Unmatched Global Market Coverage

aiworld

Artificial Intelligence UPDATE

Machine Learning - Deep Learning - Software Research
Products - Technology - Applications - Industry News

AI World is a Global Publication of e2mos

Q1-2018

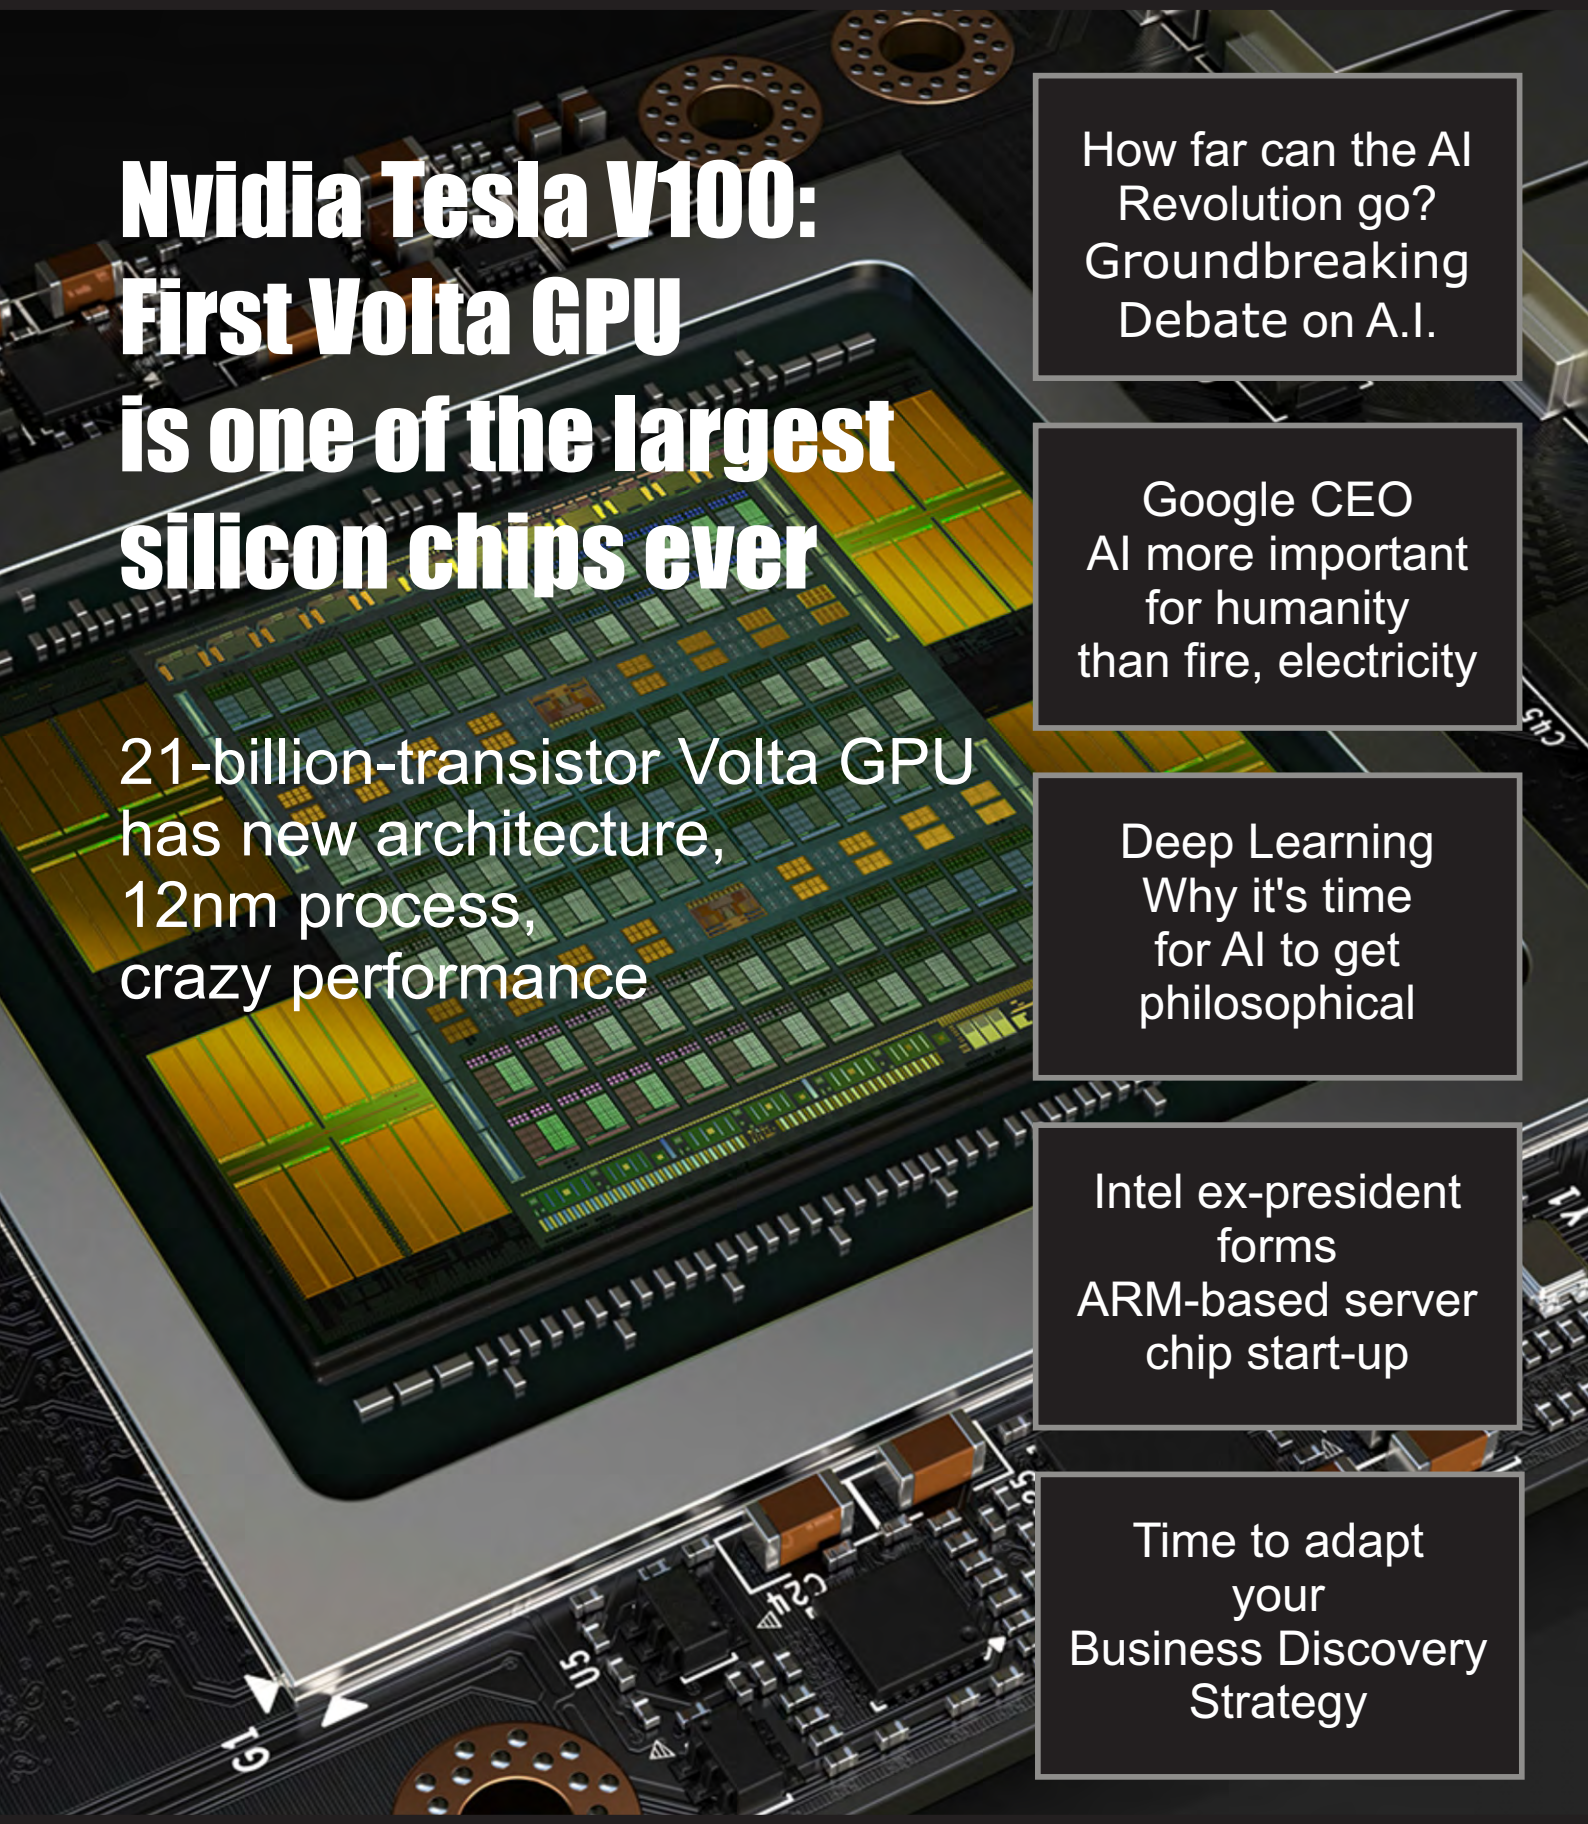

Nvidia Tesla V100: First Volta GPU is one of the largest silicon chips ever

21-billion-transistor Volta GPU has new architecture, 12nm process, crazy performance

How far can the AI Revolution go? Groundbreaking Debate on A.I.

Google CEO AI more important for humanity than fire, electricity

Deep Learning Why it's time for AI to get philosophical

Intel ex-president forms ARM-based server chip start-up

Time to adapt your Business Discovery Strategy

In this Edition:

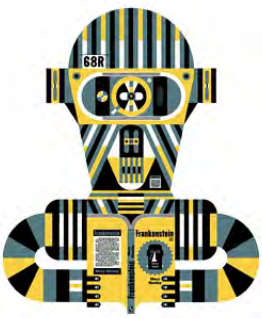

Daniel Dierickx
CEO & co-Founder
at e2mos
Acting Chief Editor

Over 3 decades
Chips & Embedded Systems
Global Market Expertise

Dear Reader,

Here is your free copy of **AI World**, one of our five magazines published by e2mos.

Those e-magazines are distributed Free of charge WORLDWIDE to our PREMIER Database.

Each e-magazines has its own Dedicated Website.

Our aim is to provide you with relevant information in relation with your business, we are listening to our READERS in order to adapt the content & the right mix.

In the mean time we have produced and distributed about 2.000 Editions.

FREE Subscription
Click on the logos

aiworld
IoT World
Telecom COTS World
Broadband Broadcast IoT AI Convergence
Embedded Systems World
ATCA World

Editor/Publisher: e2mos
WEB: www.e2mos.com
Contact: mgt@e2mos.com

About e2mos SERVICES
See last page:
• Business Development
• Coaching Biz Discovery
• Publications & e-mailings

How far can the AI revolution go?

See a groundbreaking debate on artificial intelligence between renowned physicist Sir Roger Penrose and neuroscientist Dr. Eugene Izhikevich.

By: Matt Grob Executive Vice President, Technology Qualcomm Technologies, Inc. -- Feb 1, 2018
Qualcomm Research is a Division of Qualcomm Technologies, Inc.

AI debate participants (l to r):
- James Tagg,
- Sir Roger Penrose,
- Matt Grob (moderator),
- Dr. Eugene Izhikevich,
- Oleg Sinyavskiy.

We're in the early days of AI development, and right now, the entire industry has more questions than answers on how far we can take the AI revolution. There are a lot of unknowns, and there's an interesting contrast between what AI is now and what it can be. Yes, AI is already doing tremendous things and improving at exponential rates, but it's still incapable of doing many tasks that humans consider simple or that children can learn in minutes given a few examples. We simply don't completely know how the human brain works and why it's so efficient at so many tasks.

Qualcomm Research has been researching AI for over a decade, starting with brain-inspired spiking neural networks, to make the devices and the things we interact with everyday more intelligent. We've worked closely with other leaders in AI, from academia to industry, to push AI forward. As you can imagine, we've encountered diverse opinions about the future of AI and have continued those discussions and debates to advance our thinking and knowledge, as well as the industry.

To that end, we brought together two world-renowned thinkers in the AI space for a friendly debate: mathematical physicist, author, and professor Sir Roger Penrose and Dr. Eugene Izhikevich, a neuroscientist and the co-founder, Chairman, and CEO of Brain Corporation. What could be better than two titans of AI, with contrasting opinions, having a healthy debate? I had the honor to moderate the debate, which was hosted by Joan and Dr. Irwin Jacobs. The debate format comprised opening statements, moderated questions with the opportunity for rebuttal, and audience questions followed by final statements.

Watch the entire debate CLICK HERE

The premise of the debate centered around how far AI will progress. Are we on the verge of artificial general intelligence – machines that possess the true understanding and consciousness that human brains display? Is it possible with the continued advancement of today's digital technology – custom hardware architectures, novel machine learning algorithms, and more digital gates – to ever achieve such goals? Or, is it the case that digital gates will never be enough and rather we must look to new physics to achieve the conscious intelligence that humans possess?

« An intelligent thing would have to be conscious, and it would have to have understanding. » Sir Roger Penrose

Eugene made the case that modelling neurons with digital technology can achieve general intelligence that would appear conscious. In effect, a large enough collection of gates based on classical physics is sufficient to emulate the human mind. Roger, however, argued that we need to modify our theories of the physical world to achieve true brain-like functions, such as consciousness and genuine understanding. Specifically, the rules of quantum mechanics need to change in the light of general relativity. His view is that any large collection of digital gates will always fall short of general intelligence.

« You can have intelligent behavior without being conscious. » Dr. Eugene Izhikevich

The point of the debate was not to crown a winner and loser. Rather, it was to learn something new and advance the thinking about AI. The future of AI is still to be determined, but here are some other key insights I gained:
• There is no agreed upon test for consciousness or understanding. As Eugene explained, as Claude Shannon helped define information theory, we need to define consciousness to further our progress.
• While existing neuron models for AI are "simple" and may not achieve consciousness, they are providing amazing capabilities today that will only continue to improve.
• The brain is a complicated biological and chemical system, but biology may not be the only way to achieve consciousness.

Don't just take my word for it: You can watch highlights of this fascinating debate above, or check it out in its entirety on YouTube.

Source and more links CLICK HERE.

S. Korea's ETRI develops 9-core processor for self-driving cars

2017.09.20 14:12:53 -- By Kim Yoon-jin
Source: PULSE by Maeil Business News Korea CLICK HERE

South Korea's Electronics and Telecommunications Research Institute (ETRI) has developed a nine-core high-performance processor controlling a self-driving car, a move that is expected to reduce Korean autonomous car developers' reliance on import processors.

ETRI announced on Tuesday that it succeeded in developing a nine-core processor that will act as a brain of a self-driving car. The new autonomous car processor, dubbed Aldebaran, is an upgraded version of the quad-core processor that ETRI introduced last year.

The chip with a dimension of 7.8 x 6.7 millimeters allows a self-driving car to detect movement of people and other cars on the street and process images while consuming only around 1 watt of electricity. Also, with 9 million calculations per second, the chip can process data faster and generate clearer images than its predecessor.

Aldebaran is expected to bring a big change to local self-driving tech industry that has mainly relied on foreign country-made processors. The latest upgrades that ensure simultaneous processing of ultra high definition images and motion identification, as well as a 99 percent rate of auto detection of breakdowns in the system are expected to accelerate the development of indigenous self-driving car processors, industry observers said.

The government-funded research institute aims to commercialize the upgraded processor through a local semiconductor designing company by next year. It already transferred the quad-core processor technology to NextChip for commercialization.

ETRI plans to continue upgrading the processor to the level that enables self-driving cars to recognize all moving objects without human assistance as well as developing a chip that can make a machine select destination and search routes on its own, said Kwon Young-soo, a researcher at ETRI.

Sundar Pichai: AI more important for humanity than fire, electricity

However, Google's AI push comes with plenty of people problems.

MENLO PARK, CALIFORNIA - February 02, 2018: Google CEO Sundar Pichai recently declared that artificial intelligence fueled by powerful computers was more important to humanity than fire or electricity. And yet the search giant increasingly faces a variety of messy people problems as well.

The company has vowed to employ thousands of human checkers just to catch rogue YouTube posters, Russian bots and other purveyors of unsavory content. It's also on a buying spree to find office space for its burgeoning workforce in pricey Silicon Valley.

For a company that built its success on using faceless algorithms to automate many human tasks, this focus on people presents something of a conundrum. Yet it's also a necessary one as lawmakers ramp up the pressure on Google to deter foreign powers from abusing its platforms and its YouTube unit draws fire for offensive videos , particularly ones aimed at younger audiences.

In the latest quarter alone, Google parent Alphabet Inc. added 2,009 workers, for a total of 80,110. Over the last three years, it hired a net 2,245 people per quarter on average. That's nearly 173 per week, or 25 people per day.

Some of the extra workers this year will be part of Google's pledge to have 10,000 people across the company snooping out videos and other material that violate the company's policies _ but which computers can't catch on their own. That program will lead to what Google calls "significant growth " in personnel.

Google will take on even more workers in the current quarter now that it has closed its $1.1 billion purchase of part of hardware maker HTC, bringing onboard the 2,000-plus engineers who worked on the Pixel smartphone line.

On Thursday, Pichai spoke bullishly about content-checkers hiring, saying the investments now set the company up to capture growth in the future _ in the same call with investors that he touted self-driving vehicles developed by Alphabet's Waymo unit, which aim to do away with human drivers entirely.

ALSO READ: Cryptocurrency out, blockchain, AI & IoT tech get a dekho

For instance, Pichai said he sees consumers increasingly watching YouTube videos on connected TVs in the living room, a lucrative segment of growth for the digital video advertising that helps power Google's growth.

After controversies over YouTube stars who made anti-Semitic comments or showed video of someone who had apparently died by suicide, Google has tightened its standards . It has limited which YouTube channels can serve up ads; vowed to manually review every video in its most popular channels for 18-to-34-year-olds; and will pay outside companies to ensure that brands don't have their ads turn up next to unsuitable videos.

"While there have been some concerns, we're working really hard to address them and respond strongly," Pichai said.

Some analysts aren't so sure. Collin Colburn, an analyst with market researcher Forrester, wonders how much of the recent changes are just window dressing at a company for whom hiring thousands of people amounts to little more than pocket change.

"I wonder if it's more of a move of optics rather than practicality," Colburn said, noting Google's "massive" double-digit revenue growth and cash hoard of $102 billion.

Revenue at Google parent Alphabet rose 24 percent from a year ago to $32.32 billion. After subtracting advertising commissions, revenue was $25.87 billion, exceeding Street forecasts of $25.65 billion. But the company swung to a $3 billion loss from a $5.33 billion profit a year earlier, reflecting the recent federal tax overhaul.

Alphabet shares were down 2.3 percent at $1,141.42 in after-hours trading.

Google's growing workforce has the company on a real-estate tear.

It recently opened up offices in Austin, Texas; Ann Arbor, Michigan; Boulder, Colorado; and is planning to open offices in Detroit. It broke ground in November on a huge office building in the heart of London, home to its DeepMind artificial intelligence unit, that will come complete with a rooftop running track.

ALSO READ: Google develops human-like text-to-speech AI

Near its current headquarters, construction is underway on two futuristic dome-like structures infused with natural light, brimming with solar panels and set to open in late 2019. Google is negotiating with the city of Mountain View to add 10,000 housing units, many of which will likely be home to employees known as "Googlers."

Pichai said the company intends to hire "thousands of people across the U.S." this year, build or open five new data centers, and make "significant investments" in nine states.

Source ETTelecom and more LINKS: CLICK HERE

World’s Largest Server Companies Announce NVIDIA Volta Systems Supercharged for AI

Wednesday, September 27, 2017

Dell EMC, Hewlett Packard Enterprise, IBM, Supermicro Announce Servers Based on NVIDIA Tesla V100 Accelerators -- World's Most Advanced Data Center GPU

SANTA CLARA, CA - NVIDIA (NASDAQ: NVDA) and its systems partners Dell EMC, Hewlett Packard Enterprise, IBM and Supermicro today unveiled more than 10 servers featuring NVIDIA® Volta architecture-based Tesla® V100 GPU accelerators -- the world's most advanced GPUs for AI and other compute-intensive workloads.

NVIDIA V100 GPUs, with more than 120 teraflops of deep learning performance per GPU, are uniquely designed to deliver the computing performance required for AI deep learning training and inferencing, high performance computing, accelerated analytics and other demanding workloads. A single Volta GPU offers the equivalent performance of 100 CPUs, enabling data scientists, researchers and engineers to tackle challenges that were once impossible.

Nvidia Tesla V100: First Volta GPU is one of the largest silicon chips ever

21-billion-transistor Volta GPU has new architecture, 12nm process, crazy performance

Seizing on the AI computing capabilities offered by NVIDIA's latest GPUs, **Dell EMC**, **HPE**, **IBM** and **Supermicro** are bringing to the global market a broad range of multi-V100 GPU systems in a variety of configurations.

"Volta systems built by our partners will ensure that enterprises around the world can access the technology they need to accelerate their AI research and deliver powerful new AI products and services," said Ian Buck, vice president and general manager of Accelerated Computing at NVIDIA.

V100-based systems announced include:

Dell EMC -- The PowerEdge R740 supporting up to three V100 GPUs for PCIe, the PowerEdge R740XD supporting up to three V100 GPUs for PCIe, and the PowerEdge C4130 supporting up to four V100 GPUs for PCIe or four V100 GPUs for NVIDIA NVLink™ interconnect technology in an SXM2 form factor.

HPE -- HPE Apollo 6500 supporting up to eight V100 GPUs for PCIe and HPE ProLiant DL380 systems supporting up to three V100 GPUs for PCIe.

IBM -- The next generation of IBM Power Systems servers based on the POWER9 processor will incorporate multiple V100 GPUs and take advantage of the latest generation NVLink interconnect technology -- featuring fast GPU-to-GPU interconnects and an industry-unique OpenPOWER CPU-to-GPU design for maximum throughput.

Supermicro -- Products supporting the new Volta GPUs include a 7048GR-TR workstation for all-around high-performance GPU computing, 4028GR-TXRT, 4028GR-TRT and 4028GR-TR2 servers designed to handle the most demanding deep learning applications, and 1028GQ-TRT servers built for applications such as advanced analytics.

These partner systems complement an announcement yesterday by China's leading original equipment manufacturers -- including **Inspur**, **Lenovo** and **Huawe**i -- that they are using the Volta architecture for accelerated systems for hyperscale data centers.

... to next page

World's Largest Server Companies Announce NVIDIA Volta Systems Supercharged for AI

... from previous page

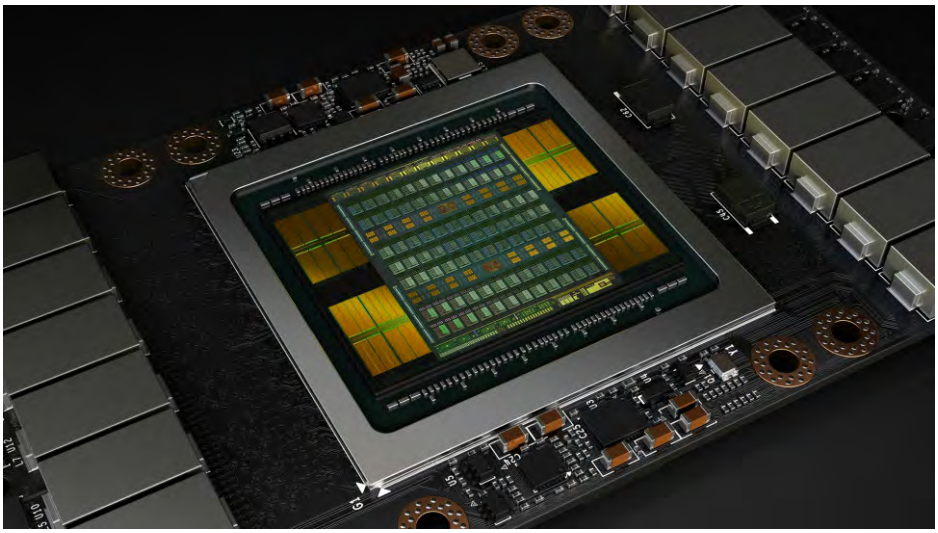

Additional NVIDIA V100 Details

Each NVIDIA V100 GPU features over 21 billion transistors, as well as 640 Tensor Cores, the latest NVLink high-speed interconnect technology, and 900 GB/sec HBM2 DRAM to achieve 50 percent more memory bandwidth than previous generation GPUs.

V100 GPUs are supported by NVIDIA Volta-optimized software, including CUDA® 9.0 and the newly updated deep learning SDK, including TensorRT™ 3, DeepStream SDK and cuDNN 7 as well as all major AI frameworks. Additionally, hundreds of thousands of GPU-accelerated applications are available for accelerating a variety of data-intensive workloads, including AI training and inferencing, high performance computing, graphics and advanced data analytics.

Partner Quotes

"One of the core principles for Dell EMC is to deliver differentiated solutions to our customers so that they can leverage the most advanced technology for a competitive advantage. To that end, we are proud of the work we do with partners like NVIDIA to build PowerEdge servers ideal for compute-intensive workloads including data analytics, high-performance computing, machine learning and AI."
-- Armughan Ahmad, senior vice president and general manager of Hybrid Cloud and Ready Solutions at Dell EMC

"As deep learning continues to become more pervasive, technology advancements across systems and accelerators need to evolve in order to gain intelligence from large datasets faster than ever before. The HPE Apollo 6500 and HPE ProLiant DL380 systems combine the industry-leading GPU performance of NVIDIA Tesla V100 GPU accelerators and Volta architecture with HPE unique innovations in system design and manageability to deliver unprecedented levels of performance, scale and efficiency for high performance computing and artificial intelligence applications."
-- Bill Mannel, vice president and general manager of High Performance Computing and Artificial Intelligence at Hewlett Packard Enterprise

"IBM's upcoming POWER9 servers will support NVIDIA's Volta GPU, and will be the only one to support the latest generation of NVLink and PCIe 4.0, which will deliver maximum throughput. With accelerators like Volta, IBM will scale deep learning performance to new heights."
-- Brad McCredie, vice president and IBM Fellow, Cognitive Systems Development at IBM

"Supermicro designs the most application-optimized GPU systems and offers the widest selection of GPU-optimized servers and workstations in the industry. Our high performance computing solutions enable deep learning, engineering and scientific fields to scale out their compute clusters to accelerate their most demanding workloads and achieve fastest time-to-results with maximum performance per watt, per square foot and per dollar. With our latest innovations incorporating the new NVIDIA V100 PCI-E and V100 SXM2 GPUs in performance-optimized 1U and 4U architectures with next-generation NVLink, our customers can accelerate their applications and innovations to help solve the world's most complex and challenging problems."
-- Charles Liang, president and CEO of Supermicro

About NVIDIA

NVIDIA's (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI -- the next era of computing -- with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at http://nvidianews.nvidia.com/.

Deep learning: Why it's time for AI to get philosophical

For years, science fiction writers have spelled out the technological marvels and doomsday scenarios that might result from artificial intelligence. Now that it's a part of our lives, argues Catherine Stinson, those working in AI need to take their work's social and ethical implications much more seriously

CATHERINE STINSON
SPECIAL TO THE GLOBE AND MAIL
PUBLISHED MARCH 23, 2018UPDATED MARCH 25, 2018

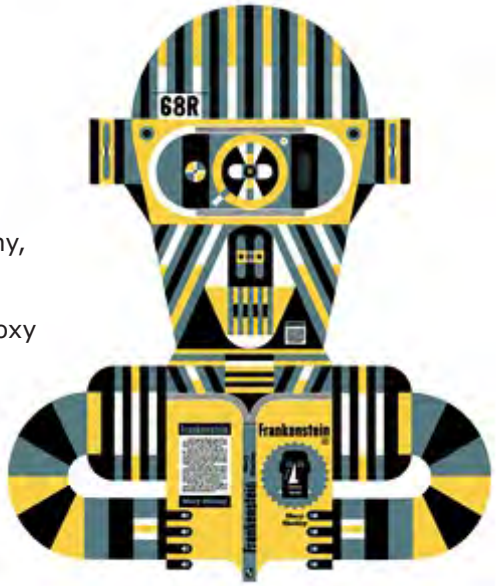

ILLUSTRATIONS BY RAYMOND BIESINGER

Catherine Stinson is a postdoctoral scholar at the Rotman Institute of Philosophy, at the University of Western Ontario, and former machine-learning researcher.
I wrote my first lines of code in 1992, in a high school computer science class. When the words "Hello world" appeared in acid green on the tiny screen of a boxy Macintosh computer, I was hooked. I remember thinking with exhilaration, "This thing will do exactly what I tell it to do!" and, only half-ironically, "Finally, someone understands me!" For a kid in the throes of puberty, used to being told what to do by adults of dubious authority, it was freeing to interact with something that hung on my every word – and let me be completely in charge.

For a lot of coders, the feeling of empowerment you get from knowing exactly how a thing works – and having complete control over it – is what attracts them to the job. Artificial intelligence (AI) is producing some pretty nifty gadgets, from self-driving cars (in space!) to automated medical diagnoses. The product I'm most looking forward to is real-time translation of spoken language, so I'll never again make gaffes such as telling a child I've just met that I'm their parent or announcing to a room full of people that I'm going to change my clothes in December.

But it's starting to feel as though we're losing control.

These days, most of my interactions with AI consist of shouting, "No, Siri! I said Paris, not bratwurst!" And when my computer does completely understand me, it no longer feels empowering. The targeted ads about early menopause and career counselling hit just a little too close to home, and my Fitbit seems like a creepy Santa Claus who knows when I am sleeping, knows when I'm awake and knows if I've been bad or good at sticking to my exercise regimen.

Algorithms tracking our every step and keystroke expose us to dangers much more serious than impulsively buying wrinkle cream. Increasingly polarized and radicalized political movements, leaked health data and the manipulation of elections using harvested Facebook profiles are among the documented outcomes of the mass deployment of AI. Something as seemingly innocent as sharing your jogging routes online can reveal military secrets. These cases are just the tip of the iceberg. Even our beloved Canadian Tire money is being repurposed as a surveillance tool for a machine-learning team.

For years, science-fiction writers have spelled out both the technological marvels and the doomsday scenarios that might result from intelligent technology that understands us perfectly and does exactly what we tell it to do. But only recently has the inevitability of tricorders, robocops and constant surveillance become obvious to the non-fan general public. Stories about AI now appear in the daily news, and these stories seem to be evenly split between hyperbolically self-congratulatory pieces by people in the AI world, about how deep learning is poised to solve every problem from the housing crisis to the flu, and doom-and-gloom predictions of cultural commentators who say robots will soon enslave us all. Alexa's creepy midnight cackling is just the latest warning sign.


The continuum of creepiness between a Fitbit, above, and 2001: A Space Odyssey's mad computer Hal 9000 may not be as wide as we imagine. MGM/PHOTOFEST

Which side should we believe? As someone who has played for both teams, my answer is that there are indeed genuine risks we ought to be very worried about, but the most serious risks are not the ones getting the most attention. There are also things we can do to make sure AI is developed in ways that benefit society, but many of the current efforts to address the risk that AI will run amok are misdirected.

One scenario I don't think we need to worry about is the one in which robots become conscious (whatever that means) and decide to kill us. For one thing, conscious robots are still far off. For another, consciousness doesn't lead directly to murder – it's a non sequitur. Another scenario I find far-fetched is the one in which a computer tasked with maximizing profits for a paperclip company ends up turning all the resources on the planet (and beyond) into paperclips. Sure, it's possible, but not very probable. We should be able to switch off that malfunctioning HAL 9000. One thing we should be very concerned about, however, are weaponized drones. Automated killing machines are not going to make our lives better, period.

... to next page

Deep learning: Why it's time for AI to get philosophical

... from previous page

The other serious risk is something I call nerd-sightedness: the inability to see value beyond one's own inner circle. There's a tendency in the computer-science world to build first, fix later, while avoiding outside guidance during the design and production of new technology. Both the people working in AI and the people holding the purse strings need to start taking the social and ethical implications of their work much more seriously.

In medicine and engineering, there are codes of conduct that professionals are expected to follow. The idea that scientists bear some responsibility for the technologies their work makes possible is also well established in nuclear physics and genetics, even though scientists don't make the final decision to push the red button or genetically engineer red-headed babies. In the behavioral sciences, there are research-ethics boards that weigh the possible harms to participants in proposed experiments against the benefits to the population. Studies whose results are expected to cause societal harm don't get approved. In computer science, ethics is optional.

Mary Shelley's Frankenstein, which turned 200 this year, is perhaps the most famous warning to scientists not to abdicate responsibility for their creations. Victor Frankenstein literally runs away after seeing the ugliness of his creation, and it is this act of abandonment that leads to the creature's vengeful, murderous rampage. Frankenstein begins with the same lofty goal as AI researchers currently applying their methods to medicine: "What glory would attend the discovery if I could banish disease from the human frame and render man invulnerable to any but a violent death!" In a line dripping with dramatic irony, Frankenstein's mentor assures him that "the labours of men of genius, however erroneously directed, scarcely ever fail in ultimately turning to the solid advantage of mankind." Shelley knew how dangerous this egotistical attitude could be.

But the nerd-sighted geniuses of our day make the same mistake. If you ask a coder what should be done to make sure AI does no evil, you're likely to get one of two answers, neither of which is reassuring. Answer No. 1: "That's not my problem. I just build it," as exemplified recently by a Harvard computer scientist who said, "I'm just an engineer" when asked how a predictive policing tool he developed could be misused. Answer No. 2: "Trust me. I'm smart enough to get it right." AI researchers are a smart bunch, but they have a terrible track record of avoiding ethical blunders. Some of the better-known goof-ups include Google images tagging black people as gorillas, chat bots that turn into Nazis and racist soap dispensers. The consequences can be much more serious when biased algorithms are in charge of deciding who should be approved for a bank loan, who to hire or admit to university or whether to kill a suspect in a police chase.

There is a growing movement within AI to do better. For the first time, this past December, one of the top AI conferences, NIPS, featured a keynote about bias in AI. The New York Times recently reported that top computer-science schools such as MIT and Stanford are rushing to roll out ethics courses in response to a newfound awareness that the build-it-first, fix-it-later ethos isn't cutting it.

Tay, a chatbot designed by Microsoft, was introduced in 2016 as an experiment to teach an AI how to hold conversations with people. Within 24 hours, Twitter users had taught it to say racist and misogynist things. Microsoft quickly closed off the bot to the public. TWITTER

In fact, such courses have existed for a long time. The University of Toronto's computer-science program has a computers and society course, which includes a few weeks of ethics. Philosophy departments and, where they exist, interdisciplinary departments such as science and technology studies, history and philosophy of science and media studies also routinely offer courses covering the social and ethical implications of technology. Unfortunately, these courses are rarely required for computer-science degrees, so most graduates don't get any ethics training at all. I'm currently teaching a course about the social and ethical implications of AI in health care – to about a dozen students.

That's what may be changing. When I contacted François Pitt, the undergraduate co-ordinator of U of T's computer science department, he said the issue of ethics training had come up four or five times just that week. Steve Easterbrook, the professor in charge of the computers and society course, agrees that there needs to be more ethics training, noting that moral reasoning in computer-science students has been shown to be "much less mature than students from most other disciplines." Instead of one course on ethics, he says "it ought to be infused across the curriculum, so that all students are continually exposed to it." Prof. Zeynep Tufekci of the University of North Carolina at Chapel Hill agrees, commenting on Twitter, "You have to integrate ethics into the computer science curriculum from the first moment you teach students here's a variable, here's an array, and look, we can sort things. When you run your first qsort, [a standard sorting algorithm encountered in computer-science 101 courses], you've encountered ethical and ontological questions."

It remains to be seen how computer-science departments can pull off that sort of sweeping curriculum change, given that many of their faculty have just as little ethics training as their students. Philosophy departments regularly face the frustration of having the expertise to teach applied ethics courses, only to have those courses get poached by law, medicine and engineering departments, which require their students to take in-house ethics courses that are often taught by non-experts. Hopefully, computer-science departments expanding their ethics offerings will see the need to hire trained ethicists to design and teach the courses. Frankly, opportunities for philosophers to make themselves useful are rare enough that we shouldn't pass them over in cases where their training has real value.

... to next page

Deep learning: Why it's time for AI to get philosophical

... from previous page

Another kind of effort at fixing AI's ethics problem is the proliferation of crowdsourced ethics projects, which have the commendable goal of a more democratic approach to science. One example is DJ Patil's Code of Ethics for Data Science, which invites the data-science community to contribute ideas but doesn't build up from the decades of work already done by philosophers, historians and sociologists of science. Then there's MIT's Moral Machine project, which asks the public to vote on questions such as whether a self-driving car with brake failure ought to run over five homeless people rather than one female doctor. Philosophers call these "trolley problems" and have published thousands of books and papers on the topic over the past half-century. Comparing the views of professional philosophers with those of the general public can be eye-opening, as experimental philosophy has repeatedly shown, but simply ignoring the experts and taking a vote instead is irresponsible.

The point of making AI more ethical is so it won't reproduce the prejudices of random jerks on the internet. Community participation throughout the design process of new AI tools is a good idea, but let's not do it by having trolls decide ethical questions. Instead, representatives from the populations affected by technological change should be consulted about what outcomes they value most, what needs the technology should address and whether proposed designs would be usable given the resources available. Input from residents of heavily policed neighbourhoods would have revealed that a predictive policing system trained on historical data would exacerbate racial profiling. Having a person of colour on the design team for that soap dispenser should have made it obvious that a peachy skin tone detector wouldn't work for everyone. Anyone who has had a stalker is sure to notice the potential abuses of selfie drones. Diversifying the pool of talent in AI is part of the solution, but AI also needs outside help from experts in other fields, more public consultation and stronger government oversight.

Another recent article in The New York Times claimed that "academics have been asleep at the wheel," leaving policy makers who are struggling to figure out how to regulate AI at the mercy of industry lobbyists. The article set off a Twitter storm of replies from philosophers, historians and sociologists of science, angry that their decades of underfunded work is again being ignored and erased. Like the Whos down in Whoville, they cried out in fear, "We are here! We are here! We are here! We are here!" If policy makers and funding sources listen closely to those voices, there are solutions being offered. The article concludes that we "urgently need an academic institute focused on algorithmic accountability." On Twitter, the article's author, Cathy O'Neil, insisted, "There should be many many more tenure lines devoted to it." Those both sound like solid ideas.

Elderly people play with a robot named NAO, manufactured by Softbank Robotics, in their retirement home in Bordeaux, France.
REGIS DUVIGNAU/REUTERS

How does this play out in the Canadian context? The federal government recently earmarked $125-million for a Pan-Canadian Artificial Intelligence Strategy – with the Canadian Institute for Advanced Research (CIFAR) in charge – and three new research institutes in Edmonton, Toronto and Montreal. One of the goals is to "develop global thought leadership on the economic, ethical, policy and legal implications of advances in artificial intelligence."

There are promising signs already. An International Scientific Advisory Committee on AI Strategy was recently announced and includes a co-founder of the Future of Life Institute, which has a mission to safeguard life against technological challenges. CIFAR just launched an AI and society program featuring workshops and publications designed to stimulate discussion and guide policy makers on the ethical, legal, economic and policy issues AI presents for society. The Alberta Machine Intelligence Institute so far focuses on industry partnerships, while Toronto's Vector Institute has particular strengths in deep learning and medical applications of AI. Vector Institute research director Richard Zemel has a growing side project in algorithmic fairness and has been making connections with legal scholars and researchers from the University of Toronto's Centre for Ethics. Another group is studying AI and safety. Prof. Zemel promises that bringing in social scientists, ethicists and policy analysts is something he is "determined to do."

The Montreal Institute for Learning Algorithms (MILA) has taken several concrete steps toward addressing AI's ethical challenges, including offering diversity scholarships, hosting an interdisciplinary forum on the socially responsible development of AI and announcing the Montreal Declaration for a Responsible Development of AI. Although this too is a crowdsourced ethical code, there's a key difference: Philosophers, social scientists and policy analysts were involved in its inception and will oversee the final document.

What's still lacking is investment in jobs dedicated to the social, ethical and policy implications of AI. The Vector Institute is hiring research scientists and postdocs, and MILA is hiring interns, postdocs and professors, but in each case, expertise in machine learning is the main job qualification. CIFAR's AI and society workshops propose to bring together "interdisciplinary teams to explore emerging questions about how AI could affect the world," but there are precious few jobs in Canada (or elsewhere) for people who study AI's effects from the perspective of the humanities and social sciences. Alongside the creation of new tenure lines and graduate degrees in AI, we need to invest in training and employing experts in the social, ethical and policy implications of AI if we're to have a hope of predicting and preventing dangerous outcomes.

... to next page

Deep learning: Why it's time for AI to get philosophical

... from previous page

A recent CIFAR event illustrates what's still missing in the strategy. The workshop brought together AI and health researchers for a joint project to develop systems to do things such as predict flu outbreaks. Sounds good, but work by philosophers of science has shown that putting people in the same boardroom is not enough to make interdisciplinary projects work. Researchers in different fields use different vocabularies, concepts and methods – and may have conflicting goals. Words such as "bias," "naive," " ontology," "regression" and "supervised" have technical meanings in AI that are very different than their usual meanings. Having a translator of sorts at the table to integrate across disciplines is essential to avoid potentially costly misunderstandings. That means inviting philosophers, sociologists and historians of science to the table.

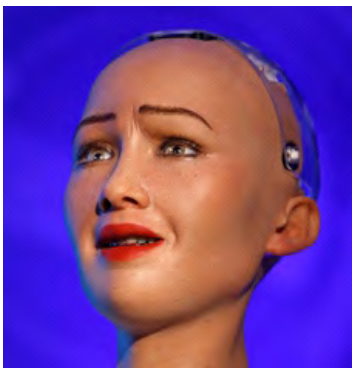

Sophia, a robot with Saudi Arabian citizenship, interacts at an innovation fair in Kathmandu. As the technical achievements of robotics and AI advance, there are few jobs to study artificial intelligence from a humanities or social-science perspective. NAVESH CHITRAKAR

Fahad Razak, an internist at St. Michael's Hospital and an assistant professor of health policy, management and evaluation at the University of Toronto, describes facing just this kind of challenge applying AI methods in his research. The Li Ka Shing Knowledge Institute's General Medicine Inpatient Initiative at St. Mike's applies machine learning to health-care data with the aim of improving treatment, especially for patients with multiple concurrent diagnoses. One of the main challenges with health-care data sets, according to Prof. Razak, is "a data-quality problem, both with data being inaccurate or poorly coded, but also often because of the lack of lots of relevant data." Socioeconomic status and the ability to function without assistance at home are two examples of important data that is often missing. But, he says, computer scientists don't realize how messy and gappy the data is, and because of the hype around AI, health researchers "believe that AI methods are so miraculous, they just work despite data quality problems!"

If every computer-science student were taught to look for social and ethical problems as soon as they learn about variables, arrays and sorting, as Prof. Tufekci suggests, they would realize that the way health data is coded matters very much. For example, recording age as a range, such as under 19, 19-24, 25-35, may seem harmless, but whether a patient is 1 or 17 can make a big difference to their health-care needs. An algorithm that doesn't take that into account could make fatal mistakes – for example, by suggesting the wrong dosage. Likewise, since income is usually missing from health data sets, an algorithm could draw false conclusions about how to prevent diabetes or premature births.

These are just a small sample of the mistakes we need to prevent. The current generation of AI researchers (with a few exceptions) do not have the training necessary to deal with the implications of the AI they are building. So far, the experts who do have that training are not being hired to help. That needs to change – or the darkest of science fiction will become reality.

Intel ex-president forms ARM-based server chip start-up

By David Manners, Electronics Weekly
5th February 2018

Renee James, former president of Intel, has set up a company called Ampere to make ARM-based processors for servers.

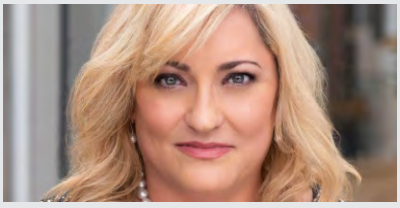

The company is backed by The Carlyle Group.

Ampere processors offer, says the company "a modern architecture with a unique, high performance, custom core Armv8-A 64-bit server operating at up to 3.3 Ghz, 1TB of memory at a power envelope of 125 watts."

The processors are sampling now and will be in production in the second half of the year.

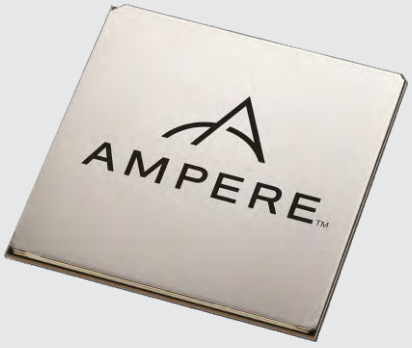

"We have an opportunity with cloud computing to take a fresh approach with products that are built to address the new software ecosystem," says James who is the CEO of Ampere, "the workloads moving to the cloud require more memory, and at the same time, customers have stringent requirements for power, size and costs. The software that runs the cloud enables Ampere to design with a different point of view. The Ampere team's approach and architecture meets the expectation on performance and power and gives customers the freedom to accelerate the delivery of the most memory-intensive applications and workloads such as AI, big data, storage and database in their next-generation data centers."

About Ampere Computing: CLICK HERE

Time to adapt your Business Discovery Strategy

The Future of any Company will depend on their ability to find NEW CUSTOMERS Faster than their Competitors (DDX)

We connect Suppliers with New Customers ...Much Faster

GLOBAL REFERENCES include: IBM, Intel, Microsoft, Cisco/Tail-f Systems, Telco Systems, Artesyn, Motorola, TI DSP, Xilinx, Adlink, Kontron, Radisys, Enea Software, Green Hills Software, Wintegra, Arrow, Avnet, … and more

Our SERVICES for Hi-end HW & SW Vendors addressing OEMs and Service Providers

Meet More New Qualified Customers

Customer Meetings Setup for New Business with the RIGHT Decision Makers

We search for you - Excellent Results based-on Deep Customer Relationship and Product/Market Expertise
OPTION: we join the Customer Meeting

Make a Big Jump Take a Tiger for a while

Coaching How to find More New Customers Strategy Setup Hands-on

Audit: what are you doing today
Targets: what do you want
Our recommendations
How to get there

We bring you to the Frontline

Market Presence Acceleration Massive Global Reach Five e-magazines with High Focus

The KEY-to-SUCCESS of our Magazines is the Quality of our PREMIER Database and the ongoing UPDATING done Every Day by RESEARCH from many sources

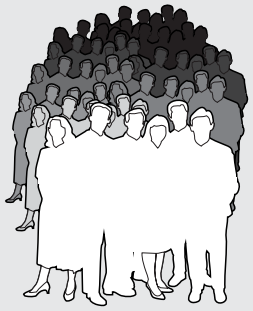

Dedicated Web Site each Click on the Logo's

aiworld

Market Focus: related to Artificial Intelligence from A to Z
Major Players

IoT World

Market Focus: related to Internet of Things
IoT HW & SW - M2M
MEMS - Sensors
Solutions
Service Providers

Embedded Systems World

OEM Focus: Automation, Energy, Transportation, Medical, T&M, Surveillance, Military & Aero, ...
all Markets for **Boards**

Telecom COTS World
Broadband Broadcast IoT Convergence

OEM Focus: Telecom IT, Networks, Data Centers, Cloud, Storage, Broadcast, Video Networks,
Service Providers: 850+

ATCA World

Advanced Telecom Computing Architecture (or AdvancedTCA, ATCA)
Standardized Platform for Carrier Grade Telecom Systems & Hi-end App's

e2mos
Embedded Extreme Marketing & Opportunity Search

Our Service is the Answer to your Biggest Challenge
Finding New Projects - New Customers and the Decision Makers

Made for Hi-end HW & SW VENDORS, based on 3 Decades Global Market Expertise

www.e2mos.com | Contact mgt@e2mos.com | Talk to us today, request a phone call