ToughTech

A publication by The Engine, built by MIT

SPRING 2020

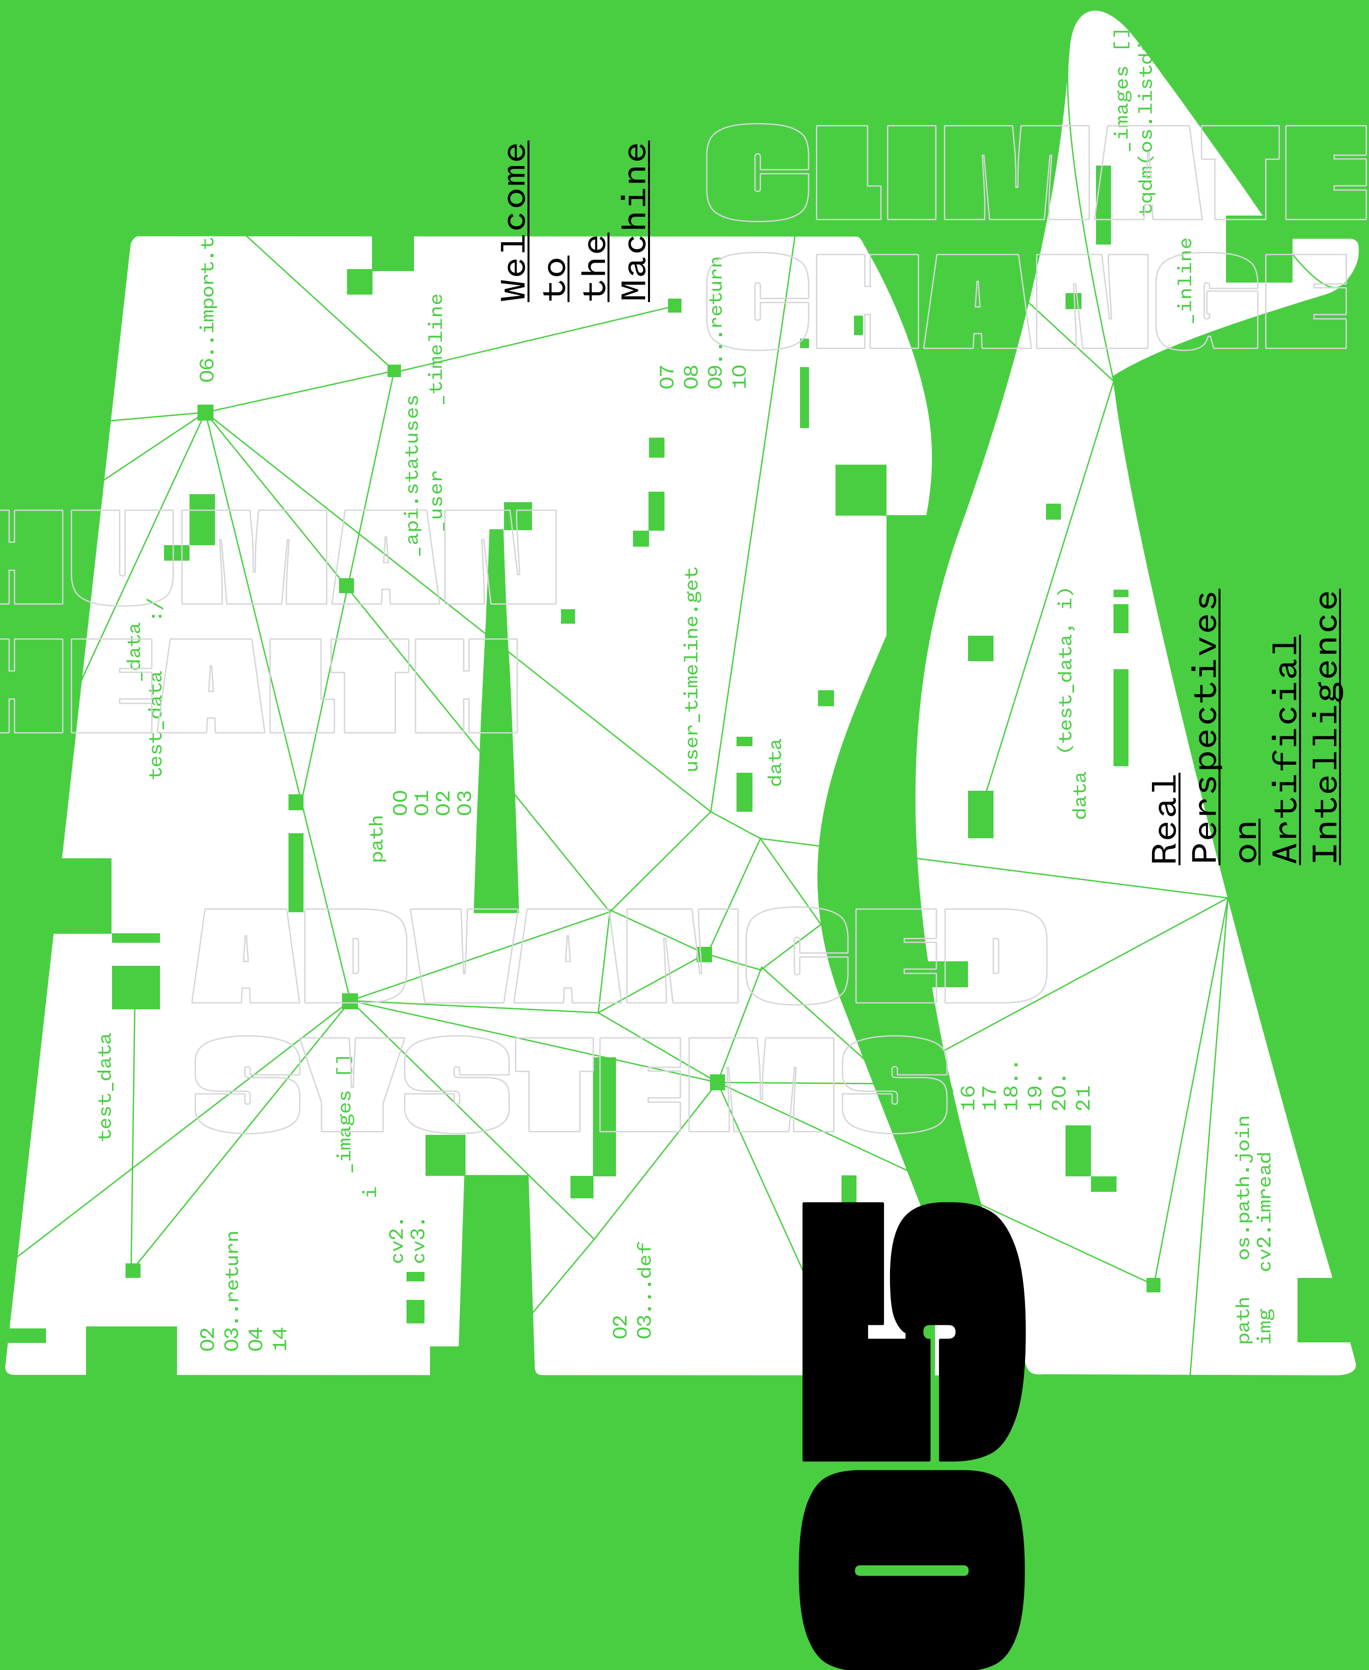

Welcome to the Machine

CLIMATE CHANGE

HUMAN HEALTH

ADVANCED SYSTEMS

Real Perspectives on Artificial Intelligence

05

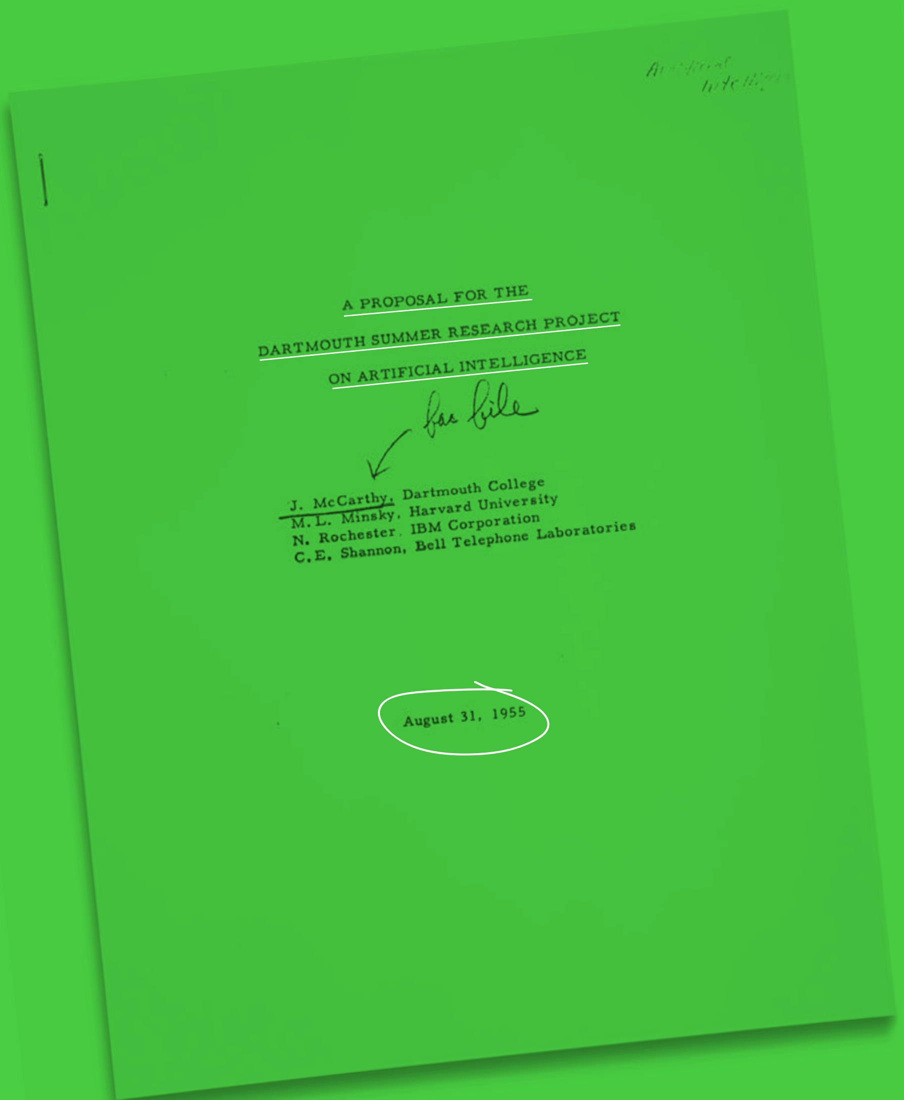

↳

**THE BIRTH OF AI:
The Dartmouth Summer Research Project
on Artificial Intelligence**

The field of artificial intelligence was born during the summer of 1956 at the The Dartmouth Summer Research Project on Artificial Intelligence. The proposal for the project, penned in 1955, introduced the term "artificial intelligence" to the world, and its core conjecture endures to this day:

"The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can, in principle, be so precisely described that a machine can be made to simulate it."

* Reproduced here are the pages from the original typewritten proposal for the project, annotated by John McCarthy. Thanks to the Rauner Special Collections Library at Dartmouth for providing scans of the proposal.

THE ENGINE
Built by MIT

A home for Tough Tech founders.

The Engine, built by MIT, is a venture firm that invests in early-stage companies solving the world's biggest problems through the convergence of breakthrough science, engineering, and leadership. Our mission is to accelerate the path to market for Tough Tech companies by providing access to a unique combination of investment, infrastructure, and a vibrant ecosystem.

Tough Tech Publication 05

May, 2020
The Engine, Built by MIT

Editor: Nathaniel Brewster
Design: www.draft.cl
Print by Puritan Capital, NH & MA

CONTENTS

The time for Tough Tech is ______ now.

THE ENGINE
Built by MIT

We started planning this publication before anyone had heard of COVID-19. The suddenness with which everything changed, however, will be one of the greatest lessons of our time.

How suddenly the world stopped moving, how suddenly our medical stockpiles got depleted, and how suddenly the death rate ticked up around the globe.

COVID-19 has taught us how little time we are left with to plan and deal with the mounting list of challenges that we face — in areas like climate, human health, and infrastructure — and exposed the systemic cracks that exist in how we respond.

It has posed certain questions, such as what do we do when an outbreak occurs in an economy that our supply chains are overly reliant on, how do we transport food and supplies when borders are closed, or how do we keep health systems running when our medical workforce itself is at risk? This time around, our systems weren't set up to adapt quickly to these pressures. Hence, we need to rethink each of them so that we are more nimble and ready to adapt when the next disruption occurs.

This is an important mandate — to create more flexible and responsive solutions to the world's growing threats — and it is Tough Tech that can get us there.

So much is already being done.

In our portfolio alone E25Bio, Biobot Analytics, Vaxess Technologies, C2Sense and Radix have dedicated their current efforts to help combat COVID-19 with better diagnostic tools and medicine delivery platforms. Commonwealth Fusion Systems, Form Energy, Lilac Solutions, Syzygy Plasmonics and others are pioneering systems to combat the threats of climate change. Sync Computing, isee, HyperLight, Cambridge Electronics and more continue to work on the advanced computing systems that will help us decode solutions to a myriad of fundamental challenges.

These examples are a few of the many from our portfolio — the companies you'll see in this publication, along with hundreds of others like them across the globe, showcase how Tough Tech can enable more resilient, distributed systems as we navigate through a dynamic and unpredictable future.

We must continue to build, look forward, and face our biggest challenges head on.

Katie Rae
CEO & Managing Partner

"Tough Tech can enable more resilient, distributed systems as we navigate a dynamic and unpredictable future."

THE AI ISSUE

Artificial intelligence is a general purpose technology with incredible possibilities. Its story winds through academic laboratories, halls of philosophy and mathematics, and our planet's most significant industries. It touches our homes and even our bodies. It is everywhere and, as some predict, will become everything.

There are vital and vigorous debates as to the societal, cultural, and moral implications of artificial intelligence. This publication largely sidesteps those conversations in favor of providing an overview of how AI has, and will, unbridle industries and technologies, namely those we call Tough Tech. Think of it as a primer — providing context for the profusion of AI approaches, technologies, and businesses.

How will AI help us develop new approaches to combat challenges like climate change, disease, and the limits of computing itself? What shape are these new Tough Technologies taking? This publication attempts to build a common understanding of AI in an effort to answer these questions.

"We can only see a short distance ahead, but we can see plenty there that needs to be done."

Alan Turing

The Tough Tech Landscape report investigates the surge in venture funding of Tough Tech since 2015. Produced by **The Engine** and **PitchBook,** the report analyzes over 4,000 companies that have received over $105 billion in funding over a period of roughly five years.

The AI Industry

AI & ML VC Deal Activity: 2012-2019

The following analysis is found on page 17 of the report.

There are significant advantages to deploying AI systems in arenas requiring significant computations to unearth actionable data and/or results with greater efficiency.

The numbers speak for themselves — from 48 financings for $253.1 million in VC in 2014, the cycle of venture investment in AI & ML has reached a new high of $4.6 billion invested across 183 transactions in 2019 to date.

More volume is found concentrated in the late stage, signifying the maturation of the space overall. Significant numbers of startups and maturing businesses are tackling arenas armed with custom applications of AI & ML, most notably in a handful of avenues that have encompassed the bulk of funding since 2012: drug discovery, resource management, autonomous vehicles and systems, and business process automation.

Autonomous vehicles and systems admittedly dominate the top 10 businesses by VC raised within this Tough Tech segment, from Cruise's multiple billions to Aurora Innovation's platform for self-driving cars. However, healthcare applications also account for a significant chunk of VC invested and volume, exemplified by businesses such as HeartFlow, which has created a product suite for healthcare professionals to identify and model diseases and treatments, or BenevolentAI, which has partnered with companies such as Novartis to aid drug discovery by leveraging patient-level data to account for patient heterogeneity.

It is likely that transportation, logistics, and healthcare applications will continue to dominate the burgeoning field of AI & ML applications given their dynamics and economics; however, multiple other avenues are cropping up, particularly in the realm of business or repetitive process automation. Take DropWise, which is leveraging ML to develop a hydrophobic coating material, or Cobalt Robotics, which makes patrolling robots that utilize ML in detection of anomalies to reinforce more accurate recognition over time.

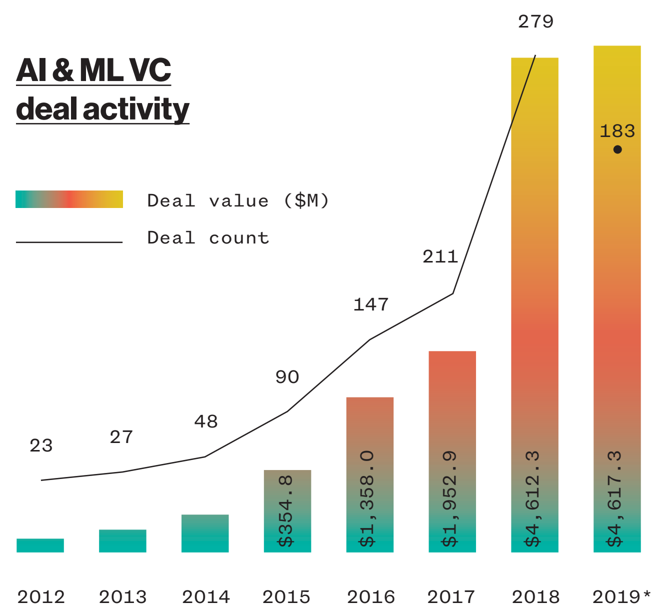

Source: PitchBook
* As of September 25, 2019.
(There were approximately 50 more deals, valued at over $3B in total, between the date of this landscape and the end of 2019.)

Strong growth rates in both VC invested and volume across all AI & ML since 2014.

Well over $9 billion invested since the start of 2018 till now.

Increased expansion of applications, both commercial and theoretical, with transportation, logistics, and healthcare dominant, yet automation of repetitive tasks and business processes also growing apace.

WELCOME TO THE MACHINE

AI is redefining how we fight climate change and combat disease. It is even revolutionizing the systems responsible for its own evolution.

by Michael Blanding
Illustrations by Harol Bustos

WE have a romantic vision of the scientific discovery process. A white-coated chemist spends long hours in the lab, titrating, pipetting, centrifuging chemicals, until — by accident or design — they stumbles upon a new molecule that might do something useful in the world. Then there are longer nights spent testing, refining, and optimizing the synthesis process, hoping that one day that chemical can be commercialized. In the luckiest of scenarios, that process might take 5 years — often it takes 10 or 20.

For Alan Aspuru-Guzik, that's 20 years too long. “Every time a human is involved, you're really wasting a lot of time,” says the University of Toronto chemistry professor, who wears big, grey-rimmed glasses and usually an impish grin. “You really want to have a machine doing as many things as possible.” Four years ago, while a professor at Harvard University, he succeeded in producing the first novel chemicals discovered by computer — materials for organic light emitting diodes[1] — selected out of 1.6 million possible candidates. From start to finish, the process just took a few months.

Since moving to Toronto two years ago, he's expanded his vision to create a self-driven laboratory, with robotic arms to do all of the hard work of discovery, development, synthesis and optimization, all at the beck and call of its human overlords, much faster than human hands ever could. Once the lab is totally up and running, it could create new chemical materials at the rate of one every hour, an incredible leap forward in productivity.

That advance is all because of one thing: the power of artificial intelligence (AI). “The chemical industry is one of the last industries that has not really been taking advantage of AI,” Aspuru-Guzik says. “Of course, many other industries, like buying stuff on the Internet have already been affected by AI, but chemistry is very hard to do because it has a very low rate of discovery, and many complicated development processes.” New advancements in power and ability, however, have put AI on the cusp of transforming chemistry discovery, as well as the innovation process of many tough technology industries, serving as a virtual assistant to dramatically speed up and improve what humans are capable of inventing on their own.

The idea that machines possess intelligence to rival a human goes back to the golems and automatons of antiquity. It didn't start becoming a reality until the 1950s, when Dartmouth computer scientist John McCarthy coined the term “artificial intelligence”[2] to describe a new class of computers capable of reasoning and making decisions for themselves when faced with novel situations. The first AI mastered simple tasks, such as playing games of Tic-Tac-Toe or Chess. However, over the decades, it has exploded into a variety of fields, including image processing, speech recognition, and natural language processing.

A subset of AI called machine learning (ML) promises to transform innovation. Simply put, ML searches through massive amounts of data to find patterns, and then uses those patterns to make future predictions. These processes power everything, from the predictions you get from shopping websites to the results you get from online search engines. “If you go to Google and type Disney, Florida, and Delta airlines, then it's pretty clear you are going to the Magic Kingdom,” says Scott Stern, a business professor at MIT Sloan School of Management and co-author of a recent National Bureau of Economic Research paper on AI and innovation. ML excels in predicting patterns too complex for humans to uncover, at least without a lot of bookkeeping. “Machine learning dramatically enhances the ability to predict things that are relatively rare in isolation, but once you aggregate them, they become clear,” Stern says.

→

THE IDEA OF MACHINES POSSESSING INTELLIGENCE TO RIVAL A HUMAN GOES BACK TO THE GOLEMS AND AUTOMATONS OF ANTIQUITY. IT DIDN'T START BECOMING A REALITY, HOWEVER, UNTIL THE 1950S, WHEN DARTMOUTH COMPUTER SCIENTIST JOHN MCCARTHY COINED THE TERM “ARTIFICIAL INTELLIGENCE” TO DESCRIBE A NEW CLASS OF COMPUTERS CAPABLE OF REASONING AND MAKING DECISIONS FOR THEMSELVES WHEN FACED WITH NOVEL SITUATIONS.

The process relies on algorithms, sets of rules, and instructions that the artificial brain can use to sift through data, supplemented by heuristics or “rules of thumb” that help narrow down the list of possibilities. This process has been significantly enhanced since the development in the mid-2000s of artificial neural networks, a set of algorithms that mirror the structure of the human brain with neurons that pass information from one to another to enable “deep learning.” Since then, there has been an explosion in the applications of AI — from stock picking to facial recognition to routing your Uber. That's been especially true in the past five years or so, as the advent of cloud computing has enabled researchers and companies to use vast amounts of computing power cheaply.

→

THE NEW FRONTIER OF AI IS REINFORCEMENT LEARNING, IN WHICH THE MACHINE BRAIN TRIES DIFFERENT COMBINATIONS BASED ON TRIAL-AND-ERROR, GRADUALLY HONING IN ON AN ANSWER BASED ON PRE-DEFINED POSITIVE AND NEGATIVE FEEDBACK.

AI generally comes in two types: the most common type, supervised learning, which happens when someone gives the AI an example of a target and then asks it to sift through data to find other examples of that target; unsupervised learning happens when the AI is given unstructured data and asked to find its own patterns by employing techniques such as “nearest neighbor search” to group similar objects into categories. The new frontier of AI is reinforcement learning, in which the machine brain tries different combinations based on trial-and-error, gradually honing in on an answer based on pre-defined positive and negative feedback. Google made a splash three years ago with its AlphaGo software, which used reinforcement learning to beat the best human players in Go, a Chinese game infinitely more difficult than chess, discovering new strategies no human had thought of.

It's the quality of discovery that makes AI so exciting not only for optimizing products and services but also for learning how to create new ones, such as the chemicals in Aspuru-Guzik's self-driving laboratory. “If you could get predictions on unstructured data to be good enough you could not simply apply them to existing things and optimize them,” Stern says, “but you could undertake a resequencing of the innovation process itself for the creation of fundamentally novel products, services, and inventions.” He compares the technology to lenses, which were used to make eyeglasses in the 1400s. “People got good at grinding lenses so that people could see,” Stern says. “But Galileo took those lenses and built the first telescope and was immediately able to resolve the existence of moons around planets. It allowed us to literally resolve phenomena we couldn't even imagine, and ask new types of questions.”

Stern, who is on the board of Atomwise, a company using AI to discover new drugs, makes the distinction between AI as a “general purpose technology” that can be used to enable other technologies, and AI as a “method of invention” that allows for the creation of fundamentally new products separate from the AI itself. Both uses could impact the development of tough technologies, allowing dramatically new breakthroughs and capabilities. Among other industries, AI is poised to transform the fields of health, climate science, and advanced computing itself.+

ALAN ASPURU-GUZIK
Professor of Chemistry and Computer Science at the University of Toronto;
Co-Founder of Zapata Computing,
Co-Founder of Kebotix

SCOTT STERN
David Sarnoff Professor of Management of Technology, Professor, Technological Innovation, Entrepreneurship, and Strategic Management at MIT

1. https://www.ncbi.nlm.nih.gov/pubmed/27500805 ; https://www.seas.harvard.edu/news/2016/08/towards-better-screen
2. https://www.technologyreview.com/s /612404/is-this-ai-we-drew-you-a-flowchart-to-work-it-out/

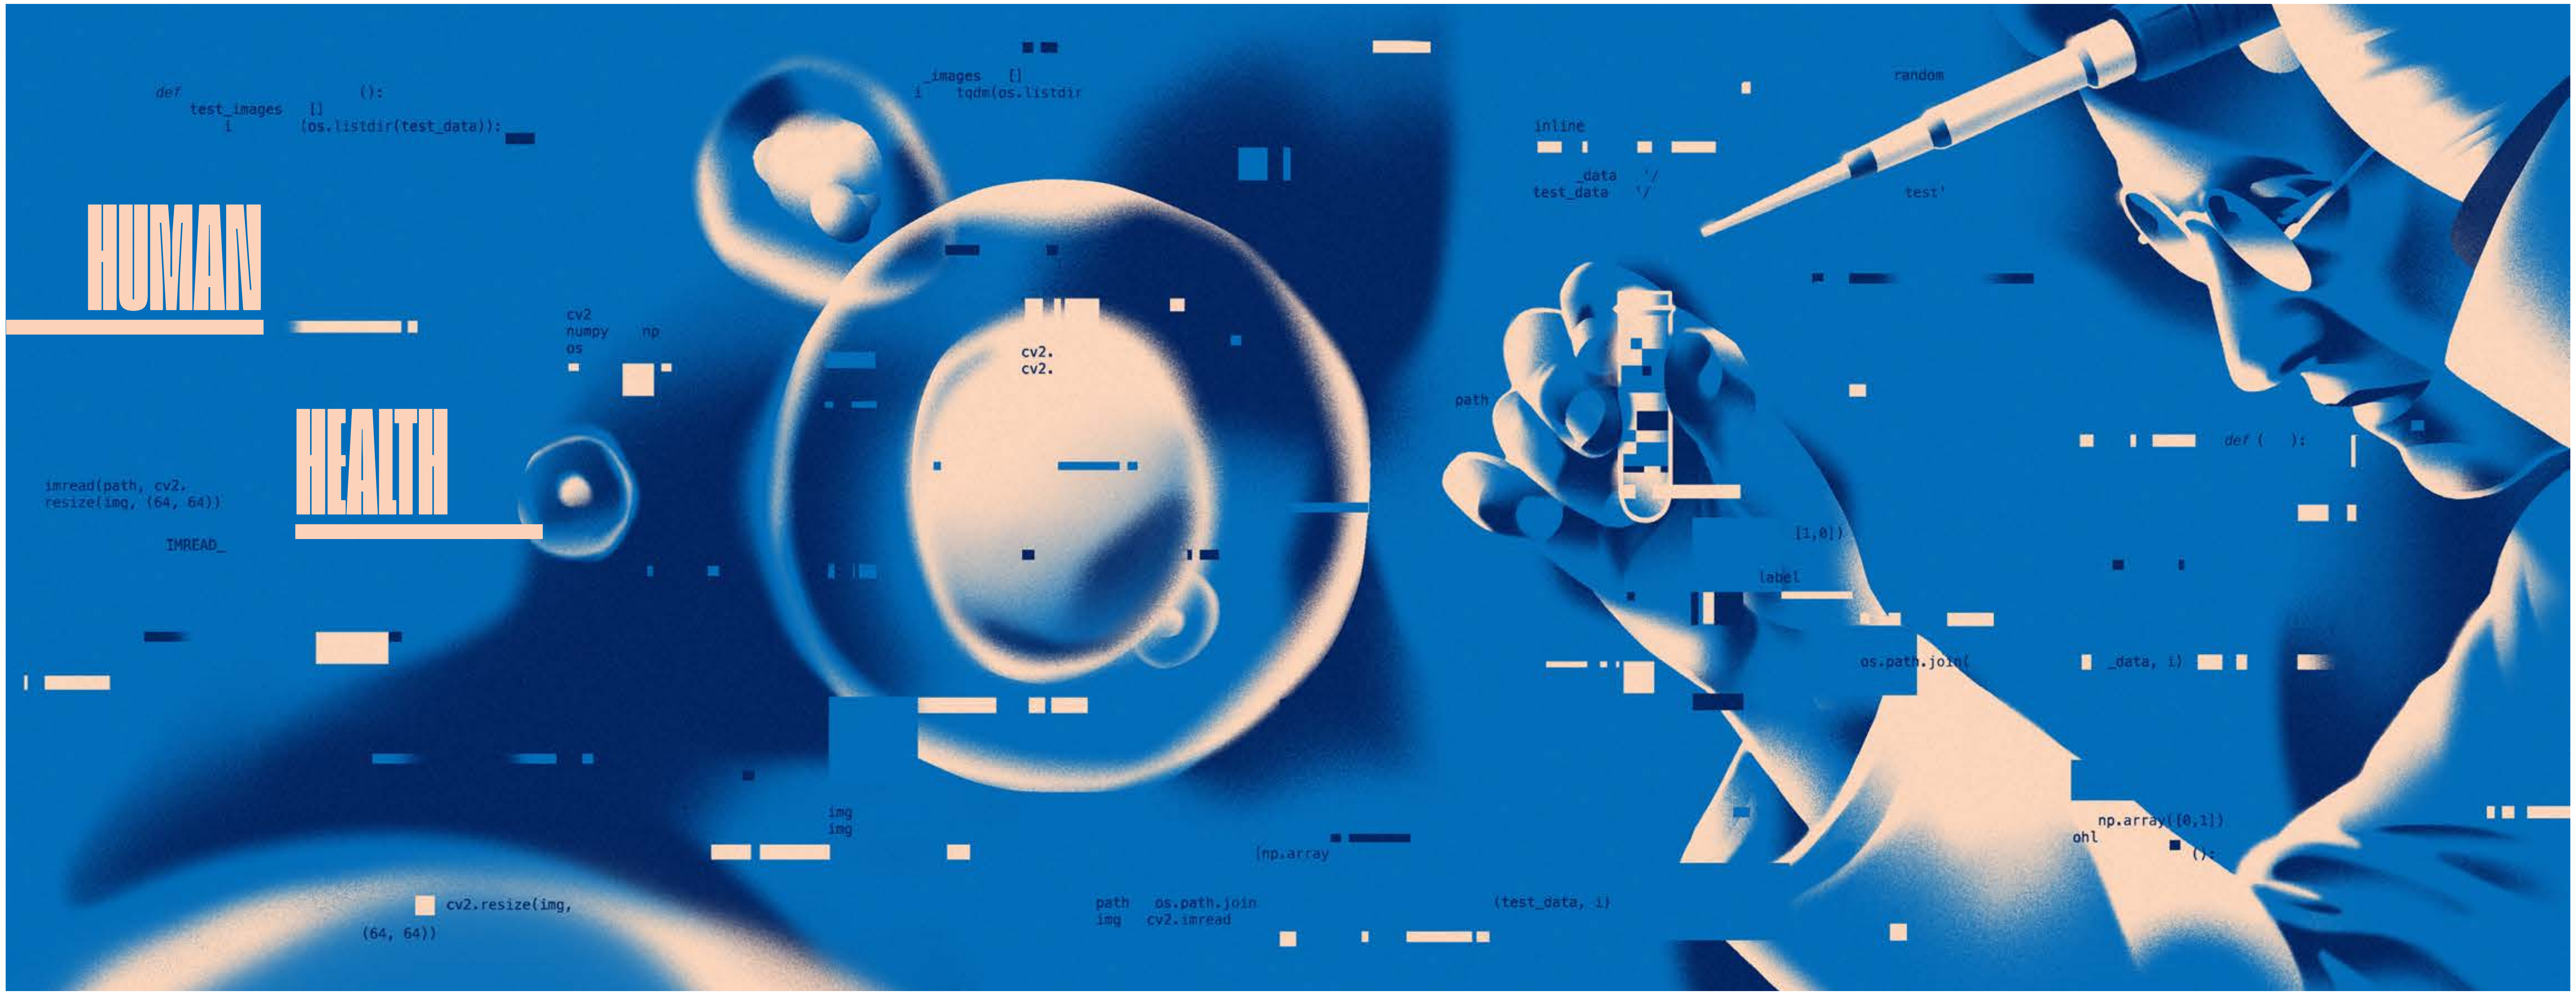

HUMAN HEALTH

Depending on whom you talk to, the number of possible chemicals in the world are anywhere from 10^{60} to 10^{80} — somewhere, in other words, between the number of atoms in our solar system and the number of atoms in the visible universe. Any one of those chemicals could hold the key to development of a new drug or other therapy. While human researchers often have a good intuition of where to explore, the sheer vastness of the space leaves ample room for AI to assist in discovering brand-new possibilities that humans might not have even considered.

"In drug discovery, you're going to have all of these ideas of what could be a good candidate molecule," says Connor Coley, an assistant professor of chemical engineering at MIT. "You know what could have good bioactivity, good potency, good solubility, and other properties, — but in order to test it, you actually have to make it." Creating drugs can be a complicated process, requiring multiple chemical reactions in a specific sequence and conditions, done at scale in order to synthesize the desired amount of drug to test.

Coley previously worked in the laboratory of Klavs Jensen, an MIT professor applying those principles to the synthesis and discovery of small molecules, which form the basis of many medicines. Using supervised learning techniques, he and his colleagues have used AI to simulate the testing of candidate molecules, varying catalysts, temperature, and concentration and predicting likely outcomes. "All of

↳ **RESEARCHERS ARE USING AI TO SPEED UP DEVELOPMENT OF NEW DRUGS AND OTHER THERAPIES. DEPENDING ON WHOM YOU TALK TO, THE NUMBER OF POSSIBLE CHEMICALS IN THE WORLD ARE ANYWHERE FROM 10^{60} TO 10^{80} - SOMEWHERE, IN OTHER WORDS, BETWEEN THE NUMBER OF ATOMS IN OUR SOLAR SYSTEM AND THE NUMBER OF ATOMS IN THE VISIBLE UNIVERSE.**

those variables will have some effect on the yield, efficiency, and cost of running that reaction," says Coley. "These adaptive algorithms help you design experiments and tell you what are the most useful experiments to run."

More recently, Coley and his colleagues have been using supervised machine learning on a huge corpus of chemical reactions collated by Elsevier and other companies to help identify new molecules to test. "We're taking information chemists have been publishing for the last couple hundred years, and then trying to apply that to new molecules," Coley says. "Then we put out a full synthetic recipe about how you might be able to make it—everything you should buy, how you should combine them, and what steps you should take to produce the molecule of interest." More recently, Coley and his colleagues, including Jensen, have worked on creating a robotic platform[3] similar to Aspuru-Guziks' self-driving laboratory to at least partially automate that process. "The idea is to make the research and development process a little bit easier and a little bit faster."

Coley is a member of MIT's Machine Learning for Pharmaceutical Discovery and Synthesis Consortium, a group that includes both MIT researchers and representatives from 14 pharmaceutical and biotech companies to help better train algorithms to discover and create new drugs. The companies provide financial support as well as feedback from their own processes to the labs, which produce models and algorithms for use of members. "We benefit tremendously from understanding what are the real pain points in their workflows, and how well the existing solutions work to alleviate those pain points," Coley says. In return, such techniques could help companies to quantitatively determine which candidates for drug discovery they should pursue for testing.

Another consortium member, MIT computer science professor Regina Barzilay, recently led the use of deep learning techniques to discover a new antibiotic. Historically, antibiotics have been discovered by screening vast numbers of soil-dwelling microbes, or creating synthetic varieties based on those discoveries that she, James Collins (a Professor of Medical Engineering and Science at MIT), and colleagues wrote about in a paper published in *Cell* in February 2020[4]. As the number of new molecules have shrunk, however,

3. http://news.mit.edu/2019/automate-molecule-production-ai-0808
4. https://www.cell.com/cell/pdf/S0092-8674(20)30102-1.pdf

CONNOR COLEY
Assistant Professor of Chemical Engineering at MIT

→

THIS WORK HIGHLIGHTS THE SIGNIFICANT IMPACT THAT MACHINE LEARNING CAN HAVE ON EARLY ANTIBIOTIC DISCOVERY EFFORTS BY SIMULTANEOUSLY INCREASING THE ACCURACY RATE OF LEAD COMPOUND IDENTIFICATION AND DECREASING THE COST OF SCREENING EFFORTS.

REGINA BARZILAY
Delta Electronics professor in the Department of Electrical Engineering at MIT

DALE THOMAS
Co-founder of mytide therapeutics

pharmaceutical companies have turned to screening large chemical libraries in search of novel drugs. Barzilay's team trained a specialized deep neural network called a variational autoencoder on more than 2,000 molecules to determine what molecular features inhibited the growth of E.coli bacteria, then applied their model to libraries with more than 100 million chemicals, ranking the results.

This process resulted in the identification of a drug with the unwieldy name c-Jun N-terminal kinase inhibitor SU3327, which the researchers renamed Halicin. Not only was it effective against E.coli, but they also found that it inhibited other pathogens, including A. baumanniias, a bacterium that has developed antibiotic resistance and is one of the World Health Organization's top pathogens to target. In addition, they discovered two other antibiotics effective against antibiotic-resistant E.coli. "The machine learning model can explore, in silico, large chemical spaces that can be prohibitively expensive for traditional experimental approaches," Barzilay told MIT News[5]. "This work highlights the significant impact that machine learning can have on early antibiotic discovery efforts by simultaneously increasing the accuracy rate of lead compound identification and decreasing the cost of screening efforts," the researchers wrote.

Another veteran of Jensen's lab, Dale Thomas, has spun out a company to synthesize molecules for commercialization. Called MyTide, the startup focuses on peptides, short chains of amino acids including insulin and other hormones that drive reactions inside the body. Among other uses, they can help in regulating obesity, developing vaccines, and targeting cancer though personalized peptide therapeutics. "You can develop a peptide specifically tuned to a cancer within a patient's body and trigger an immune response that allows the body to recognize a tumor as something that is not natural."

With 20 different amino acids to choose from, however, the number of potential peptides that could be used in therapies is overwhelming. "There are over a trillion different peptides that can be manufactured—and that's an understatement," Thomas says. In addition, each amide bond between amino acids is unique, requiring different conditions to form, and with dangerous side reactions if you get it wrong. "You have to understand not only the conditions for linking one amino acid to another, which is 400 different types of reactions, but also how previous couplings affect the current coupling. If you do not have the correct data collection and ability to use in-line analytics, you don't have the ability to capture the data of this building of peptides."

In order to solve the problem, the company has built a synthesizing machine that can move amino acids and other reagents through a reactor at the rate of 3,000 reactions a day. As the reactions occur, the system collects second-by-second data, including traces of UV light absorbed at different wavelengths, to show the extent of coupling of molecules. "That enables us to leverage machine learning, because we are building high-quality datasets that feed into machine learning models," says Thomas. Using the output from those models, the machine operators can then adjust variables such as the type of reagents, reaction time, and temperature, to optimize molecule synthesis, as well as to create new molecules for new purposes.

"After we run 5,000 molecules through this platform, we have 150,000 data points on the output. And we can feed that back in, in this closed loop manner, in order to redefine the next set of peptides that we should manufacture, and then carry through again to find the right target." Using these techniques, the company has been able to reduce the time required to synthesize non-naturally occurring peptides from nine weeks to four days — and use less reagents overall. "Our company's real aim is to

5. http://news.mit.edu/2020/artificial-intelligence-identifies-new-antibiotic-0220

HOW AI IS PERSONALIZING MEDICINE

Doctors are already starting to use AI in limited ways as a virtual medical assistant with special training for pattern recognition. Radiology has been an early pioneer in this area, using AI's ability to recognize complex patterns in images.[6] The FDA has already approved more than 50 algorithms[7] for use with AI platforms to help detect early signs of lung, breast, and brain cancer. In Uganda, an AI Lab at Makerere University uses a smartphone app and a microscope to detect malaria in blood samples, compensating for a lack of trained pathologists in the country.[8]

In the coming years, however, AI has the potential to transform every aspect of clinical diagnosis and treatment. "AI is not the decision-maker," says Finale Doshi-Velez, assistant professor of computer science at Harvard, who has pioneered such efforts, "but it can be the tool that provides the information and potentially the recommendations or anti-recommendations based on a patient's history to help doctors make decisions under conditions of uncertainty."

With new legislation on Electronic Medical Records (EMR) under consideration to allow patients better share their medical information digitally, new apps could help doctors sift through a patient's histories quickly in order to arrive at more effective diagnoses and treatments. "You have a lot of data about the patient in the health record but the doctor just doesn't have the time to look at it," says Doshi-Velez. An app could flag, for example, when a cardiac medication a doctor prescribes might produce unforeseen side-effects. "Errors may get made simply because the doctor didn't know enough about the patient."

Doshi-Velez is currently working on more advanced algorithms to personalize treatment modalities for treating diseases such as HIV, which relies on complex cocktails of drugs given in individualized combinations, dosages, and sequences. Using data on efficacy in past patients, an algorithm could analyze a new patient based on the patient's history, including resistance mutations in the virus and previously prescribed drugs, to recommend an optimal cocktail for that patient. "The trick is to look at the data and then propose alternate strategies that have not been tested, and be able to assess how well you think your new proposal will do," says Doshi-Velez.

Not all conditions may require such highly individualized treatments, however. Doshi-Velez' lab is also working with algorithms to produce more generalized treatment protocols using past data to sort patients into categories based on the symptoms they present. "My hunch is that a lot of times that stratification is going to get us a long way; if you could split people among the right variables, then you could put them into one of 300 bins, or 10 bins," Doshi-Velez says. Once the AI has helped create those bins, she says, then humans could help determine which treatment bins patients fall into. In an ICU setting, AI could even develop protocols for treatment that might fit on just a few pieces of paper, differentiating patients based on temperature, blood pressure, and other symptoms. "They could say, whenever you're making this decision ask a patient about these five variables and follow the flow chart."

6. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6268174
7. https://www.acrdsi.org/DSI-Services/FDA-Cleared-AI-Algorithms
8. https://www.cnn.com/2018/12/14/health/ugandas-first-ai-lab-develops-malaria-detection-app-intl/index.html

start building in more complex chemistries than have previously been able to be manufactured," Thomas says. For example, the company is experimenting with building using D-amino acids, mirror images of naturally occurring L-amino acids that show increased bioactivity. Other projects include adding polymer side chains to peptides to create cardiac medicines and oils to treat eczema.

In addition to drug development, researchers are using AI to create other new therapies as well. Cellino, one of The Engine's portfolio companies, is developing cell and tissue therapies by guiding the development of stem cells using lasers. Induced pluripotent stem cells (iPSCs) are grown on a well plate with a 20-nanometer undercoating of titanium. A laser shot at the bottom of the plate can target an individual cell, poking a tiny bubble in the cell membrane through which a liquid-borne cargo of growth factors or other components can enter, or — if the bubble is big enough — killing the cell outright. Using

FINALE DOSHI-VELEZ
John L. Loeb associate professor in Computer Science at the Harvard Paulson School of Engineering and Applied Sciences

MARINNA MADRID
Co-founder of Cellino

ARNALDO PEREIRA
Principal Machine Learning Engineer at Cellino

these controls, company scientists can precisely control the differentiation of cells into therapeutic cell types to treat various disorders.

"We are positioning ourselves as a cell and tissue foundry," says Marinna Madrid, cofounder of Cellino. "But what's unique about our technology is that we're able to do that with single-cell control because of the laser-based technology we use." The company hopes to partner with pharma and biotech companies that are differentiating cells for therapeutic applications but are having difficulty in achieving uniform, high-quality cells for implantation.

One therapy the company has been working on, for example, is the implantation of retinal pigment epithelial cells in order to treat age-related macular degeneration (AMD). As it grows these cells, the company has worked with biologists to identify which cells or regions of tissue to keep or enhance, killing the others with the laser. "There are a few things that correlate with the cells being good at doing their job in the body," says chief science officer Arnaldo Pereira, "for example, the amount of melanin granules in them, or how easily they can phagocytose the garbage that the photoreceptors sitting atop them slough off."

Ordinarily, it would take a Ph.D-level biologist painstaking hours with a microscope to identify these cells. Using AI, however, Pereira and his team have been able to design an algorithm on microscope images to quickly identify which cells to save. "If you can remove the underperforming cells and leave space for good cells to grow in that space, then you can increase your yield, and also increase the variability of cells that you get in the end," Pereira says. The company is working on a way to totally automate the process, including cell detection and lysing with the laser in order to streamline the process.

The same technique could be used for other cell therapies, for example, transplanting dopaminergic neurons to treat Parkinson's disease. Some studies have shown that transplanting such cells at an earlier progenitor stage produces better outcomes, but it can be difficult to determine at that stage which cells will succeed in fully differentiating into mature cells. "We plan to develop machine learning algorithms that can predict at earlier and earlier times whether or not the cell is going to be functional down the road," Madrid says.

In addition to just killing underperforming cells, the process is also able to manipulate cells by controlling the strength of the laser to poke just a small hole, through which liquid-borne cargo can enter, without killing the cell. The process could be used to insert growth factors or gene modulators to turn on or off genes within cells to individually control their growth with a high degree of precision. "So if you see that a cell is starting to differentiate down an incorrect pathway, but it's not completely lost, you could deliver a gene modulator into that cell to turn the right genes on or off to kind of course correct," Madrid says.+

IN

addition to applying AI in the energy industry, a burgeoning new field is the implementation of machine learning in materials production, similar to the way that it has been used for drug design. Instead of clinical effects of molecules, however, scientists are looking to create materials with specific chemical properties. MIT professor of chemical engineer, Tonio Buonassisi, for example, has increasingly devoted resources in his lab towards creating materials for use in solar panels and environmentally sustainable applications. "I think of machine learning as a cognitive assistant that helps the researcher do their job more efficiently and faster."

In 2018, he started a program in Singapore through the Singapore MIT Alliance for Research and Technology (SMART) with $18 million (USD) in funding from the Advanced Manufacturing & Engineering (AME) domain of the National Research Foundation (NRF) of Singapore, specifically to use AI to speed up the creation of solar cells and other environmental materials. Much of the challenge in producing more efficient solar cells lies in optimizing several variables at once. "It's not enough that it absorbs a lot of light, it's not enough that it has the ability to extract a lot of carriers, it's not enough that it's cheap, it has to be all of those things at the same time — and more," Buonissisi says. "In mathematics, we call it a heavily constrained optimization problem."

The human brain isn't good at keeping track of so many different variables at once, so machine learning can help solve those complex problems. "It's a matter of bookkeeping," Buonassisi says. "Our brains can't really handle these very high dimensional spaces, very well, but the machine learning algorithms." In order to test the design of the experiment, Buonissisi applied it to a "toy problem," which sought to tune the size and shape of silver nanoparticles so that they would change to a desired color at a certain temperature. "With only 100 points of training data, we could predict this output parameter within just a few degrees Celsius," he says. "It was a beautiful case study of creating materials with user-defined properties that nobody had ever made before. The algorithm told us what to make."

For their next experiment, Buonissisi's lab created a convolutional neural network that could predict stable crystals made from a set of 28 precursor solutions that could be used in a particular type of solar cell. "It is a very slow technique," Buonissisi says. "It took us three to five hours per sample to fire the data and for the human to analyze it. And we developed a machine-learning algorithm that cut that time down to five minutes per sample." In the end, the process discovered two new materials, including novel alloy, and four crystals which had only been grown in bulk before, but could now be spread in a thin film. "In the old way of doing science, this would've been six independent research papers that each would have taken three or four months to produce, but we're doing it here in two months. The rate at which you are learning is increasing thanks to the availability of these tools."

The crystal films might be useful for optoelectronic applications, such as light-emitting diodes, Buonissisi says. In order to create more efficient materials for solar panels, however, you'd need to optimize not just for stability but for long-range carrier transport as well. "To really change the world, you need to move from single objective optimization to multi objective simultaneous optimization and that's where some of our experiments are going at the moment," he says. "Now we are starting to push stability and performance at the same time." The AI zooms in on the region of multi-variable parameter space that's most promising and then churns through samples in that region and does the final optimization from there. That emergent toolset should knock on wood, yield many materials with user-defined properties, and hopefully a lot of them will succeed in addressing climate needs and sustainability needs within the next two to five years.

On a long term basis, the process might even be able to create mass-customized materials, for example, making solar materials that would work irrespective of the different locations. "We know, for example, that a material with a given band gap would perform better in Singapore than it would in Boston because the temperature and humidity are different," Buonissisi says. "If you had the power to tailor a solar cell for one region versus another, you could squeeze out a few extra percent energy yield, which is a really big deal in terms of profitability for companies.

Another organization tackling similar problems of designing new materials to solve climate challenges is Mission Innovation, a worldwide nonprofit funded in part by Bill Gates to help spur research towards the goals of the Paris Climate Accords. Among its eight "challenges" are such fields as batteries, smart grids, and biofuels, as well as one on Clean Energy Materials[9] specifically dedicated to using AI to create new environment-friendly materials. "Alan was the brainchild behind it," says Mark Kozdras, a program manager at Natural Resources Canada and the co-lead of the challenge. After proposing the initiative, Aspuru-Guzik helped lead the first international conference dedicated to the concept in Mexico City in September 2017.[10]

Since then, Kozdras has led the effort with an initial $10.8 million grant from the Canadian government and research support from several Canadian universities to create a self-driving laboratory, specifically engineered for sustainable materials called a Materials Acceleration Platform, or MAP. One of its most successful efforts to date is a robotic platform created by researchers at the University of British Columbia (UBC) along with Aspuru-Guzik called Ada, which uses machine learning to design and synthesize thin-film materials. In a video on the Project Ada website,[11] a white robotic arm swivels around a platform, selecting test-tubes full of reagents, twirling them in a spin coater, and pipetting them to combine on a well plate. As a proof of concept, the UBC lab showed how they could use the platform to create a new hole transport material, a component of certain types of solar cells, optimizing it on two parameters: stability and connectivity. In a working paper published in March of this year, the team described how the AI platform shrunk the number of experiments needed to discover the material from 1,000 to 30.

"They're taking experiments that could take a year and doing them in a week," Kozdras says. "and with a 10th of the chemicals." The Canadian government has now dedicated in excess of $28 million to build a new two-floor laboratory dedicated specifically to MAPs that can solve a variety of chemical problems. In addition to photovoltaic materials, Kozdras says, the lab is planning to explore batteries and thermoelectric materials that could transform waste heat to energy. Another candidate is concrete, which is similarly created by mixing materials that vary in composition from place to place. An AI brain could optimize production for different areas based on their unique raw materials. "If you could increase the strength of cement, then you could use less material," Kozdras says. "Even a 1 percent reduction could have a huge environmental impact."

Aspuru-Guzik also co-founded a company called Kebotix designed

↳

"THEY'RE TAKING EXPERIMENTS THAT COULD TAKE A YEAR AND DOING THEM IN A WEEK," KOZDRAS SAYS. "AND WITH A 10TH OF THE CHEMICALS."

THE UTILITY OF AI

Another major challenge that AI is beginning to tackle is the threat posed to our climate from global warming. Already AI is being used in "smart grids" that help better distribute the energy load to require less standing energy at power plants and in "smart buildings" to draw electricity for heating and cooling systems at off-peak times to put less pressure on the system as a whole. A new company, Urbint, is using a vast number of variables to help utilities cut down on direct leakage of greenhouse gases into the atmosphere. Among other energy infrastructure, Urbint targets natural gas pipelines, which transport methane gas — 84 times more heat-trapping than carbon dioxide — and are prone to corrosion and breakage.

The country has over 2 million miles of natural gas pipelines and countless more local pipes, many over 100 years old. "Just last summer, a local utility replaced a gas pipe in Chicago that had been installed the year before Lincoln became president in 1859, and had leaked 30 times in the past 30 years," says Urbint founder and CEO Corey Capasso. A successful software company entrepreneur, Capasso, helped found Urbint two years ago to help companies manage the whole infrastructure using AI. Most of the threats to energy assets comes from outside, such as fallen trees from wildfires, corrosion from flooding, or a backhoe hitting an underground pipe. "If we could predict and pinpoint these very specific risks, then the utility could anticipate those risks and take action," Capasso says.

The company was launched in 2018 with $25 million in funding, mostly from Energy Impact Partners and National Grid Partners, and now serves 40 clients, including the 10 largest natural gas companies. For each company, Urbint gathers data on the location of assets and historical data on damages. "We get three or four rows of data from the utility, and then we add hundreds and hundreds of columns of our own to that," Capasso says. Among the data it includes are soil conditions, weather conditions, construction activity in the area. "We even realized that moon phase has something to do with corrosion," Capasso says. "In some areas, the moon affects the tides which affects underground aquifers, which can lead to increased corrosion in certain types of soil conditions."

The company then applies AI models to identify the patterns and sequences of events in real time that could lead to specific types of leaks or other incidents and then help the company incorporate that prediction into its workflow. "If we predict, 'Hey, this excavator is going to hit a gas pipe, the utility can drive on site and talk to the excavator in person and tell him why his job is extra-risky." So far, he says, the company calculates that it has been able to help large utilities reduce third-party damages by more than 30 percent. "The whole goal is to make a prediction and then take an intervention to reduce the risk," Carpasso says.

TONIO BUONASSISI
Professor of Mechanical Engineering at MIT

COREY CAPASSO
Founder and CEO of Urbint

MARK KOZDRAS
Program Manager at Natural Resources Canada

9. http://mission-innovation.net/our-work/innovation-challenges/clean-energy-materials/
10. http://mission-innovation.net/wp-content/uploads/2018/01/Mission-Innovation-IC6-Report-Materials-Acceleration-Platform-Jan-2018.pdf
11. http://www.projectada.ca/

specifically to use AI and self-driving laboratories to custom-design new materials based on previously undiscovered molecules. “Our tagline really is ‘Materials for tomorrow today,’” says Jill Becker, the company’s CEO, a Harvard Chemistry Ph.D and formerly CEO of the successful company Cambridge NanoTech. “What we want to be is the DuPont of the 21st century.” Currently, the Cambridge-based company employs 17 people and recently closed an $11.4 million Series A round.

Along with developing novel chemicals to suit specific needs of clients, Kebotix is consciously focused on environmental sustainability. “We have aligned ourselves with several of the United Nations sustainability goals, and are really interested in making chemicals and materials green,” says Becker, an irreverent entrepreneur taken to wearing a leather jacket onstage. “Like one day I’m envisioning the future where I can walk up to one of our self-driving labs and say, ‘Hey, Kebotix platform, make me a chemical that will make me a truly recyclable coffee cup and won’t leach out nasty shit that gives people stomach cancer.”

Along with neural networks and other techniques, the company uses a type of unsupervised AI for unstructured data called a generative area network (GAN), refined by Aspuru-Guzik, which can explore a space of hundreds of millions of potential chemical candidates in search of viable molecules. Shown the structure of a particular molecule, say ibuprofen, the GAN crawls the space, generating other molecules it thinks will behave similarly, then automatically challenging and rejecting them as it refines its choices. Alternatively, the GAN can explore the space between multiple molecules at once. Over video chat, Becker puts up a slide with an image of a woman wearing eye-shadow and another woman smoking a cigarette, along with a half-dozen images between them that show a graded combination of the two.

“We call this the road from aspirin to caffeine,” Becker says. “Aspirin’s on the left—that would be the lady with the cigarette—and caffeine is on the right—that’s the lady with the eye-shadow, and instead of a roadmap of pictures of women that this narcissistic society doesn’t need, it dreams up molecules in the latest space. Then we can go back in there and say, hey, I’ve never seen you before, can I make you? What else do you do? Can I patent you? What are your properties?”

So far, the company has both worked with clients as well as dreaming up and patenting its own unique molecules. Kebotix gave one company, for example, some 2,300 possible chemical compounds all ranked to meet certain properties, including boiling point, vapor pressure, and melt viscosity. “They liked four of them, and made three of them,” Becker says. “That saved them years and hundreds of thousands of dollars.” In its own explorations, the company searched through a dataset of 7 million chemicals in order to find a new electrochromic material. “So it should be transparent, and when you send a current through it, it should give off light, but not heat,” Becker says. The company has identified one such material it is in the process of patenting for smart windows, which can change color and transparency to conserve heat. “You can imagine the savings in terms of cost of electricity,” Becker says.+

JILL BECKER
CEO of Kebotix

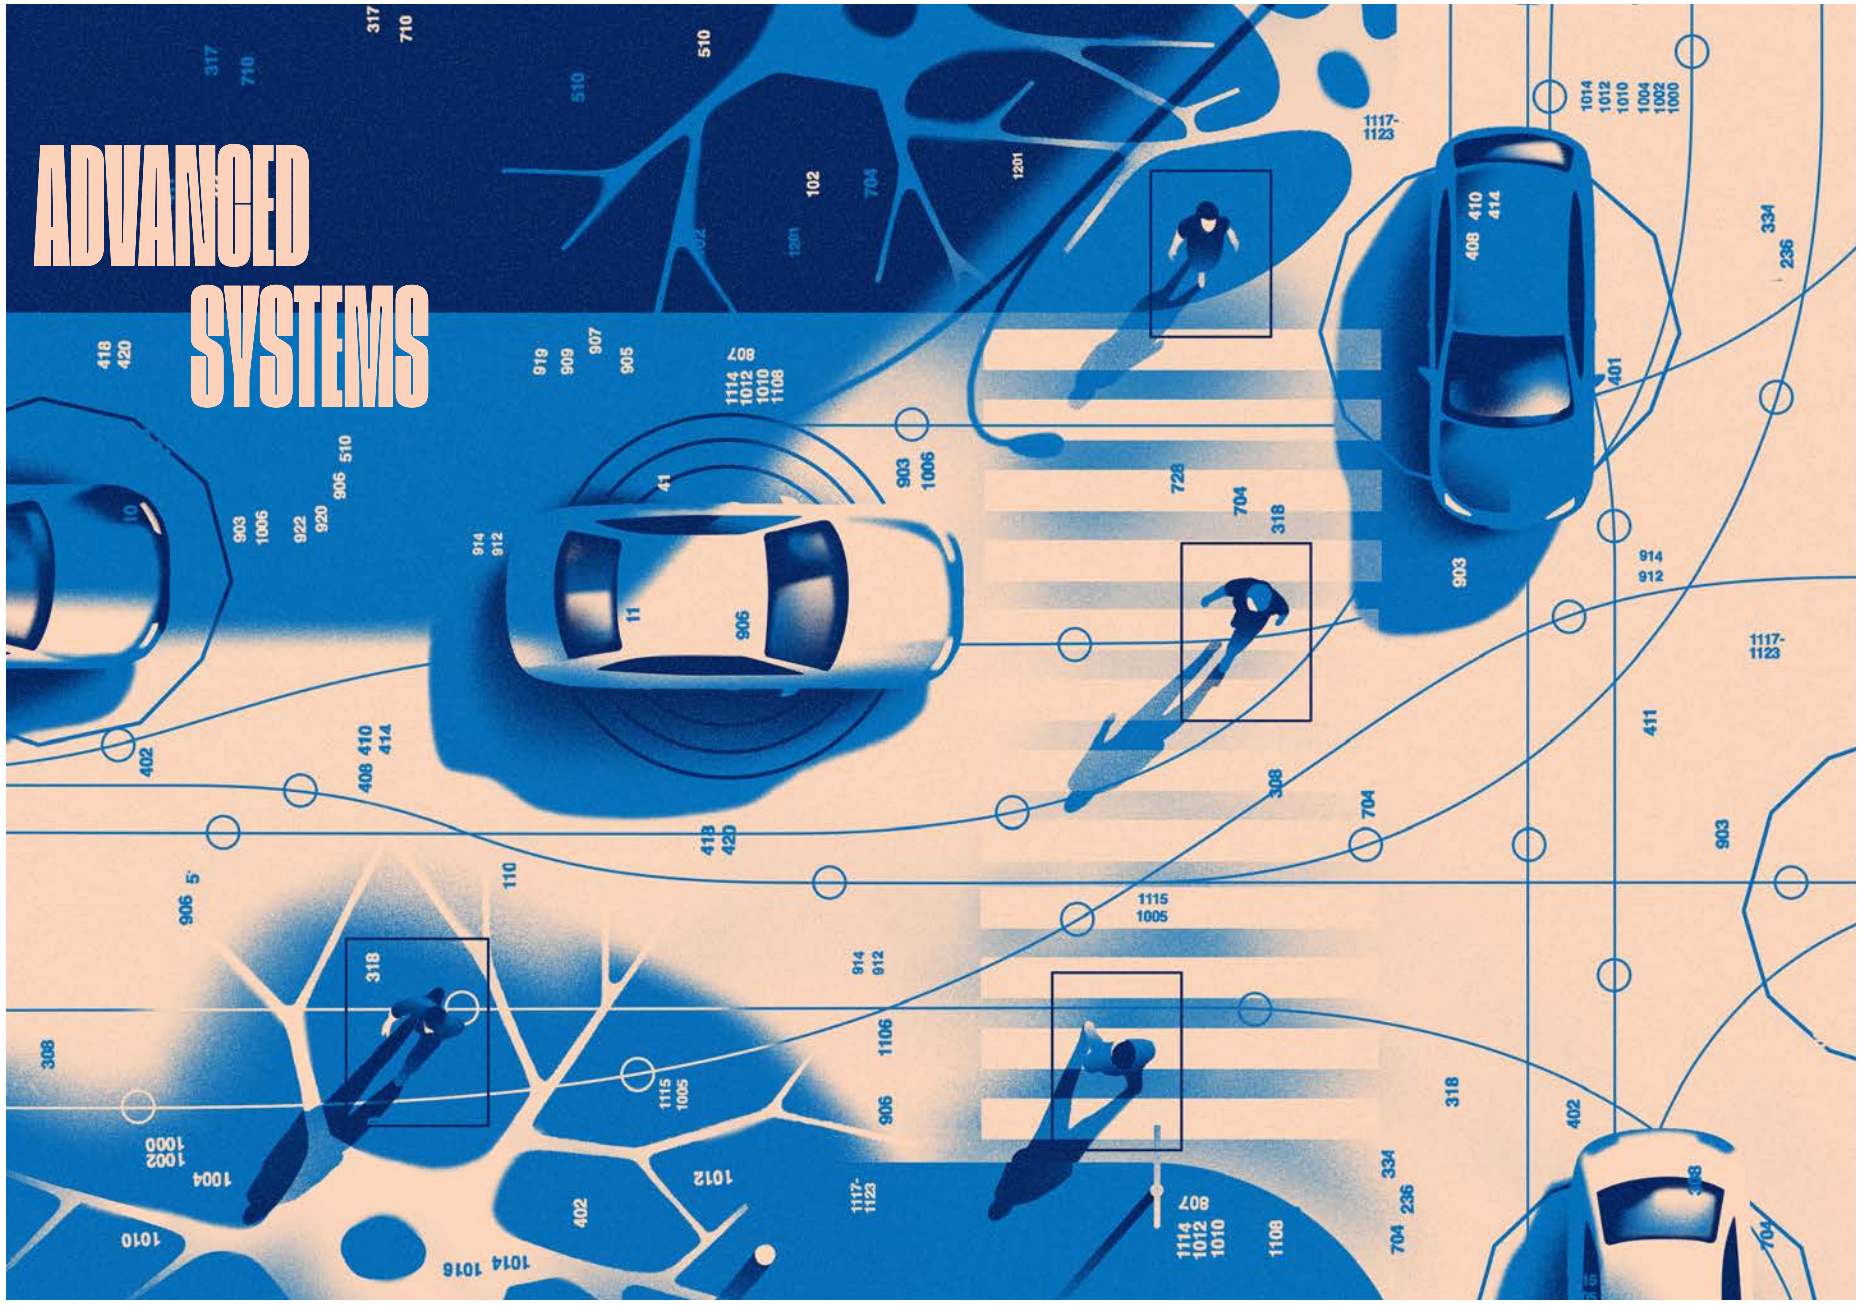

ADVANCED SYSTEMS

Artificial intelligence is using advanced computing to transform every other industry, so perhaps it was inevitable that AI would be used to optimize computing itself. Earlier this year, Google announced that it was using AI to help determine the optimal placement of components on chips for its Tensor Processing Unit, a processor designed to optimize Google’s AI program TensorFlow — in other words, using AI to develop a better chip for AI.[12] Ordinarily, human designers determine chip placement using a set of ingrained heuristics to reduce the amount of wire needed for hundreds of components.

In a paper released in March of this year, the company describes how it is using a different type of machine learning called deep reinforcement learning — the same kind of AI that Google used to create its AlphaGo program — to solve the problem. Instead of predicting results based on a pre-trained set of data, deep reinforcement learning picks solutions to problems based on how closely they meet a pre-defined “reward function” — in this case, a composite of different measures of performance. As it generates different solutions, the model uses positive and negative reinforcement to guide it towards the best solution.

While Google engineers acknowledge the difficulty of accurately programming such models, they have been able to use an algorithm to generate tens of thousands of possible solutions to chip placement, gradually honing in on the best solution. At the International Solid State Circuits Conference in San Francisco this February, Google’s head of AI, Jeff Dean, described how within 24 hours, one algorithm came up with a more efficient solution than designers who had spent up to eight weeks on the task.

Intel, meanwhile, has been using the principles of neural networks to design the very architecture of a new type of “neuromorphic” computer. Two years ago, the company released its Loihi neuromorphic research chip, which instead of having traditional architecture of parallel processors attached to memory, is arranged around artificial neurons and synapses that learn and store information in the same way as the brain. “This is about finding a new computer architecture that is inspired

12. https://www.zdnet.com/article/google-experiments-with-ai-to-design-its-in-house-computer-chips/; https://www.technologyreview.com/f/615419/google-ai-chip-design-reinforcement-learning/

JEFF DEAN
Lead of Google AI

MIKE DAVIES
Director of the Neuromorphic Computing Lab at Intel

CHRISTOPHER SAVOIE
CEO and Founder of Zapata Computing

YUDONG CAO
CTO of Zapata Computing

more directly from our modern understanding of neuroscience of how brains process and compute information, which is dramatically different than our standard way of computing," said Mike Davies during a press call in March to announce the release of a new neuromorphic research system, Pohoiki Springs,[13] which integrates 768 chips containing nearly 100 million artificial neurons — the number in the brain of a small mammal. Notably, the quirky system names come from geographical features on the Hawaiian Islands.

For its first application, Intel has trained an olfactory sensing system using machine learning to recognize smells. As an odor arrives, sensory neurons "spike" in a telltale pattern that is then compared to stored patterns based on training data. Unlike vision, which relies on three colors — red, green, and blue — smell integrates thousands of different olfactory components, making it a formidable problem. With only a single sample of a particular chemical, however, the system was able to outperform a deep learning AI system by 40%. It takes 3,000 samples, says Davies, for the deep learning system to catch up to the neuromorphic system, even though eventually he concedes that the traditional system surpassed the new design. The speed with which it can make judgments could be an advantage in certain applications in which a quick answer is important.

In addition, says Davies, the system uses significantly less energy than traditional computing systems, since unlike them, the neuromorphic system doesn't store a series of 1's and 0's. Individual neurons only spike with 1's, while the 0 state takes no energy to maintain. "Once you have an architecture that is actually exploiting the fact that one of these two binary states can save energy, then you can skew your whole algorithmic formulations to prefer zero states, and get gains that way," Davies says. "The downside is you have to rethink the algorithms because all of our algorithms don't have that property by default." In November, the company announced the formation of a Neuromorphic Research Community, with 500 corporate members, including Accenture, Airbus, GE, and Hitachi, to experiment with the system. The company hopes that in the future, it might show promise in other machine learning problems, such as image recognition and route optimization.

Other designers are using AI to help better design next-generation quantum computers, which use quantum superpositions to short-cut thousands of cumbersome computations to arrive at solutions to problems exponentially faster [see *Tough Tech* 01 "The Future is Quantum"]. While a variety of different techniques exist to create quantum computers, including superconducting circuits and trapped ions, they share a common difficulty: random errors that must be corrected to get an accurate result. Until quantum computers are able to surpass the power of conventional silicon computers, AI can help bridge the gap.

"There have been a few proposals about using machine learning to detect and classify the errors in these near-term devices," says Chris Savoie, CEO of Zapata, a quantum software company founded in part by Aspuru-Guzik. "AI could help correct those errors without even necessarily knowing the source of them. Even without correcting errors, machine learning could help improve the accuracy of quantum computers in the near-term," says Zapata chief technology officer Yudong Cao. "A quantum computer is essentially a sample generator," he says. "A machine-learning model could capture the distribution of samples and use that as a proxy for the actual distribution."

Among the applications that quantum computing could potentially revolutionize is artificial intelligence, using quantum's superior computing power for ML to come up with better solutions to optimization problems. In other words, using machine-learning techniques to optimize quantum would once again be like using AI to create better AI. No matter how "intelligent" artificial intelligence gets, the dystopian fantasies of sentient machines taking over the planet are likely to stay forever in the realm of science fiction. Instead, AI is likely to remain a powerful complement to human intelligence, allowing us to create new breakthroughs in health, climate, computing, and other industries, we could never dream up alone.+

→

FOR ITS FIRST APPLICATION, INTEL HAS TRAINED AN OLFACTORY SENSING SYSTEM USING MACHINE LEARNING TO RECOGNIZE SMELLS. AS AN ODOR ARRIVES, SENSORY NEURONS "SPIKE" IN A TELLTALE PATTERN THAT IS THEN COMPARED TO STORED PATTERNS BASED ON TRAINING DATA.

13. https://newsroom.intel.com/news/intel-scales-neuromorphic-research-system-100-million-neurons/

WHAT IF DRIVERLESS CARS COULD THINK LIKE A DRIVER?

Autonomous driving couldn't exist without artificial intelligence; the complicated nature of driving a car on a public street requires that a driving system is able to make split-second decisions about when to speed up, slow down, and turn to stay on the proper route and avoid obstacles. So far, however, the promise of autonomous driving has not lived up to the hype, with existing systems unable to show the control needed for real-life conditions on public roads.

The Engine portfolio company isee has a theory for why that is so – it's not the self-driving cars that are the problem, it's the other cars on the road. "If we could completely remove other humans, and just have the autonomous driver out on the streets, that should already be possible," says Yibao Zhao, founder and CEO. "But we can't move those humans, so we have to have an autonomous vehicle co-exist with other human-driven vehicles and pedestrians, and they are the most unpredictable part of the environment."

Zhao and his cofounder Chris Baker are implementing a different kind of AI based on their previous work at MIT's Computational Cognitive Science Lab, which focused more on making an artificial brain aware of the intelligence of others. "When humans make decisions, we have a lot of common sense knowledge that helps us to have a much better understanding of the environment in order to understand events that are happening," Zhao says. One of those senses, known as "theory of mind" allows us to put ourselves in the mind of another person to better understand what they are thinking and feeling and predict what they might do because of it.

"It's fundamentally a kind of probabilistic reasoning, where we realize we don't have certainty about the causes of events, or know what is going to happen next," says Baker. "But as humans uncertainty is pervasive, and yet we can make intelligent decisions and gauge risk and live healthy, successful lives." In essence, he says, that's because we are always building probabilistic models of the world, based on the possible causes and future effects.

The company is working to train autonomous driving brains not only to anticipate behavior of other drivers, but also to understand the causes of their behavior. "If we see someone change lanes, we ask why they did that," Zhao says. Are they getting ready to exit, for example, or is there something in the road. Answering that question could better determine whether we change lanes as well. It's important that causal reasoning be a part of this." Building a sophisticated model of the environment can help us understand uncertainty and reason about the 'what-if' in the future."

That's a dramatic departure from other autonomous driving companies, which are pouring their energies into creating better sensors, cameras, and handling in order to help cars better control their environment. By contrast, isee's focus on using deep learning to teach its autonomous autos "common sense" helps them better understand their environment, and anticipate and adapt to changing circumstances on the road.

After raising a first round of funding from The Engine and a second round last year of $15 million from Founder's Fund, isee has been able to design an artificial brain with enough common sense to be able to safely put cars onto public roads, hauling freight in Texas and California. "If we can move freight more efficiently, and reduce two days to one, or one day to five hours, that's a huge benefit to customers," Zhao says. "The supply chain is full of opportunities that we can use AI to optimize."

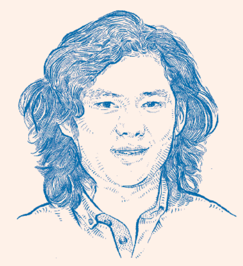

YIBAO ZHAO
Co-Founder and CEO of isee

CHRIS BAKER
Co-Founder and Chief Scientist of isee

—Real Perspectives on Artificial Intelligence—

For extended interviews, visit medium.com/@the_engine

Texts edited for length and clarity
Illustrations by Harol Bustos

Rick Calle
Artificial Intelligence
BD @ M12

Rohit Prasad
VP & Head Scientist @ Alexa AI

Dan Huttenlocher
Dean @ MIT Schwarzman
College of Computing

Dayna Grayson
Founder @ Construct Capital

Daniela Rus
Professor Electrical Engineering
and Comp Sci & Director @ MIT CSAIL,
Deputy Dean @ MIT Schwarzman
College of Computing

Randy "Laz" Gordon, Colonel USAF
Director @ USAF/MIT Artificial
Intelligence Technology Accelerator

Interviewed 3.30.2020

DAN HUTTENLOCHER

Dan is the inaugural dean of the MIT Schwarzman College of Computing. Previously he helped found Cornell Tech, the digital technology oriented graduate school created by Cornell University in New York City, and served as its first Dean and Vice Provost.

■ *No matter how responsibly developed AI may be, its generality seems at once its greatest asset and its greatest danger. How do you reconcile this duality?*

■ I think it's a great question. Much of this — and I don't know if it's fear of AI, exactly — but much of the sense of the potential dangers of AI comes from a misunderstanding of what the technology really is.

There's often a tendency to anthropomorphize technology and with AI this tendency is much more extreme. In the end these are still just algorithms. When we're using machine learning, for example, they are algorithms that we can teach instead of algorithms that we have to code. But they're still algorithms. They're not going to become evil. There's no rational basis for that worry at the present time (I'm not saying it's impossible for that to be true in some future with technologies one cannot foresee today, but for the path we are on it is science fiction).

If these algorithms are used for something bad, it's because people decided to use them for something bad. They don't decide to do things that are bad on their own.

■ *I've heard similar sentiments expressed throughout the AI community. On one hand this sentience like Data from Star Trek or the Terminator robot is far fetched, but it's also so ingrained in popular culture, that it's hard to think about AI any other way.*

■ It's easy to anthropomorphize AI when you have an Alexa or Google Home or something similar in your house answering your questions. And just imagine when these tools become robots that you're conversing with — which won't be long, I'm sure. It's natural that people will then assume that these machines have the full range of human intelligence. People already anthropomorphize their Roombas.

How these intelligences will evolve, especially considering how machine learning has grown lately, poses some interesting questions. Many of these platforms are now learning from us rather than being programmed — they learn what we do, not what we aspire to do, not what we say we're going to do. They are not necessarily learning our best attributes, they are learning our actual selves. And that can be embarrassing.

■ *And these attributes can be interpreted as biases.*

■ Yes. Biases we may not even know we have. Or biases we do know we have some of which we legislate against. So when you start observing behaviors and decision making behaviors by humans to train machine learning, there's all kinds of bias there, which may be undesirable.

Once we recognize that these platforms are surfacing biases then we can get more systematic in combating those biases. It's what good actors are doing now. And what legislation will eventually require. I think we're in a time period now where people are rightly agitated — and that's good because it's highlighting things that need attention in the world, not only attention in AI. I think AI will continue to evolve in a positive way — we'll end up in a much better place than today.

■ *I like the optimism! A bit of a pivot here — will AI, like calculus or geometry, become a prerequisite for every engineering program?*

■ Yes, I think so. And I would broaden this beyond AI to encompass "computing" in general. Even now it is something that every student needs to learn. As AI gets to be a more and more important part of computing, it too will become required. MIT already has a large percentage of undergraduates studying machine learning. In many senses, computing is already a de facto requirement.

■ *So much has changed since your time as a student in MIT's AI Lab. What's surprised you the most since you left?*

■ The big surprise is how quickly AI became practical. As a grad student in the 1980s, we joked around that AI was the perfect research problem, because it would always look like it was almost ready. Now here we are, a few decades later, and AI is used everywhere.

■ *The final question: in 20 or 30 years, do you see any industries or sectors being AI holdouts? Those in which AI will be difficult to implement for technological or cultural reasons?*

■ For technological reasons? I can't think of anything. Maybe if there's some big AI backlash, maybe we'll see some sort of Neo Luddites who are specifically anti-AI, that is certainly not impossible.+

Interviewed 3.30.2020

RANDY "LAZ" GORDON

Colonel Gordon is the Director of the USAF/MIT Artificial Intelligence Technology Accelerator. His team is building a unique pipeline of AI technologies to help give the U.S. a competitive advantage in AI capability breakthroughs. He is also a Presidential Fellow, Harvard Business School Alumnus, DARPA Fellow, and Massachusetts Institute of Technology Fellow. He served the F-22 Combined Test Force as its commander. He has experience as a combat officer and pilot, with experience in over 74 military and civilian aircraft.

■ *Should we remain optimistic about a future increasingly reliant on AI?*

■ When I came into this, I was very much focused on the notion of "AI as a savior". Whatever problem you have, whatever it is you're trying to do, there's an AI for that. But ultimately it's a reflection of we as humans and the society that we come from. So rather than getting hyper-focused on the technology of AI, what I've learned is that it's actually a behavioral science — it's a reflection of the society that we come from.

So should we be optimistic? Absolutely. In the United States, we have a society built on some very fundamental values. During my time in the military, I traveled around to a lot of different places in the world where those fundamental values just aren't there. In those nations in which AI technology develops, it will be a reflection of those times, those places. The U.S. is by far, in my view, the place to be optimistic about such development.

■ *When you meet with senior Air Force leaders, how do you explain the reality of AI? To what degree do you have to temper expectations?*

■ I often have to temper expectations in my conversations with senior leadership — state of the art is nowhere near what you're thinking and nowhere near what you picture in science fiction. My team's been trying to teach people about the reality of AI versus the rhetoric of AI. And that gulf is pretty huge.

I think a lot of that is driven by just the name itself. Many people interpret "artificial intelligence" as trying to say that a machine has some level of consciousness, that we share the same level of understanding and appreciation. In reality, the technology is essentially just very good statistical modeling. There is no "understanding."

AI is subject to all the biases that we as humans ultimately feed it. It is subject to the quality of the data and then the construction of the algorithm itself. So it's extraordinarily brittle.

In the Air Force we tend to think if there's a problem "I'm going to throw an AI against it and AI is going to fix everything". Then you have a deeper conversation about quantity and quality of data. Is that data shareable and releasable? And even if you have all those things, it's only really good at a narrowly defined statistical model.

It's not like it's going to show up and suddenly fix the entire Air Force personnel management system. It's not going to suddenly create artificial flying machines. There's a great difference between that rhetoric and the reality of where things are at.

■ *Even so, AI is ubiquitous and incredibly powerful. What do you see as the most significant breakthroughs in the evolution of AI in the past 20 years? And how are those breakthroughs impacting our lives today?*

■ The growth of compute power and sheer amount of data that's available now has created a weaponization of trust. What do you believe anymore?

When you see something online, our natural human tendency is to trust it because the engines that run those algorithms are designed to understand your preferences and continuously feed you more of the thing that you've already expressed a preference for liking. It's almost like a silent curator that's operating behind the scenes. A secondary effect of this is that it splits people such that we don't even have agreement on the same level of facts.

Compounding that, when people see something written, they think that that comes from another human being when in reality it could have been created by an AI. Add in the power of generative adversarial networks, and you can create realistic images of people that don't actually exist. How do you believe what you are seeing and how does it not become a contagion? This is a psychological dimension of technology that we haven't really had to deal with as a society yet.

■ *So how do we combat this "weaponization of trust"? What should the US and its academic institutions be doing to ensure not only an enduring competitive advantage in AI, but also responsible development of the technology?*

■ When it comes to any emerging tech that rapidly changes — like what we're seeing with respect to AI — there really is no roadmap because you're at the tip of the spear. The only way to ensure lasting success is to have open, direct dialogs between disciplines and sectors. That's something that's very unique to the United States that will keep us on the forefront.

Take the current COVID-19 pandemic, for example. If there's a means with which to use the statistical power of AI and machine learning to better diagnose the disease, to better understand where breakouts are occurring, to better deliver telemedicine — that's something you'll see developed here. And when we do that — when we lead with our values — then that becomes an example to the rest of the world. +

Interviewed 2.19.2020

DAYNA GRAYSON

Dayna is an expert in industrial software, automated systems, and additive manufacturing with nearly 15 years of experience as an early-stage investor. Her firm, Construct Capital, funds early-stage companies enabling massive new markets through breakthrough technologies. Previously, she spent seven years at NEA as an investing partner. She is currently a director at Desktop Metal, Tulip, and others.

■ *You've been involved in software and systems engineering for over two decades. What's contributed most to the progression of AI during that time?*

■ The accumulation of data. We have access to so much data now. And it's growing exponentially in the industrial space as more systems are instrumented for measurement and monitoring, in particular.

■ *Where is this data best put to use?*

■ I think the next immediate application of this data will be within logistics space and supply chain. There is so much data to learn from in the warehouse alone. Then there's the entire autonomous systems sector and how that can leverage this data to automate functions in the warehouse. Many of us think only of autonomous driving, but the more interesting applications are within industrial robotics, UAVs, minor navigation tasks, and more.

And of course there's manufacturing, which is just a huge, huge sector that hasn't really changed since the industrial revolution. Manufacturing is just being digitized — less than 1% of the market has been digitized. We haven't even scratched the surface. The data that it produces will power the next big wave of innovation.

■ *Let's take manufacturing specifically. Where do you see the bulk of that innovation happening?*

■ Robotics and additive manufacturing. It's only been over the past five or six years in additive manufacturing, after the sector made it prime time, that the level of data that can be adjusted in that space made it right for AI implementation. There are some exciting advancements with generative design and hardware utilization, for example. We're also seeing AI being applied to robotics in traditional manufacturing environments, especially when it comes to human-machine collaboration and automating human-like skills — using AI to teach new bodies to behave like old bodies.

The CPG and other consumer brands that leverage this for vertical integration and decentralized (closer to the consumer) productions will be the leaders of the future. Very few brands fully grasp this yet — it's a huge opportunity, especially for startups disrupting incumbent brands and driving down costs.

■ *There are so many facets of additive manufacturing — from materials to hardware, processes, to design of the objects to be printed — where will AI play the most significant role?*

■ There will certainly be advancements related to precision and resolution thanks to AI, as well as advancements in the print materials themselves.

But generative design is the most exciting for me. The future of additive manufacturing will be not just printing new parts but architecting new parts. For example, you design a particular part the way you think it should be, and then a program can take that design, along with information about its intended application, and redesign it based on optimal strength, optimal function, and optimal fit. The part may come out looking like something you would never imagine to design yourself, but it is totally optimized for material savings, strength, and performance. This is where I think AI & ML will be the most powerful.

Couple that with faster and smaller production facilities closer to the customer and the whole manufacturing and supply chain can be reinvented. The cost of creating this and integrating it within a new products company from the start is surprisingly low.

■ *What's surprised you most about the evolution of AI?*

■ Everyone thought that AI adoption would be much faster than it's actually been. Many digital businesses exist without an AI underpinning. Even existing companies that we think are AI-powered really are not. Things like marketing platforms, sales software, and similar businesses — AI is gradually being layered in; it is not central to their business model yet.

■ *So we don't have to worry about the robots taking our jobs…tomorrow?*

■ I don't think we'll suddenly see a full, wholesale change. AI adoption is gradual — it will not happen overnight. We will have time to adapt. Just like the digital or industrial revolution, changes to the economy and jobs are slower than we all may think.+

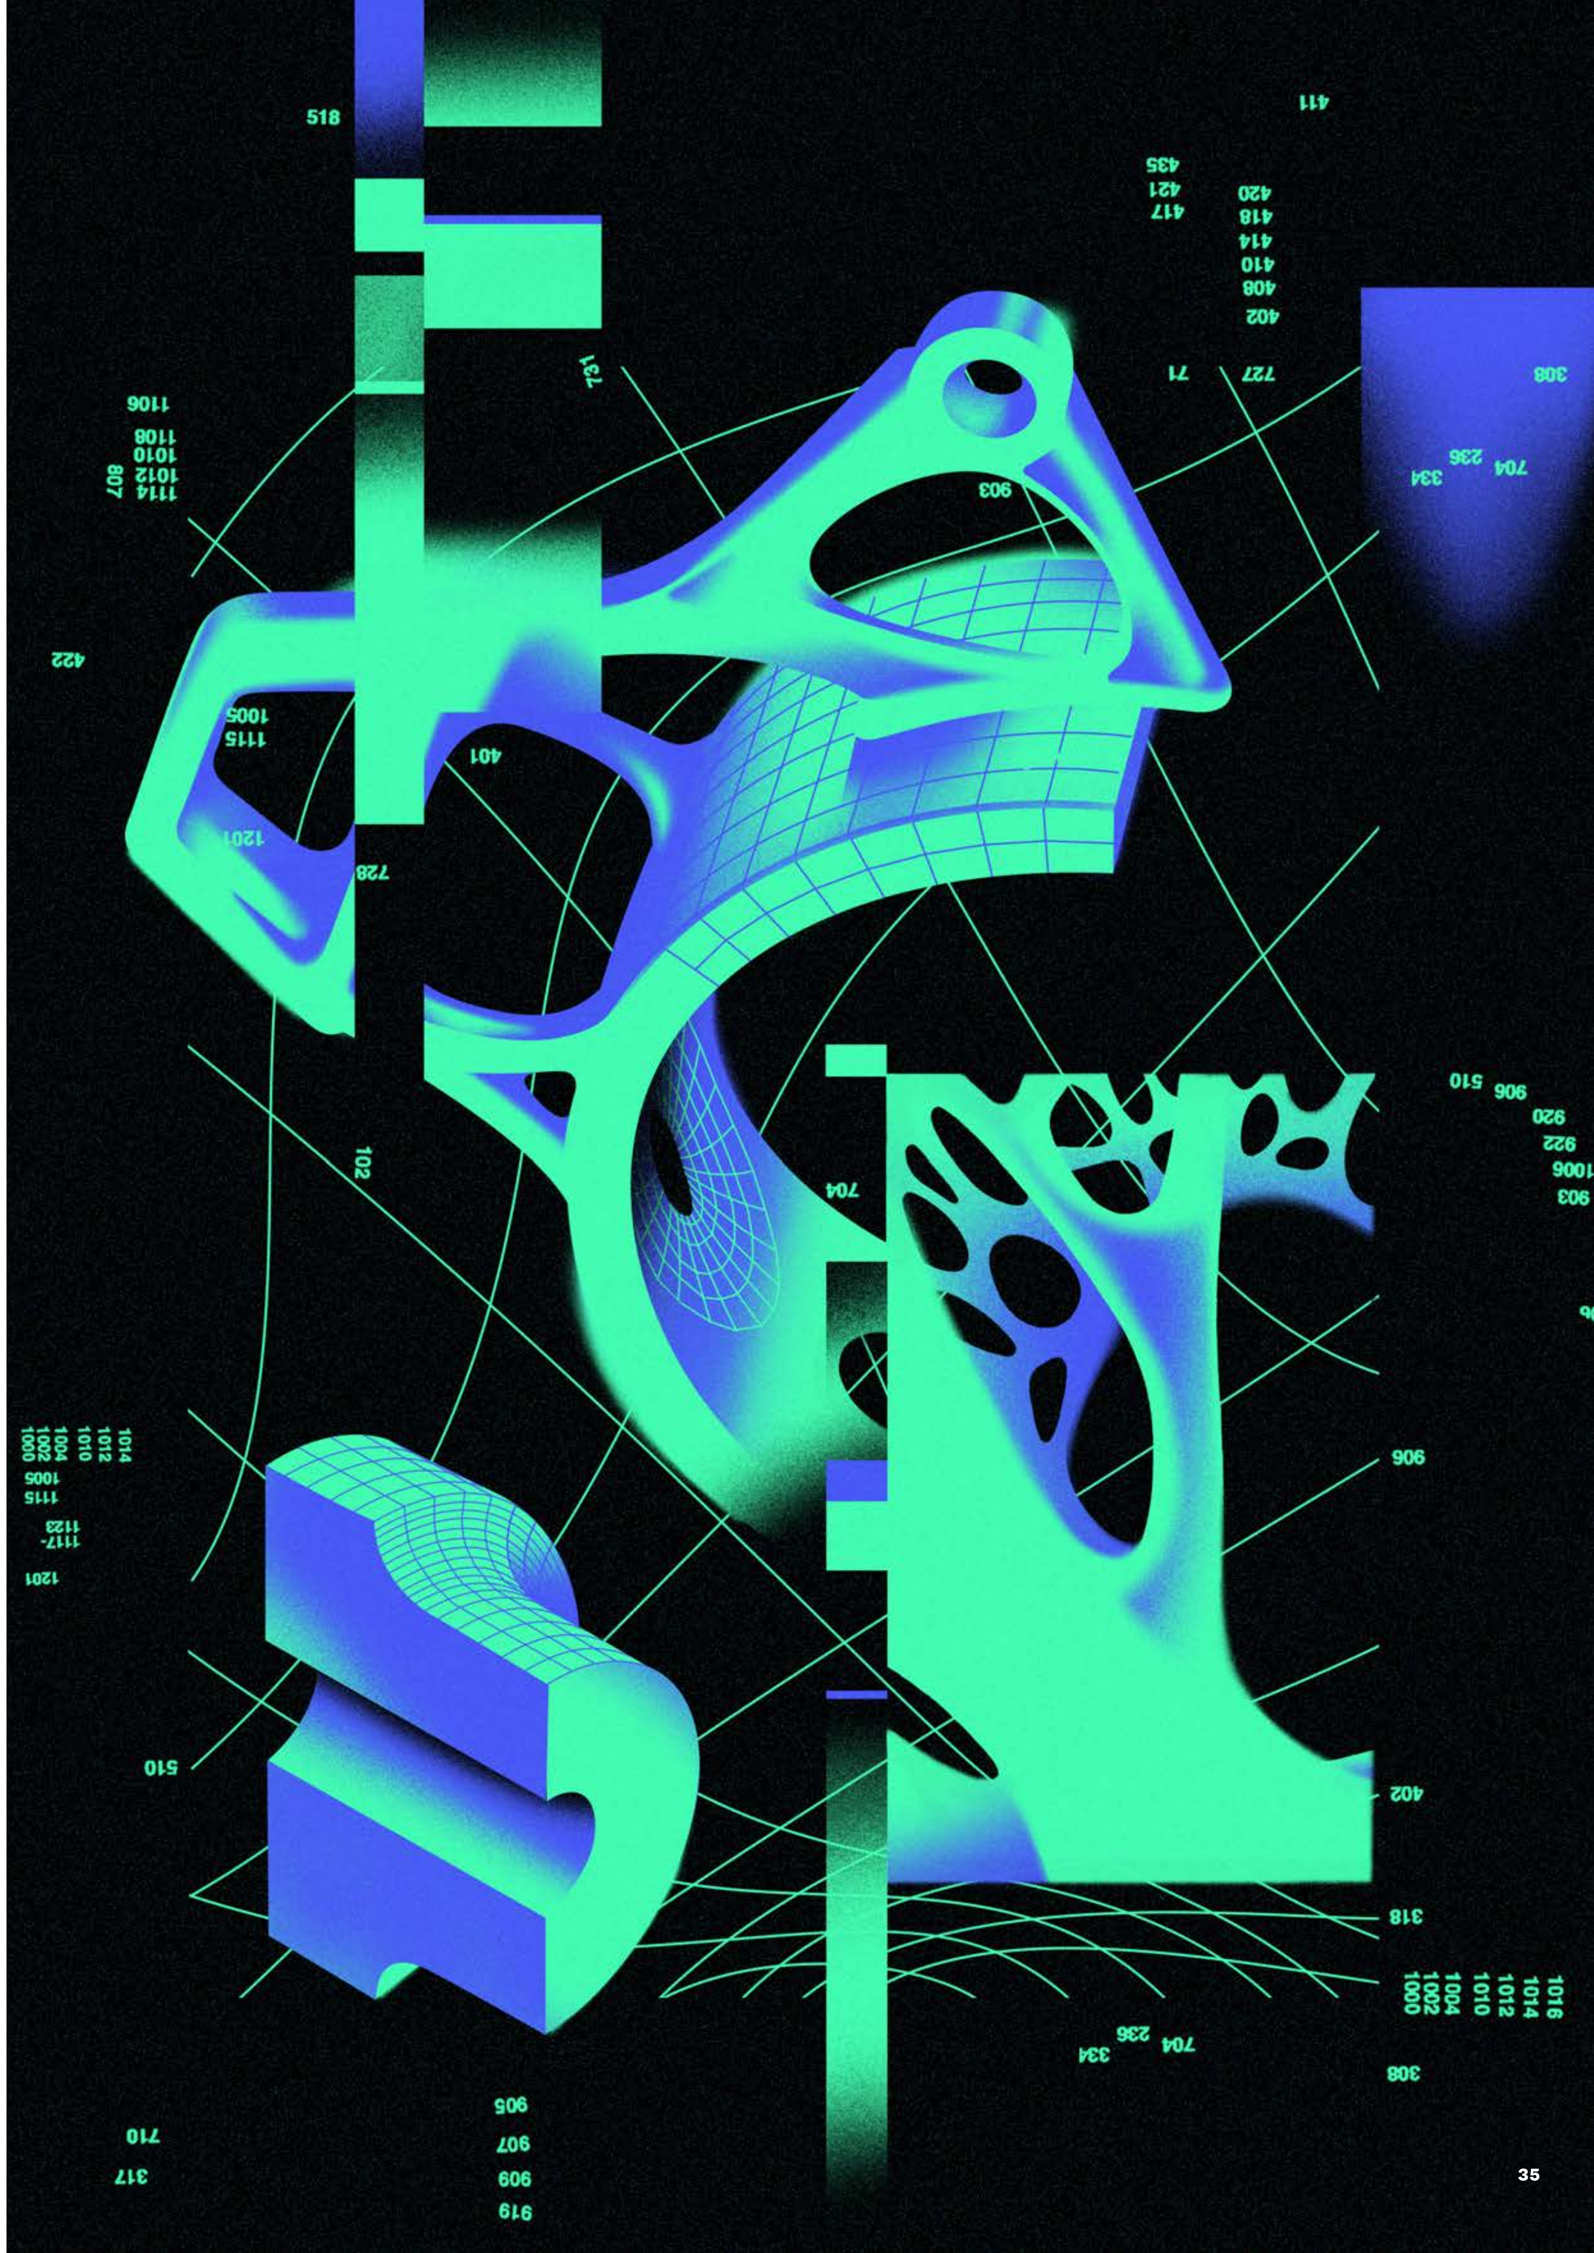

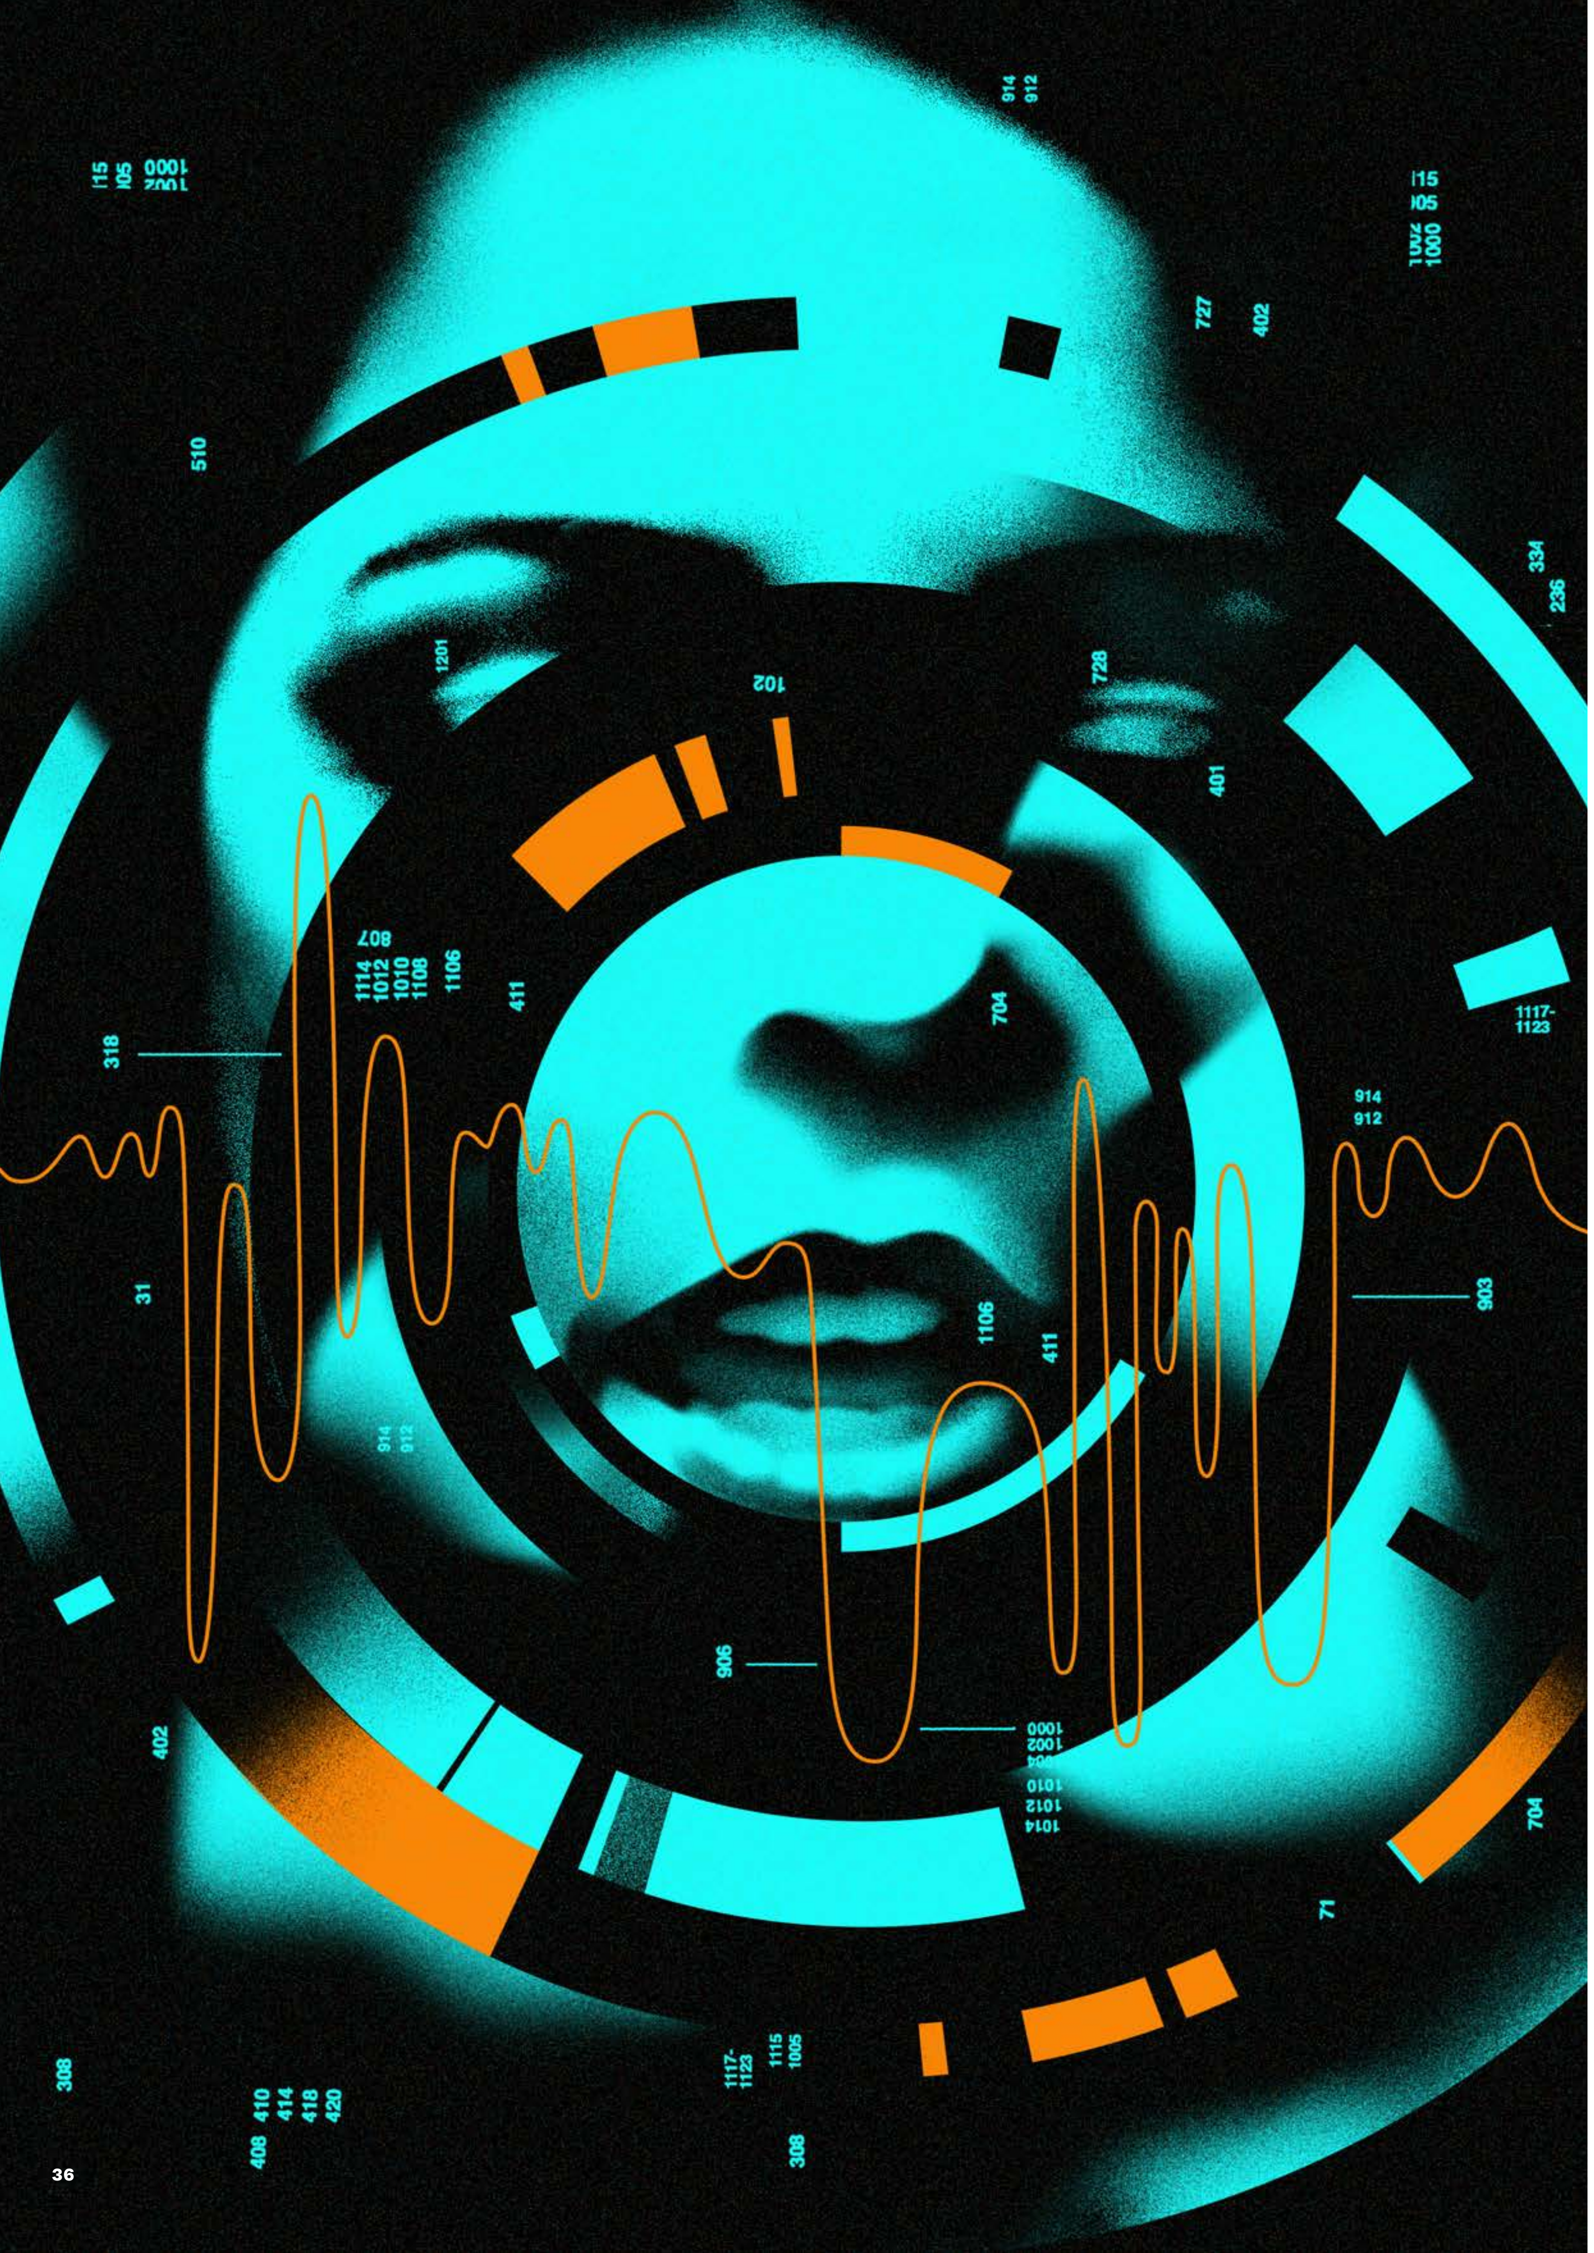

Interviewed 4.7.2020

ROHIT PRASAD

Rohit is the Vice President and Head Scientist for Alexa AI. He directs conversational AI R&D for Amazon Alexa, with a focus on natural language understanding, machine reasoning, and underlying machine learning capabilities.

■ *I recently read a note by President Reif of MIT regarding the school's new College of Computing. He stated that its students will be "navigating an algorithmic future." That is a provocative concept — how do you envision such a future?*

■ It's definitely a provocative statement — I agree with it. We are living in the early years of the golden age of AI. AI algorithms in hardware and software are becoming more and more indispensable and ubiquitous. They are part of the social fabric.

In some cases, I think algorithms will recede into the background where they assist with decision making. In other cases, I think algorithms will be embodied in physical form factors — self-driving cars, robots; these are manifestation of AI in the physical world. Alexa, as it currently stands, is an example that blends these two schools of thought.

The pervasive adoption of AI will be so profound that it will have a deep impact on all types of education. The barrier to programming AIs will also be lowered with the ongoing push in democratizing AI. Just a fascinating future from a student perspective.

To close on this — being in this field for a long time — I've noticed the conversation in industry is completely shifting. Leaders in enterprise are not asking the questions: what is AI? Is it real? Does it have any benefits? They're asking: how do I apply AI right away? Or how can I go faster with AI? It is just inconceivable for me to imagine a future where AI will not touch every person in the world.

■ *So have we closed the perception gap — the gap between the capabilities of AI and the perception of those capabilities?*

■ That gap has closed so much in the last decade, especially in the last five years. This isn't an issue we should worry about anymore. Take Alexa as an example — if there was a gap between perceived capabilities and actual abilities, then customers wouldn't use it.

Five years back when we were working on launching Echo and Alexa, even the academic community was extremely skeptical about the possibility of far-field speech recognition in a mainstream consumer product. There were lots of studies in labs, but to have a product that millions of customers with different accents and noisy households could interact with from a distance...that was unheard of.

So why did this change? We've all dreamt about the Star Trek computer, and we've all thought about making it a reality for a long time. The expectations were always high. The reason we were able to bridge that gap of perception versus capabilities was because Alexa did a few things quite well early on, and then became more useful, more accurate, more conversational, and more natural in understanding and decision-making. Alexa is also striving to be a fast self-learner and is another reason for bridging the perception gap in human-machine interactions.

■ *I'm sure you're a little biased here, but which AI innovations excite you most?*

■ I get most excited by AI innovations that transform our daily lives in a positive way. I also cherish AI advances that make complex tasks simple.

I am super excited by AI's role in medicine; I'm very optimistic that AI will have a profound impact on personalized medicine. There are already early signs of that happening. Self-driving cars will be revolutionary, but it's still very early. And the integration of AI into robotics, in both industrial environments and other forms, will be transformational.

Of course, my favorite is conversational AI, which means you interact with the AI as you do with a human. It is, to me, the hardest AI challenge — you are seeking to model human intelligence and then transcend it by adding superhuman capabilities.

■ *As you noted, you've been in the industry a while. What's surprised you the most regarding the evolution of AI over those years?*

■ The first is Alexa related — if you asked me five years back, "In five years, will you see billions of customer interactions every week with Alexa, and more than 100,000 skills on the Alexa service?" I would say, "nope." I had planned for success but not of this kind. We're now comfortable speaking to a device at a distance and even anthropomorphize the AI. Such adoption is just mind boggling. I may have a bit of proximity bias, but I really think it is a testament to the AI community that there has been such mainstream adoption.

The second positive surprise is the democratization of AI, where it is not a privileged technology any more. Developers can integrate AI in their environment via open source ML frameworks or AI tools from cloud providers such as AWS and even build an AI experience on a voice assistant using tools like the Alexa Skills Kit without worrying about how AI works. Such democratization has been the most fulfilling, positive surprise to me in the past 20 years.+

Interviewed 4.3.2020

RICK CALLE

Rick leads AI business development for M12, Microsoft's venture fund. He works at the intersection of AI algorithms, hardware computing efficiency, and novel AI use cases. During his time with Qualcomm's AI Research, he worked with the team that launched Qualcomm's AI Engine into over 100 different models of AI-enabled mobile phones.

■ *How energy intensive is the AI infrastructure today? And what does that mean for the future of discipline?*

■ Today's AI algorithms, software and hardware combined are 10X to 100X more energy intensive than they should be. In light of Microsoft's recent announcement of its carbon negative commitment, my challenge to the industry is clear: let's improve AI hardware and software so that we don't overheat our planet.

The computing industry is always optimizing for speed and innovation, but not necessarily considering the lifetime energy cost of that speed. I saw an inflection point around 2012 when the progression of AI hardware and algorithmic capabilities began to deviate from Moore's law. Prior to that, most AI solutions were running on one, maybe two processors with workloads tracking to Moore's law. A steady progression of workloads from the Perceptron in 1958 to systems like Bidirectional LSTM neural networks for speech recognition in the mid-2000s.

Training AI models with multiple GPUs changed everything. After Alex Krizhevsky and team designed the AlexNet model with two GPUs in 2012, the computing power and electrical energy involved in training AI models took off at an entirely different pace: over 100X compounding every two years. Theirs was certainly not the first Convolutional Neural Network (CNN), but their "SuperVision" entry swept the field, winning the 2012 ImageNet competition by a huge margin. The next year nearly all competitors used CNNs and trained with multiple processors!

Fast forward to 2019, and quickly developing innovative neural networks for Natural Language Processing may require hundreds or thousands of distributed GPUs — like self-attention encoder-decoder models which employ Neural Architecture Search (NAS) methods. According to a recent University of Massachusetts Amherst study, the amount of CO_2 emitted from energy generation plants to power the computation involved in creating a new state-of-the-art AI model, was the equivalent of five automobile lifetime's worth of CO_2 emissions. If that's what it takes to train only one new AI model, you can see that it is just not compatible with a prioritization of sustainability.

I believe we can incentivize the AI industry to make a change in the overall lifetime energy budget for AI workloads, and identify startups that are already committed to this cause.

■ *Where do you see the biggest opportunities for the highest impact energy savings?*

■ My colleagues and I think it's joint optimization of three things: energy-efficient AI hardware, co-designed efficient AI algorithms and AI-aware computer networks.

The challenge is that the energy consumption of AI models is likely the last thing an AI algorithm developer is thinking about (unless they're focused on mobile phones). Usually the early optimizations are foremost around performance. AI engineers often think: "what's my peak accuracy" and "how fast can I train the model" — both of which need faster computing and more energy.

I support a new success metric to help incentivize the AI industry and startups to reduce energy and CO_2 emissions at data center scale. We need to shift the focus to higher throughput and lower lifetime total cost of ownership of a system for given computing workloads. I stress "system" because often hardware marketing metrics forget to mention the energy cost of extra processors, memory, and networks required for an AI training system.

Success Metric=Workload Throughput ÷ [($ Cost of System) + ($ Cost of Lifetime Energy of System)]

Throughput measures how fast we can compute the required AI algorithms. In the phraseology of the late Harvard Business School Professor Clayton Christensen, workload throughput is the "job" that matters at the end of the day. Not peak Floating Point Operations Per Second (FLOPS) which are magical, mystical marketing numbers only loosely related to getting the computational "job" done.

The denominator of this ratio is the computing hardware cost plus the lifetime energy cost of operating that hardware including cooling and any extra network and processors required. With this new ratio, AI designers have far more degrees of freedom to optimize software, hardware and algorithms. For example, the power consumption of an AI chip itself — whether it is 50 watts or 450 watts — doesn't matter as much. The lifetime energy consumption of many chips to deliver a certain workload throughput is what matters most. If we can maximize this success ratio, then by definition energy and CO_2 emissions are reduced as well.

■ *Why change the "performance" mindset that has been the status quo for so long?*

■ AI has an existential problem. As its models continue to get larger, more computationally complex, and more accuracy is desired to reach human performance levels, the energy required to train those models increases exponentially. At some point if things continue as they have, researchers won't be able to get enough computers or energy to create the new AI algorithms we want. I'm really worried about that potentially stalling AI innovation. Not many research labs can string together 4,000 leading-edge processors and run them for weeks. They just don't have the resources to deploy exascale computers. So at some point — without change — we have the potential to reach a ceiling of innovation. I'd hate to see another AI winter.

If our AI industry innovates around the success metric, then we will benefit from AI that is more compatible with sustainability, yet meets performance goals with lower lifetime energy hardware, more efficient AI algorithms and lower energy infrastructure. +

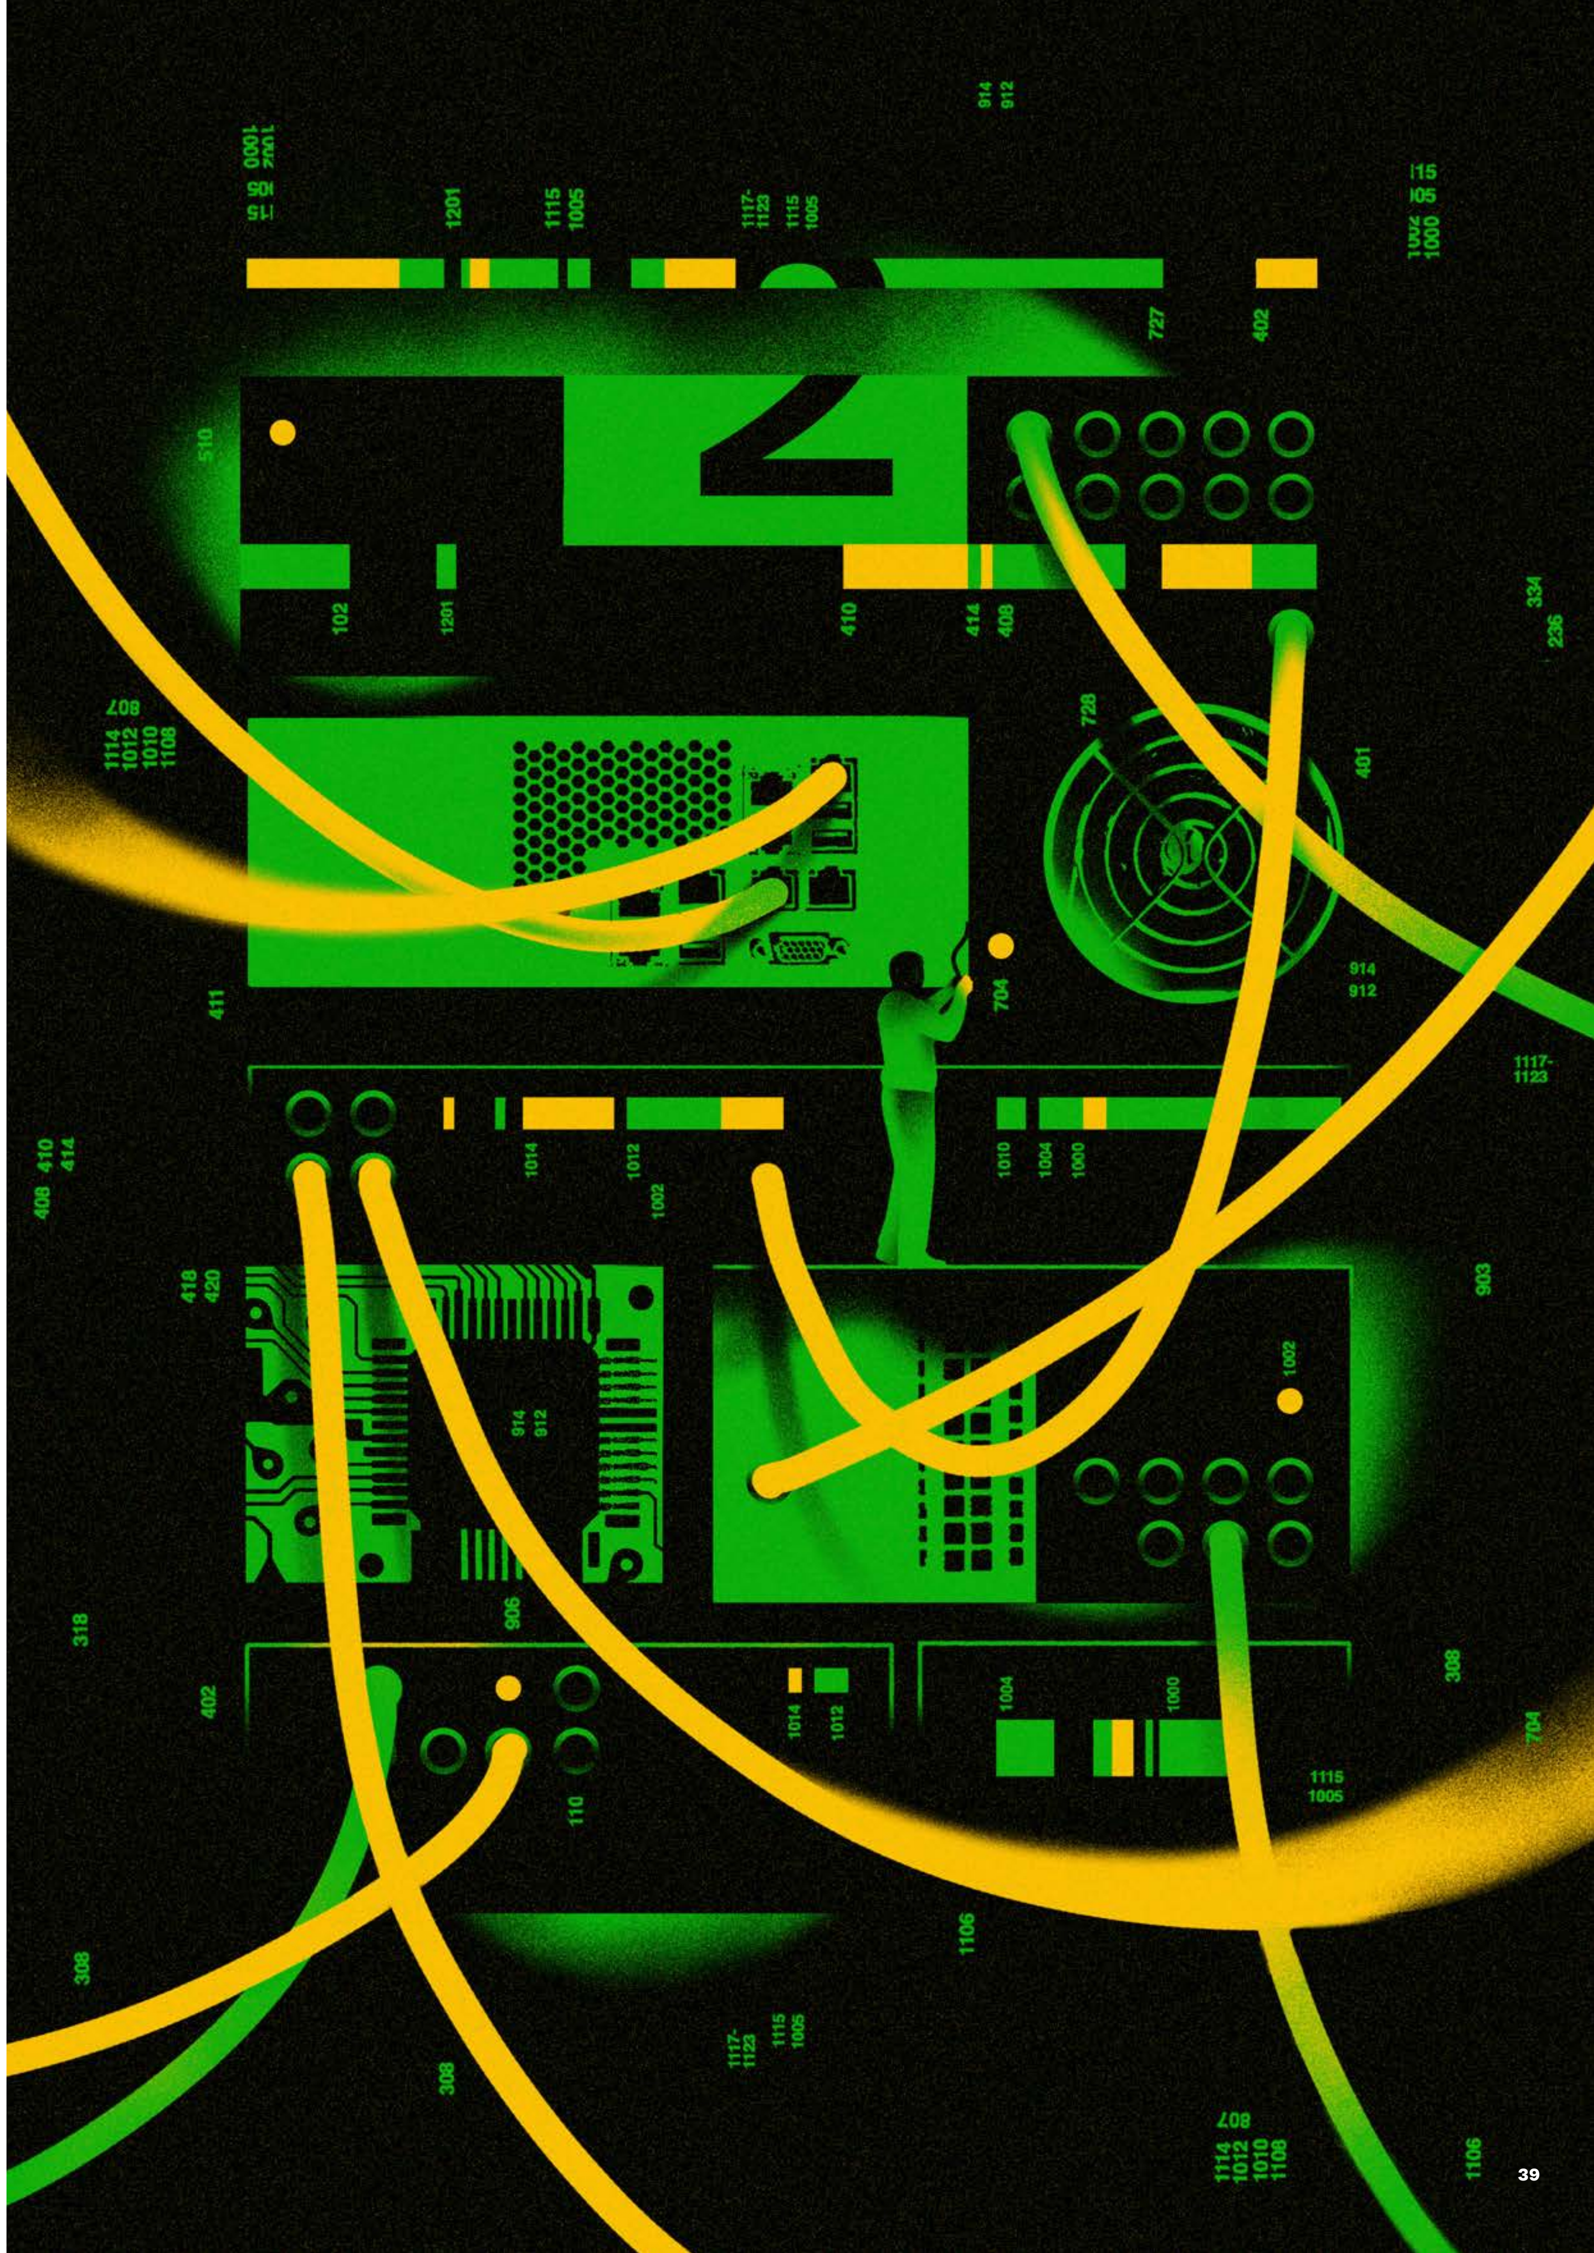

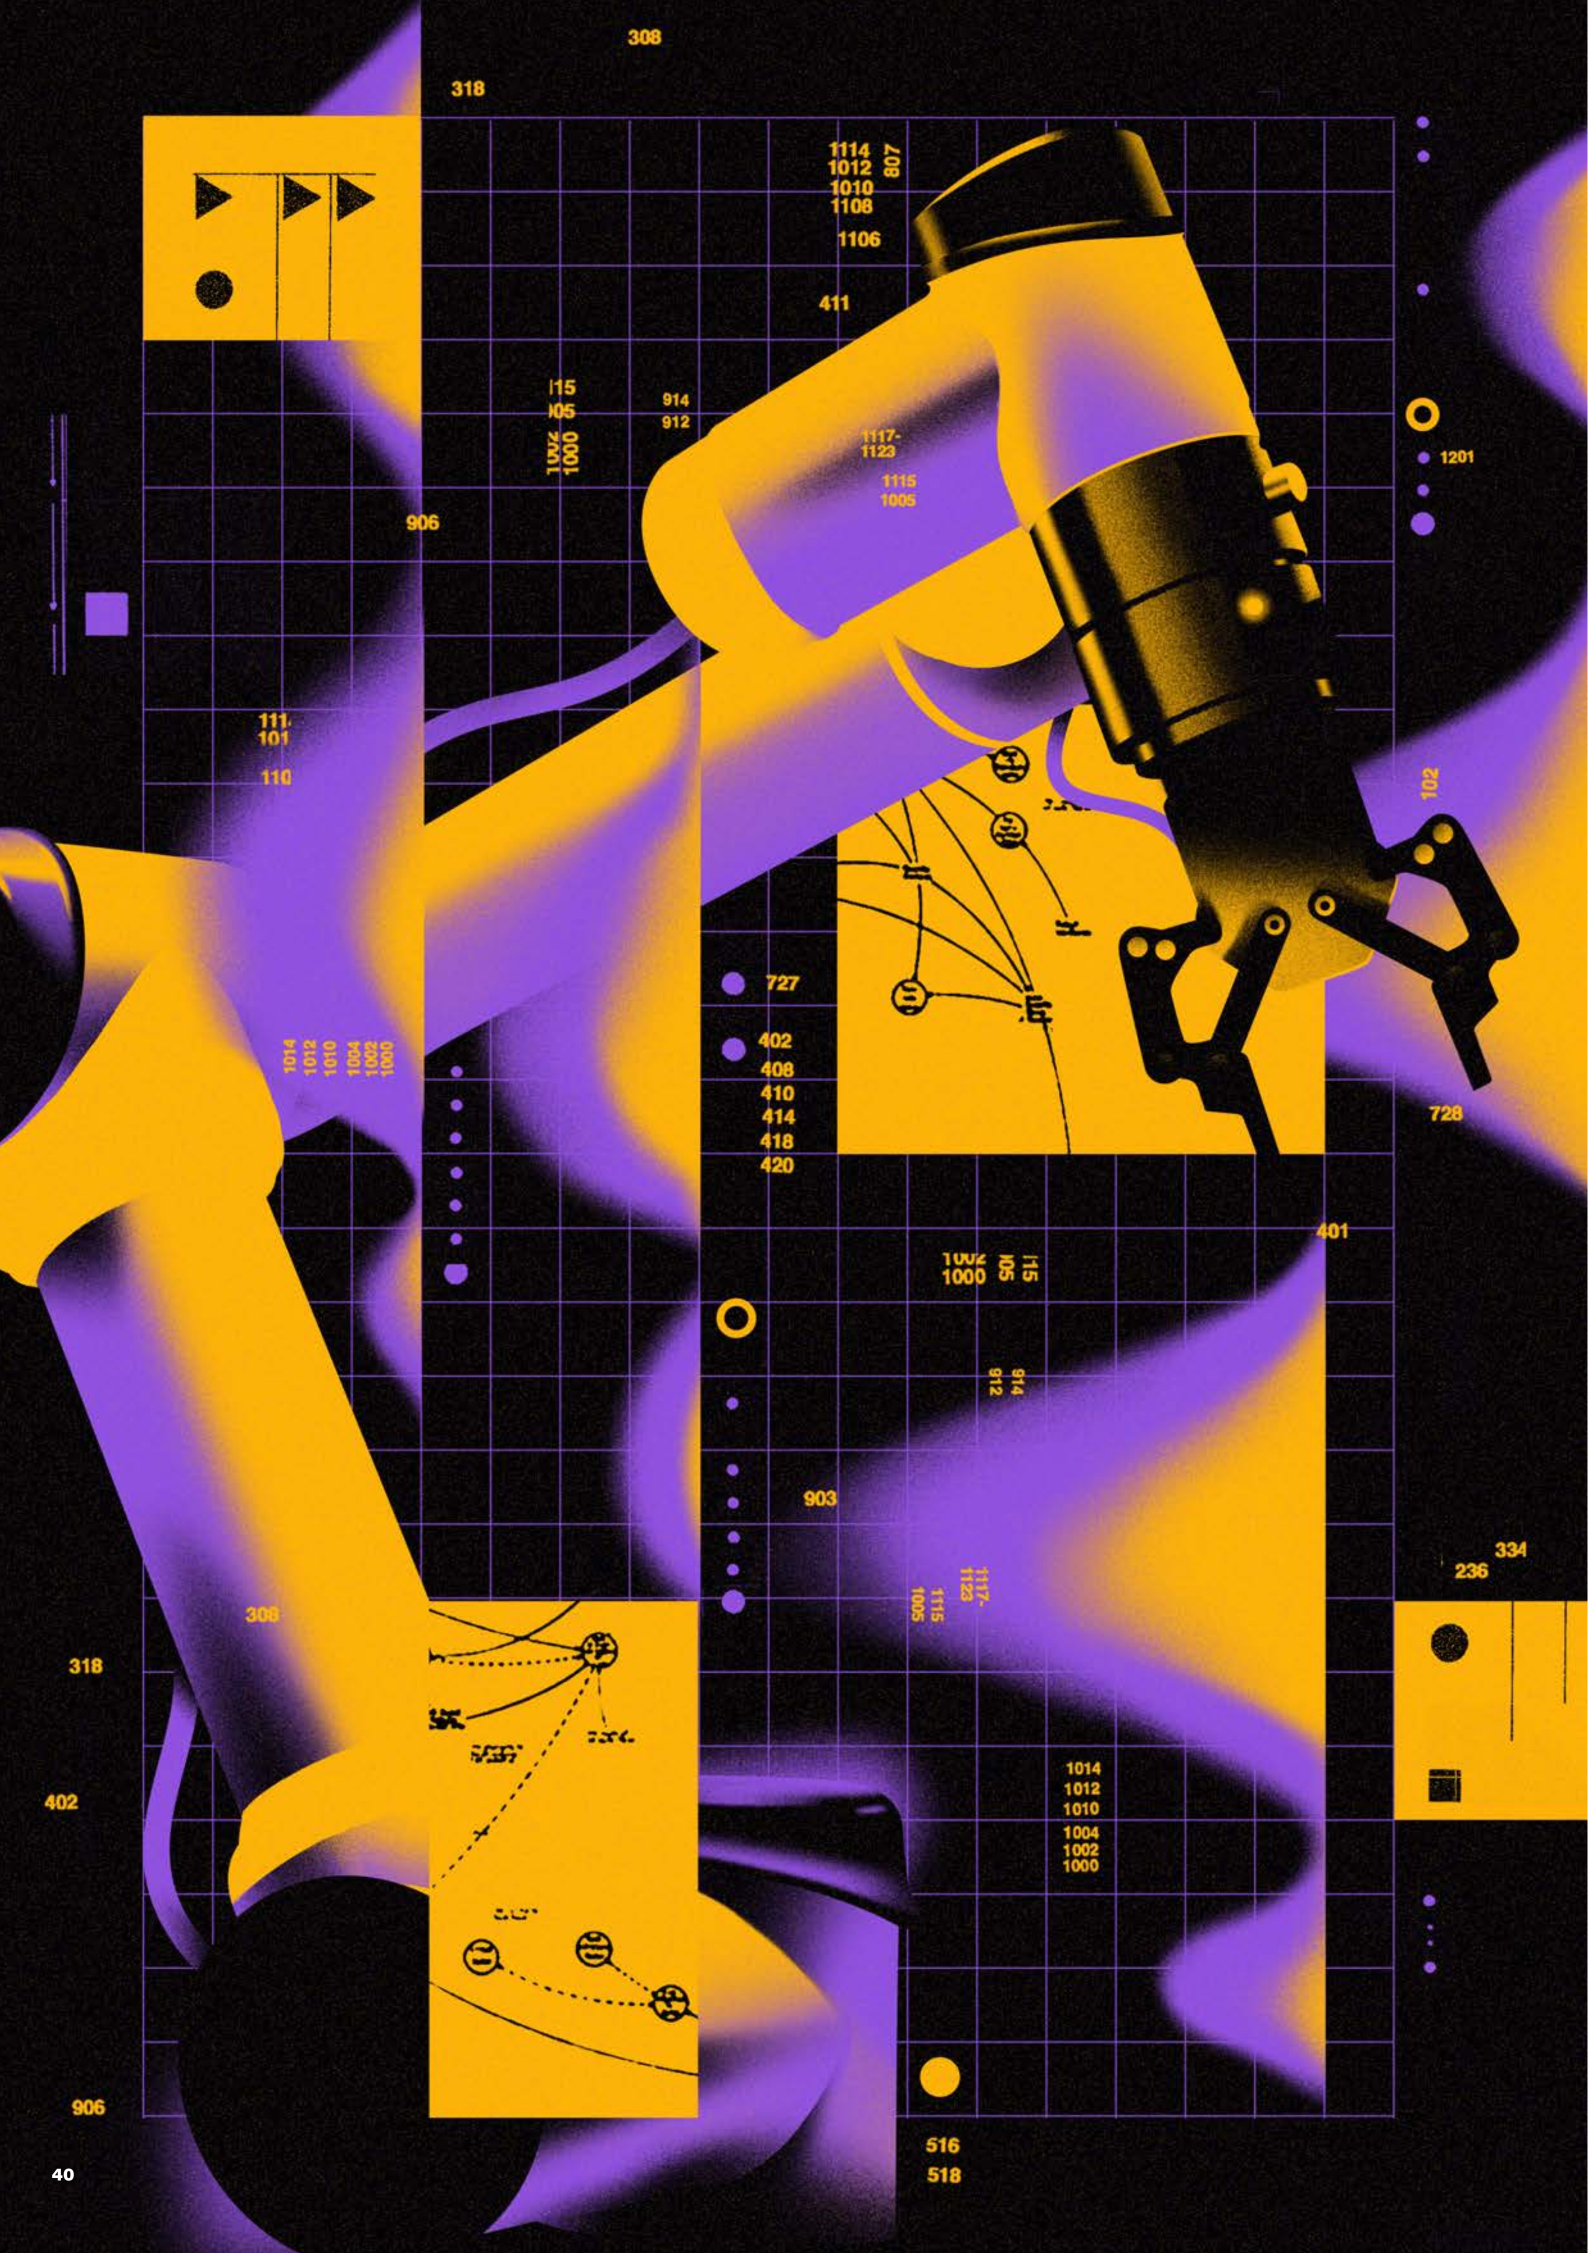

Interviewed 4.14.2020

DANIELA RUS

Daniela is the Andrew (1956) and Erna Viterbi Professor of Electrical Engineering and Computer Science and Director of CSAIL at MIT. She is a 2002 MacArthur Fellow, a fellow of ACM, AAAI and IEEE, and a member of the National Academy of Engineers, and the American Academy of Arts and Sciences.

■ *AI is polarizing — you see it referenced in pop culture and the media as both a savior and a threat. Should we remain optimistic about a future increasingly reliant on AI? Why?*

■ When I tell people that I am a robotics researcher, I tend to get one of two reactions. Some people get nervous. They make jokes about Skynet and ask when the robots will take over their jobs. They worry about autonomous vehicles because they don't trust computers to avoid crashes as well as they could themselves.

The other group gets excited. This is the group that buys Roombas to vacuum their living room and dreams of the day when their car will drive them home from work. They can't wait to hear about the latest research and discuss all the ways that robotics and AI will change our lives and our world for the better.

As the Head of CSAIL at MIT, I'm sure you won't be surprised to hear that I'm in the latter group. But I also know it's an important part of my job to understand the fears of that first group, to listen, to ask questions, and to give them some perspective on why I see things differently. That starts with understanding that AI is just a tool. It's an incredibly powerful one, but, like other tools, it isn't inherently good or bad. It is what we choose to do with it. And I believe we can do some truly incredible things.

On a global scale, AI will help us generate better insights into addressing some of our biggest challenges: understanding climate change by collecting and analyzing data from vast wireless sensor networks that monitor the oceans, the greenhouse climate, and the plant condition; improving governance by data-driven decision-making; eliminating hunger by monitoring, matching and re-routing supply and demand, and predicting and responding to natural disasters using cyber-physical sensors.

It will help us democratize education through MOOC offerings that are adaptive to student progress. It means customized healthcare, built using knowledge gleaned from enormous amounts of data. And counter to common knowledge, it means more satisfying jobs, not less, as the productivity gains from AI and robotics free us up from monotonous tasks and let us focus on the creative, social, and high-end tasks that computers are incapable of.

■ *What does the future look like?*

■ Picture a world where routine tasks are taken off your plate. Fresh produce just shows up on your doorstep, delivered by drones, garbage bins take themselves out, and smart infrastructure systems support automated pick-up. AI assistants, whether embodied or not, act as our guardian angels, providing advice to ensure that we optimize our lives to live well and work effectively.

Technology in general—and AI in particular—can be an incredible vector for positive change as we work together to figure these things out. It can give us ways to navigate the truth in a world inundated with alternative facts. It can take care of the routine stuff, and that gives us more time to focus on solving the big challenges that require creativity, collaboration, and strategy.

I know this sounds optimistic. AI is not going to solve all of our problems any more than it is going to destroy the world. It isn't either/or humans or machines, consider instead humans and machines. Working together with AI systems, humans can augment and amplify many aspects of work and life.

■ *Which sector(s) do you see being most affected by autonomous robotics and systems in the short term (2-5 years) and the long term (10+ years)?*

■ Transportation, logistics, and warehousing are great examples for 2-5 years impact. It's much easier to move a robot through the world than it is to build a robot that can interact with it. Technology has taken us to a point in time where we can realistically discuss the idea of autonomy on the roads. But not full autonomy. Robots don't share our recognition capabilities. How could they? We spend our whole lives learning how to observe the world and make sense of it. Machines require algorithms to do this, and data—lots and lots and lots of data, annotated to tell them what it all means. To make level 5 autonomy possible, we have to develop new algorithms that help them learn from far fewer examples in an unsupervised way, without constant human intervention.

■ *What innovations in robotic hardware excite you most? What makes them transformative?*

■ I am particularly excited by the transformative potential of soft robotics. While the industrial robots of the past 60 years have mostly been inspired by the human form, the next stage will be soft robots inspired by the animal kingdom and by our own built environments, with broader applications potential. With the development of soft materials, machines and materials are coming closer together—this raises an interesting question: what is a robot? Traditionally, we have considered industrial manipulators and robots on wheels, but what about robots made out of food or paper or ice, or giving everyday objects the ability to move and compute so they become intelligent and autonomous?

■ *Will AI, like calculus or geometry, become a prerequisite for every engineering program?*

■ I believe that computational thinking should be part of literacy in the 21st century.+

The Portfolio Companies

The Engine invests in founders solving the world's biggest problems through the convergence of breakthrough science, engineering, and leadership.

We've seen our investments coalesce into three areas of impact: those companies whose core technology will help solve climate change; those that will create new human health solutions; and those that will usher in a new era of advanced systems.

HUMAN HEALTH

- Biobot Analytics
- Cambridge Crops
- Cellino
- E25Bio
- Kytopen
- Lucy Therapeutics
- Seaspire Skincare
- Suono Bio
- Vaxess Technologies

CLIMATE CHANGE

- Boston Metal
- Cambridge Electronics
- Commonwealth Fusion Systems
- Form Energy
- Lilac Solutions
- Syzygy Plasmonics
- Via Separations

ADVANCED SYSTEMS

- Analytical Space
- C2Sense
- HyperLight
- isee
- Radix Labs
- RISE Robotics
- Sync Computing
- Zapata Computing

Biobot Analytics

Founders	**|1|** Mariana Matus, **|2|** Newsha Ghaeli
Background	MIT
Industry	Biotech & Life Sciences, Data Science, AI & ML

We can learn a lot about a community from its waste. And until now, such portraits would disappear with every flush.

Biobot Analytics analyzes urine and stool collected in sewers and unites that data into a dynamic picture of a community's wellbeing. By mapping this anonymous data, Biobot helps communities and private health partners to proactively tackle public health issues. Imagine detecting the signs of a potential pandemic before the community exhibited symptoms and, as a result, implementing protective measures before the disease had time to spread.

The vision of co-founders Mariana Matus and Newsha Ghaeli, Biobot is working to tackle the COVID-19 pandemic as one of the first large-scale implementations of its platform. With the help of communities across the country, the startup hopes to sidestep inefficiencies inherent in the existing testing platforms to more efficiently combat the historic outbreak.

Matus and Ghaeli founded Biobot by asking, "Can we look at sewer systems the same way that biologists study the human microbiome? Are our sewer systems essentially the 'guts' of our communities?"

Matus, a computational biologist, studied wastewater epidemiology at MIT. The concept of "sewage as our collective microbiome," was one she pursued during her doctoral studies. Ghaeli, with training in urban planning and architecture, has long known that the future of humanity is urban — and uncovering smarter ways of harnessing that closeness is necessary if we are to thrive in that future. The pair met thanks to a joint research venture between Matus' lab and Ghaeli's department.

Some of Biobot's first pilot initiatives were concerned with treating one of this country's most insidious non-viral epidemics — opioids. By analyzing the waste of hard-hit communities, the startup helped local governments identify the type of opioid being used, the general quantity, and the timeframes associated with these variables. Its platform can identify, for example, the difference between prescription opioids and heroin. Notably, it is sensitive enough to detect one dose of particular chemicals within a sample size of 5,000 people.

Biobot conducted its first commercial implementation of its technology with an opioid analytics program in Cary, NC. The company's analysis gave local officials accurate information on the use of opioids, so they could better lead productive interventions, reducing overdoses by 40% and lowering their associated cost to the healthcare system. Today, seven cities in Massachusetts are conducting initial studies with the Biobot opioid product. The startup is also exploring Hepatitis C pilot programs.

Biobot is built to respond to the public health priority of the moment, without changing the physical sampling device within the sewer systems. It can just as easily compile and analyze data to combat addiction as it can to thwart the next pandemic. And a single of the startup's samplers can account for varying rates of sewage flow — from groups of 5,000 to up to 20,000. If a community has a larger population, simply deploy more sensors.

As Matus, Ghaeli, and the Biobot team gather more community data, they will begin to build proprietary predictive models tailored to specific public health applications. By merging these predictive models with the dynamic picture of real-time community wellbeing, the team can empower all facets of those communities, from citizens to public health experts and private sector partners, to make decisions that will help save lives today — and tomorrow. +

Photo: Webb Chappell for the Boston Globe

Data analysis that will transform wastewater infrastructure into public health observatories

Lilac Solutions

Founders	**|1|** Tom Wilson, **|2|** Dave Snydacker, **|3|** Nick Goldberg
Background	Northwestern University
Industry	Advanced Materials

The lithium industry stands at the edge of a revolution. As the batteries which power our electronics are being adapted and deployed at a massive scale to power electric vehicle fleets, lithium demand will grow exponentially. But lithium is difficult to produce — nobody has figured out how to do it efficiently or in a way that is environmentally sustainable. Today's battery makers face a critical question: how will they secure the huge volumes of low-cost lithium they need to remain competitive in the electric vehicle era?

Lilac Solutions is the answer. The company, helmed by Dave Snydacker, Nick Goldberg, and Tom Wilson, is commercializing a new lithium extraction system that recovers twice the lithium at half the production cost, with 1000x smaller footprint and 5000x faster processing speed than conventional evaporation-pond extraction methods. These incredible numbers are supported by extensive test work on more than 30 lithium brine resources from around the world.

Lithium brines hold approximately 75% of the world's lithium resources and are considered easier to access than hard-rock ore, but lithium extraction from brines represents a significant technological challenge.

In traditional lithium brine extraction, brine is pumped to the surface into giant evaporation ponds where it evaporates over the course of months, at which point the remaining concentrated brine is pumped out and sent to a recovery plant and then further purified. The process requires enormous tracts of land and has a devastating impact on local water resources. Also, traditional brine extraction is shockingly inefficient — it can only recover 40% of available lithium.

Lilac Solutions will eliminate the need for evaporation ponds. Its system extracts lithium directly from brine through a patented ion exchange process that is highly efficient (70%-98%, depending on brine quality). It also uses 99% less land, 90% less water, and reduces greenhouse gas emissions by 80% relative to current brine-extraction methods. Moreover, these projects can reach commercial production in just two years instead of the usual ten.

The Lilac team is uniquely qualified to champion this new lithium extraction process. The startup's core technology was born out of Dave Snydacker's obsession with solving climate change as well as his expertise in materials engineering. Snydacker met co-founder Goldberg when the pair were students at Wesleyan University, eventually reconnecting after Goldberg had established himself in a top corporate law firm in New York. Goldberg joined Lilac's advisory board where his role quickly evolved to Chief Operating Officer.

In 2019, the pair was introduced to Chief Development Officer Tom Wilson at Stanford's Graduate School of Business as Wilson was completing a Sloan Fellowship. Wilson had spent 14 years developing projects in the oil and gas industry. Lilac presented an ideal opportunity to apply his significant oil and gas experience to an emerging clean energy industry.

After observing major industry investments in traditional lithium ion technology and repeated failures of competing battery technologies, Snydacker came to believe that further major breakthroughs in battery technology were unlikely. For automakers looking to modernize their fleets, the bottleneck was no longer battery technology — it was sourcing the raw materials for those batteries, especially lithium.

Snydacker knew that ion-exchange — a proven technology used in water treatment, mining, and other industries that demand purified liquids — had the ability to transform lithium extraction. But those who had previously pursued the technology had failed to create a material that could absorb lithium without dissolving in the acid bath required to ultimately extract the element. The process had potential, but faced a fundamental materials science problem.

For a period of 30 weeks in 2017, the Lilac team made small changes to the formulation of their ion exchange bead. As Snydacker tells it, "It was an epic 30-week process. Every week we'd make small changes to the material and test it, and every week it would fail. And at week 30 it finally clicked." Lilac had created a material that could absorb lithium from brine, then be reused over and over again without dissolving.

With this breakthrough, Lilac began scaling the core ion-exchange process. The company has now successfully demonstrated its system at 1,000 L/hr scale, and the efficiency of the process has not been affected by this scale-up across four orders of magnitude. Commercial units on the horizon are planned to have a capacity of 10,000l /hr.

Lilac Solutions has the potential to be the technology of choice for all new lithium brine projects. The company is partnering with resource developers and lithium producers around the world to unlock sub-economic brine resources and efficiently expand lithium production. Lilac is harnessing its proven technology to transform an industry and usher in a renewable, electrified future. +

RISE Robotics

Founders	**|1|** Blake Sessions, **|2|** Toomas Sepp, **|3|** Kyle Dell'Aquila, **|4|** Arron Acosta
Background	MIT
Industry	Robotics

Heavy machinery moves our world. It lifts the steel for our buildings. It carves roads out of mountains. It plants and harvests our foods. It is powerful, durable, and ubiquitous. While there is tremendous diversity of scale and application within the heavy machinery industry, the lifeblood of the vast majority of these machines is diesel fuel, which, in turn, is needed to run complex and power intensive hydraulic systems. While relatively long lasting and capable of brute force, these platforms lack fine control and produce 22 lbs of CO_2 per gallon of diesel, resulting in 55,000,000 tons of CO_2 annually.[1,2]

RISE Robotics has invented a replacement for hydraulic systems that will enable the next era of fully electrified heavy machinery — one that is at once sustainable, robust, and precise. The startup's core technology is an electrically-powered mechanical linear actuator with all the abilities of a hydraulic cylinder, but vastly improved efficiency and control. RISE also supplies electrification systems through partnerships with heavy machinery OEMs, helping maximize the impact of its hardware.

The demand for electrically powered heavy machines is rapidly increasing. The industry is already seeing electric and hybrid retrofits of existing hydraulic platforms. But these conversions are grossly inefficient and significantly more expensive than their diesel counterparts. The batteries required to run the current hydraulic systems of heavy machinery are massive, unwieldy, and offer less runtime than the status quo. Some warehouse-scale machinery like forklifts are transitioning to alternative fuels such as hydrogen in an effort to be more sustainable. All of these solutions have shortcomings. The reason? They rely on an intrinsically inefficient system — hydraulics.

Born from their relationship formed while at MIT, founders Arron Acosta and Blake Sessions started RISE to create human-amplifying machines, also known as exosuits. These machines give their wearers superhuman motion. But creating an exosuit required a motion component — one that was powerful, lightweight, efficient, and precise — that no linear actuator on the market could provide. RISE had to invent a solution.

The team, with the addition of Toomas Sepp and Kyle Dell'Aquila, discovered that high-strength steel cables in an electrically powered and digitally controlled pulley system offered a powerful combination of efficiency, size, and precision. They quickly realized that the potential of their platform stretched far beyond exosuits — with their new linear motion solutions, RISE could help usher in a new era of sustainable heavy machinery.

While its platform is broadly capable, RISE is first targeting the lifting machine market. Think forklifts of every scale (from warehouse to port) and construction lifts. There is significant interest in electrifying these machines in order to meet imminent oil-free and zero-emission mandates.

The evolution of heavy machinery is marked by technological step changes, rather than a gradual optimization of platforms. First, there were large steam-powered, cable driven machines — think of the steam shovel from the classic 1939 children's book Mike Mulligan and His Steam Shovel. Then there was the transition to diesel and hydraulics, enabling the compact heavy machinery we see today. RISE will lead the next revolution, enabling fleets of fully electric, sustainable, and precise heavy machinery to build the world of tomorrow. +

1. http://www.patagoniaalliance.org/wp-content/uploads/2014/08/How-much-carbon-dioxide-is-produced-by-burning-gasoline-and-diesel-fuel-FAQ-U.S.-Energy-Information-Administration-EIA.pdf fluid_power_workshop_master_presentation.pdf
2. https://www.nrel.gov/transportation/assets/pdfs/mobile_fluid_power_workshop_master_presentation.pdf

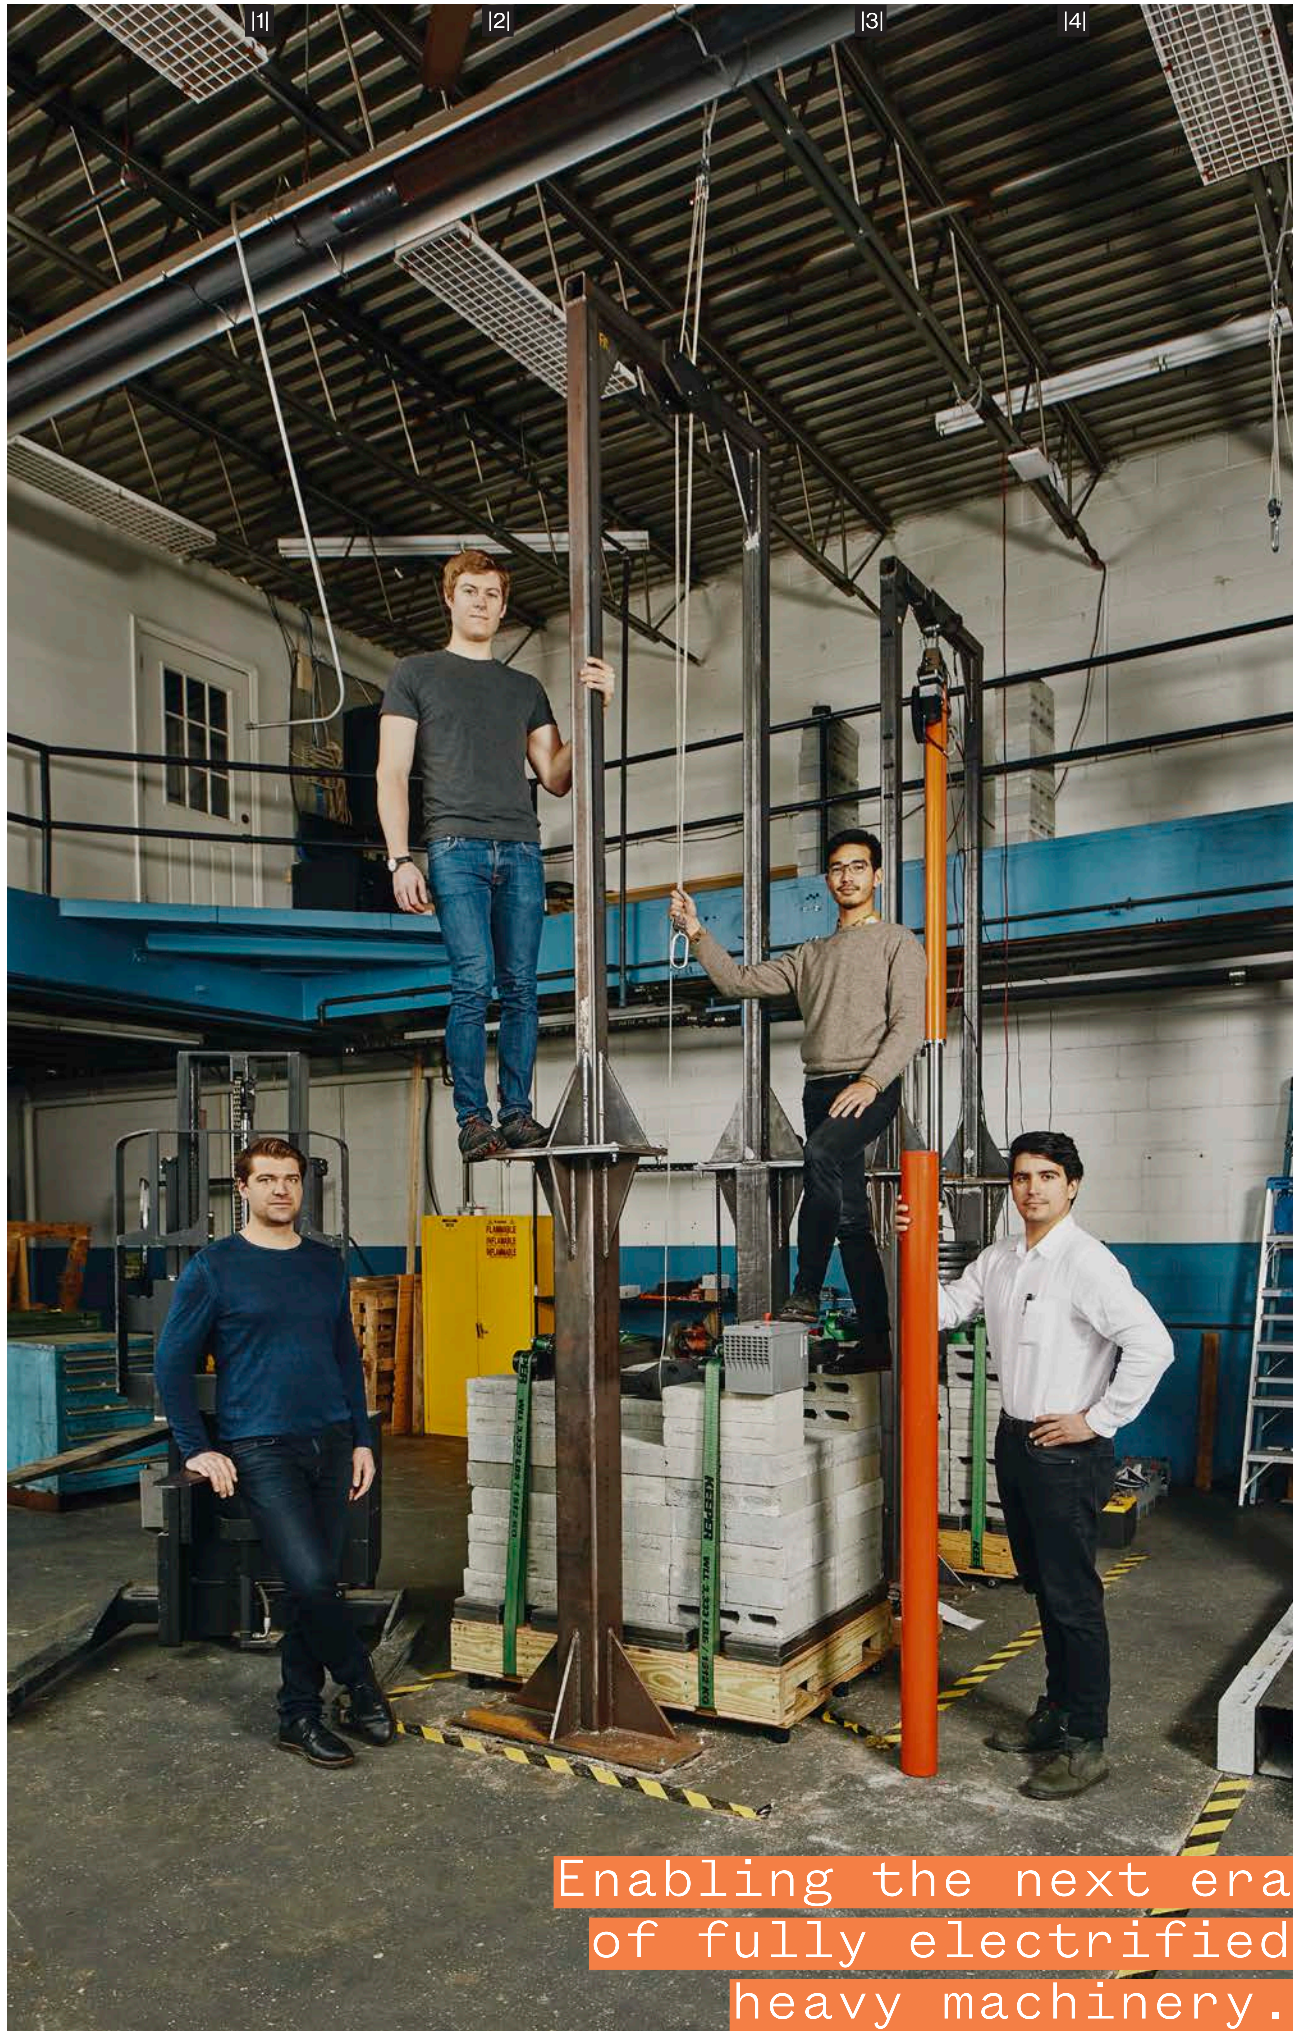

The future of sustainable skincare.

Photo: Ruby Wallau for Northeastern University

THE FOUNDERS

Seaspire Skincare

Founders	**|1|** Leila Deravi, **|2|** Camille Martin
Background	Northeastern University
Industry	Advanced Materials

Seaspire Skincare, a startup born out of Northeastern University, is pioneering a new category of multifunctional materials with extensive implications for human health and environmental safety. The team has unlocked the ability to recreate and package the chemical machinery of the chromatophore, a pigment-containing organ found in the skin of marine life such as squid and other cephalopods. Chromatophores enable these animals to adapt to their environment by changing color almost instantaneously.

Camille Martin and Leila Deravi, Seaspire's co-founders, will use this chromatophore-inspired class of pigments — coined xanthochrome by the team — to replace the harmful active ingredients that provide UV filtering in existing sunscreens.

Xanthochrome has the potential to be among the first UV filters that can protect against solar irradiation and be used in preventative skincare, while reducing toxicities to marine organisms. The compounds have the ability to scatter and absorb light and can protect against a broader spectrum of UV radiation (from UV to near IR) compared with current physical and chemical UV filters.

Recent research indicates a link between ubiquitous UV chemical filters and the health of marine ecosystems like coral reefs. Chemical filters are easily washed off into the oceans, contaminating the seabed. Major tourist destinations like Key West, the US Virgin Islands, Palau, and the entirety of Hawaii, have banned the use of sunscreens with these chemicals.

With growing ecological concerns regarding regular sunscreen use, efforts have also been made to re-evaluate the safety and efficacy of sunscreens for human use, an initiative spearheaded by several regulatory agencies worldwide. There is no doubt, however, that sunscreen is one of the easiest and most proven ways to prevent skin cancer. This disparity demands a solution — something that offers all the benefits of traditional chemical UV filters, without adverse effects on human and marine health. And with a global market size of almost $10B for suncare products, such a material presents a massive business opportunity.

Martin, Seaspire's CEO, had long known that she would pursue some type of novel cosmetics and skincare. A chemist by training, she completed her PhD in co-founder Deravi's lab, where she helped synthesize and process xanthochrome. While at Northeastern, Martin was also inducted into the Huntington 100 — one of the university's most prestigious student awards — in recognition of her dedication to entrepreneurship.

Deravi, a cephalopod specialist, is an assistant professor in the Department of Chemistry and Chemical Biology at Northeastern University, where she focuses on understanding the role of pigments and proteins involved in the dynamic color change of living systems. Her team's research has received support from the US Army, US Navy, and the National Science Foundation. It has also been published in leading academic journals, and featured in The New York Times and Chemical and Engineering News.

Such recognition is not without warrant. Martin, Deravi, and team are decoding the secrets of compounds that took millennia to evolve and until now have remained under the stewardship of some of the planet's most striking animals.

It seems only fitting that xanthochrome has the potential to transform the health of the marine ecosystems from which they arose. As components of suncare products and cosmetics, the pigments could transform these industries as sustainable, high-performance, and nontoxic alternatives to the status quo. +

THE FOUNDERS

Sync Computing

Founders	**|1|** Suraj Bramhavar, **|2|** Jeff Chou
Background	MIT Lincoln Lab
Industry	Computing

From the nanoscale confines of a computer chip to the labyrinthian patchwork of roads, rails, and flight paths, our world is overrun by waste. Not grime and garbage, but rather underutilized and misused time and energy. These trashed seconds and watts may be invisible, but they affect our lives in very real ways.

Think of your commute. Do you go down one road to avoid traffic? Or wait it out, hoping that the line of cars in front of you gets moving? Do you follow directions from an app? Or do you go off script and take a shortcut? Every choice you make, and every choice your fellow commuters make, has an impact on the time you spend on the road and the amount of fuel you burn.

Now think of the same scenario for a logistics company delivering millions of packages worldwide, every day. The challenge to sort packages, find the most efficient delivery route, use the least amount of fuel, and ensure that packages arrive within a prescribed time window, is extraordinary. It is an overwhelming optimization problem — so overwhelming that these companies exist knowing that each day is inherently wasteful.

These types of problem, known as combinatorial optimization problems, exist in industries like telecom, pharma, finance, logistics, those that rely on machine learning, and more. Solving them is notoriously difficult because of the sheer number of possibilities that must be sifted through to determine the most efficient solution. They represent a challenge that digital computers, by their very nature, will never master.

Researchers have long searched for a viable computing approach to combinatorial optimization problems.

That search may now be over. Sync Computing, a startup born out of the MIT Lincoln Laboratory, is pioneering the world's first Optimization Processing Unit (OPU) that could provide a solution to these stubborn optimization challenges.

Think of it as an "algorithm in hardware form," says Sync's co-founder Jeff Chou. Inside the OPU is a system that naturally wants to find a path to optimal energy use. This path, when found, can be interpreted as the solution to the problem.

This differs from a digital computer which must test every single scenario before reaching a reliable conclusion, an approach that is both time and energy intensive. The system that Chou and his co-founder Suraj Bramhavar have invented uses simple electronic oscillators to take advantage of nature's unfailing ability to optimize energy.

Similar approaches to solving combinatorial optimization problems have been attempted with quantum computers. But scaling quantum computations to anything beyond experimental levels is an exceptional challenge. The machines are temperamental and expensive, requiring exotic materials and engineering to keep the computer operating reliably at close to absolute zero. The Sync platform, in contrast, uses inexpensive off-the-shelf electronics components and can operate at room temperature.

The startup's elegant use of everyday componentry means that it can experiment and scale quickly and with extraordinarily low financial risk. It will help the team create new solutions, faster.

Much of what Sync is pioneering runs counter to decades of problem-solving strategies dominated by a standardized digital architecture. Their computing solutions that power their OPU are not taught in most computer science classes. And that is precisely what makes it so exciting. Bramhavar sums it up well, "We are embarking on a completely alternative form of computing that does not rely on the standard digital architecture. If we can prove it out, it will open up an entirely new way of solving some of our trickiest problems."+

Building the world's first Optimization Processing Unit to help unlock solutions to some of society's most profound problems.

CLIMATE CHANGE

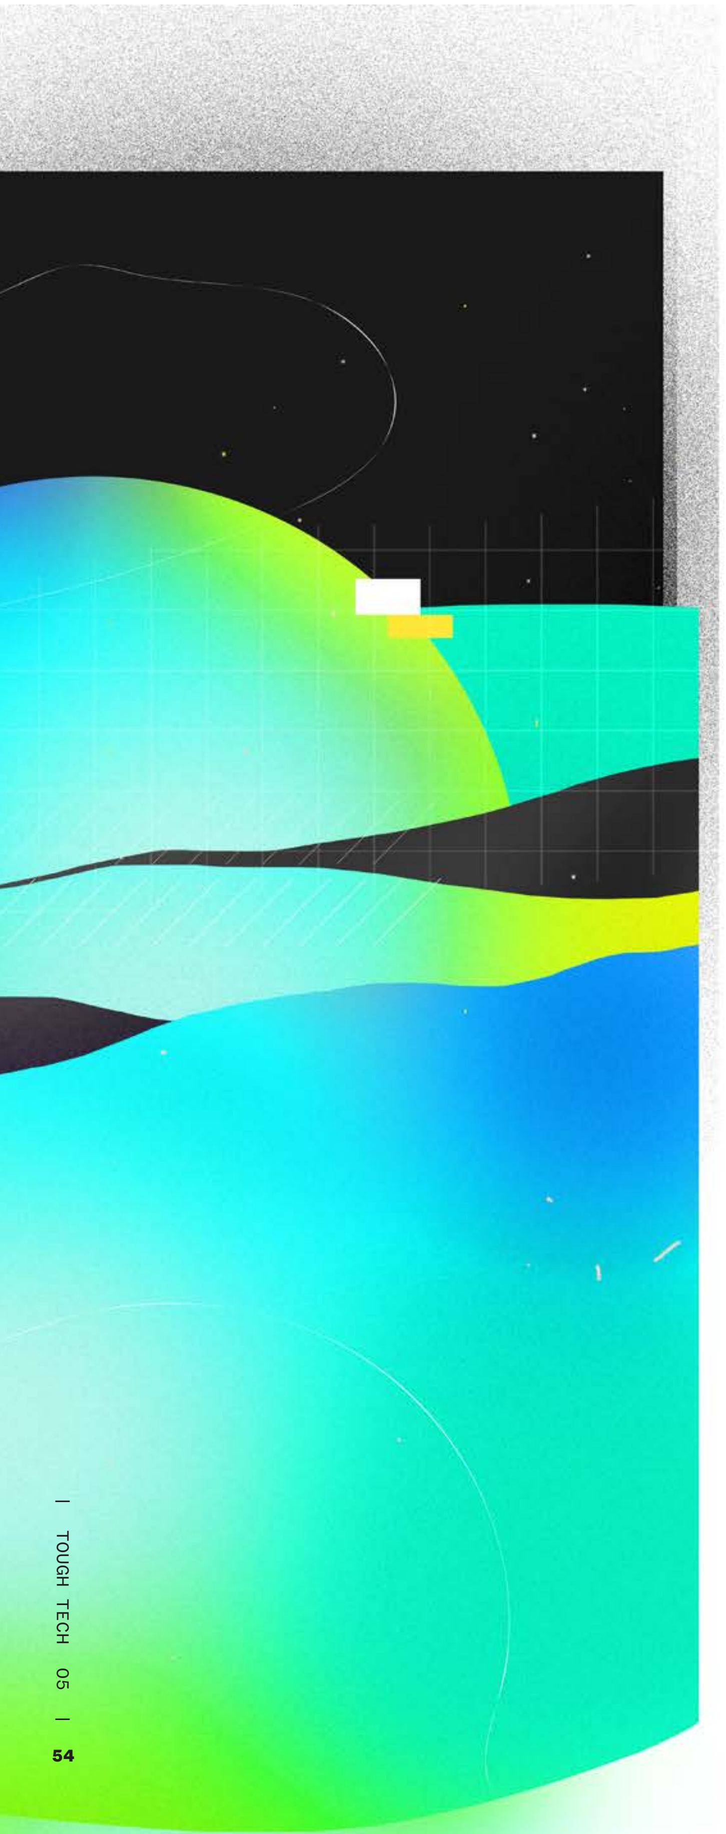

Boston Metal

Founders & Leadership
Tadeu Carneiro, Rich Bradshaw, Adam Rauwerdink, Donald Sadoway, Antoine Allanore, Jim Yurko, Bob Hyers

Background
MIT Department of Materials Science and Engineering

Industry
Advanced Manufacturing, Energy

Boston Metal has invented a coal-free, emissions free, modular method of industrial steel and ferroalloy production using electricity. It's called molten oxide electrolysis (MOE) and combines transformative materials engineering and novel systems engineering with elements from industrial aluminum production, traditional blast furnaces, and arc furnaces to produce steel and ferroalloys more efficiently, at lower costs than traditional methods, and with zero greenhouse gas emissions.

Significance
Today, the steel industry is the largest industrial source of CO2 emissions because of a reliance on coal. Boston Metal removes coal from the process, driving CO2 emissions to zero, while also providing substantial OPEX and CAPEX savings.

Efficient, lower-cost production of steel and alloys with zero emissions.

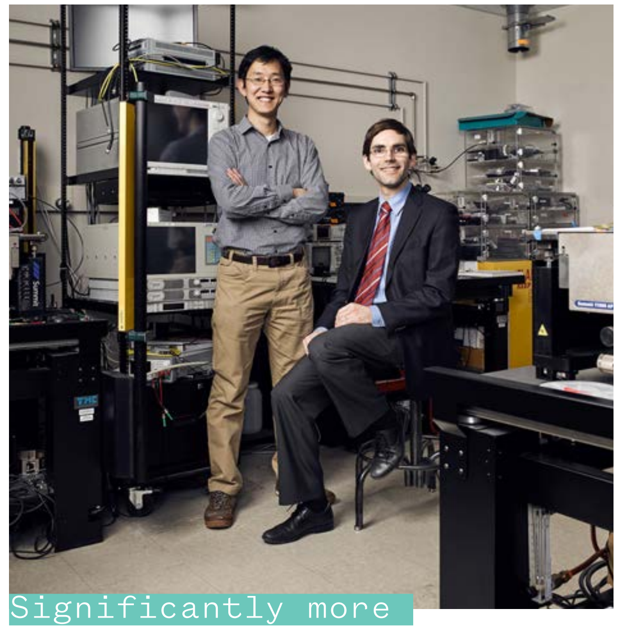

Significantly more efficient electronics – from data centers to electric vehicles.

Cambridge Electronics

Founders
Bin Lu, Tomás Palacios

Background
MIT Microsystems Technology Laboratories, MIT Department of Electrical Engineering and Computer Science

Industry
Semiconductors, Advanced Materials

Today's electronics rely on silicon processing. From data centers, electric vehicles, and consumer electronics, the ubiquitous material is used to control and convert power. As these technologies advance, industries are challenged to build increasingly efficient (and increasingly small) power electronics. In many cases, we have reached the limits of silicon. Cambridge Electronics has invented a proprietary gallium nitride (GaN) technology that is less expensive and exponentially more efficient than silicon, while also having a smaller footprint.

Significance
Cambridge Electronics' technology will bring significant energy savings to diverse and power-reliant industries like data centers, renewable energy, manufacturing, automotive, and consumer electronics.

Commonwealth Fusion Systems

Founders & Leadership
Bob Mumgaard, Brandon Sorbom, Dan Brunner, Dennis Whyte, Martin Greenwald, Zach Hartwig

Background
MIT Plasma Science and Fusion Center

Industry
Energy, Advanced Materials

Commonwealth Fusion Systems (CFS) aims to provide a new path to fusion power by combining proven fusion physics with revolutionary magnet technology to deploy the first working, economic fusion reactors to the world. The team will develop high field magnets based on a new class of high temperature superconductor materials that will allow fusion reactors to be 10 times smaller, economically feasible, and operational in the next 10 years.

Significance
Fusion energy is the holy grail of clean energy: limitless, no greenhouse gases, baseload, concentrated, no meltdown, and no proliferation. If successful, the world's energy systems will be transformed.

Creating safe, unlimited, carbon-free fusion power for the grid in 10-15 years.

Form Energy

Founders
Mateo Jaramillo, Ted Wiley, William Woodford, Yet-Ming Chiang, Marco Ferrara

Background
MIT Department of Material Science and Engineering, 24M Technologies, A123, Tesla Energy

Industry
Energy, Advanced Materials

Form Energy will solve large-scale renewable energy's most fundamental limitation — reliability — through energy storage. Rather than thinking of batteries in the traditional sense, simply as storage vessels, Form is designing bidirectional power plants. Built to displace fossil fuel baseload generation plants, Form Energy's core technology will store and supply hundreds of megawatts via the existing energy grid.

Significance
Form Energy will help bring renewables to the masses at an affordable price by storing energy from sources like wind and solar to power thousands of homes and businesses.

Engineering a bidirectional power plant to make renewable energy available 24/7.

Chemical manufacturing driven by light, enabling cheaper scalable on-site production.

Syzygy Plasmonics

Founders
Trevor Best, Suman Khatiwada, Naomi Halas, Peter Nordlander

Background
Rice University, Baker Hughes

Industry
Advanced Manufacturing

Syzygy Plasmonics is pioneering a new type of chemical reactor driven by light rather than heat, eliminating the greenhouse gas (GHG) emissions associated with burning fuel to power a reaction. At the heart of the reactor is a novel photocatalyst with 10,000x greater efficiency than competitive examples. The company has focused its first efforts on hydrogen production, but the underlying technology platform can be tailored to produce other chemicals as well.

Significance
Syzygy Plasmonics' technology platform allows for the production of chemicals on site, in a modular, scalable, and cost-effective way, with reduced GHG emissions. This will revolutionize the entire chemical manufacturing industry not in the least because the decentralization can open new markets by avoiding the need to rely on costly or inefficient transportation chains.

Via Separations

Founders
Shreya Dave, Brent Keller, Jeff Grossman

Background
MIT Department of Materials Science and Engineering

Industry
Energy, Advanced Materials, Advanced Manufacturing

Separation processes are the building blocks for materials, chemicals, and consumer goods — they are core to the industrial ecosystem. Currently, most separations are done with thermal processes such as evaporation and distillation, which are very energy intensive. Via Separations is commercializing novel membrane materials and manufacturing processes to replace evaporation and distillation with filtration.

Significance
The company's technology has the potential to replace thermal separation processes, saving the energy equivalent used by the entire gasoline industry every year in the U.S.

Up to 90% energy savings in separation process in the pulp & paper, chemical, and dairy industries.

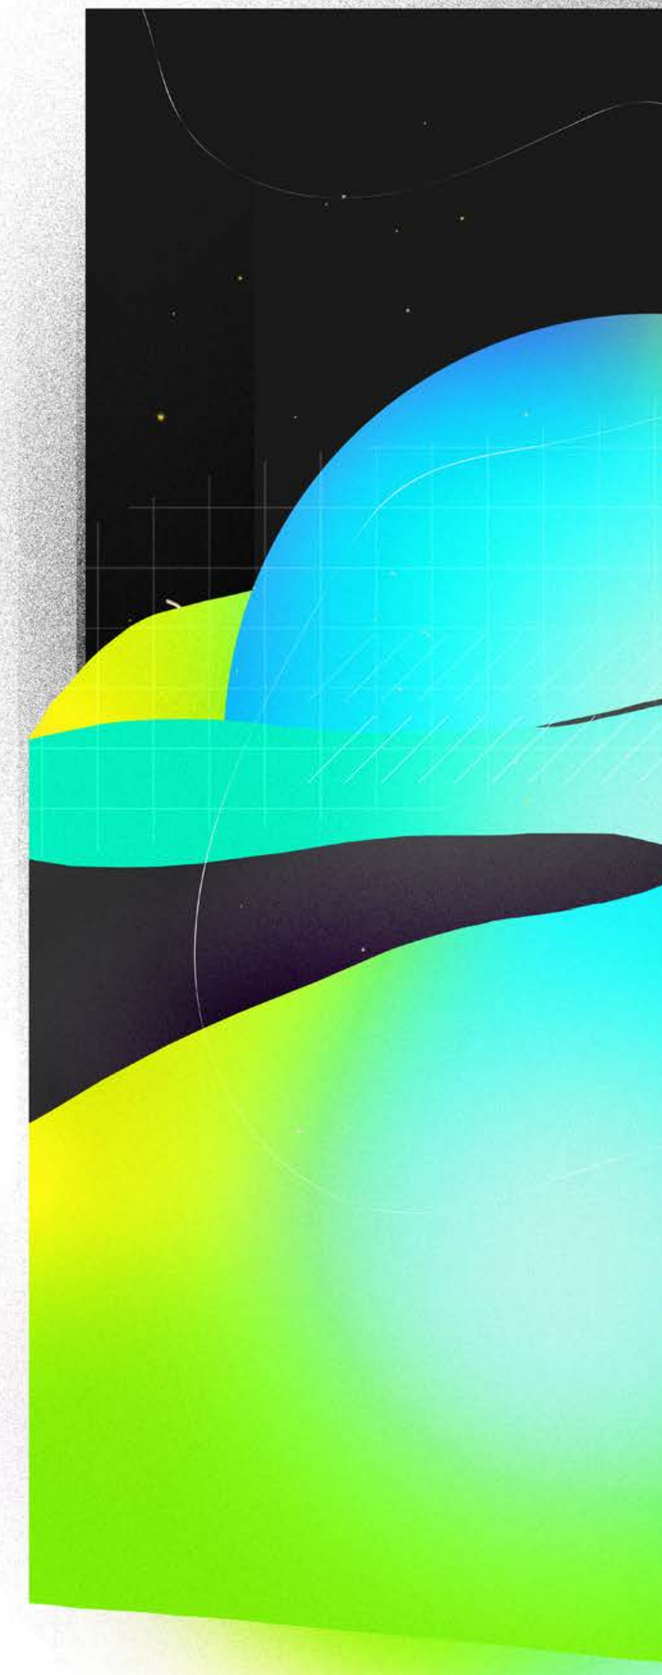

HUMAN HEALTH

A natural coating that reduces food spoilage and packaging waste.

Cambridge Crops

Founders
Adam Behrens, Sezin Yigit, Benedetto Marelli, Livio Valenti, Fiorenzo Omenetto

Background
MIT Laboratory for Advanced Biopolymers, Tufts University SilkLab

Industry
Food & Agriculture, Advanced Materials

Cambridge Crops is addressing the problem of food spoilage and waste by pioneering a natural, ultra-thin water-based coating that preserves the freshness of food longer. It's tasteless and invisible and can be applied to everything from fresh and cut produce to proteins like meat and fish. The coating dramatically extends shelf life by slowing the exchange of gases that cause decay, making food accessible to more people for longer times. In addition, the coating has the potential to support enhanced nutrients for food and also help reduce packaging.

Significance
One third of the food produced in the world is wasted. Cambridge Crops' technology helps to reduce food spoilage across the supply chain, decreases logistics costs, and makes healthy food more accessible.

Cellino

Founders
Nabiha Saklayen, Matthias Wagner, Marinna Madrid

Background
Harvard Physics Department, Harvard School of Engineering and Applied Sciences (SEAS), Harvard Medical School

Industry
Biotech & Life Sciences, Advanced Manufacturing, AI & ML

Cellino has built the first platform that enables precise control over iPS cell fate in their natural environment. The Cellino Tissue Engineering Platform manufactures high-quality, impurity-free tissues for new regenerative medicines. Cellino will use its platform to manufacture tissues at scale, delivering the highest quality human tissues made to date. Such tissues are poised to lead to significant gains in therapeutic benefit to the patient.

Significance
Cellino's approach for high-throughput, computer-guided engineering human cells will create new tissues with significant gains in therapeutic benefit to the patient and further transform the biotech industry.

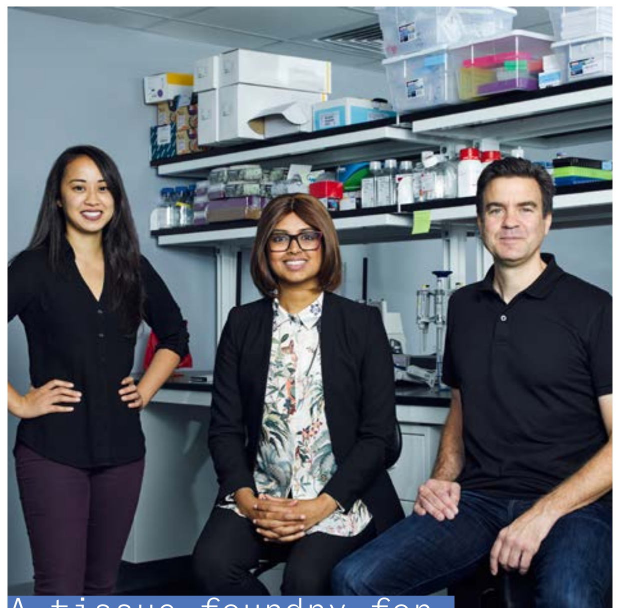

A tissue foundry for regenerative medicine.

Rapid, accurate diagnosis of infectious disease at the point of care.

E25Bio

Founders
Irene Bosch, Bobby Brooke Herrera, Lee Gehrke

Background
MIT Institute for Medical Engineering & Science, MIT Tata Center

Industry
Biotech & Life Sciences

E25Bio is pioneering an at-home rapid fever panel for mosquito-borne diseases. With its first-in-class antibodies identified with a novel screening method, E25Bio's diagnostic test is the first of its kind to distinguish between dengue (as well as all four subtypes of the disease), chikungunya, and Zika.

Significance
E25Bio is reducing a specialized central medical testing facility to a single over-the-counter test. Initially, the company's rapid fever panel will empower patients, healthcare workers, and public health officials in Latin America. But the company's ability to quickly discover and produce effective antibody pairs means that it has the potential to help patients across the globe.

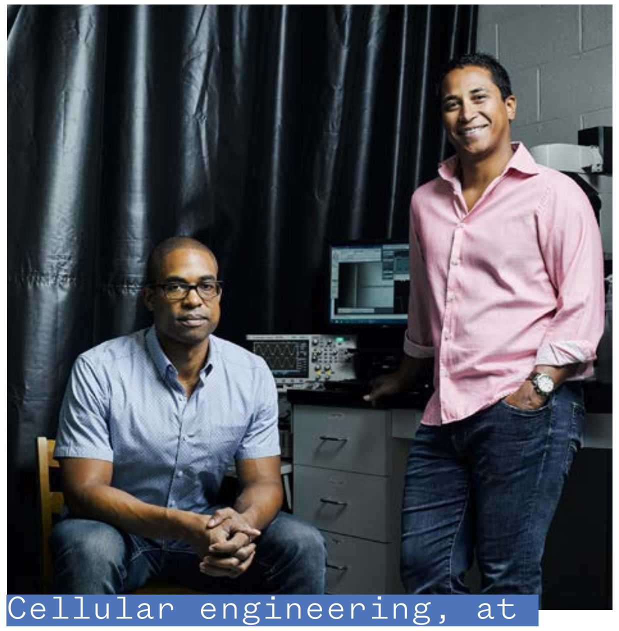

Cellular engineering, at unparalleled speeds, from discovery to the clinic.

Kytopen

Founders
Paulo Garcia, Cullen Buie

Background
MIT Department of Mechanical Engineering

Industry
Biotech & Life Sciences, Advanced Manufacturing

Kytopen aims to transform the cell and gene therapy industry with its microfluidics and electric field-based platform that can automate and manufacture the genetic engineering of cells 10,000x times faster than current methods. With continuous flow of cells during genetic manipulation, the products in development address both small and large sample volume, and enable both drug discovery and manufacturing at scale.

Significance
Cell and gene therapies currently suffer from major challenges in efficiency, reproducibility, and cost. Kytopen's solution can solve a huge bottleneck in the development and manufacturing process, reducing costs and accelerating time to market for these therapies.

Lucy Therapeutics

Founder
Amy Ripka

Background
University of Wisconsin-Madison, The Scripps Research Institute

Industry
Biotech & Life Sciences

Lucy Therapeutics is pursuing more effective clinical results in neurological diseases like Rett Syndrome, Parkinson's, and Alzheimer's by targeting dysfunctional mitochondria in neurons. The insights that underpin Lucy Therapeutics' drug discovery platform may also lead to a biomarker that would enable early, presymptomatic diagnosis of these diseases.

Significance
Imagine a world in which doctors can diagnose and treat patients before the tremors, the dementia, or the seizures from neurological diseases like Rett Syndrome, Alzheimer's, and Parkinson's take control. This is a world that Lucy Therapeutics is working to realize.

Breakthrough mitochondrial-based therapies for neurological diseases.

Ultrasound drug delivery for difficult-to-treat diseases.

Suono Bio

Founders & Leadership
Carl Schoellhammer, Robert Langer, Amy Schulman, Gio Traverso

Background
MIT Department of Chemical Engineering

Industry
Biotech & Life Sciences

Suono Bio has reimagined ultrasound as an effective and elegant delivery mechanism for the most delicate therapeutics. Its technology can push molecules like DNA, RNA, and proteins directly into cells without disrupting the surrounding tissue or harming the molecule itself. The flexibility and efficacy of the Suono Bio therapeutic platform brings with it the potential to treat and cure diseases with targets once deemed undruggable.

Significance
Suono Bio will more effectively treat challenging chronic gastrointestinal diseases and enable new therapies for other pressing health challenges like diabetes, cancer, and viral infections.

Vaxess Technologies

Founders
Michael Schrader, Kathryn Kosuda, Livio Valenti, David Kaplan, Fiorenzo Omenetto

Background
Harvard Business School, Tufts University SilkLab

Industry
Biotech & Life Sciences, Advanced Materials

Vaxess Technologies is pioneering a technique it calls Infection Mimicry to help increase the effectiveness of immunotherapies for infectious diseases and cancer. The company's first product, named MIMIX, is inspired by the body's natural immune response to infection. MIMIX is a smart-release therapeutic patch that, after only minutes of wear-time, can release treatments into the skin at precise rates for up to months after the initial application.

Significance
The same biology that allows MIMIX to activate the immune system against infectious diseases like influenza can also be used to activate the immune system against cancer cells. When a MIMIX patch loaded with a chemo agent is applied to a certain tumor, for example, it kickstarts a natural immune response, eventually eliminating metastases throughout the body.

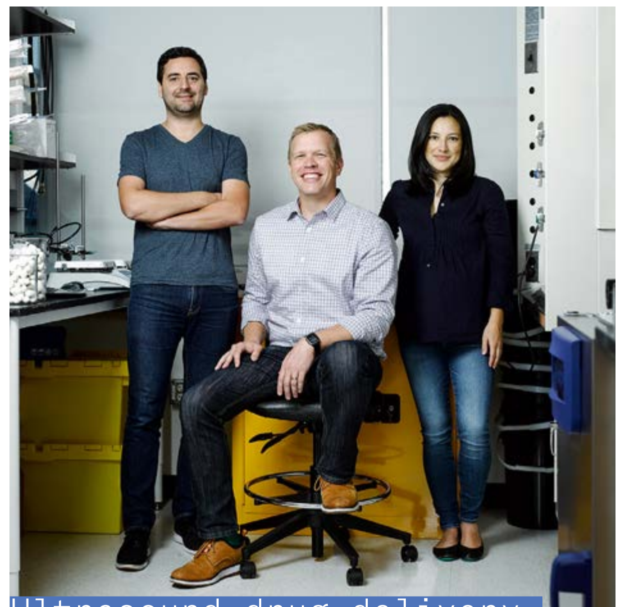

Ultrasound drug delivery for difficult-to-treat diseases.

ADVANCED SYSTEMS

Real-time satellite network connecting space to Earth 24/7.

Analytical Space

Founder
Dan Nevius

Background
NASA, Planetary Resources, White House, HBS

Industry
Space, Internet of Things

Analytical Space (ASI) is building a network of in-orbit communication relay satellites that offers expanded connectivity for data transfer, without any change to existing hardware. This results in faster data downloading, more access to download windows, lower latency, and improved cost structures, while being compatible with heritage satellites and new satellites alike.

Significance
Analytical Space will liberate and deliver terabytes of untapped data gathered by hundreds of satellites, helping industries from agriculture to defense operate with greater precision, efficiency, and insight.

C2Sense

Founders & Leadership:
George Linscott, Tim Swager, Eric Keller, JT Mann

Background
MIT Department of Chemistry

Industry
Advanced Materials, Internet of Things

A digital olfactory sensor platform for industry, C2Sense's technology transforms smell into real-time data that can be accessed remotely. With high-fidelity electrochemical sensors at a low price point, C2Sense will empower a broad array of industries, including those involved in food supply, product authentication, and chemical production, to take control of their environments.

Significance:
By making gases detectable and trackable on an industrial scale, C2Sense reduces food waste, improves safety and health of employees, and builds a more efficient and productive world.

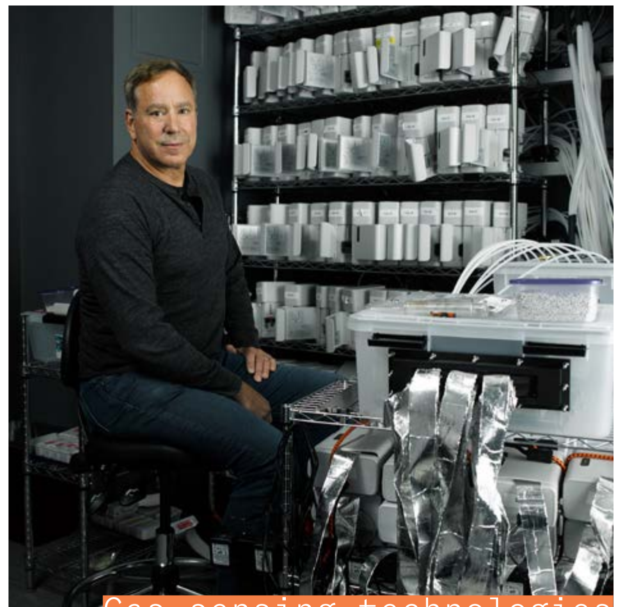

Gas sensing technologies to track and quantify the invisible.

Ultra-efficient optical circuits to de-bottleneck data centers and telecom networks.

HyperLight

Founders
Mian Zhang, Marko Loncar, Cheng Wang

Background
Laboratory for Nanoscale Optics at Harvard University

Industry
Semiconductors, Advanced Materials, Advanced Manufacturing

HyperLight has invented unique processes and designs for fabricating integrated, chip-scale Lithium Niobate (LN) modulators with extremely low signal loss. These integrated optical circuits hold the potential to reshape the world's relationship with optical data and enable novel functionalities from communication to spectroscopy. The startup's technology was developed at Harvard University and is featured in multiple publications in the journal, "Nature."

Significance
The information age relies on billions of devices converting signals between electricity and light waves. These integrated light circuits are the backbone of telecommunication, data centers, and even secure quantum communications. HyperLight's devices will force industries to rethink and reimagine their standards.

isee

Founders:
Yibiao Zhao, Debbie Yu, Chris Baker

Background
MIT Computational & Cognitive Science Group

Industry
Deep Software & AI

isee is engineering next-generation, humanistic AI to automate the logistics industry from dock to door. Their technology is built for complex environments with high uncertainty (shipping yards and congested highways) and can integrate into an existing logistics workflow without infrastructure change. isee was the first to achieve exit-to-exit autonomous highway driving, the first to merge onto a highway in heavy snow, and the first to handle congested traffic better than a leading autonomous driving startup.

Significance:
isee plans to first automate the shipping yard, reducing yard costs by 50% and increasing yard throughput by 30%. The same AI that will power yard trucks can be used to transport freight across our highways; it will add value and increase safety throughout the logistics supply chain.

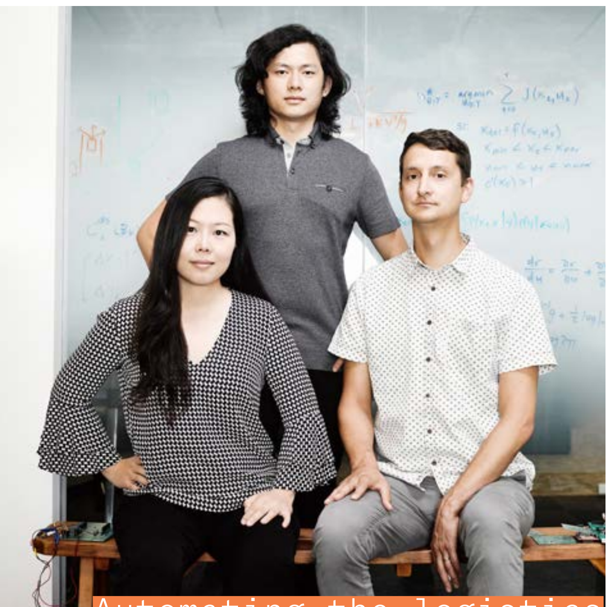

Automating the logistics industry with a humanistic AI-powered autonomous driving system.

Automating biology lab processes from experiments to mass production.

Radix Labs

Founder
Dhasharath Shrivathsa

Background
Olin College, MIT Media Lab

Industry
Robotics, Deep Software & AI, Internet of Things, Biotech & Life Sciences

Radix Labs has built a programming language that unites biologists and their lab machinery in one automated unit. This programming language is the heart of software that manages both human and machine tasks. It is the first time disparate lab machinery can communicate with one another under the control of one centralized platform — it is, for all intents and purposes, an operating system for biology labs.

Significance
Designed around an approachable user interface, this software solution intentionally distances the specification of the program — in this case the lab protocol — from the execution. It does this with the hope that biologists would spend less time in the lab, and more time focusing on experimental design and analysis.

Zapata Computing

Founders
Christopher Savoie, Alán Aspuru-Guzik, Jonathan Olson, Peter Johnson, Yudong Cao, Jhonathan Romero Fontalvo

Background
Harvard Department of Chemistry, University of Toronto Department of Chemistry

Industry
Quantum Computing

The team at Zapata Computing writes algorithms that harness the power of quantum computing to help predict and simulate some of the universe's most complex interactions, such as the behavior of molecules at an atomic level. When used in tandem with quantum hardware, these algorithms have practical industrial applications, like the optimization of supply chains and travel routes, or the prediction of drug efficacy before compounds are synthesized in the lab.

Significance
By creating algorithms that bridge advances in quantum computing hardware and commercial applications, Zapata has the potential to discover new life-saving molecules, energy efficient materials, and much more.

Quantum software and algorithms to solve industry's hardest problems.

THE ENGINE NETWORK

The Engine Network facilitates the creation of long-term, mutually beneficial relationships between founders, startups, strategic corporates, policymakers, and investors across the capital stack — in short, all of the stakeholders necessary to build successful Tough Tech companies.

Strategic Members

Platinum Members

U.S. Army · DEVCOM Army Research Laboratory · Breakthrough Energy Ventures · Hyundai Cradle · IMA Sustain Ability · Mitsubishi Electric · PNC Bank

Gold Members

Arrow Five Years Out · Brex · Bayer Life Hub Boston · Camber Road Partners · First Republic Bank · Foley Hoag LLP · Gunderson Dettmer · Mars · Mintz · SVB Silicon Valley Bank

SAVE THE DATE

The Engine Tough Tech Summit 2020
Build 10.27. | Invest 10.28.

At Tough Tech Summit 2020, you'll have the chance to discuss challenges and share ideas in energy and climate, human health, computing, automation, and much more, with entrepreneurial engineers and scientists, investors, academics, policymakers, and business leaders from across the country.

“the sun is a beautiful thing
in silence is drawn
between the trees
only the beginning of light”

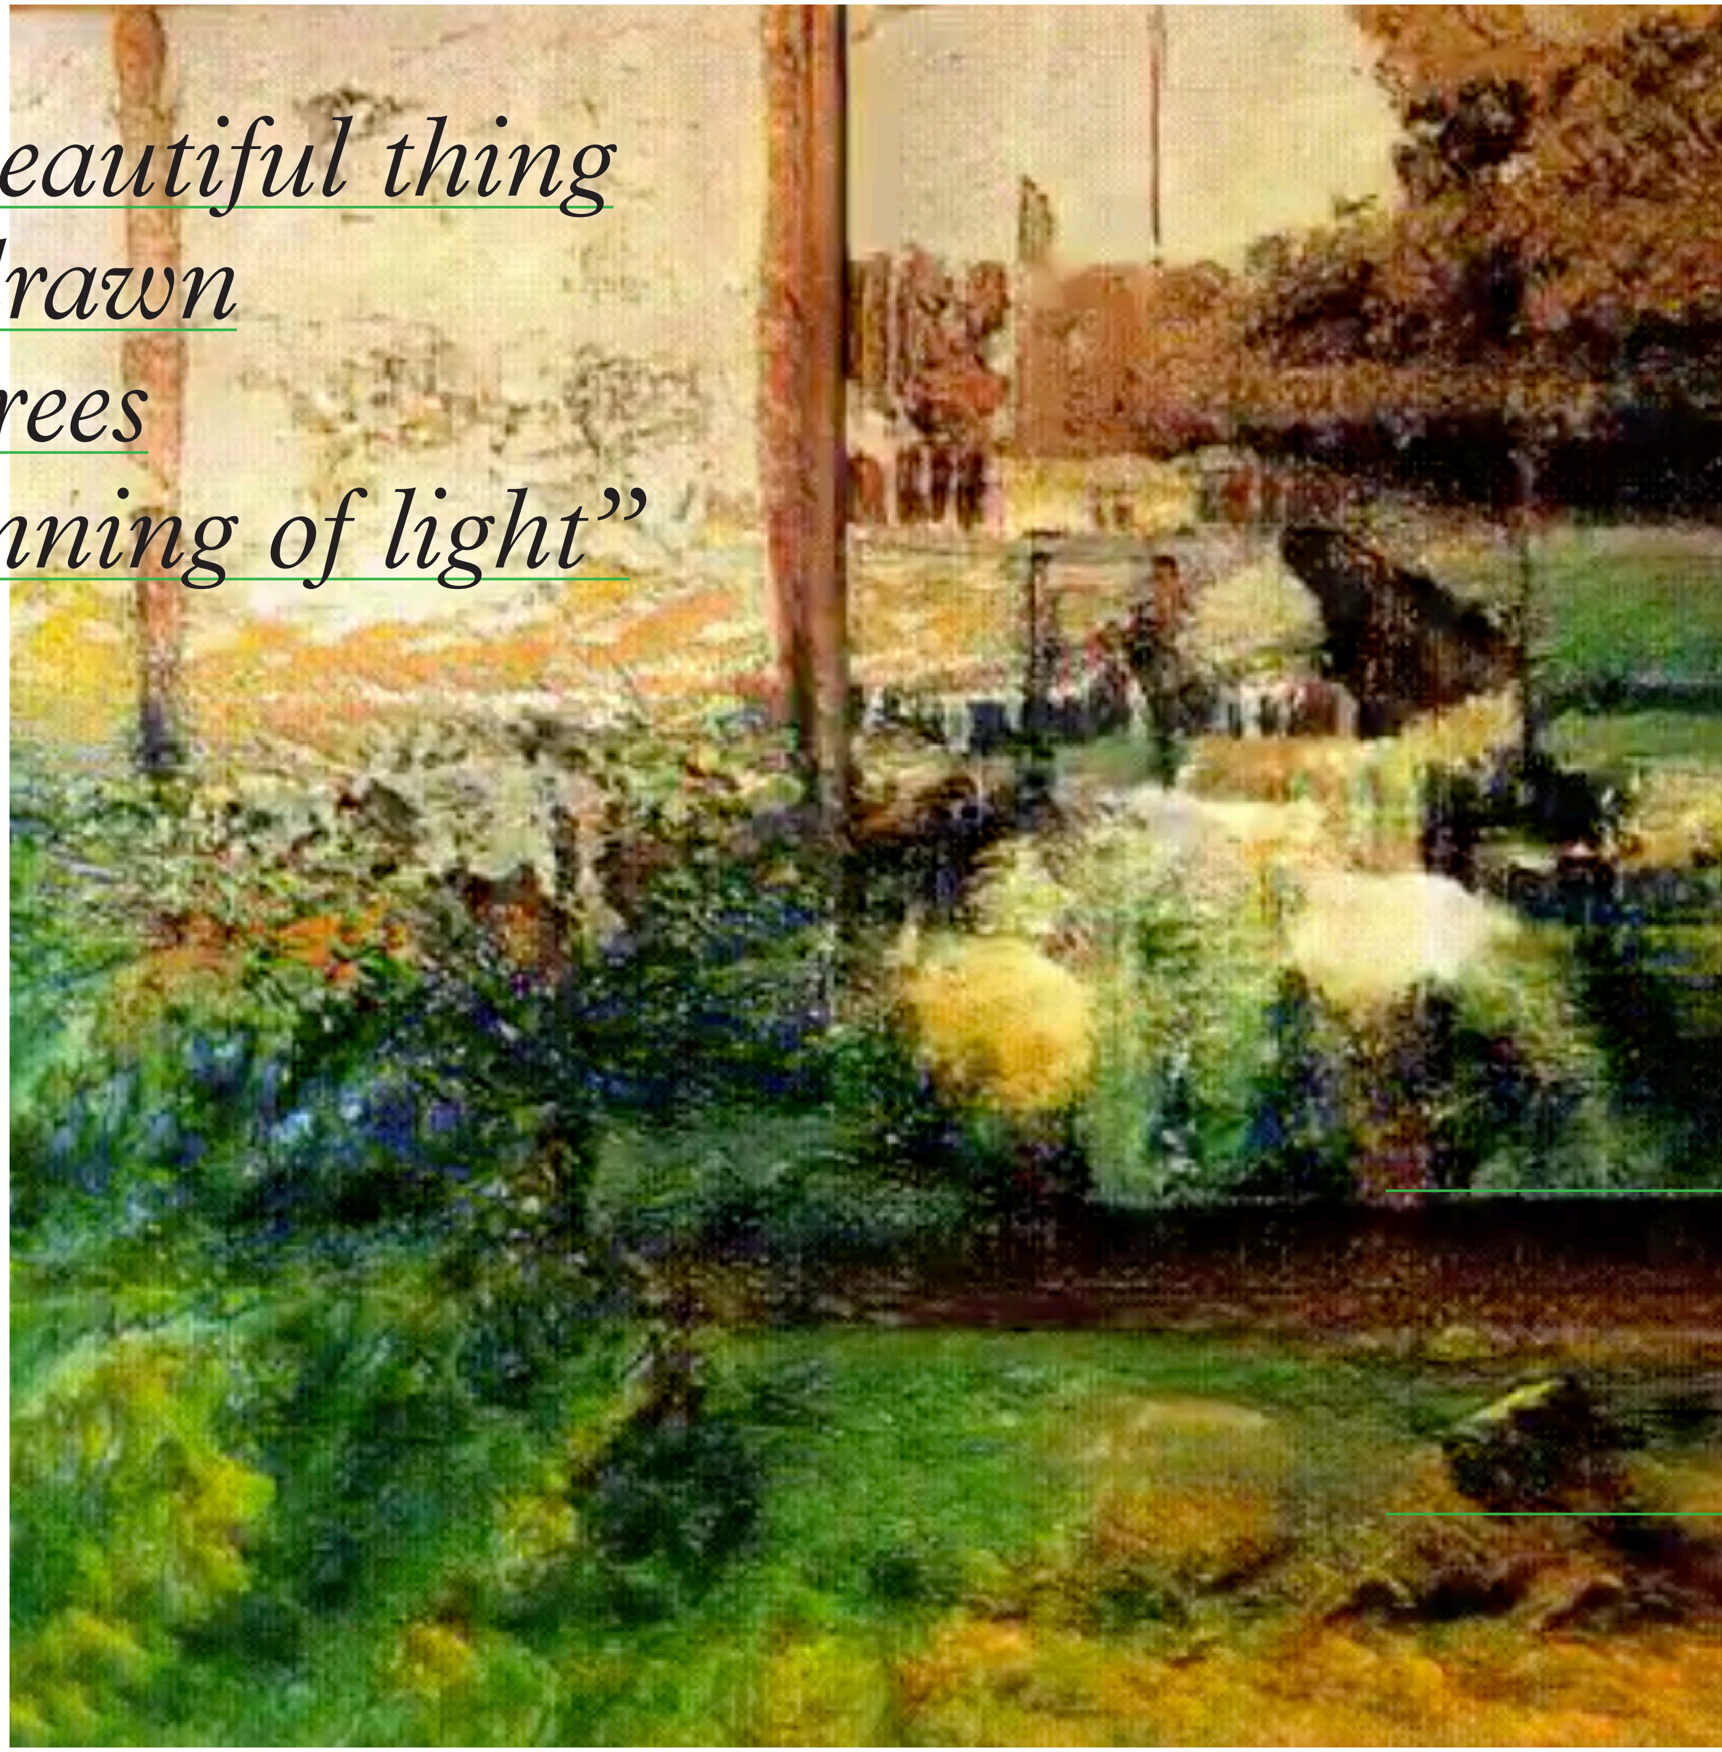

The image you see here is the result of a progressively-grown generative adversarial network (GAN) trained on approximately 80,000 paintings. It was released, along with dozens of others, by Gene Kogan in 2018. www.genekogan.com

The poem was also created using AI. It is featured in the 2018 paper Beyond Narrative Description: Generating Poetry from Images by Multi-Adversarial Training authored by Bei Liu of Kyoto University, Jianlong Fu of Microsoft Research Asia, Makoto P. Kato of Kyoto University, and Masatoshi Yoshikawa of Kyoto University.

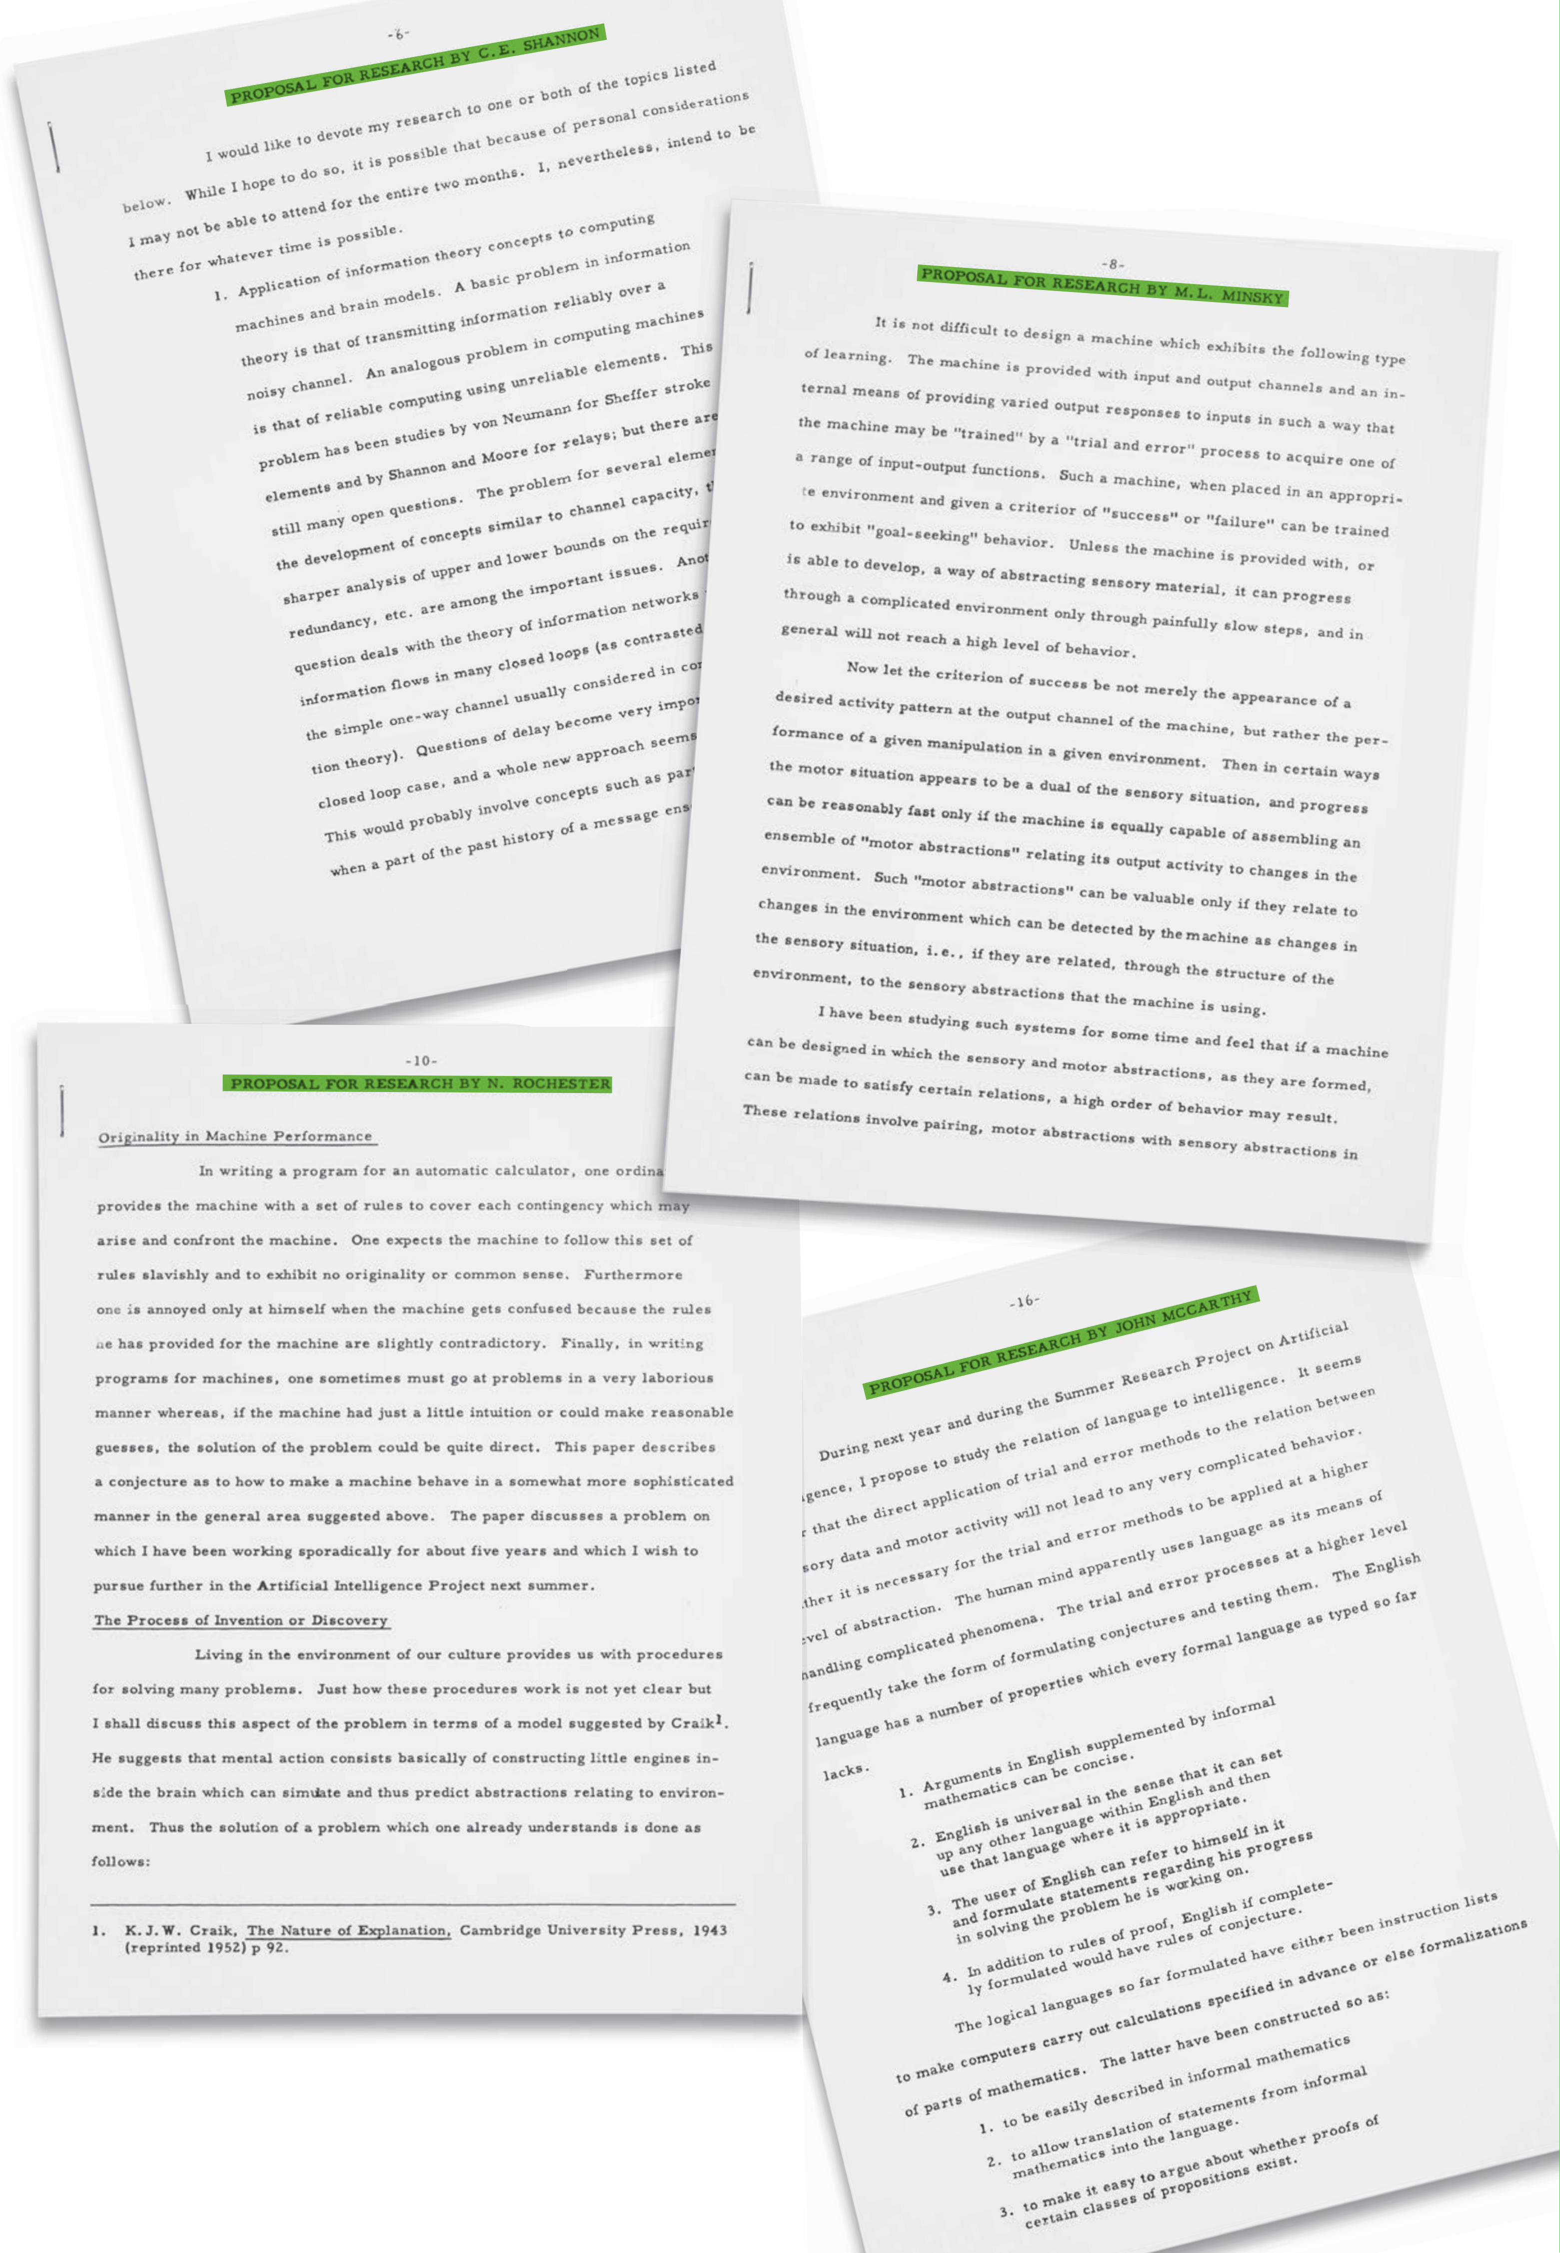

-6-

PROPOSAL FOR RESEARCH BY C.E. SHANNON

I would like to devote my research to one or both of the topics listed below. While I hope to do so, it is possible that because of personal considerations I may not be able to attend for the entire two months. I, nevertheless, intend to be there for whatever time is possible.

1. Application of information theory concepts to computing machines and brain models. A basic problem in information theory is that of transmitting information reliably over a noisy channel. An analogous problem in computing machines is that of reliable computing using unreliable elements. This problem has been studies by von Neumann for Sheffer stroke elements and by Shannon and Moore for relays; but there are still many open questions. The problem for several elemen the development of concepts similar to channel capacity, t sharper analysis of upper and lower bounds on the requir redundancy, etc. are among the important issues. Anot question deals with the theory of information networks information flows in many closed loops (as contrasted the simple one-way channel usually considered in co tion theory). Questions of delay become very impo closed loop case, and a whole new approach seems This would probably involve concepts such as par when a part of the past history of a message ens

-8-

PROPOSAL FOR RESEARCH BY M.L. MINSKY

It is not difficult to design a machine which exhibits the following type of learning. The machine is provided with input and output channels and an internal means of providing varied output responses to inputs in such a way that the machine may be "trained" by a "trial and error" process to acquire one of a range of input-output functions. Such a machine, when placed in an appropri-te environment and given a criterior of "success" or "failure" can be trained to exhibit "goal-seeking" behavior. Unless the machine is provided with, or is able to develop, a way of abstracting sensory material, it can progress through a complicated environment only through painfully slow steps, and in general will not reach a high level of behavior.

Now let the criterion of success be not merely the appearance of a desired activity pattern at the output channel of the machine, but rather the performance of a given manipulation in a given environment. Then in certain ways the motor situation appears to be a dual of the sensory situation, and progress can be reasonably fast only if the machine is equally capable of assembling an ensemble of "motor abstractions" relating its output activity to changes in the environment. Such "motor abstractions" can be valuable only if they relate to changes in the environment which can be detected by the machine as changes in the sensory situation, i.e., if they are related, through the structure of the environment, to the sensory abstractions that the machine is using.

I have been studying such systems for some time and feel that if a machine can be designed in which the sensory and motor abstractions, as they are formed, can be made to satisfy certain relations, a high order of behavior may result. These relations involve pairing, motor abstractions with sensory abstractions in

-10-

PROPOSAL FOR RESEARCH BY N. ROCHESTER

Originality in Machine Performance

In writing a program for an automatic calculator, one ordina provides the machine with a set of rules to cover each contingency which may arise and confront the machine. One expects the machine to follow this set of rules slavishly and to exhibit no originality or common sense. Furthermore one is annoyed only at himself when the machine gets confused because the rules he has provided for the machine are slightly contradictory. Finally, in writing programs for machines, one sometimes must go at problems in a very laborious manner whereas, if the machine had just a little intuition or could make reasonable guesses, the solution of the problem could be quite direct. This paper describes a conjecture as to how to make a machine behave in a somewhat more sophisticated manner in the general area suggested above. The paper discusses a problem on which I have been working sporadically for about five years and which I wish to pursue further in the Artificial Intelligence Project next summer.

The Process of Invention or Discovery

Living in the environment of our culture provides us with procedures for solving many problems. Just how these procedures work is not yet clear but I shall discuss this aspect of the problem in terms of a model suggested by Craik[1]. He suggests that mental action consists basically of constructing little engines inside the brain which can simulate and thus predict abstractions relating to environment. Thus the solution of a problem which one already understands is done as follows:

1. K.J.W. Craik, The Nature of Explanation, Cambridge University Press, 1943 (reprinted 1952) p 92.

-16-

PROPOSAL FOR RESEARCH BY JOHN McCARTHY

During next year and during the Summer Research Project on Artificial igence, I propose to study the relation of language to intelligence. It seems that the direct application of trial and error methods to the relation between sory data and motor activity will not lead to any very complicated behavior. ther it is necessary for the trial and error methods to be applied at a higher evel of abstraction. The human mind apparently uses language as its means of handling complicated phenomena. The trial and error processes at a higher level frequently take the form of formulating conjectures and testing them. The English language has a number of properties which every formal language as typed so far lacks.

1. Arguments in English supplemented by informal mathematics can be concise.
2. English is universal in the sense that it can set up any other language within English and then use that language where it is appropriate.
3. The user of English can refer to himself in it and formulate statements regarding his progress in solving the problem he is working on.
4. In addition to rules of proof, English if completely formulated would have rules of conjecture.

The logical languages so far formulated have either been instruction lists to make computers carry out calculations specified in advance or else formalizations of parts of mathematics. The latter have been constructed so as:

1. to be easily described in informal mathematics
2. to allow translation of statements from informal mathematics into the language.
3. to make it easy to argue about whether proofs of certain classes of propositions exist.

THE ENGINE
Built by MIT

@enginexyz
www.engine.xyz
501 Massachusetts Ave, Cambridge, MA 02139