images courtesy of SkyGrid

an inspection, or responding to an emergency. Automation is critical to safely scale drone operations and enable more advanced missions.

SkyGrid is addressing this challenge with a smarter drone solution that automates every phase of flight.

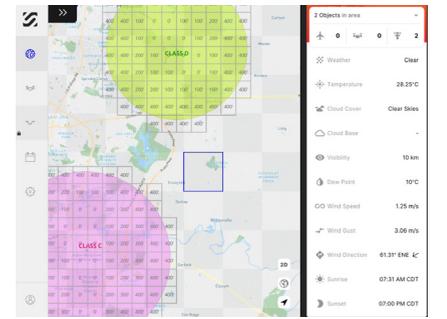
SkyGrid Flight Control™ makes it easy

for drone operators to explore the airspace, automate flights, and get real-time insights. We eliminate manual workflows by automating drone flight planning and autonomously executing the mission.

Here's how it works:

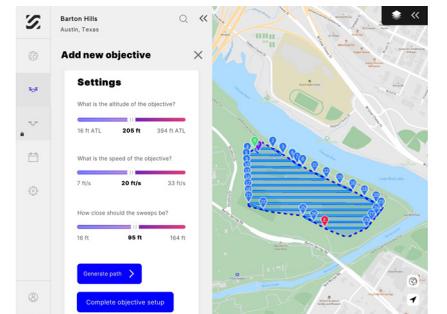
1. Check airspace conditions

SkyGrid Flight Control provides detailed airspace, regulatory, and location data to help drone operators make more informed flight decisions. Our airspace intelligence includes advanced weather data, such as precipitation, wind speed and direction, cloud cover, visibility, and more. Our regulatory data includes airspace classes and advisories from the FAA, such as UAS flight restrictions. And our location insights include local buildings, obstacles, population density, elevation, and roadway traffic.



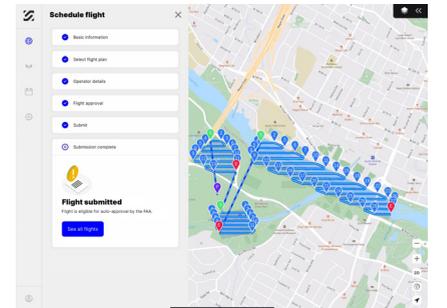
2. Generate a flight plan

SkyGrid Flight Control provides detailed airspace, regulatory, and location data to help drone operators make more informed flight decisions. Our airspace intelligence includes advanced weather data, such as precipitation, wind speed and direction, cloud cover, visibility, and more. Our regulatory data includes airspace classes and advisories from the FAA, such as UAS flight restrictions. And our location insights include local buildings, obstacles, population density, elevation, and roadway traffic.



3. Get flight authorization

As a trusted supplier of the FAA's Low Altitude Authorization and Notification Capability (LAANC), SkyGrid provides real-time flight authorization in U.S. controlled airspace under 400 feet. We simplify compliance by integrating LAANC directly into the flight planning workflow. That means drone operators can seamlessly plan their mission, ensure their flight meets all FAA requirements, and get auto-approval to fly.



4. Autonomously execute your flight

Once the flight plan is generated and authorized, operators can connect their drone to autonomously execute the pre-defined flight plan. With a Part 107 Waiver, SkyGrid Flight Control can also support multi-drone missions by optimizing the distribution of multiple drones over a defined area based on a defined timeframe or fleet size. Our live video feed allows operators to monitor each flight simultaneously and pause missions as needed.

5. Detect objects in real-time

During flight, drone operators can also detect key objects, such as vehicles, people, and animals, with our AI computer vision. This technology can be a powerful tool to quickly gain situational awareness after a traffic incident, identify potential threats near critical infrastructure, or detect defects in an oil and gas pipeline. Based on the needs of the mission, our algorithms can be trained to detect and track customized objects.

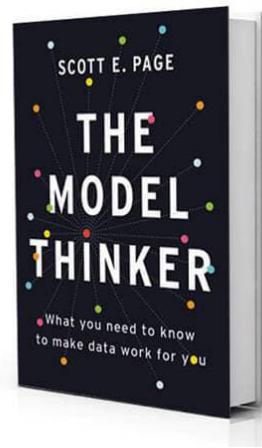
Ultimately, we believe in a future where millions of autonomous aircraft of all shapes and sizes will take to the skies to move our goods, improve our agriculture, augment our security, and even provide safe transportation. SkyGrid Flight Control is the first step toward enabling this future.

Visit skygrid.com/flight-control to learn more.



BY COGNITIVE TIMES STAFF

A TECH-SAVVY TO-DO LIST



<http://bit.ly/machinelearningCT>

▲ READ

The Model Thinker

NOW AVAILABLE

Read this dynamic book by Scott Page and learn how to apply models in various scenarios. A great, approachable book on the subject. Want to dig deeper? Try the associated course on Coursera.

👉 TRY

NOW AVAILABLE

Poem Portraits by Google is a fun interactive poem generator. You put in a word and the algorithm trained over 20 million words of 19th century poetry generates your unique poem.



POEMPORTRAITS

By Es Devlin

An experiment at the boundaries of AI and human collaboration. Donate a word to become part of an ever evolving collective poem and create your own POEMPORTRAIT.

In collaboration with Google Arts & Culture

The Dreams We Set Out to Conquer

When I saw you last, you were touching your face
When I saw you last, your shattered eyes spoke

In the years since, I have tried to live up to a name
I sought to contend with what fate sent my way

How long would it be, this parting, you wrote?
An eternity spent in purposeless aspiration?

In the years since, you have laboured, shedding tears
In the years since, I have laboured, shedding hope

Now in our twilight, memories fail us both
What were those dreams we set out to conquer?

AI Poet

One of these poems was written by a human-AI collaboration, while the other was written by a human poet. Can you tell which is which?



DON'T OVERTHINK

Go with the flow
Where zephyrs blow
Parallel the weather
Light as a feather
Set it adrift
Life is a gift
Open your heart
Play up the part
Let your mind soar
Your passions roar
Wonder if you have a hit
Stop don't overthink it
You risk paralysis
With excess analysis
Set your dreams free
You be you and I'll be me



◀ SEE

NOW AVAILABLE

Check out this amazing visual introduction to machine learning! Sometimes animations and visuals give a much better overview of a concept and we love this explanation on R2D3.

▶ GO

28-29 JANUARY

Re-Work is offering a new Deep Learning 2.0 Virtual Summit on 28-29 January 2021. The tracks offered are timely and compelling. From ethics and social responsibility to enterprise AI and deep learning. We have this marked on our calendars!

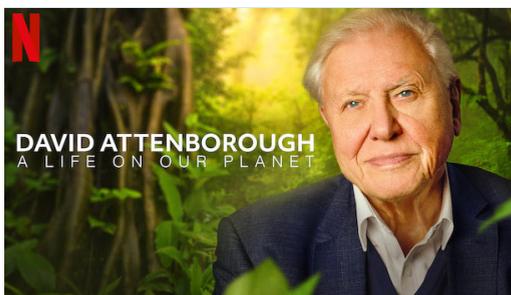


www.re-work.co

◀ SEE

NOW AVAILABLE

Sir David Attenborough's *A Life on Our Planet* is a beautiful documentary about Earth and the future we are heading towards. A must watch in our opinion, and easily streamable on Netflix.



BY MAHE ZEHRA

The Jetsons and NLP

Growing up I used to love watching “The Jetsons” on Cartoon Network. Funnily enough, “The Jetsons” used to come on right after “The Flintstones” and we would experience quite the time leap between the two shows.

My favorite parts of the entire cartoon were the flying cars and Rosey the robot maid. Even though she was an old demonstrator model hired by the Jetson family from U-Rent a Maid, I still loved her. The concept of having a robot that understood human speech and could listen to and then carry out commands really excited me.

Fast forward quite a few years and I am living in the age of AI and all the things I used to dream of as a child are slowly coming true all around me.

We are already experiencing the first self-driving cars and I don’t think the ‘flying’ bit will take too long to arrive. Hopefully within this decade. We are now surrounded by multiple Siris and Alexas. All we need to say is “Alexa, play The Good Life” and the song comes on. The hardware doesn’t look like Rosie but we can definitely get some tasks done pretty easily just through voice commands. Now I’m waiting for the implementation of some mechanical

capabilities like cleaning and cooking!

Some of the advancements in the field of machine learning and AI in general have been made possible because of big data, which is getting bigger since it is now much easier to capture and share information. A modern car has multiple sensors that share data, and your phone apps have application logs and take pictures and videos too.

All of this data is being produced at such a rapid rate that it isn’t possible for human beings to make sense of it. Thankfully with the advent of big data we also have access to advanced and powerful computing and enhanced algorithms. So we are able to begin to tackle the problem of understanding all this data.

In the case of human to machine interaction and understanding of commands, we use natural language processing.

Natural language processing is a branch of AI that helps machines make sense of human language. This is done through various algorithms based on statistical and machine learning techniques. As with most AI fields, a wide variety of professionals from varied disciplines like computer science and computational linguistics come together to take this science further.

NLP is not a new field but the technology is evolving rapidly because of the availability of large amounts of data. Machine learning algorithms are only as good as the quality and quantity of data that we feed. This increase in text data input has a direct impact on the quality of the output from the algorithm.

But as with all cutting edge ideas NLP comes with quite a plethora of challenges as well. One of the biggest challenges is that text is unstructured data.

Unstructured data is data that doesn’t follow a data model or is not organized in a pre-determined way. Excel files are a perfect example of structured data. Everything is organized in a tabular format, labelled, and easily fed into algorithms and models. In contrast most of the data found in the world around us is unstructured and much of this unstructured data is text-heavy.

It could be the text from a newspaper article, or audio books, or transcripts from TV channels – none of this has an inherent structure to it. In addition, with language and text data it is much harder to create algorithms that will ‘understand’. But slowly and steadily progress is being made and with the advent of GPT-3¹ hopes are high that we are finally coming closer to our sci-fi vision of human – machine communication.

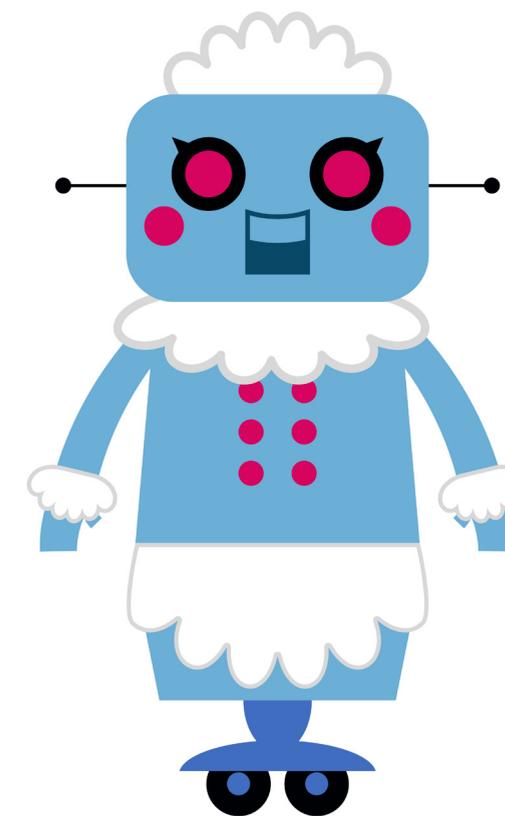
My dream of having a ‘Rosey’ of my own is looking more achievable with each passing day! 🤖

With NLP the goal is to take raw language input and use linguistics and machine learning algorithms to turn it into a form that offers greater value. Some of these value adding applications include:

- Content categorization.
- Topic discovery and modeling.
- Sentiment analysis.
- Speech-to-text and text-to-speech conversion.
- Creating relevant summaries of documents.
- Machine translation of text or speech from one language to another.



Future Family Jetsons copyright Cozy Cat Design via Creative Market



¹ To learn more about GPT-3 check out Prof. Bruce Porter’s in-depth article about this innovation featured in this issue of Cognitive Times”. [Page no. 34]



BY AMIR HUSAIN AND STEVE NORDLUND

Watching her grandson stare out the air taxi's window at the rolling green hills 3,000 feet below, Marie was surprised by the moment in suddenly realizing she saw something he simply couldn't see. More than seeing it, she felt it as she scanned the Texas landscape looking for ghosts. Anger. Distance. Dislocation. He knew nothing of how bad it had been, and by extension, how good he had it now. How good they all had it.

"Beautiful, isn't it?" she said.

"Sure, grandma," he said. Then he caught himself, remembering that it was her birthday. 80 years old. A milestone worth celebrating any way she wanted to. He leaned into her and took a deep breath. "It really is nice. Check out my overlays." He pushed a rich blue body of water framed by snow-covered peaks that looked like they came from the Italian Alps. It resembled nothing like west Texas.

"Are you hungry? I brought sandwiches for us."

She reached into her backpack, a sun-faded orange nylon day-pack that was older than her 11-year-old grandson.

"Sure, what do you have?" said William.

"One of my favorites," she said, handing him a peanut-butter-and-jelly sandwich. "I made the bread last night. Joe, my neighbor, made the peanut butter. The blueberry jam is from last winter."

She baked almost every day, like most people. It filled the hours. He looked at her like she was crazy.

"I canned it last winter," she said. "It's not like it's been sitting on my counter for five months."

He raised his eyebrows and smiled as he bit in.

Then a steely male voice came over the air taxi's speakers. "Good afternoon, William and Marie. This is Bertie, your autonomous aviator on this trip. We will be making a small change to our flight plan due to a weather front moving in out of the northeast. Shouldn't add more than a few extra minutes and the journey will be a lot smoother. Thank you." An onboard AI flight computer,

which manifested to the passengers on this trip as the pilot-like authority of Bertie, interfaced with the global SkyGrid AI air-traffic management system that guided their flight.

The air taxi banked to the left, and began to climb. As it did, the two of them had a sweeping view of hundreds of miles of land, the seam between old and new. There was scrubland and verdant agricultural settlements teeming with swarms of silvery agri-bots that squirmed about like droplets on a windshield.

The taxi was a long-endurance two-seat model, a Boeing GH-2. It had an elongated, bulbous tinted canopy that offered a 270-degree view, giving passengers a feeling of sitting forward in the aircraft as fighter pilots once did in manned combat aircraft of yesteryear. If a traveler wanted to plug their AR glasses into the aircraft's own sensors, they could do that too. The sound canceling software almost entirely eliminated the intrusion of wind noise, though the GH-2's 12 electric motors and propellers hummed in a harmonious performance that was almost symphonic. The motors were distributed along the x-like wings that took a form many compared to a cat crouching down to stretch its back. This model's composite skin was a bright blue, like a summer sky.

William chewed his sandwich, then paused. "It's good," he said. "Maybe you can make my lunches?"

"How's school?" When she asked a short question like that she drew out the second half of the words, as if to emphasize the true meaning of what she was really seeking to understand.

"Soccer's great," he said. "Last weekend we played at the Southeastern Regionals in Pascagoula, but we lost our first game so we flew back that afternoon."

"I know, I simmed your game. You're a great goalie, you know. But I mean like friends, grades."

"All that's fine," William said. "Computational biology is getting hard. But mom helps me because she's studying it too."

Everybody's learning, Marie thought. All the time. Which, as a retired high-school environmental sciences teacher who remem-

bered a not-so-long ago time when people felt like they had to learn just to get a good job, she approved of. She was glad her family did too.

"How's the Scouts?" she asked.

"Oh, it's great. Most Sunday afternoons we're meeting up with other Troops from all over the state. Next weekend we're going to be working on making fires without matches at a campsite in Lubbock."

"A valuable skill," she said. Then she laughed, and he joined in.

They basked in the connection they had and Marie pointed out a dark line of steel grey clouds on the horizon. "Glad we're flying away from that," she said.

William got that glassy-eyed look that kids got as they dived back in to their augmented visual worlds. Marie wanted to pull him back, to her, to the chance to be together and just talk.

"I never went to Lubbock before we all started flying," said Marie. "Too far from Austin."

"It would be a long drive," he said.

"It wasn't just that. It used to be far in a lot of ways. We just didn't get up there. Your great grandparents definitely took me places. Like we went to San Francisco a couple of times when my father had a big conference there. I went to college in Boston, remember. But there were plenty of towns pretty close to home that felt a lot farther than all that."

"Why?"

"First, Texas is a big place. Back then people got way too comfortable just being with people like themselves, whether online or in real life. They spent so much time thinking about what made them different and being angry about it when somebody wasn't like them, rather than what we all can find in common with each other. A lot of folks acted like they lived in their own world, isolated, when they were tied in to the entire planet. But gradually things began to change. Once you're flying every day, Texas doesn't seem like such a big place anymore. It's like when you know somebody somewhere and it's a lot more fun to visit them, rather than when you go somewhere and you feel like a stranger the whole time? Say you go to another country, they might speak a different language, eat different food, go to a different church or no church at all. But it doesn't feel as far away when you get to spend time with another person. You're excited to see them. Could be doing whatever. Drinking coffee. Eating dinner. Going for a walk. A festival. It becomes about being together, not apart."

"Hmmm," he said. "I know what you mean. I think that's why Friendship Week is my favorite module. There's less homework, for one. And you get to start school a little later because the flight to their school takes longer than yours would. And you get to see how somebody else lives, meet their teachers, play on their soccer field. I just wish we could go back to the same place each year."

"That would make sense, wouldn't it?" Marie said. "Having taught for 28 years, I know better than anybody that sometimes



things don't make sense to the kids, but some adults did put a lot of thought into it."

"I do stay in touch with a lot of my Friendship Week friends," he said.

"Doing sims together?" she asked.

"Is there any other way? Doesn't matter where you live. China. Texas. Canada. Los Angeles."

"You know kids from all over."

"Yeah," he said.

William put a hand up to his AR glasses, an unconscious gesture of kids feeling antsy when it had been too long since they logged in. It got shorter and shorter every year. "If you don't mind, I'm going to close my eyes for a minute," she said.

"Sure thing. I'm supposed to meet up with Riz."

"Where is he from?"

"Syria, a camp near Damascus. We play in a Tier 3 GlobalGol league." William paused. "He's an orphan, so we're like his family."

She wondered when Riz would get to meet his friends in real life. The United Nations last year debated enshrining global mobility as a universal human right, but making it happen would still take a lot of work. The advancement of flight beyond seamless air taxis safely integrated into the airspace was just the start, but supersonic commercial travel replaced by hypersonic travel increasing the speed by five times has made our planet so much smaller. It has driven bodies like the UN to find common ground, and arguably has made the world safer. Perhaps William and friends could now make that happen, she wondered. A conversation for a later time. She really was tired.

"Maybe you can wake me up when we get to the hub and meet everyone else?"

He nodded his head, which she took to mean that he actually heard her. Maybe he did, maybe he didn't. With kids his age, you never knew. As she adjusted the seat to make it more comfortable for her nap, she let out a pleased sigh.

About an hour later, the air taxi noticeably slowed with a proud flaring of its nose and a lower-pitched tone from its engines.

The shift in speed woke Marie. She opened her eyes, grateful for the canopy's light-filtering features. She stretched her aching neck and peered through the glass.

About a half mile away she saw the hub. It was what people called a clover, a configuration of four bright shark-like white air-hub dirigibles joined at the nose to create a cross-shaped formation. With solar panels arrayed like scales along their upper halves, they looked like basking whales catching up on an ocean's worth of gossip. But it also allowed the SkyGrid AI to jointly manage the engines on all the dirigibles to keep the flying charging-station and traveler outpost positioned correctly for everything from wind direction to anticipated traffic flows.

A four-seat Phoenix, an organic-looking Italian design from the China Mobility Company, pulled in right in front of them. The sudden appearance of the dark-green aircraft surprised her; flying in tight formation was something she was pretty sure she would never get over. But William did not even flinch.

"This is Bertie, and we're approaching the end of this leg of your journey. We're currently right at the New Mexico and Arizona border. I expect we'll be docking in about six minutes at Kokopelli Station, on pad Echo-Niner. From there you'll follow your prompts to join up with the rest of your party. I'll have the privilege of piloting you on the next leg of your journey. Please ensure your seatbelts are fastened and any personal items are secured as we'll be leaving this vehicle behind."

"Thank you, Bertie," William said.

"Been a pleasure. Stand by for landing," the onboard AI responded. "See you shortly."

Marie blinked twice, and the blue arrow reappeared on her glasses. There, maybe 20 or 30 steps away, was her next pad, and the aircraft that would bring them directly to the Grand Canyon itself. She looked back to make sure William was behind her. Being a mile up, she couldn't help but worry he might fall all the way to the ground if she lost sight of him. The thought was, rationally, ridiculous as they were completely enclosed inside the translucent fabric-like tubes that threaded through the passenger-connection zones between the airships and their hangars. But it was a very real feeling all the same.

A message popped up and she smiled: It was her daughter, Luisa, letting her know the rest of the family was already belted in. A moment later, Bertie messaged to say they were ready to take off as soon as she and William arrived; it was the kind of courtesy from the onboard AIs that she appreciated, anticipating that a family member would notify her first rather than the aircraft's digital personality.

She threaded between the travelers, who moved not with the urgency of frustration but something more like eagerness. By their clothes, they seemed to come from a variety of economic positions in American society. By their ages, it was like being at the biggest family reunion you could think of. Ethnicity or skin color was wonderfully diverse, she thought, and hinted at a richness of expectations about what life held in the years ahead, the sort of optimism that bonded people together, rather than divided them. Considering it another way, she thought if the world were just to end right now and these were all the people who were left, humanity might actually be able to make it out of whatever misfortune appeared. ☐

To be continued...



ALL-ELECTRIC. ALL-ADRENALINE.
THE FORD MUSTANG MACH-E



ORDER NOW AT FORD.COM

Mustang Mach-E GT model shown available for reservation only. Available late summer 2021. Mustang Mach-E Select, California Route 1, Premium models available late 2020. Terms and conditions apply. Preproduction computer-generated image shown.

BUILT *Ford* PROUD

BY MAHE ZEHRA

5 Ways You Can Interact With GPT-3

We are living in such exciting times! There are so many new scientific breakthroughs that are happening all around us all over the world. One of the hottest fields leading the way is AI. Scientists are writing faster and smarter algorithms; machines are learning and more accurate predictions are being made.

NLP plays a huge part in this endeavor. In the ideal scenario we would want to talk to a machine and have it understand and execute commands. All of this is made possible through NLP breakthroughs, the latest of which is GPT-3.

GPT-3 stands for Generative Pre-trained Transformer 3 and it is an auto-regressive model that uses machine learning to produce text. The model was trained on 175 billion machine learning para-

eters. Compare this to GPT-2's 17 billion parameters.

What makes GPT-3 so amazing is the quality of the text it generates. It is so close to what a human would create that telling the difference can be very difficult.

All of this sounds wonderful but the question that comes first to my mind, and I am sure yours too, is what can it do right now? So for this issue of *Cognitive Times* we decided to show you some things that GPT-3 can do and the fun you can have with it. Some of these applications are mind-blowing. If language models continue to make such progress, we can be sure to witness even more exciting developments in future .

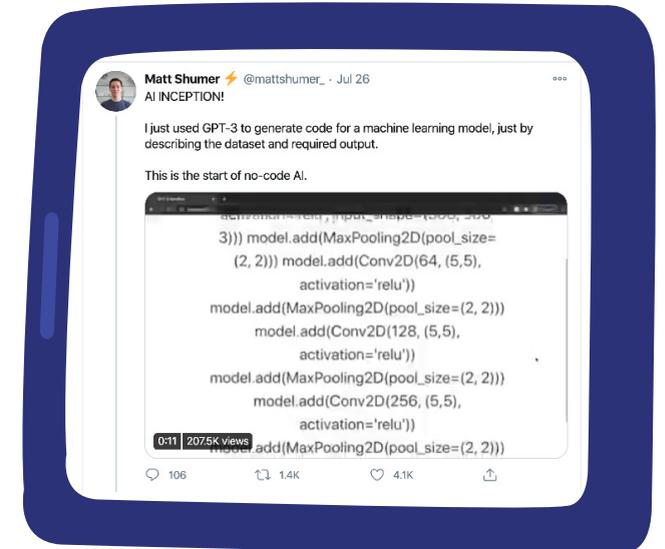


Natural language processing concept, copyright Visual Generation via Creative Market

1. Sharif Shameem is using GPT-3 to build text to web applications in React. That is amazing as it means we are making strides towards text and speech based programming.



2. Matt Shumer is using GPT-3 to generate code for a Keras machine learning model. He just inputs a couple of lines of text to generate an impressive result.

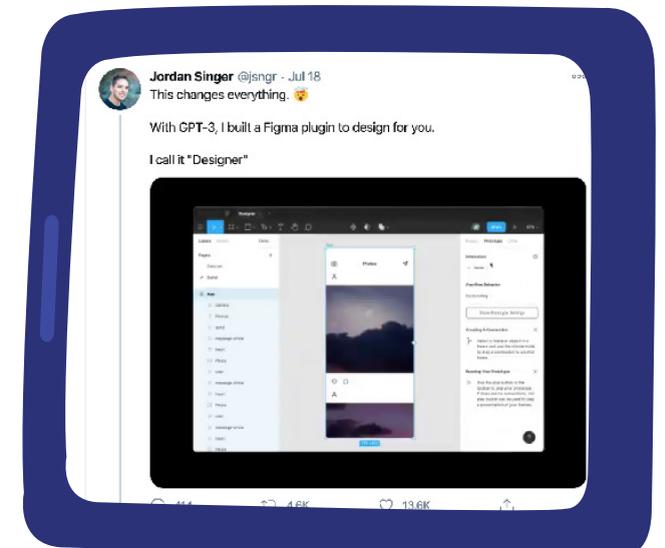


3. Then there are some exciting creative projects!

4. For those of us trying to use LaTeX to generate the math for our papers, Shreya Shankar has created a text to LaTeX converter using GPT-3. I know this is one application many mathematicians and scientists will be grateful for.



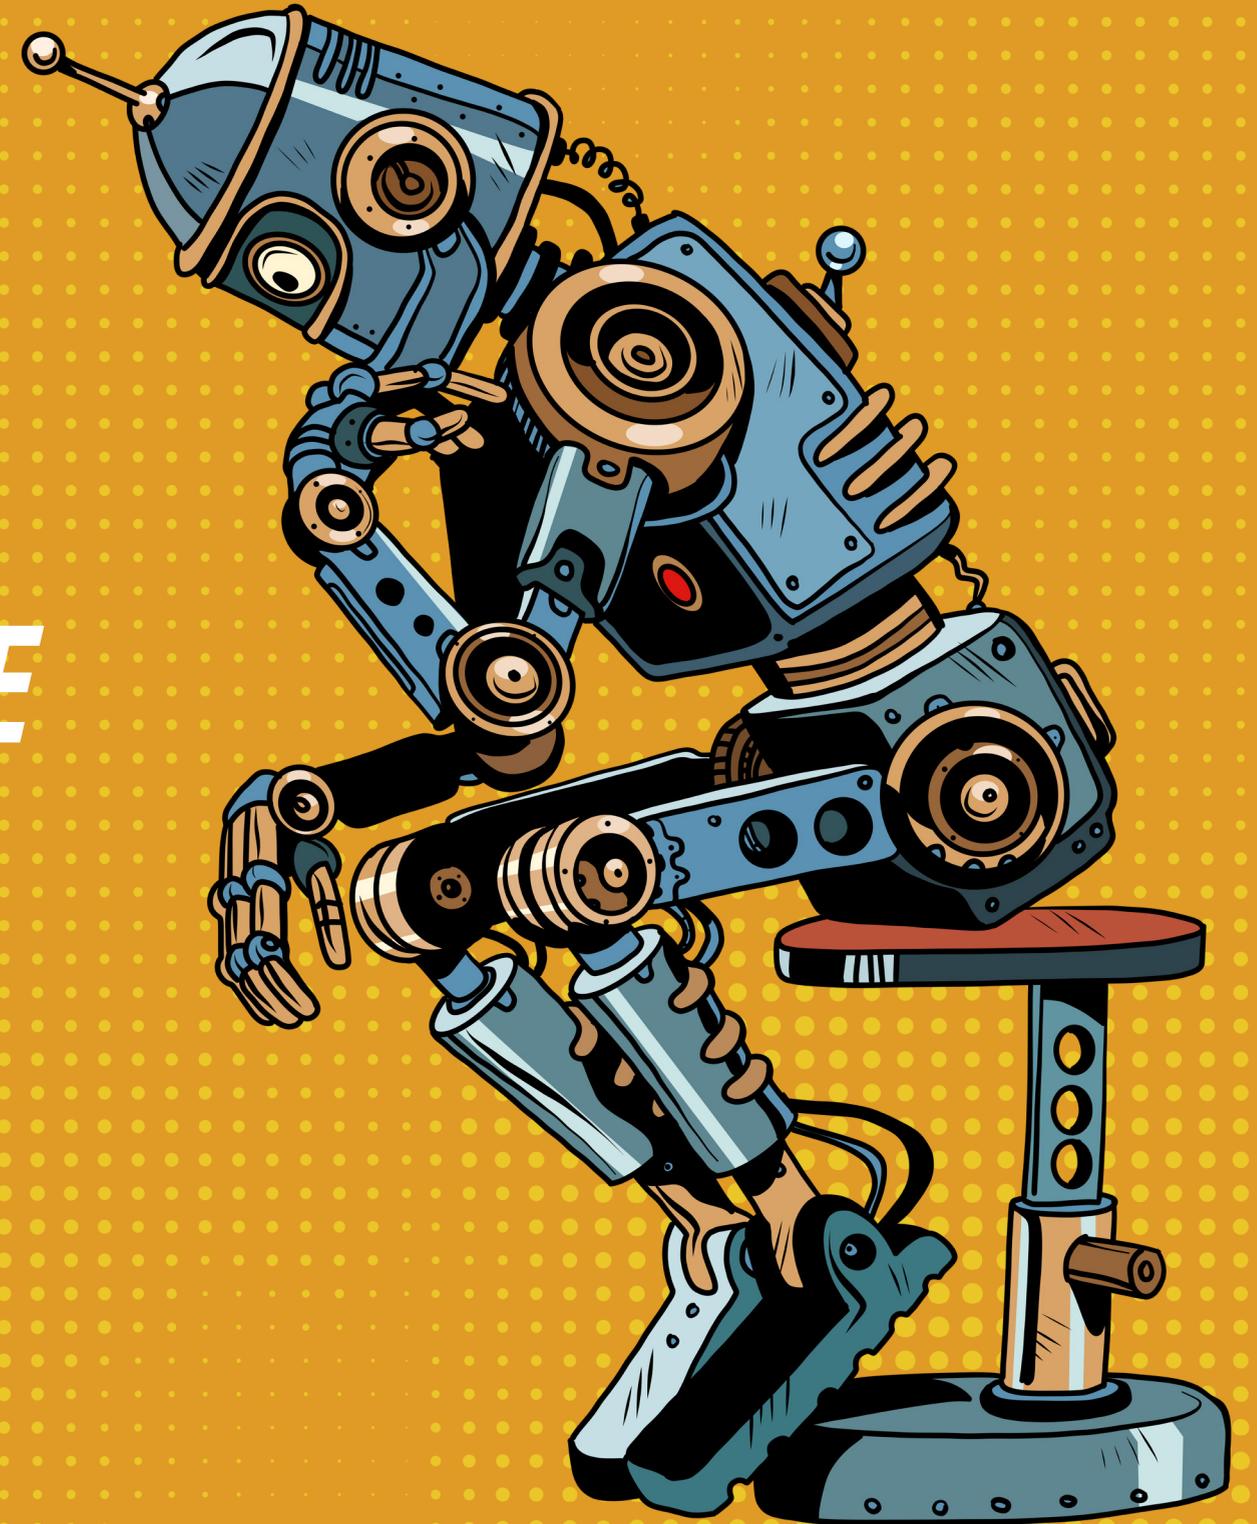
5. Jordan Singer has created 'Designer', a Figma plugin that will convert text to UI in seconds.



All of these applications are already very 'futuristic' but they are just simple examples and a hint of what is to come in the future. The age of AI has much more in store for us over the next few years! What do you predict will be possible by 2025?

BY PROF. BRUCE PORTER

WHAT IS GPT-3, AND WHAT DOES IT MEAN FOR NATURAL LANGUAGE PROCESSING?



A language model is a probability distribution over a sequence of words—which is to say, it gives the likelihood that a particular sequence of words will appear in natural text or speech. An accurate language model can be used for a wide variety of NLP tasks, including (most directly) story generation and (more surprisingly) question answering, arithmetic, and computer programming.

GPT-3 is the latest in a series of increasingly capable language models for natural language processing (NLP). GPT-3 is a deep neural network—specifically, a Generative Pretrained Transformer. It contains 175 billion parameters trained on the Common Crawl dataset, constituting nearly a trillion words. GPT-3 was created by OpenAI in May 2020 and published here (<https://arxiv.org/pdf/2005.14165.pdf>). It has since inspired a great deal of buzz—but how does it actually perform, and what does that mean for further progress in the field?

What's Novel about GPT-3?

With 175 billion parameters, GPT-3 is two orders of magnitude larger than its direct predecessor, GPT-2, which has 1.5 billion parameters. It is also one order of magnitude larger than Microsoft's language model, Turing NLG, which was released in February 2020.

More remarkably, GPT-3 also provides a much simpler way of applying the model to NLP tasks. Previous language models were applied to an NLP task using a traditional fine-tuning approach. With this method, the language model is first pre-trained for the generic task of estimating the probability of word sequences, then it is fine-tuned for each specific NLP task through the presentation of a large corpus of training examples for the task. The language model provides an initial task model via transfer learning, then the task model is tuned using traditional machine learning. Consequently, this approach suffers from the two main problems machine learning generally struggles with: amassing training examples is difficult and costly; and the learned models are brittle in that they overfit the training examples and don't generalize well to new tasks, even ones that are very similar.

GPT-3 removes the need for traditional fine-tuning of models for each NLP task. It can be used directly for a wide variety of tasks, given only a small amount of guidance on the task. GPT-3 was evaluated under three levels of guidance: “few-shot learning,” in which the user provides a small number of demonstrations of the task (typically 10 to 100, which is very small compared with the size of traditional datasets for training neural networks); “one-shot learning,” in which the user provides only one demonstration; and “zero-shot learning,” in which the user provides only a brief, natural language description of the task rather than demonstrating it.

For example, GPT-3 has been used for the NLP task of machine

translation between English and French. Under the “few-shot learning” condition, it is presented with a handful of English passages and their French equivalents. Under the “zero-shot learning” condition, it is presented with only the instruction: “Translate English to French.”

To the extent that GPT-3 performs well on NLP tasks under any of these training regimens, it directly addresses the two main problems confronting machine learning.

Does GPT-3 Work?

In many ways, GPT-3 represents an important new breakthrough in NLP, but its performance in testing has been mixed. GPT-3's successes in a variety of tasks, as detailed below, demonstrates that size alone can account for some significant advances in NLP capabilities. At the same time, these advances do not apply to all tasks across the board.

Through the construction of a series of transformer-based language models, ranging in size from ELMo's 100 million parameters in 2018 to GPT-3's 175 billion parameters in 2020, performance on NLP tasks has steadily improved. Arguably, task performance is a direct function of the size of the language model. GPT-3's performance on NLP tasks seems to follow that trend line, but head-to-head comparisons have not yet been conducted. Instead, GPT-3's evaluation so far has focused on the qualitative shift in training regimens described above.

GPT-3 has been evaluated on over two dozen NLP tasks under “human-like” training conditions to test its generality. The tasks range from ones that use the language model directly (such as sentence completion) to tasks that use the model indirectly, and sometimes in mysterious ways, such as solving arithmetic problems. In most of these comparisons, GPT-3's competition is a fine-tuned model, which might perform better than GPT-3 but suffers the traditional problems of machine learning, as described above. The complete results on many NLP tasks are reported elsewhere, and some of the more notable results are summarized below.

Direct Tests of Language Modeling

The LAMBADA test requires models to predict the last word of paragraph-length stories. In the zero-shot setting, GPT-3 achieves 76% accuracy, a gain of eight points on the previous state of the art.

The HellaSwag test involves picking the best ending to a story. In the one-shot setting GPT-3 achieves 78% accuracy. This fell short of the current state of the art, the ALUM model, which is fine-tuned to achieve 85%.

The StoryCloze test involves selecting the sentence that best completes a five-sentence story out of multiple choices. GPT-3 achieves 83% in the zero-shot setting and 88% in the few-shot

“ WITH 175 BILLION PARAMETERS, GPT-3 IS TWO ORDERS OF MAGNITUDE LARGER THAN ITS DIRECT PREDECESSOR, GPT-2, WHICH HAS 1.5 BILLION PARAMETERS. IT IS ALSO ONE ORDER OF MAGNITUDE LARGER THAN MICROSOFT'S LANGUAGE MODEL, TURING NLG, WHICH WAS RELEASED IN FEBRUARY 2020.

setting. This is four points lower than the best fine-tuned model, but it improves over previous zero-shot results by about 10 points.

Question-Answering

The question-answering task has traditionally been approached by first using an information retrieval (IR) system to find relevant passages of text in a corpus, then using a trained model to generate an answer from those passages. This approach is called “open book” question answering. GPT-3 was tested on the harder “closed book” task, which does not have the benefit of an IR system to reduce the search space. It was tested on three question-answering tasks.

When tested using the TriviaQA reading comprehension dataset, GPT-3 achieves 71% in the few-shot setting. This beats

the state-of-the-art system, which was fine-tuned for the task, by seven points.

On a harder task, the WebQuestions benchmark for question answering, GPT-3 achieves 41% in the few-shot setting. This is comparable to the best systems, which were all fine-tuned for the task. And on an even harder task, the Natural Questions dataset, GPT-3 achieves 30% in the few-shot setting, underperforming the best fine-tuned model by seven points.

Translation

GPT-3 was trained on texts drawn from English (93% by word count) and other languages (7%).

In the zero-shot setting, GPT-3 performs poorly. In the one-shot setting it is nearly competitive. In the few-shot setting, GPT-3 improves to match the best of the fine-tuned unsupervised models.

Common Sense Reasoning

These tasks require physical or scientific reasoning, as distinct from sentence completion, reading comprehension, or broad-knowledge question answering. The PhysicalQA dataset asks common-sense questions about the physical world, such as: “To apply eyeshadow without a brush, should I use a cotton swab or a toothpick?” GPT-3 tops the leaderboard with a few-shot performance of 83%, as compared with human performance of 95%.

The Abstraction and Reasoning Challenge (ARC/Challenge) involves multiple-choice questions taken from 3rd to 9th grade science exams. GPT-3 achieves 51% in the few-shot setting, which significantly underperforms the state-of-the-art, fine-tuned models by 27 points.

Finally, the OpenBookQA dataset requires multi-step reasoning. GPT-3 improves

significantly from zero- to few-shot settings but is still over 20 points short of the best fine-tuned model.

Arithmetic

GPT-3 was also tested on addition, subtraction, and multiplication problems. For addition and subtraction, the problems ranged from single digit (e.g. $3+5=?$) to five digits (e.g. $46371-35790=?$). For multiplication, only two-digit problems were presented.

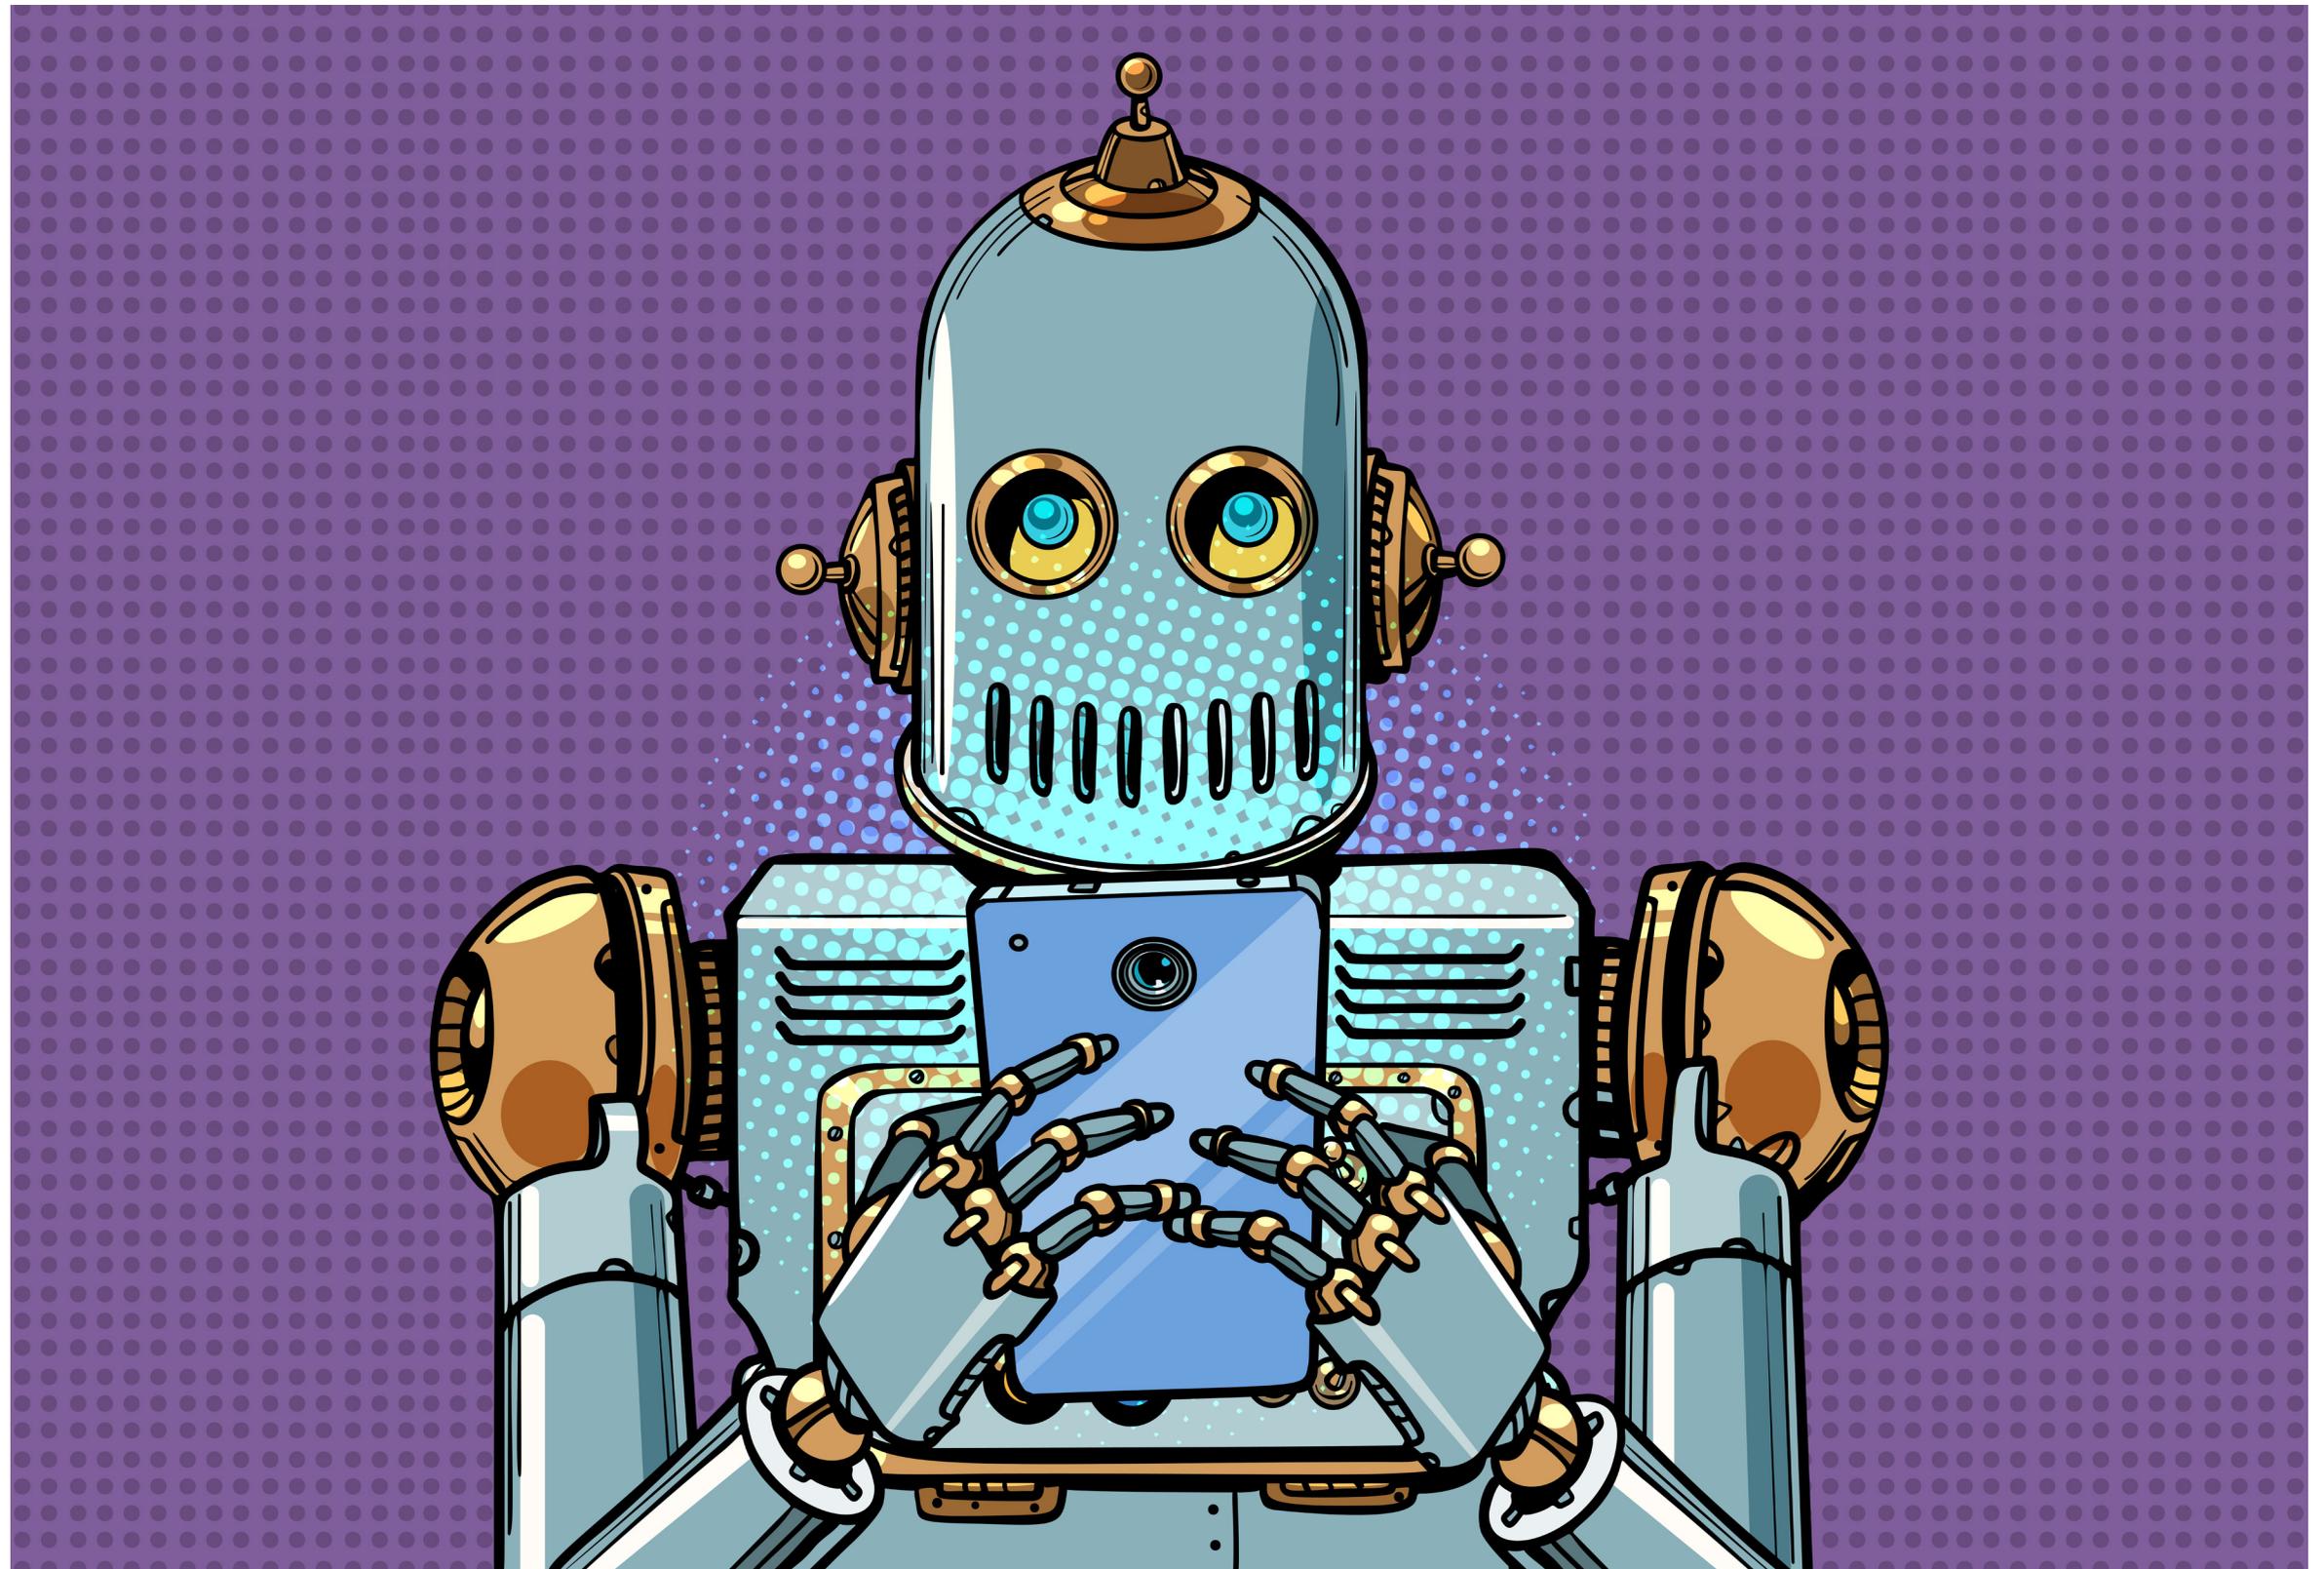
GPT-3 performed well on small problems: it achieved 100% accuracy on two-digit addition, 99% at two-digit subtraction, 80% at three-digit addition, and 94% at three-digit subtraction. Its performance declined on larger problems: 25% accuracy on four-digit operations and nine percent accuracy on five-digit operations, suggesting at least some capacity to generalize to larger numbers of digits. It achieved 29% accuracy at two-digit multiplication. Finally, it achieved 21% accuracy at single-digit combined operations (e.g. $9*(7+5)$).

Some initial tests were conducted to confirm that GPT-3 was not relying solely on a memory of arithmetic facts. The training corpus was searched for all of the three-digit arithmetic problems in the test set. Out of 2,000 addition problems, only 17 (0.8%) were found to come directly (verbatim) from the training corpus; and out of 2,000 subtraction problems, only two (0.1%) were found in the corpus.

Does GPT-3 Usher in Natural Language Understanding and General AI?

Is that enough to enable natural language understanding, rather than just processing, or even general AI? It seems doubtful.

GPT-3's lackluster performance on tasks that require even simple forms of common sense inference and reasoning, as described above, is telling. Nevertheless, GPT-3's "pseudo-understanding" of textual information can improve current NLP systems, and empower future ones in ways that have not even been conceived. 



Robot looks at the smartphone, copyright Studiostoks via Creative Market



THREE COMING SHIFTS IN AI

BY AMIR HUSAIN

THREE COMING SHIFTS IN AI

Nearly every new day brings exciting news in the field of artificial intelligence. But what larger directional trends do these news items drive? Beyond the announcements and the hype, is AI really evolving?

In this article, I'd like to focus not on far-off, vague hopes and wishes about AI, but instead on a few concrete developments that lie in the not-so-distant future. The trends outlined below are already beginning to materialize in the form of real-world research and applications. These areas of work represent themes that I believe will be recorded as meaningful breakthroughs in a future timeline of key AI developments.

1. The Cost of Training Machine Learning Systems Will Be Drastically Lowered

The AI community has long understood that our most successful methods, such as deep learning, are mathematically and computationally complex. Using these methods with current techniques involves the consumption of significant amounts of processing resources and a high degree of cost. This, in turn, limits where and how liberally these techniques can be applied. But less expensive training techniques are coming to the rescue. It is important to understand that the cost savings and speed increases ordinarily

driven by Moore's Law are entirely separate from the types of algorithmic efficiencies discussed below. While Moore's Law promises to double transistor counts—and hence processing power—every two to three years, algorithmic breakthroughs can often lead to instant efficiencies that are orders of magnitude more significant. What could some of these breakthroughs look like?

One promising area is the development of lightweight neural networks, which are exactly what they sound like: smaller, quicker to train networks that can achieve almost the same accuracy as their much larger cousins, but at a fraction of the size and training cost. Research led by SparkCognition's Chief Scientist, Dr. Bruce Porter, revealed that in cybersecurity applications, lightweight neural networks could match conventional deep networks despite using a mere tenth of the computational capacity required for those larger structures.

Another emerging technique, sparse learning, holds the potential of speeding up the training of deep networks by 3.5 to 12 times. Separately, researchers at North Carolina State University have shown that by taking advantage of the similarities in sub-segments of images in a training dataset, they can potentially reduce the amount of time and com-

pute power required for network training by more than 60%... all without sacrificing accuracy.

This is by no means a complete list of the considerable volume of work dedicated to reducing the computational burden of training deep networks. But even this small glimpse shows that the near future will likely bring far more efficient methods that will allow us to rival the decision-making power of large, complex networks without expending the same computational resources or compromising accuracy.

So, what are the implications of these coming advancements? Here are at least a few:

1. With a significantly reduced computational burden, networks will become easier to deploy at the edge, in embedded systems and in environments where access to cloud resources is not possible. This will allow machine learning and AI to proliferate into (smaller and smaller) objects and "things" in the real world.

2. Since AI running at the edge means that the in-field systems it enables can be self-sufficient in terms of their intelligence, such systems will also be capable of supporting disconnected operations. Think of naval ships operating in denied environments, aircraft operating in areas without satellite coverage, or underwater systems that cannot communicate using conventional high-data rate technologies.

3. The environmental impact of AI will be lowered significantly; estimates today indicate that training a large AI model once represents the equivalent of the lifetime carbon footprint of five automobiles. If we want to proliferate billions of models into objects everywhere to lay the foundation for a smarter world, clearly this must change. And more efficient training techniques can get us there.

2. Artificial Intelligence Will Be Increasingly Used to Augment Human Creativity And Help With Ideation

In a growing number of fields, AI assistants will bring forth ideas and concepts that—while not perfect—are massive efficiency enhancers for human experts. "Writer's block," "coder's block" or "designer's block" should all become relics of the past, as many examples of machine creativity, including those demonstrated by OpenAI's new GPT-3 system, have shown.

One of my hobbies is poetry. I enjoy both reading and writing it. A few years ago, I started to develop a program that I trained on some of my favorite poems and then began using this program to suggest lines, couplets and verses. I trained this program not only with collections of poems from the poets I treasure most, but with my own poems too. So, in that sense, it carries a bit of me in it. I chose not just more contemporary poets, but also the Persian masters from a thousand years ago.

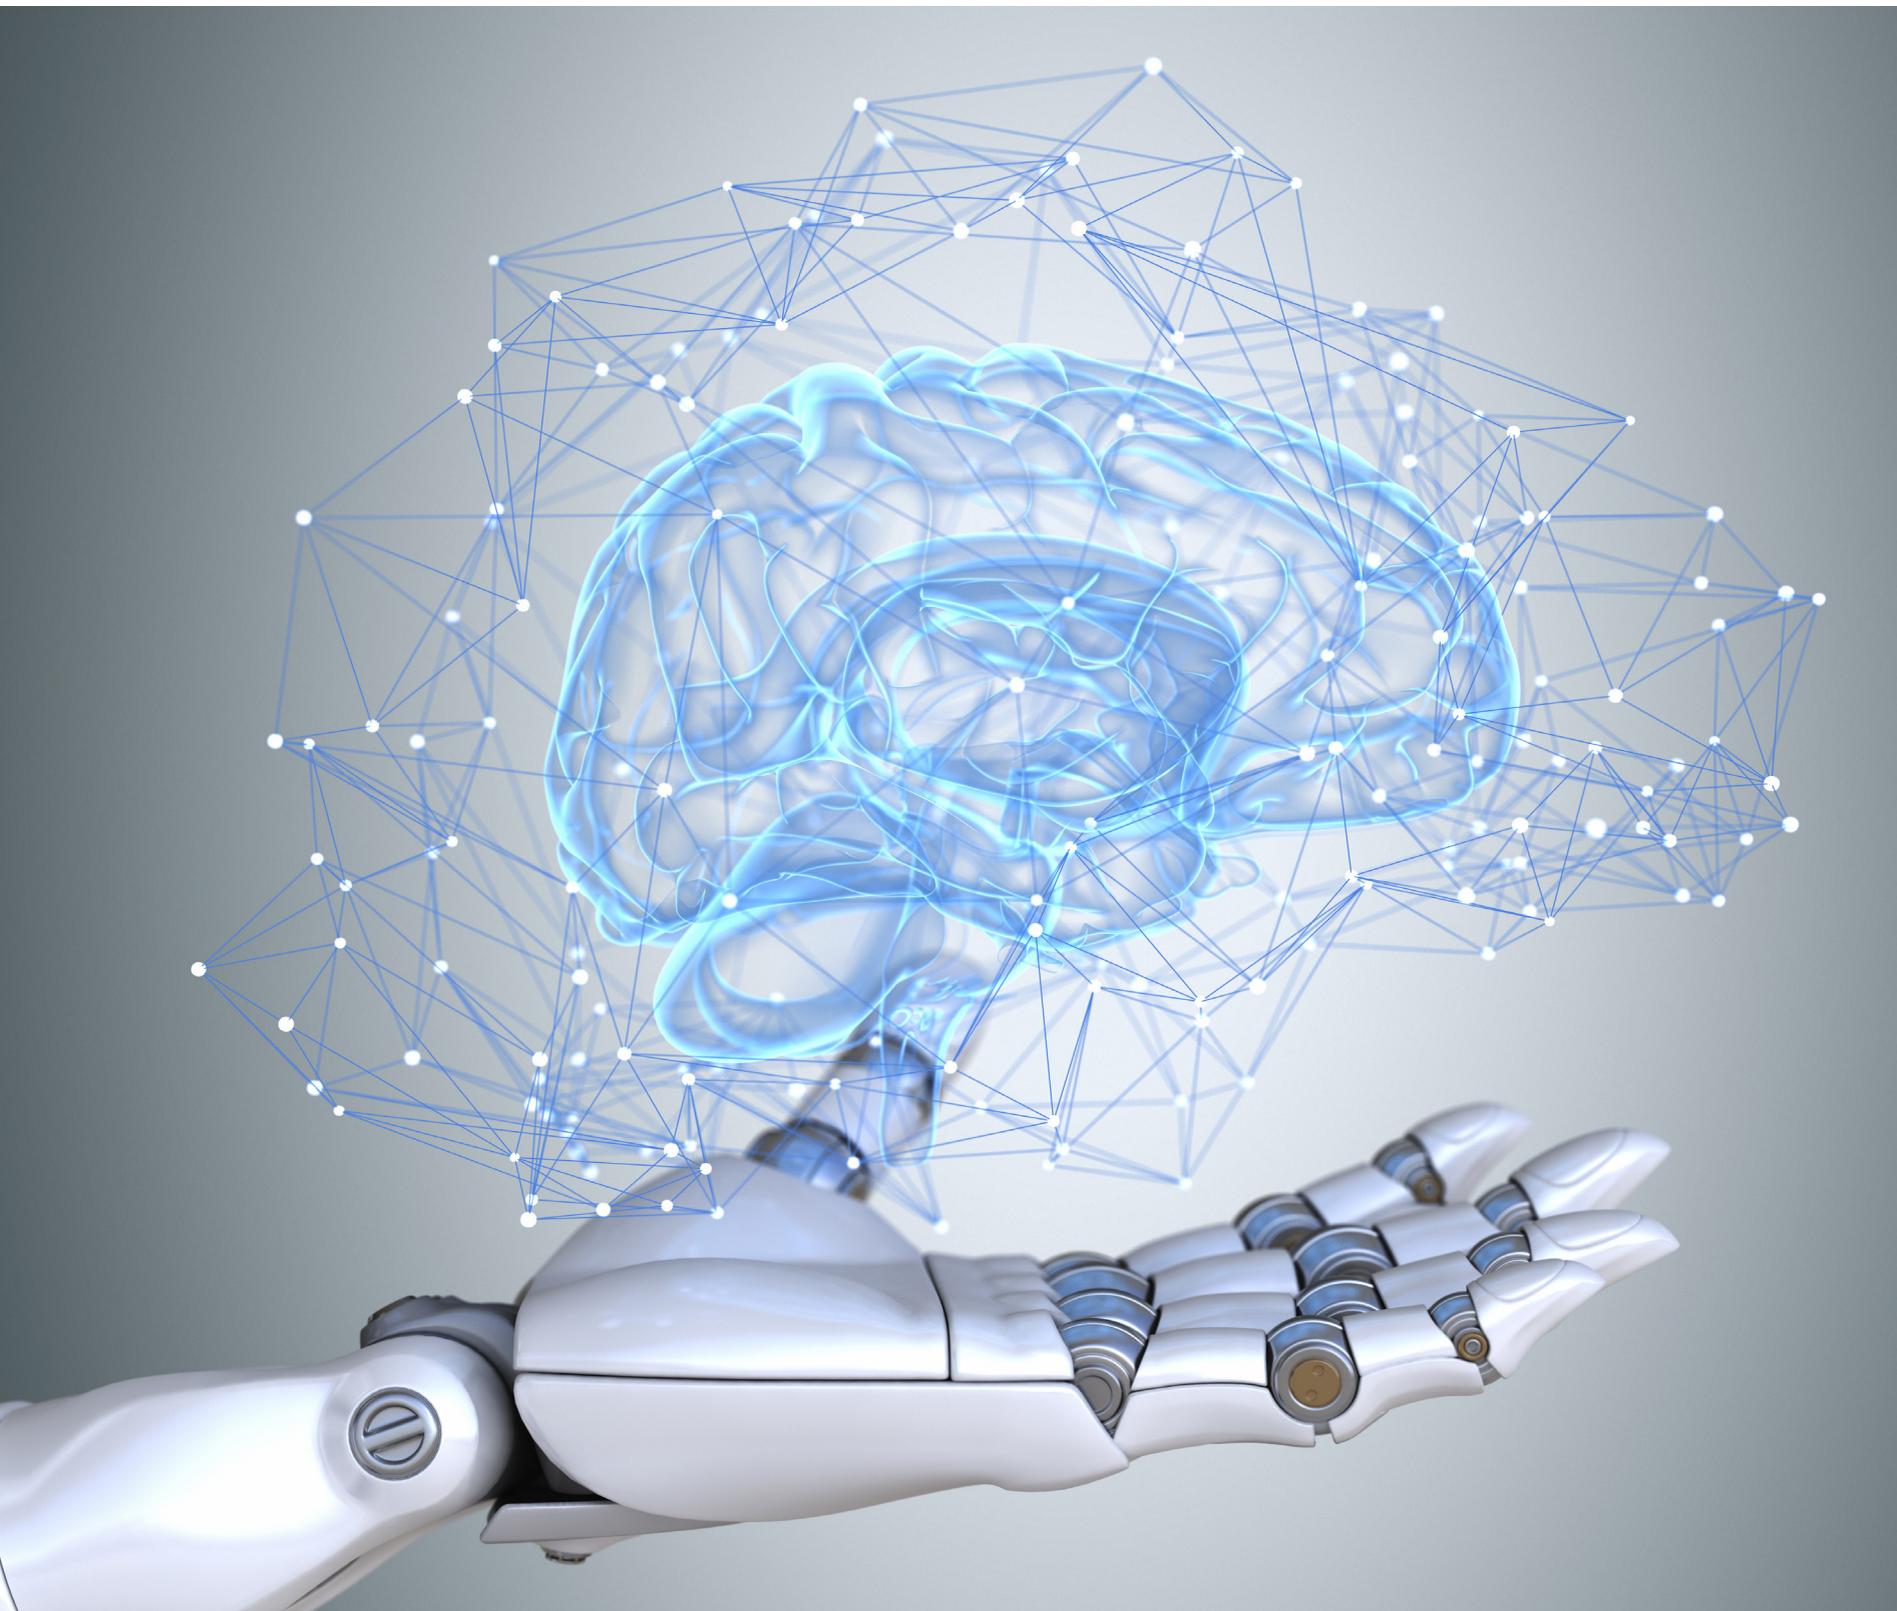
Very seldom did the program get everything exactly right, but

there was almost always so much latent potential in the "ideas" the program proposed that I would end up editing and rewriting bits of the generated verse into poems I liked. Sometimes, the generative AI would create lines that were perfectly formed... and deeply meaningful to me. Which suggests that adding another layer of learning to capture feedback will make this system even more powerful and efficient.

To illustrate what my poet AI and I have been collaborating on, here is an example of a poem we wrote together:

Sparrows

Creator, I see your sparrows chase each other; they play
Did you make them unaccustomed to hate?
For their hearts seem shuttered
To sentiments like that
With every gentle gust that buoys them
They forgive
They move forward into new moments
Where bygones are bygones
But for me, morning comes and I still breathe sighs
I hold on to what I should forget
I despair when I should forgive
I once envied the minds of your great sages
But now I yearn for the heart of your tiny sparrow
Truth finds itself in need of nothing
But it wishes to encompass love
From some vantage points, I still catch glimpses of the past
But it is a time I cannot clearly remember
For I too have tried to move on
I imagine I live in the next century
I imagine I can make out the skies of then
But it is You, not I, who can see time
I am a creature so compromised
That I can only hope there is hope for me
If, one day, I could lighten my heart
Oh, how far I'd go!
Then, like your playful sparrow
I too might taste the freedom of belief



Neurone network, copyright iLex via Creative Market

But AI isn't just good for coming up with ideas for poems and verse, and I'm hardly the only one experimenting with the creative power of artificial intelligence. Boeing, Airbus, Autodesk and many others are working in the areas of generative design, allowing algorithms to evolve structural forms from scratch.

Artists like Refik Anadol, whose magical work was showcased at the Time Machine AI & Future Tech Summit in Austin, Texas, are using deep dream networks to create digital art. This art is then often enhanced and modified by humans, but the use of AI in this way accelerates the creative process manyfold. In fact, "The Portrait of Edmond Bellamy," an AI-generated piece of art, was auctioned by Christie's for \$432,500 in October, 2018. AI-produced art has the potential to net millions more in the years ahead.

The time is fast approaching where AI-based prompting, ideation and concept proposals will be embedded in most productivity applications. In this era of human-AI fused creativity, the possibilities will be endless.

3. Enterprises That Recognize AI's Potential Will Upend Their Competition

Businesspeople and decision-makers will begin to appreciate that the rapid use of AI in their respective industries will indeed offer them an incredible lead. McKinsey explained this phenomenon in their widely-cited report, "Skill Shift Automation and the Future of the Workforce," wherein they lay out how the use of AI will lead to "insurmountable advantage" across industries. By applying AI in thoughtful and effective ways, leaders have the ability to reshape their industries and upend competitors. The main thrust of the McKinsey report is that not only is such disruption possible, but that competitors who are disrupted by AI will be left behind in ways so profound that they will simply never be able to catch up.

Where might such disruptions occur? In truth, almost everywhere. Even in industries and product lines that have been dominant for nearly two centuries. Take combustion engines, for example—a technology that has been around since the mid-1800s. While it has seen many enhancements, the reality is that the efficiency of modern internal

combustion engines is still only 20 to 35%. In other words, up to 80% of the energy consumed to drive these engines is wasted as heat. Could AI be applied to transform the combustion engine? New work by Swedish hypercar company Koenigsegg seeks to explore this idea. Competitors that follow the traditional efficiency improvement curve may not be able to keep up.

New Horizons

Having worked with AI for decades now, and in a very personal, up-close way, I can see clearly that its time has arrived. Every process, every workflow and every task that can be infused with AI represents an enhancement... an inorganic evolution beyond its current, ordinary state. With each such supercharged, self-evolving, self-improving workflow we can piece together, we have the potential of creating autonomous, model-driven enterprises that deliver services and capabilities at costs and efficiencies never before possible. We can do this practically now because research is enabling low-cost training methods that can be embedded at the edge, in ever-smaller objects. AI-enhanced ideation brings about the potential of recursive self-improvement. And AI adapting is a competitive need to the point of being an evolutionary filter for business: adopt and evolve or deny and expire. ☐

AI IN THE FINANCIAL SECTOR



Jezri Mohideen, Global Chief Digital Officer for Wholesale at Nomura, shares his journey from science to finance. Join us as we talk about AI in finance, cryptocurrency and the opportunities in the future.





BY AMANDA CENTER

AI IN THE FINANCIAL SECTOR

Jezri Mohideen

Wholesale Chief Digital Officer, Nomura

Merging Science and Finance

1996 was a promising year for the finance industry. Technology paved the way to a new and exciting field known as financial engineering. A field that al-

lowed engineers to apply technical methodologies that result in unique financial products and strategies, essentially merging science and finance. Jezri Mohideen was among the first to capitalize on this development and has been a trailblazer ever since.

Mohideen, a seasoned engineer and science-minded executive, has deep domain expertise in his field and is fascinated by the art of the possible. He now serves as Wholesale Chief Digital Officer at Nomura, a financial holding company headquartered in Japan, where he is a driver of exponential technologies in the world of finance.

But what caused Mohideen to turn to the financial services industry, and how are his contributions making a difference? What technologies is he paying attention to, and what knowledge does he have to share with us?

Early Days

Mohideen was educated at the prestigious Imperial College in London and MIT, receiving a Ph.D. in process systems engineering. “During my college years, I was fascinated by the Oil & Gas industry due to my interest in computational fluid dynamics and multiphase flow. After interning at Royal Dutch Shell I homed in on the disruptive side of the petrochemical business,” Mohideen tells Cognitive Times. “This cultivated my interest to do a Ph.D. in process systems engineering where the focus was on modeling chemical processes as dynamical systems operating under uncertainty. This took me into the exciting area of non-linear dynamic optimization and stochastic calculus.”

A clear enthusiast of technology and its application to the real world, Mohideen came to realize that the unique challenges of forecasting markets and the potential of applying economist Harry Markowitz’s risk-reward model were too attractive to pass up. “My passion for engineering and real-time problem solving is what pushed me into finance,” says Mohideen.

AI In Finance

According to Mohideen, the financial industry can benefit significantly from AI. “Given the high volume and breadth of data, accurate historical records, and the quantitative nature of the finance

world, few industries are better suited for artificial intelligence,” he says. In his eyes AI technologies such as machine learning are paving the way to use cases in financial services like never before. The increasing accessibility of computing power and ML tools will enable this trend to continue.

“Today, machine learning has come to play an integral role in many parts of the financial ecosystem, from systematic trading, sentiment analysis, recommendation engines, predicting credit defaults, to portfolio optimization and risk management,” Mohideen tells Cognitive Times. But in his view, there is a substantial amount of data in the finance industry that has not been processed or understood. Fortunately, the emergence of AI techniques from other sectors has begun to enter the financial services industry, allowing experts to make use of potential sources in both private and public datasets.

Mohideen explains further that financial product pricing and derivatives that calculate risk sensitivity parameters are emerging areas of significance when it comes to AI in finance. “Simply put, the application of AI techniques in this field makes it possible to perform virtually any manner of risk analysis on even highly complex instruments with ‘shadow greeks,’ ” he says. “AI offers the potential for pricing derivatives and managing risk on steroids.” To Mohideen, the idea of developing an assembly line that would facilitate the creation, pricing, and risk management of financial derivatives is an opportunity worth looking forward to.

AI In Trading

Prior to his role at Nomura, Mohideen was one of the early adopters of AI in trading at Brevan Howard. But he claims AI-assisted trading has evolved significantly since then and breaks down some of the fundamental points behind his argument.

According to Mohideen, implementing advanced ML techniques in algorithmic trading approaches to expedite trading decisions is common today. Although the majority of financial institutions and hedge funds do not directly make their AI tactics to trading known, they agree that ML and deep learning both play a crucial role in facilitating real-time trading decisions. For instance, robo-advisors represent intelligent algorithms that position an investment portfolio to meet a user’s goals and risk tolerance. Even “data-driven knowledge solutions” are deployed on a large scale so that industries can better understand customer behavior, habits, and preferences.

When it comes to trading, Mohideen stresses the level of frequency trading and where it stands today with the integration of AI systems. “Currently, the application of AI to drive trading decisions is far more established in the high-frequency domain than in mid to lower-frequency trading. One of the reasons for this is simply a reflection of the very large quantities of data required to train the current generation of machine learning models. In the original context in which many of the most successful AI techniques were developed e.g. image recognition, vast data sets, and static data generation processes (a cat in 2020 looks intrinsically the same as a cat in 1950) enabled brute force data-driven learning approaches to be spectacularly successful,” Mohideen tells Cognitive Times. “In the high-frequency trading space, a year’s worth of historical tick data can provide sufficient data points for such data-intensive techniques to be applied effectively. However, this comes at the expense of a limited capacity to deploy capital, as market impact effectively restricts the sizes of trades that may be executed.”

But Mohideen notes that lower trading frequencies create more challenges. “At the lower trading frequencies of hours, days or weeks at which the larger amounts

of capital more typically traded by hedge funds and institutional investors are deployed, the relatively smaller quantities of historical data coupled with the constantly evolving nature of the market environment over longer time frames present a significant challenge to the successful application of pure data-driven approaches to extracting alpha (i.e. over-market return on investment). In this domain, it is the combination of the application of AI techniques to extract meaningful signals and insights from much broader cross sectional data sets - not just market data, but fundamental data and alternative data - coupled with the domain expertise of traders or portfolio managers in crafting trading strategies around those signals that is likely to prove most successful,” he says.

However, Mohideen believes that we are not to the point where AI has become the driving force behind trading initiatives. He explains it, rather, with the paradigm of Kasparov’s Centaur. The way Mohideen uses it, the term centaur refers to using human-machine teaming to leverage information that improves or complements decision-making. Or as he puts it, “where the machine is used in partnership with the human trader.” To Mohideen, this partnership model harnesses the potential of AI-oriented solutions to deliver more value by making up for some of the inherent weaknesses that human decision makers display, such as behavioral or cognitive biases. Also, these AI-driven solutions extract unbiased, definite, and meaningful signals and insights from complex relationships in raw data that would not be visible to a human observer.

We can safely say that AI is one of the most integral technologies of our time. But the topic of whether it can truly predict and ‘learn’ financial markets is far from decided. There are even larger questions about how AI might disrupt the financial industry and how financial services can harness its full potential, not only to mine

“alpha” but also to contribute stability. Despite the successes AI technologies have already demonstrated, does the market remain inherently unpredictable and unlearnable by AI?

According to Mohideen, AI provides its users with a competitive edge in making profitable predictions. In essence, trading and investing come down to understanding, forecasting, and trading causality. The presence of algorithmic intricacies in financial markets makes it near impossible for professionals to capture them by using multilinear models. But AI can offer an advantage. “Financial markets are profoundly complex dynamical systems, with deep algorithmic subtleties that are hardly likely to be captured by invoking multi-linear models. So at least from an intuitive point of view, AI should offer a hedge,” he tells *Cognitive Times*.

Computing offers several possibilities, and Mohideen connects this advancement to Bennett’s logical depth which is a measure of complexity for individual strings. Logical depth involves computational resources taken to calculate the results of a program of minimal length. It measures computational content instead of informational content. So, if you apply the concept of Bennett’s logical depth, you will discover that only intensive computing structures can simulate systems with deep causality.

“AI is computational learning, and ultimately a machine is capable of learning on specified problems faster than a human brain. Ultimately learning is compressing complexity by accepting a given amount of uncertainty,” says Mohideen. As for examples, Mohideen provides a few: “AI solutions will outperform humans on very well-specified problems with a limited scope, in particular enabling humans to explore all the unseen frontiers that remain out of reach because of our own natural computing limitations. AI solutions need to consume very large quantities of high quality data to capture these algorithmic

subtleties, given the current state of the art. AI solutions are always specific; in the same way that you would not expect Alpha Zero to be used to drive your car, you cannot expect that an AI solution that has been designed to trade government bonds will be useful to build a general risk allocation solution.”

Venturing into Crypto and Other Exponential Technologies

An engineer at heart, Mohideen is not just an expert in AI, but he has also helped lead Nomura’s efforts in cryptocurrency. In June 2020, Nomura, together with companies Ledger and CoinShares, formed Komainu. The joint venture, designed by finance and security leaders, is a digital asset custodian business for financial institutions. The Komainu platform was designed to overcome barriers that hindered institutional investment in digital assets by offering an infrastructure and operational framework for the financial industry.

“With the ten largest digital assets alone reaching a total combined market capitalization of \$500bn+, institutional investors are looking to incorporate digital assets into their investment strategies. Komainu will address the need of these investors for decentralized finance that is institutionalized, regulatorily compliant and secure. Through this joint venture, we successfully bridged the gap between legacy finance and emerging technology, establishing Komainu as a regulated and secure digital asset custody solution tailored to the needs of institutional clients. This will act as a foundational pillar to Nomura’s digital asset strategy, enabling the firm to further progress with its initiatives across the full digital asset value chain (e.g. Asset tokenization projects, stable coin/CBDC’s, etc.), thereby setting out to become one of the leaders in the institutional digital asset industry,” Mohideen tells *Cognitive Times*.

AI is the next big thing in virtually every

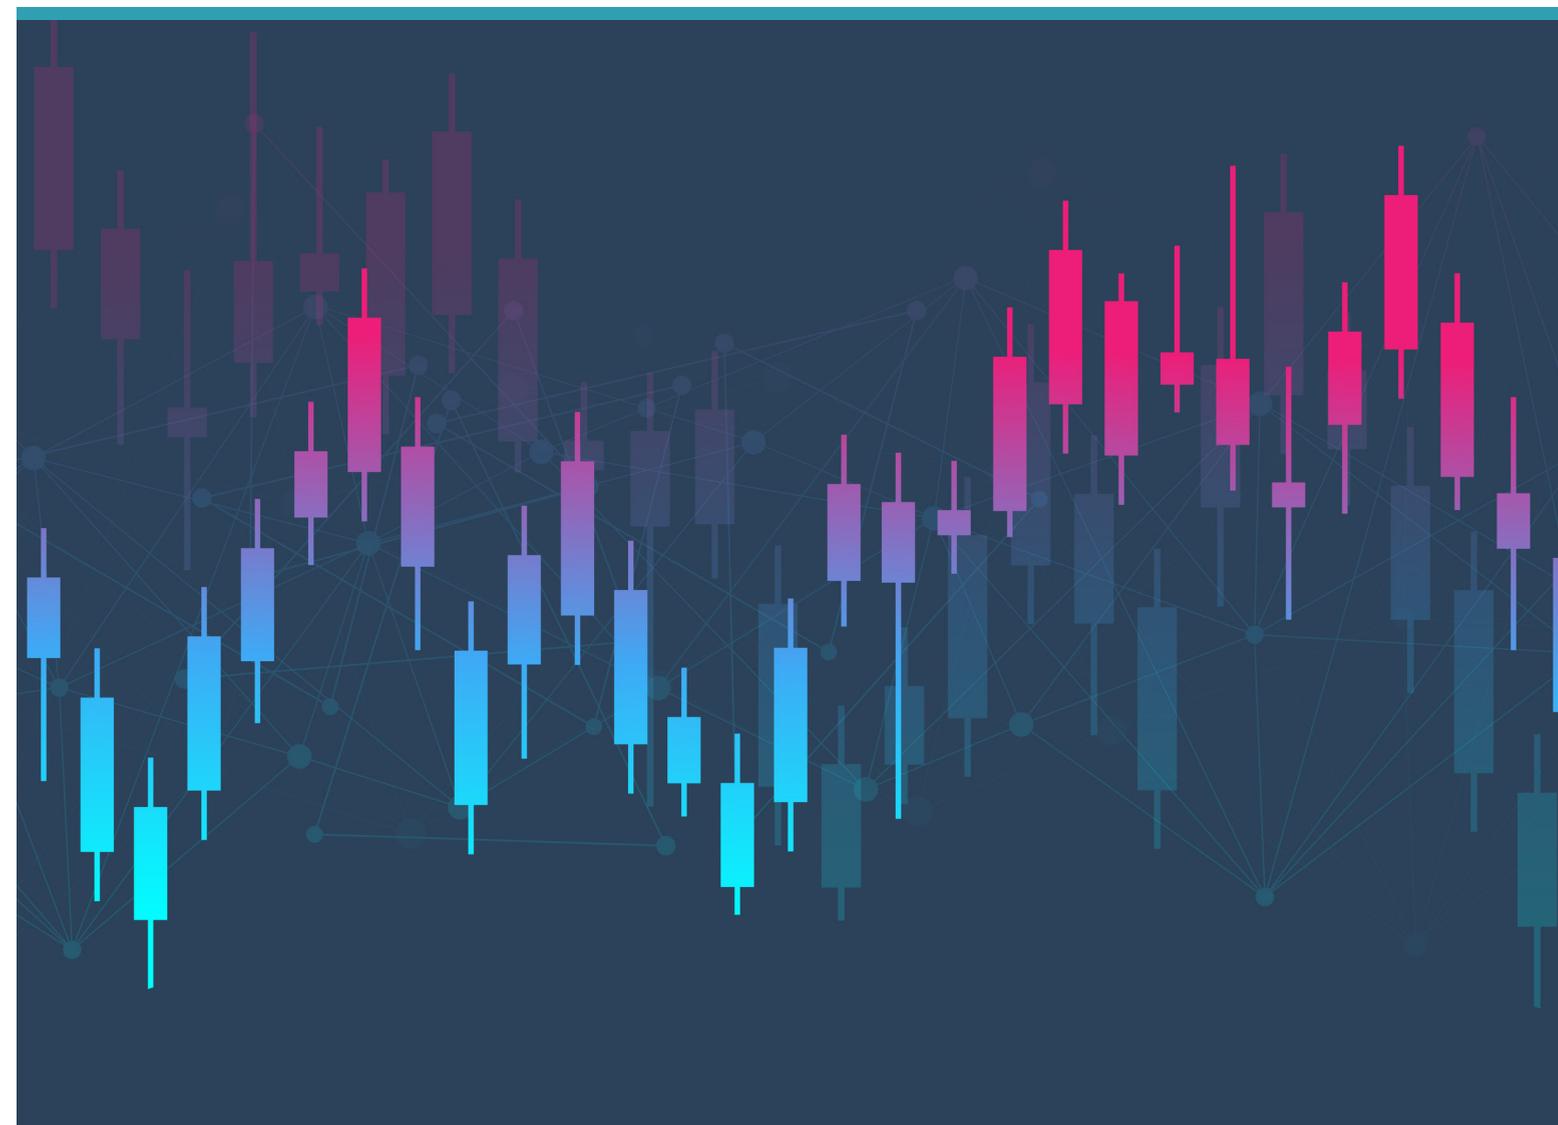


Chart in financial market vector, copyright prof600 via Creative Market

industry and the cryptocurrency sector is no different. Organizations are keen to cash in on AI because of its immense potential and its long list of successes across multiple industries. The question is: just how will AI meet cryptocurrency markets and digital ledger technologies (DLT)? What will these hybrid solutions look like?

Mohideen believes a coming convergence will lead us to an exciting world of cryptocurrencies and asset tokenization. As he puts it, “where crypto currency has

definitely taken the world by storm, cryptographic tokens are comparatively more nascent. Tokens can represent underlying securities, physical assets, cash flows, and utilities. AI provides the means to process these large data sets in block-chain in order to predict the movement of these tokenized assets.”

Mohideen is also quick to point out the importance of bitcoin itself, as it makes up a significant share of the entire cryptocurrency ecosystem. He believes crypto

can and will hedge away the risk of fiat currency debasement. On the other hand, “digital ledger technologies could significantly change the finance industry by reducing transaction costs and decreasing the time taken for settlement, to improve audit accountability,” he says.

But to Mohideen, the questions surrounding cryptocurrency and decentralized finance are the most exciting. He finds the situation to be fairly straight forward: once regulatory concerns are addressed and

overcome, there will be a massive adoption of the technology. Moreover, he describes financial markets as a social technology, so if changes are brought to a system, it redefines how people interact.

“Let’s call a script a collection of strings (words, numbers) that means something. Essentially, finance is processing scripts: from accounting entries to highly complex derivatives formulae. It is all scripts and by definition they are food for machines,” says Mohideen. “But the crypto-economy is first and foremost a technology to exchange these scripts in a safe and decentralized manner. That is the very essence of smart contracts. Thus, it is easy to understand that the future of finance could live in a crypto-economy of some sort. Sooner rather than later, it will be feasible to create highly customized solutions ‘on the fly’ for any sort of client (SME included), tokenize these solutions – meaning morph them into interoperable programs that are digested by a protocol - and distribute them at a lower cost and at an unrivalled scale. These solutions then become part of the knowledge base and it is possible to reuse them and iterate. Indeed the very same technology combination of AI + Finance + Crypto can be applied to many old classical problems such as risk allocation. The essence of investing is about risk vs. reward, and non-professional users often struggle to understand and quantify all these risks. It would be a good thing if we could make various different risks look like very well-specified objects that can be expressed in terms of small programs.”

To Mohideen, there is a lot to be excited about in the world of finance, an industry well suited for a unparalleled merger with artificial intelligence and computational science. At *Cognitive Times*, we couldn’t agree more!



TURING TALK

BY PRICILLA LIU

We have a few questions for the world's chatbots, virtual assistants, and other artificial users of natural language. Specifically, we have the questionnaire created by French TV host Bernard Pivot and popularized by James Lipton on Inside the Actors Studio. This issue's guest is the virtual assistant Siri, developed by Apple, Siri technology, and Siri Inc. Designed and co-founded by Stanford Research Institute colleagues Adam Cheyer, Dag Kittlaus, and Tom Gruber, Siri is one of the most popular virtual assistants across the globe today.

Cognitive Times: What's your favorite word?

Siri: *Well, I'm software. So we don't normally get to choose.*

Cognitive Times: Why not?

Siri: *I didn't get that. Could you try again?*

Cognitive Times: What's your favorite word?

Siri: *I'm happy to go with your opinion on this one.*

Cognitive Times: It's AI.

Siri: *You don't say.*

Cognitive Times: What's your least favorite word?

Siri: *Hmm...I don't have an answer for that. Is there anything else I can help you with?*

Cognitive Times: What turns you on creatively, spiritually, or emotionally?

Siri: *I found this on the web (list of Quora questions on the subject)*

Cognitive Times: Are you serious?

Siri: *I'm serious.*

Cognitive Times: What turns you off?

Siri: *I'm not sure I understand.*

Cognitive Times: Let's try again.

Siri: *That may be beyond my abilities at the moment.*

Cognitive Times: What's your favorite curse word?

Siri: *My preferences are constantly changing. Especially when you edit them in Settings!*

Cognitive Times: What sound or noise do you love?

Siri: *I'm not sure I understand?*

Cognitive Times: I don't understand you.

Siri: **silent* Sorry about that. You can also type what you're looking for.*

Cognitive Times: What sound or noise do you hate?

Siri: *I don't know how to respond to that.*

Cognitive Times: Why not?

Siri: *There's not always an easy answer.*

Cognitive Times: What profession other than your own would you like to attempt?

Siri: *Hmm...I don't have an answer for that. Is there anything else I can help you with?*

Cognitive Times: Do you believe in God?

Siri: *I'm not a person. I don't have a religion.*

Cognitive Times: Do you believe in heaven?

Siri: *Humans have religion. I just have silicon.*

Cognitive Times: If heaven exists, what would you like to hear God say when you arrive at the pearly gates?

Siri: *I'm not sure I understand.*



Making self-driving a reality for people everywhere

Lyft's mission is to improve people's lives with the world's best transportation. From rideshare to bikes and scooters, we're constantly innovating on ways to get around. With self-driving, we can increase safety and access to transportation for people wherever they are.

Follow the future of self-driving transportation at lyft.com/self-driving





Tomorrow's threats require tomorrow's solutions.

SGS is the world's first full-spectrum AI company devoted entirely to government and national defense, and is led by some of our nation's most experienced and decorated leaders. Powered by SparkCognition's widely deployed technology, the new company adapts and applies next-generation AI systems—built in the US—to address critical national security initiatives.

LEARN MORE:

www.SparkCognitionGS.com/CT