



CB DIGEST FOR TECHNOLOGY

FEB 29,
2020

Every week CB DIGEST scours many of newspapers, magazines, and websites, searching for the most intriguing tech stories and the most thoughtful things – left, right, and in-between. The CT DIGEST also reports on what the smartest people are saying about the world.

CHAMBIZ 

Find the Article You Need

[Technology and Industrial News](#)

[Acquisition & Funding News](#)

[Robotic piece picking gets more flexible with AI-enabled adaptive tooling](#)

[A boom, a bust, a reckoning, a race: Four takes on today's startup market](#)

[Public markets fall yet again as venture deal counts appear to slip](#)

[Meet ByteDance's Biggest Homegrown Rival](#)

[Prophesee's Event-Based Camera Reaches High Resolution](#)

[Uber's Old Rivals in Southeast Asia Hold Merger Talks](#)

[Fake images can fool autonomous cars, posing risks, Israeli researchers warn](#)

[What's the Difference Between a Classic Curve Tracer and SMU with Curve-Tracer Software?](#)

[Instead of IPOs and acquisitions, exiting to community is one alternative](#)

[Inside 'Amazon Go Grocery': Tech giant opens first full-sized store without cashiers or checkout lines](#)

[Swarming robots avoid collisions, traffic jams](#)

[Anki was developing security robots before shutdown](#)

[Partnership Aims to Speed Delivery of Scalable, 3-in-1 eDrive Systems](#)

[Operators are best placed to serve Industry 4.0 with 5G slices, says Huawei](#)

[Image courtesy of Intel. Intel all in on 5G infrastructure—eyeing \\$25 billion opportunity](#)

[CommScope: '2020 will deliver the 5G promise'](#)

[Andreessen makes Ribbon Health the first investment from its \\$750 million new healthcare fund](#)

[Hustle Fund's Elizabeth Yin discusses 2020's fundraising landscape](#)

[DoorDash, Airbnb Face Wary Investors in Plans to Go Public](#)

[Open Images V6 — Now Featuring Localized Narratives](#)

[Inside the mind of an autonomous delivery robot](#)

[All Processing Bends Toward AI](#)

[California's self-driving car reports are public. Here's what they don't mean](#)

[Buyer Beware: Why Retailers Should Think Hard About Amazon 'Go' Partnerships](#)

[Engineers develop miniaturized 'warehouse robots' for biotechnology applications](#)

[Einride ready to recruit first remote operators for autonomous trucks](#)

[Quaternium drone HYBRiX sets a new World Record of Endurance with a flight of 8 hours and 10 minutes](#)

Google to invest more than \$10 billion in offices and data centers across the United States in 2020

Alphabet Inc's Google announced on Wednesday it would invest more than \$10 billion in offices and data centers across the United States this year. The tech giant, which currently has a presence in 26 states across the country, will now focus the new investments in 11 of them: Colorado, Georgia, Massachusetts, Nebraska, New York, Oklahoma, Ohio, Pennsylvania, Texas, Washington and California. "These investments will create thousands of jobs – including roles within Google, construction jobs in data centers and renewable energy facilities, and opportunities in local businesses in surrounding towns and communities," Chief Executive Officer Sundar Pichai said in a blog post.

Sundar Pichai added: "Everywhere we invest, we strive to create meaningful opportunities for local communities. A powerful example is our data center in Pryor, a town in Mayes County, Oklahoma. Last year, I visited Pryor to announce a \$600 million investment, our fourth expansion there since 2007. It felt like the whole community came out to welcome us, from small business owners to teachers to Google employees. Pryor Mayor Larry Lees told the crowd that Google's investments have helped provide local schools with the resources they need—including the latest textbooks and STEM courses—to offer a world-class education. He talked about the small businesses we have helped train and the mentorship Googlers have provided to Pryor's students."

Google hopes the investments in new office and data center projects in 2020 will create thousands of jobs—including roles within Google, construction jobs in data centers and renewable energy facilities, and opportunities in local businesses in surrounding towns and communities. This effort builds on the momentum of the \$13 billion investment in communities from South Carolina to Nevada we made in 2019. Combined with other R&D investments, Google's parent company Alphabet was the largest investor in the U.S. last year, according to a report from the Progressive Policy Institute.

Sony launches first 5G smartphone in a bid to prop up its struggling smartphone business. Sony has unveiled its first 5G smartphone as it hopes to turn around its struggling mobile business. With the Xperia 1 II flagship handset, Sony has leaned in to its heritage in camera and display technology. The device has a triple lens setup with the ability for users to adjust some of the settings found in professional cameras. The device has a 6.5-inch 4K high dynamic range (HDR) display. 4K refers to the resolution while HDR is a technology that helps make the contrast between black and white more distinct to make the image on the screen look more realistic.

Microsoft announces public availability of Azure Sphere, a solution designed to combat escalating threats to billions of IoT devices. About two years ago, Microsoft introduced Azure Sphere, a hardware and service designed to better secure 41.6 billion internet for things (IoT) devices. The program, which includes a Microsoft-developed Linux operating system for microcontrollers, evolved out of a Microsoft Research project, called Project Sprints. Today, Microsoft [announced](#) public availability of Microsoft's Azure Sphere after several years of testing and previews.

Cruise can now transport passengers in self-driving cars in CA

Cruise recently received a permit to transport passengers in its autonomous vehicles in California. Granted by the California Public Utilities Commission, the permit is part of the state's Autonomous Vehicle Passenger Service pilot. As part of the program, Cruise must provide data and reports to the CPUC regarding any incidents, the number of passenger miles traveled and passenger safety protocols. Cruise must also have a safety driver behind the wheel and not charge passengers for rides.

Cruise is now one of five self-driving companies allowed to participate in this program. The others are Zoox, Waymo, Pony.ai, Aurora and AutoX. This program is a bit different from the one run by the California Department of Motor Vehicles, which has granted 66 companies permits to test their respective vehicles in the state. Cruise originally

aimed to launch its robotaxi service last year, but it scrapped those plans last July due to technical challenges, and the creation of a vehicle platform and the ride-hailing app itself. At this point, it's unclear as to when we can expect to see Cruise deploy any type of public service.

Samsung Could Be Next to Face Production Issues Due to Coronavirus

Samsung might be the next company to face supply-chain challenges related to the coronavirus, which is spilling over into neighboring countries such as South Korea and Vietnam. Reuters reported over the weekend that one employee was infected with the coronavirus at one of Samsung's mobile device factories in South Korea and that the company could also face issues producing smartphones in Vietnam. Samsung shut the South Korean facility over the weekend and said it wouldn't reopen the floor where the employee worked until Tuesday, adding that it put coworkers who came into contact with the infected employee in self-quarantine. Meanwhile, Vietnam's Ministry of Industry and Trade said Samsung could face production delays in the country related to smartphone production, according to Reuters. The ministry said many of Samsung's smartphone components are still made and shipped from China, Reuters said.

Walmart Expands Its Amazon Prime Competitor

Walmart has been preparing to rebrand its subscription program that charges \$98 annually for unlimited same-day fresh grocery deliveries as Walmart+, Vox reported on Thursday. The program could eventually add other perks, such as discounts on prescription medicine and gas at Walmart stations, and appears to be an effort to compete with Amazon Prime. The price may later change, as the company adds more perks, according to Vox. Walmart also might add the ability for customers to place orders via text message. It's unclear whether a rebrand is enough to position Walmart's offering as a serious competitor to Amazon Prime. Amazon Prime includes other benefits and already has more than 150 members worldwide, who may already be locked in and unlikely to pay to try out Walmart's service. As we've reported, Walmart also has a much smaller selection available for one-day and two-day shipping.

Lufthansa gets spectrum licence, deploys Nokia private 5G for remote engine checks

German airline Deutsche Lufthansa (Lufthansa) has gained a spectrum licence for operating private LTE and 5G in the 3.7-3.8 GHz band in Germany. The licence is with Lufthansa subsidiary Lufthansa Technik, which provides technical aircraft services to airline carriers. Lufthansa Technik has built a 5G network with Nokia at its Hamburg facility to enable remote inspection of engine parts for its civil aviation customers. The company is managing the network internally. Nokia described the 5G deployment as "fully functional". The 3.7-3.8 GHz band in Germany is dedicated to enterprises for to operate private industrial campus networks. Siemens and Bosch have private licences at 3.7-3.8 GHz in Germany; Volkswagen is also believed to have a licence.

Department of Energy invests \$74m in building and construction technologies

The U.S. Department of Energy (DoE) announced \$74 million for 63 selected projects to research, develop, and test energy-efficient and flexible building technologies, systems, and construction practices to improve the energy performance of the country's buildings and electric grid. In a statement, the DoE said that awardees include national laboratories, universities, small businesses, and industry partners.

America's 125 million residential and commercial buildings use more energy than any other sector in the United States, accounting for 40% of the country's energy use and nearly 75% of its electricity consumption, according to the department. The research partnerships announced that they will pursue new technologies to enhance the energy productivity of buildings and improve the capacity of buildings to operate more flexibly, the DOE said.

Nokia launches end-to-end 4G and 5G NR slicing

Nokia announced the launch of its new end-to-end slicing network functionality for 4G and 5G New Radio (NR), claiming to be the first vendor to offer this capability. In a release, Nokia said that the solution will support connectivity from 4G and 5G devices over the sliced network to applications running in private and public clouds. The vendor said that the solution will be available later this year.

Ericsson launches its dynamic spectrum sharing technology

Ericsson announced that its Dynamic Spectrum Sharing (DSS) solution is commercially available, stating that the new technology is already live in a number of 5G networks around the world. Ericsson said that its spectrum sharing solution allows both 4G and 5G technologies to be deployed in the same band and on the same radio through a software upgrade – dynamically allocating spectrum based on user demand on a 1 millisecond basis.

Ericsson also noted that dynamic spectrum sharing is the most economically feasible way to deploy 5G on existing bands – enabling wide 5G coverage from day one – making more efficient use of spectrum and delivering superior user performance.

Eric Schmidt Warns Silicon Valley Could Lose to China

The U.S. stands to lose to China in technology wars without more help and input from the federal government, according to Eric Schmidt, former CEO of Google, writing in a New York Times op-ed.

Schmidt writes that government grants allowed him to finish his graduate work in computer science, and that the government should invest in workforce training, developing a competitive alternative to Huawei, and double funding for quantum computing and AI research. Schmidt's plea is indicative that technology leaders are worried about China gaining edge in areas like artificial intelligence and 5G, especially since the Chinese government has invested greatly in those areas and has major input in new developments.

The government has too much faith in the private sector to innovate to beat China, Schmidt writes. "Ultimately, the Chinese are competing to become the world's leading innovators, and the United States is not playing to win," Schmidt writes.

Dell-backed edge AI firm FogHorn gets \$25m fund, signs Stanley B&D, Honeywell

California-based industrial AI developer FogHorn (** Chambiz DF 11 Aug 2018*) has closed \$25 million in a series C round of funding, led by LS Corp, part of \$25 billion South Korean industrial conglomerate LS Group. It said the new funds will accelerate growth among industrialists seeking digital change in East Asia, in particular. Forte Ventures also joined the round, along with existing FogHorn investors, including Dell Technologies Capital, Intel Capital, Saudi Aramco Energy Ventures, Honeywell Ventures, GE Ventures, Robert Bosch Venture Capital, March Capital Partners, and Darling Ventures.

FogHorn, which develops edge analytics software for industrial and commercial IoT, has raised \$72.5 million in total, to date. In 2019, it achieved nearly 10-times growth in annual license bookings over the previous year, it claims, buoyed by demand for edge analytics in discrete and process manufacturing, oil and gas, transportation and smart buildings.

Softbank leads \$100 million investment in Behavox, the world's only end-to-end data operating platform. [Behavox](#),

a provider of an AI-driven platform that helps organizations to mitigate compliance, cyber and conduct risk, while identifying revenue opportunities in large volumes of communications data, announced it has raised a \$100 million investment from SoftBank Vision Fund 2.1. Founded in 2014 Erkin Adylov, Alexander Glasman, Kiryl Trembovolski, Roman Zelov, and Slav Slavinski, the New York City-based Behavox is an enterprise platform that unifies all types of data into a single data lake, allowing users to query the data and apply machine learning to identify capital markets misconduct and insider threats. Powered by machine learning and advanced analytics, Behavox is an organization's single entry point for all internal data.

Pet Plate raises \$9M Series A funding to provide fresh pet food directly to the consumers. [Pet Plate](#), a subscription service startup that provides direct-to-consumer fresh pet food, announced it has closed \$9 million dollars in Series A funding to expand the company's product offering, including new recipes, organic treats, and nutritional supplements to offer a holistic solution for online shoppers. The round was led by DFE Capital Management and 301 INC, the venture capital arm of General Mills. Additional investors include Marco Polo and Fernbrook Capital Management LLC, along with existing investors, The Yard Ventures and Castor Ventures. Pet Plate was founded in 2016 by Renaldo Webb, an MIT graduate, who made it his mission to provide better quality dog food in a more convenient fashion.

Karius raises \$165 million for its liquid biopsy technology identifying diseases with a blood draw

While the new money may have been raised under the looming threat of Covid 19, the company's technology is already being used to test for infection-causing pathogens in immunocompromised pediatric patients, and for potential causes of complex pneumonia, fungal infections and endocarditis, according to a statement from the company. Liquid biopsy technology has been widely embraced in cancer treatments as a way to identify which therapies may work best for patients based on the presence of trace amounts of genetic material in a patient's bloodstream that are shed by cancer cells.

Karius applies the same principles to the detection of pathogens in the blood — developing hardware and software that applies DNA sequencing and machine learning techniques to identify the genetic material that's present in a blood sample. As the company explains, microbes infecting the human body leave traces of their DNA in blood, which are called microbial cell-free DNA (mcfDNA). The company's test can measure the cell-free DNA of more than 1,000 clinically relevant samples from things like bacteria, DNA viruses, fungi and parasites. These tests indicate the types of quantities of those pathogens that are likely affecting a patient.

Revolut raises \$500 million at a \$5.5 billion valuation

Fintech startup Revolut is raising a large Series D round of funding. TCV is leading the \$500 million round, valuing the company at \$5.5 billion. Over the past few years, Revolut has raised \$836 million in total. Some existing investors are also participating in today's funding round, but Revolut isn't sharing names. Previous investors include DST Global, Index Ventures, Balderton Capital and many others.

If you're not familiar with Revolut, the company is building a financial service to replace traditional bank accounts. You can open an account from an app in just a few minutes. You can then receive, send and spend money from the app or using a debit card. On top of that, Revolut has added a ton of features that it has built in-house or through partnerships. You can insure your phone, get a travel medical insurance package, buy cryptocurrencies, buy shares, donate to charities, save money and more. Revolut currently has more than 10 million customers, mostly in Europe and the U.K.

Molekule inhales \$58M to purify air

Air purification startup Molekule landed \$58 million in a Series C round led by RPS Ventures. The San Francisco company plans to use the cash to invest in research and development, and scale the business.

UK fintech startup Revolut raises \$500M

British fintech startup Revolut has raised \$500 million in a fresh round of funding led by TCV. The financing sets a value of \$5.5 billion for the digital banking upstart, more than triple the valuation from its Series C round in 2018.

M33 closes on \$260M to scale bootstrapped businesses

Boston-based venture capital firm M33 Growth has raised \$260 million for its second flagship fund. The firm focuses on health care and technology companies that have bootstrapped their way to profitability and are looking to accelerate growth.

Agrifood investment tops \$19B

Investors put \$19.4 billion to work in seed through growth stage food and agriculture deals in 2019, according to a new report from venture firm AgFunder, based on Crunchbase data. The total represents a wide swath of industries, ranging from meat alternatives to indoor farming to robotic food delivery and cloud kitchens.

Pipe secures \$6M for SaaS financing

Los Angeles-based Pipe, a provider of cash advances for SaaS companies based on anticipated annual revenue, raised \$6 million in a seed funding round led by Craft Ventures.

Toyota invests \$400M in driverless tech startup Pony.ai

Toyota is investing \$400 million in Pony.ai, a developer of self-driving technology with headquarters in Silicon Valley and Guangzhou, China, at a valuation of just over \$3 billion. The financing comes amid a period of rising startup investment by automakers, who have been putting record sums to work in a bid to stay competitive.

AI chipmaker Graphcore secures \$150M

Graphcore, a UK-based startup that develops processors for artificial intelligence applications, has raised another \$150 million at a reported valuation of \$1.95 billion.

Salesforce nabs Vlocity for \$1.3B

Salesforce announced it is spending \$1.3 billion to acquire Vlocity, a developer of industry-specific cloud and mobile software built natively on the Salesforce platform. Founded in 2014, San Francisco-based Vlocity previously raised more than \$160 million from backers including Salesforce Ventures.

Tempo lands \$17.5M for home workouts

Connected fitness startup Tempo raised \$17.5 million for its Series A round and will start taking pre-orders for its AI-enabled system, which will start shipping this summer. Investors in the round include Founders Fund, Khosla Ventures, DCM and Bling Capital.

Everstream To Buy Detroit-Based Rocket Fiber

Everstream has agreed to acquire Rocket Fiber LLC, a fiber-based network service provider in Detroit. The acquisition completes Everstream's coverage in all major Michigan markets and gives Everstream the largest fiber footprint in greater downtown Detroit, according to a news release. Terms of the deal were not disclosed.

Exclusive: AI chipmaker SambaNova nearly triples valuation to \$2.5B

SambaNova Systems has been valued at \$2.5 billion in a new funding round, according to regulatory filings reviewed by PitchBook, making the company second only to China's Horizon Robotics as the most valuable semiconductor specialist in its field.

The valuation, 2.8x higher than what the company claimed in April, follows a \$250 million Series C led by asset manager BlackRock. Existing investors GV, Intel Capital, Walden International, WRVI Capital and Redline Capital also participated in the round.

Chipmaker Graphcore adds to war chest to take on semiconductor rivals

UK-based chipmaker Graphcore has raised an extra \$150 million as it goes toe-to-toe with industry rivals buffeted by the US-China trade war. The new funding round—comprising investments from Mayfair Equity Partners, Baillie Gifford and Amadeus Capital Partners—values the company at \$1.95 billion. The deal is an extension of its \$200 million Series D, which was co-led by Atomico and Sofina in 2018.

Based in Bristol, a semiconductor hub, Graphcore is a lead developer of AI microchips designed for use in a variety of applications, from voice recognition to self-driving cars. Its clients include Citadel Securities and French search engine Qwant, as well as Microsoft, which launched Graphcore's chips on its Azure software last year. This latest investment will go toward research and development, as well as helping to expand the business' sales and engineering centers.

The market for AI semiconductors is expected to grow exponentially in the next couple of years. Global revenue in the space is estimated to surpass \$30 billion by 2022, up from 2018's \$6 billion, according to a report released last year by PwC. Established chipmakers such as Intel, Nvidia and Qualcomm have been pursuing their own AI chips for years, along with tech giants Apple, Google and Huawei.

TPG-backed McAfee lines up acquisition after nixing IPO

Cybersecurity pioneer McAfee has agreed to acquire Light Point Security, a Baltimore-based creator of security software built to protect web browsers from attacks.

TPG Capital acquired a 51% stake in McAfee at a \$4.2 billion valuation in 2016, with previous owner Intel retaining a 49% stake. Thoma Bravo bought a minority interest in the security company the next year.

Last summer, McAfee brought on Morgan Stanley and Bank of America to underwrite an IPO that could have valued the business at \$8 billion or more, according to Bloomberg. But the company joined a swelling number of businesses to delay IPOs late last year, a response in part to public-market investors pushing back against lofty private valuations.

Lendio closes \$55M Series E funding to accelerate the development of next-generation loan marketplace. [Lendio](#), the nation's largest marketplace for small business loans, today announced it has secured \$55 million in capital, including \$31 million in equity led by Mercato Partners' Traverse Fund and a \$24 million debt facility from Signature

Bank. The equity funding came from existing investors including Napier Park Financial Partners, Comcast Ventures, Blumberg Capital, Stereo Capital and Runa Capital. The funding enables Lendio to increase the scope and precision of its loan marketplace while expanding new bookkeeping and lender services functions. Founded in 2011 by Brock Blake, Levi King, and Trent Miskin, Lendio provides a platform that helps small business owners find lenders and secure loans.

Sequoia India leads \$16M funding in Qure.ai, a healthcare AI startup that helps physicians with routine diagnosis and treatment using artificial intelligence. [Qure.ai](#), a Mumbai, India-based AI tech startup on a mission to make healthcare affordable and accessible using the power of artificial intelligence, announced today it has raised \$16 million in a funding round to drive geographical expansion, expand product portfolio and support regulatory clearances. The round was led by Sequoia India, with participation from MassMutual Ventures Southeast Asia (SEA). Founded in 2016 by Prashant Warier and Dr. Pooja Rao, and incubated by Fractal, Qure.ai today has presence in over 20 countries and has impacted over 600,000 lives. Using artificial intelligence developed on a database of over 7 million scans, Qure.ai solutions provide fast and accurate interpretations of radiology scans within seconds, reducing time to treatment in the critical moments of care.

Wheely, Europe's 'luxury' competitor to Uber, plans on raising \$30 million to fuel global expansion. [Wheely](#), a London-based premium ride-hailing startup is planning to raise \$30 million to expand into either the U.S. or Asia, according to a report from CNBC, citing the company's CEO. The London-based company's platform connects riders with chauffeurs who drive high-end Mercedes-Benz cars. The announcement comes as Uber faces the prospect of being blocked from the U.K. capital, which is its biggest European city. Founded in 2010 by Swiss/Russian CEO and entrepreneur Anton Chirkunov, Wheely is a luxury ride-hailing service, allowing customers to book chauffeur-driven journeys both on-demand and pre-booked through an app in London, Moscow and Paris.

Healthtech startup NeuroVasc raises \$34M funding to support clinical trials and product development. [NeuroVasc Technologies](#), a healthtech startup focused on novel catheter-based technologies to treat neurovascular disease, announced that it has closed \$34M in funding through a strategic partnership with the Wego Group. The new capital infusion will be used to support the company's product portfolio development and global clinical trial program. NeuroVasc has developed a comprehensive portfolio development strategy to bring differentiated tools to the INR community to facilitate broader treatment options for patients suffering stroke and other neurovascular diseases.

Roblox raises \$150M for gaming platform

Gaming platform Roblox raised \$150 million in a Series G round led by Andreessen Horowitz, at a reported valuation of around \$4 billion. Along with the new funding, the Silicon Valley-based company has a secondary offering of up to \$350 million to provide liquidity for early employees and stakeholders.

K Health raises \$48M for primary care

K Health, a primary care consultant powered by artificial intelligence, raised a \$48 million Series C round led by 14W and Mangrove Capital Partners. The New York company uses AI and anonymized health records to augment diagnoses and allows users to chat with a doctor through its app.

By Eugene Demaitre



To achieve maximum throughput in robotic piece picking, it's not simply a matter of adding machines to the line. Robotics designers, suppliers, integrators, and users need to identify the best combination of robot arms, sensors, and end effectors for a particular payload or task. In addition, many robots need the right intelligence and ability to choose the right grippers. XYZ Robotics Inc. (**Chambiz DF 13 July 19*) is an example of a company that has developed systems to address these needs.

Allston, Mass.-based XYZ's piece picking system uses machine vision, but it does not rely exclusively on artificial intelligence models. A combination of mechanical and machine learning approaches is necessary, said Peter Yu, chief technology officer at the startup.

"With both [approaches] and our tool changer, a robot can pick nearly anything, which is useful in logistics and manufacturing," he told The Robot Report. "AI is important for tool selection. Changing between a large cup to a small cup or a bag cup gripper — that's a challenge from both the tool side and the AI side."

XYZ Robotics' grippers pick consumer electronics, apparel, cosmetics, and other objects for e-commerce order fulfillment. With AI guidance and the ability to change end-of-arm tooling, one robot can handle a wide variety of items with speed and accuracy.

"For example, if the SKU is a plastic bag, our system will know and choose a suction cup to pick it up," Yu said. "But if it's mesh or a thin pencil or screwdriver, there's not much area, so the robot can choose a two-fingered gripper."

Choosing the right grippers for piece picking

XYZ's vision-guided tool changer can swap out end effectors in about half a second. "For vision, the time to identify the piece picking points is 0.1 sec. with VGA, and at 720p, it is 0.3 sec.," Yu said. "In addition, the tool changer has locating pins so the robot knows the gripper is engaged in a specific orientation."

XYZ Robotics uses a combination of standard and custom end effectors with its tool changer. "A suction cup may come from off-the-shelf vendors, but the bag cup is designed, made, and patented by us, as well as the tool changer," said Yu.

"There are many kinds of grippers — some, like Schunk, are electric, and others are pneumatic, like SMC's," he added. "Usually, we use vacuum to do suction. We want to use the same source to drive everything. That's why we have our grippers driven by vacuum, so we don't need to add lines for electricity or compressed air."

"We're still working on a two-fingered soft vacuum gripper for picking items that are not graspable by suction cups, such as lipsticks, measuring spoons, and screwdrivers," Yu said. "We want to provide a holistic piece-picking solution to our customers."

Engineering reduces the need for big data

Machine learning typically needs large, clean data sets, and humans play a major role in annotating training data, Yu acknowledged.

"Everyone needs data and observation of that data," he said. "Our overall approach was statistical-based machine learning, but we don't throw raw data to the algorithms to train it. That's an end-to-end approach — it's cool, but it requires a bigger magnitude of data."

"Machine learning scientists come in, as humans give some heuristics to model, which then needs less data to train," Yu said. "For example, for piece picking, intuition is that we observe the geographical features to grasp. Context helps help reduce complexity."

How long does it take to train a piece picking robot on a novel object? "The algorithm generalizes on

Key Features

- Model-free:**
Generalize to novel objects
- Fast Prediction:**
< 0.1s in VGA mode and
< 0.3s in 720p mode
- Adaptable:**
Handles dark, shiny and
semi-transparent objects
- Quick Setup:**
Plug and play in 15 minutes
- User-friendly:**
Browser-based UI for
calibration and inspection
- Reliability:**
Remote software upgrade,
support & monitoring

Mixed random items

Sample axis-aligned pick region (yellow circles) visualization

Mixed random items Black items with reflective packaging Tightly packed items

Features of XYZ's vision processor. Source: XYZ Robotics

most novel objects,” Yu said. “But if certain objects need a specific tool or to be grasped in a certain way, we need to teach the robot, and it can learn pretty fast — one hour after inputting the data.”

“There are more ways, like self-supervised learning,” he added. “We throw in an item, and the robot tries different tools and puts item at different locations. If we let the robot explore, it could take five minutes. Then the label is fed into the training algorithm, and it takes an hour or so to get a new model.”

By taking multiple approaches, robotic grasping can improve over time, said Yu. “Before, we started at 80% of items that were graspable or suctionable by our tools,” he recalled. “It was an engineering effort to push from 90% to 99%. That was some innovative engineering going on, which means a lot when translating to accuracy, reliability, and speed.”

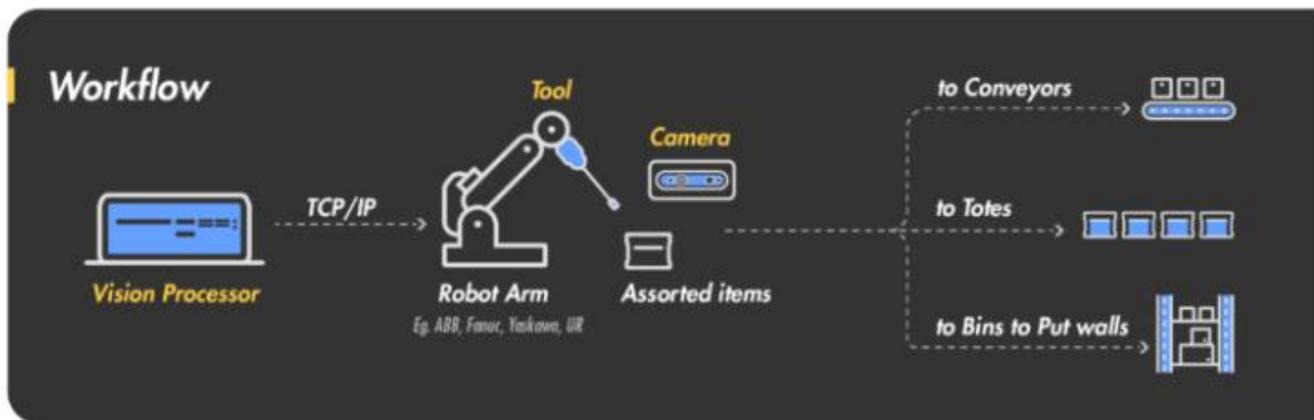
Solving the double-picking problem

“Accurately detecting double picking — that problem is hard in industry,” Yu noted. “If two boxes with little textures are tightly packed together, it’s difficult for a vision system to see the seam in between. We take a combined approach using vision and weight sensing.”

“In the e-commerce space, a customer might otherwise get two iPhones. The rate of double picking is typically 1% of items,” explained Yu. “Using the combined verification approach, we can avoid 95% of that 1%, so the overall rate of double piece picking goes down to 0.05%.”

Yu said that XYZ is constantly pushing its technology for more accuracy, reliability, and cost savings. “We keep pushing speed for the tool changer, and broadening the variety of items we want to pick is a big thing,” he said.

“Accuracy is important to reducing double picking, and we’re working on the motion of the robot to make it faster,” said Yu. “We’re pushing the system this year to do 1,200 picks per hour. To meet our goals, we have to incorporate everything.”



A vision processor workflow. Source: XYZ Robotics

Remote assistance and 5G

Like other robotics companies, XYZ expects to offer remote assistance to help deal with rare cases. “Our method is machine learning-based,” Yu said. “Sometimes, if an item has an odd shape, a vision system will not work well. A robot can send a request for help to the cloud, and a human monitor can see what’s on pick point.”

Another aspect of remote assistance is the ability to train robots in less time, thanks to 5G. “If we know a certain item isn’t performing well, the system will collect data and then go training to improve grasping within half a day and send to robots,” said Yu. “We already bought a 5G access point in China and have started testing it. Speed

improves 10 times in general, usually in urban areas. Most areas have only 4G access, and the speed may be affected a bit.”



XYZ bundles machine vision and piece picking technologies. Source: XYZ Robotics

Connecting with the rest of the warehouse

XYZ Robotics’ piece-picking and sorting system can connect with automated storage and retrieval systems (AS/RS) and place items on a shelf, on a conveyor belt, or in a box on an automated guided vehicle (AGV). Machine vision is also relevant to bar-code scanning, noted Yu.

“For warehouse applications, XYZ’s system is connected with the WCS [warehouse control system] to keep track of all the orders.”

“In terms of piece picking, the capability for the robot depends on visual data. The most data shared is the vision data, which is collected, and then the model is deployed to all the robots.”

“We do not share that much data with the AS/RS or AGVs, which are under the control of the WCS. It controls all the components, including robots.”

“On the business side, we’re working on partnering with AS/RS systems and AGV vendors,” Yu said. “We’re also working with partners like SF Express and JD.com.”

Demand for vision-guided piece picking

Although many startups are chasing the piece-picking market, Yu said they need to focus on solving the right problems.

“When we go to a customer sites and ask them what they need — there’s no product in the world that has really made a huge impact yet because humans are still pretty good in terms of speed, accuracy, and variety,” he said. “There is rarely a huge deployment of piece picking robots. My company feels that the competition is with human dexterity, not other companies.”

XYZ said it is gradually reducing the need for fine-tuning the model for specific products from specific customers. “Usually when we deploy, our normal model works out of the box,” Yu said. “But as we collect more data, we can better solve various rare situations.”

Example XYZ applications in logistics and e-commerce industries



Picking Station

- 1 Identify pick points
- 2 Pick with the best tool head
- 3 Place gently without rough dropping



Rebinning Station

- 1 Incoming SKUs via conveyor belts
- 2 Identify and pick
- 3 Sort onto put walls via barcode scanning

By Alex Wilhelm

Today let's try to figure out where the startup and private markets stand, as there are a few different takes out there that I can't reconcile. Our efforts to better understand how young companies are faring comes, of course, in the shadow of the impending \$7 billion Intuit-Credit Karma deal — the second, multi-billion dollar fintech exit so far in 2020.

So where are we today in the startup business cycle? We'll summarize a few different perspectives on the question, and then come up with our best synthesis of the group.

If you observe the behavior of the venture class, it's a full-speed-ahead market. This is contrasted by a summary of recent private-market tech stumbles compiled by New York Times's Erin Griffith. Bolstering Griffith's take are a set of long-running complaints by select tech leaders regarding the health of some of the private market's most valuable companies.

But as Credit Karma looks set to exit at a huge valuation, it's hard to tease out which perspective is the most correct. So let's try to answer that for ourselves.

Perspectives

It's tricky to summarize others. But let's be bold and do so, breaking the world into a few different viewpoints.

The following are constructed based on dozens of interviews over the past few months with founders, investors and operators. Weaved through them are the news of the last six or twelve months, replete with wins, losses, disappointments, mistakes and turnarounds. Here are our takes, explained:

- **The optimist:** Many founders and venture capitalists remain optimistic about the private markets and startups in general. There are hotspots around the technology world, including fintech, finservices, vertical SaaS, work with regulated industries, and more. The best example of the in-market impact of this particular view of the market is the galloping pace at which venture capitalists and startups continue to raise capital. While venture activity has slowed in China, there's still so much activity that it's hard to see much deceleration amongst startups building outside of that country.
- **The mid-cycle pragmatist:** There are some venture capitalists sounding notes of caution. Perhaps we're in for a cloud growth slowdown; perhaps some valuations are a bit high; and perhaps some startups got a bit silly. This perspective loves baseball analogies, telling TechCrunch over the phone or in person that we're in the fourth inning of key trends like the cloud takeover of enterprise software. This perspective often feels like our optimistic take, but with just enough caution included as to cover the speaker's backside in case venture really does go off the rails.
- **The caution-minded:** Some private market players are nigh-shouting that some things aren't going well. Bill Gurley is one of those folks, arguing in favor — over the years — of more healthy, stable, secure startups, and against unicorn excess. This perspective is not always markedly conservative, but the cautioned-minds are focused on the health of their startup or portfolio and worried about the rest of the market.
- **The post-peak scold:** While many private market players want to paint the media as generally censorious, no-fun and rude about startups, I can tell you that it's not just journalists rolling their eyes at startup excess in recent years. Some founders and investors are spooked and annoyed as well. The WeWork implosion has scared more than a few folks, and they will tell you about it if you agree not to quote them. There are not mobs of individuals in the startup world who are irked at how some silly companies made the party worse by torching too much money to avoid notice and ridicule, but they exist. And, yes, the media is chronicling the

downfall of a lot of folks who talked big on the way up and now want clemency on the way down. But the post-peak scolds aren't a single group, and they don't have a unified gripe.

It's worth noting that this column winds up embodying each of those perspectives at times, depending on what the news is telling us. No single perspective is entirely correct or completely wrong.

In fact, and I am sure that you saw this coming, each has a pretty good partial argument about where we are today.

It's fine to be an optimist in today's startup market. Founders and venture capitalists are paid optimists. They can't do their jobs unless they have that perspective; there are no short-selling private market investors. The mid-cycle pragmatists and cautioned-mind types are probably pretty close when we come down to it, each blending a mixture of selective positive expectation with the view that some other people are schmucks. Our final category is mostly filled with industry observers, focused on covering the market from the perspective of trends, along with some in-market players who are seeing their own efforts brought down by the mistakes of others.

I wanted to write all this to help the both of us remember that we're never utterly correct, and that there is never a time in which all things go well or all things go poorly. Sure, there are startup layoffs, but there is also a startup talent crunch. At the same time, there are lots of new, highly-valued rounds, but we've also just seen a high-flying unicorn get its valuation cut to shreds in its debut. And while some valuations are coming back to earth, One Medical shows that there is still appetite for rich valuations for tech-enabled entities.

Enlightened centrism is crap, but I'd hazard that the most accurate take is probably a little of all the takes. A Franken-perspective that says something like this:

We're well past the halfway mark in the current startup cycle, and there are fewer great places to place bets. That's leading to stiffer VC competition, faster rounds for winners, and rising valuations. VCs are raising more money, even as there are signs that the early-stage market is slowing down. And while there are some impressive exits, others are shit. So things are ok for now, clouds on the horizon notwithstanding. But the shakeout of weaker players in the market is causing some folks to put a bit too much faith in their perspectives, positive or negative.

I'd say that nothing much changes until the global macro climate falls apart. Then, all bets are off.

By Alex Wilhelm



All around, this has been a tough week. The coronavirus is spreading and worry is running high as infections mount. In economic terms, global markets were repeated declines last night, and the U.S. indices are off again this morning (Feb 28, 2020).

There's been plenty of bad news to read, even in our private market, startup-focused world. Yesterday the impact of COVID-19 on earnings became more apparent, bringing what has, for months, been an external concern to domestic technology companies. The problems are now. The past

week's market collapse into correction territory hasn't helped,

But the story so far has largely been *public-market* focused and with good reason: You can see the public markets contract in real-time. It's far harder to see into the shifting dynamics of the private market. Today, however, we are going to try, all the same, by digging into some preliminary venture capital data.

I realize that the last few days have been awful. So, at the end of this piece, I've excerpted a quote from a recent interview I held with the CEO of [Smartsheet](#), [Mark Mader](#), about tech cycles, downturns, and getting through tough times. It's perhaps useful today as the downward trend appears to continue.

Let's start with a brief reminder of how elevated stock prices remain and what that means for tech multiples, and then look at early February VC results from the U.S., China and Europe. With that, in Sanskrit: *अभिमुखी करोति*.

Multiples, Markets

Before we dig into the venture capital data, a reminder that, even with recent declines, we're still in warm waters as far as tech valuations go.

Don't believe me? Here are some headline stats: Two of the tech Big 5 are still worth over \$1 trillion today, and two more are worth \$900 billion or more. Also, SaaS and cloud stocks have declined, sure, *all the way back to their summer 2019 levels*. This means that things are about as good as they were last summer for the key cohort of tech companies (public SaaS firms are critical drivers of many private software startups' valuations), with the Bessemer group [posting revenue multiples](#) (calculated using enterprise value) of 11.9x as of this morning.

So the pullback has been not good, certainly, but tech shops are still richly valued. Hell, we've even seen some out-of-favor players bring it back in recent quarters, posting better-than-expected growth to the surprise and delight of their shareholders.

And that's where the good news stops, and the bad news begins.

Falling IPO activity, falling VC totals

The 2020 IPO market has been weak so far. Casper and One Medical put up some points, but with just Asana and DoorDash promising to go public eventually, we could wind up done with Q1 2020 with just two venture-backed, non-biotech IPOs. That's *far* under the pace that unicorns need to clear enough of their inventory to create material investor liquidity.

Perhaps in reaction to that issue, or perhaps for a host of others, venture activity looks weak. (Looking at recent venture capital totals is tricky, given the inherent lag in venture data; deals are often reported long after they are closed.)

Here's the following February totals via the Crunchbase dataset:

- US, equity only recorded private investments, February 2019: 999
- US, equity only recorded private investments, February 2020: 461

The story is similar in China:

- China, equity only recorded private investments, February 2019: 189
- China, equity only recorded private investments, February 2020: 32

Those figures are staggering. Keep in mind that China's 2019 totals were depressed *already* from earlier years. What we can see in the February data is that Chinese venture capital activity has all but stopped during the outbreak

How about Europe?

- Europe, equity-only recorded private investments, February 2019: 631
- Europe, equity-only recorded private investments, February 2020: 349

That's the least bad decline, notably, in percentage terms. The 2020 results *will* improve as more data comes in, but I'd bet you a full lunch that we see material, double-digit percentage declines in venture activity in February. What we can say is that it appears that the private markets are tracking the public markets: the public downturn is a private correction. We will naturally keep tabs on this during March as well and run our own dives into Q1 data when it shakes out. Consider this an *early and preliminary* look at the venture market in 2020.

Good news?

How about some encouragement to close? I could use some, I think. You too.

This Monday, as the market's terrible, horrible, no-good, very bad week got underway, I chatted with **Smartsheet's** CEO Mark Mader. Smartsheet, a member of the 2018 IPO class, is a big SaaS company worth a bit more than \$5 billion. Here's a snippet from our conversation that I've only modestly cleaned up to preserve tone:

TechCrunch: *The market is falling apart today, as I'm sure we're all aware, due to a lot of kind of little fears, and I think Smartsheet's down like six points. I am kind of curious how you approach that as a CEO? Do you just not think about short-term fluctuations that are driven by macro events in the value of your company? Or do you actually tell people internally like, "Hey, guys, you know, chin up, this is going to be okay. It's nothing to do with us. We're just caught up in this broader market chop?"*

Mark Mader: Yeah, I think it's a delicate balance. You can't be oblivious to it. But I think I answer the question of how you're doing very differently these days, Alex. If you asked me three years ago, I would have said, "Hey, Alex, this amazing deal just came in. It's really cool and here's what it means to our company," or, if it was a really shitty day, I would say, "Alex, you know, I didn't get that candidate I really wanted." I never answer that question that way anymore. I always say, on a net basis, Alex, life's great. It's great. And it's the true SaaS answer. And in the sense that we're judged on a net basis, we're not judged on "did you get a win yesterday or today [or a] loss yesterday or today?"

So I would say, in this case, you know, having lived through 2000 at a software company; having lived through '07, '08 when we just released a new version of our software, when the financial markets collapsed, to what's going on today, I've been around the block couple hundred times over 25 years, and you realize that you will shorten your life considerably, I think, if you overreact to these types of days.

The reason you should be aware and you should be sensitive is because not everyone maybe takes that approach, right? So if your customer is nervous, or an employee who just started with you is like, “*Oh my gosh, what’s happening today? I thought everything was perfectly stable in the world,*” you need to have a listening ear. But in terms of how you drive your company and make investments and decisions, I think you have to have a long term view. You have to.

Onward

So 2008 this is not. At least not yet. And tech and venture and the private markets got through that. This too shall pass, even if we don’t know yet how bad things will get.

Meet ByteDance's Biggest Homegrown Rival

By Yunan Zhang

While China's ByteDance and its TikTok video app are getting a lot of buzz in the U.S., a smaller rival called Kuaishou has quietly become one of the most highly valued tech startups in the world.

Kuaishou (KWHY-shoh) has built a social network of more than 300 million daily active users in China—an audience three-quarters the size of that for its archenemy Douyin, ByteDance's domestic version of TikTok. While Douyin has grown by feeding funny videos to bored urbanites, Kuaishou has instead fostered a devoted fan base in smaller cities and countryside with an emphasis on building an online community.

Backed by investors including Chinese internet giants Tencent and Baidu, as well as Sequoia Capital and DST Global, Kuaishou was valued at \$28 billion in its most recent fundraising round late last year. The company recently told its existing investors that it plans to go public next year, although it hasn't decided whether to list in Hong Kong or New York. ByteDance's most recent valuation was \$75 billion.

Kuaishou's and ByteDance's apps represent a new type of social media entertainment that, at least in the case of ByteDance's TikTok, has proved to have legs outside China. The apps from both companies, including the China-focused Douyin app, differ from U.S. video apps in their emphasis on very short videos, sometimes just 15 seconds. Other video apps, like Snap and YouTube, offer a broader range of content and video lengths.

The speed of the Chinese apps' growth has forced U.S. firms to respond. Facebook is testing a short-form video-sharing app in Latin America called Lasso. Google is also mulling the introduction of more short-form video features to YouTube, according to a person familiar with the situation.

To be sure, there are differences between Kuaishou's and ByteDance's apps. TikTok and Douyin tend to give prominence to slicker content with mass appeal, such as funny dance routines. Kuaishou works more like a social app. It has a rustic feel and emphasizes videos posted by friends or other users, showing people going about their life on the farm or cooking at home.

In some ways, Kuaishou is like a hybrid down-market Instagram serving the majority of China's population, which lives outside the biggest cities. It's notable that Kuaishou has yet to establish itself outside China. It has an overseas video app called Kwai, which it has tried to promote in Southeast Asia and Brazil. But compared to ByteDance, its presence beyond China is insignificant.

Kuaishou's approach is paying off. Its revenue doubled to about \$5.7 billion last year and the startup even turned a profit, according to investors who have seen its financials. The company earns the majority of its revenue by taking a cut of the tips that some users give to performers, who are often amateurs much like themselves. A smaller portion, about \$1 billion last year, comes from advertising. E-commerce is a potential new business: Kuaishou fans who have gotten used to sending cash tips to their favorite live-streaming performers are now starting to buy the products the performers promote.

But competition between Kuaishou and ByteDance is intensifying, raising questions about Kuaishou's future growth rate. Millions of Chinese, homebound because of restrictions imposed by the Chinese government to prevent the spread of the novel coronavirus, are looking online for entertainment. ByteDance and Kuaishou, along with a slew of other Chinese online entertainment providers including Tencent and Baidu, are battling to grab their attention.

Even before the outbreak of the virus, Kuaishou had made a big bet to increase its viewership by sponsoring the Chinese Spring Festival Gala, a five-hour annual television variety show more than a billion people watch on the night before the first day of the Lunar New Year holiday. That audience is about ten times that of the Super Bowl in the U.S.

Kuaishou got less of a boost from the sponsorship than it expected, people familiar with the matter said, because the coronavirus spoiled the celebratory mood. Kuaishou also hasn't been able to retain all of the new active users it gained due to the sponsorship, according to employees. ByteDance, meanwhile, grabbed some of the spotlight when it announced it had secured streaming rights to broadcast one of the most hotly anticipated movies of the Lunar New Year holiday season, "Lost in Russia."

"Healthy competition is good for the market," a Kuaishou spokeswoman said. "Kuaishou has been growing as well when Douyin grows. It's not a winner-take-all market. Both parties have been providing users with high-quality content, satisfying different demands, which is positive for both audience and market."

Fast Hand

Kuaishou, which means "fast hand" in Chinese, started life in 2011 as GIF Kuaishou. Founder Cheng Yixiao built it as a tool to help users turn short videos into GIFs, targeting customers with cheaper low-data smartphone plans. In 2014, the company hired Su Hua, a former Google and Baidu software engineer who was working on his own startup, as its new CEO, while Cheng became chief product officer.

Su dropped GIF from the company's name and turned it into a platform for users to broadcast short videos. Shortly after he joined, the app's daily active users topped 1 million for the first time.

Su, described by his colleagues as idealistic, had told friends that he wanted to create a down-to-earth online community where ordinary Chinese people can feel welcome. On Kuaishou, they can make their voices heard. Compared to other viral video apps, Kuaishou tends to allocate more traffic to unsophisticated videos posted by average people—instead of driving more traffic to those users with the most followers.

Kuaishou executives didn't design the app exclusively for users in less developed parts of China. But the app's product philosophy worked especially well in smaller towns and the countryside, with users sharing snippets of their lives with each other. Kuaishou users in rural regions say their videos get "likes" even though their accounts don't have many followers.

Kuaishou "grew fast organically through word-of-mouth, like most big ideas do," said XVC partner Hu Boyu, who invested early in Kuaishou when he was at DCM Ventures.

Kuaishou's early success made other Chinese tech companies realize that the potential market for short-video apps is huge and growing. In 2016, ByteDance, which until then had mostly focused on its news aggregator Toutiao, launched Douyin. A year later, ByteDance bought the lip-syncing app Musical.ly, which was folded into what later became TikTok.

Compared to Kuaishou, Douyin is more like a personalized TV than a social network. ByteDance's founder, the cerebral Zhang Yiming, leaned heavily on AI-driven recommendations to decide what to show users. Douyin predicts which video the user will want to see next. The app also actively recruits a professional cadre of entertainers who produce high-quality videos.

Douyin's approach quickly became popular, especially in richer cities such as Shanghai, Beijing or Shenzhen. By 2018, the number of Douyin's daily active users was 200 million, surpassing Kuaishou's 160 million users.

Kuaishou hit the 200 million user mark in May 2019. But Su wasn't happy. While his staff was celebrating, he fired off a companywide email complaining, "Being slow has become our company's tag." Su set a new goal to add 100 million more users by Jan. 23, 2020, the start of the Chinese Lunar New Year holiday. Kuaishou was in "battle mode," he wrote.

A month after Kuaishou's declaration of war, Douyin announced it had 320 million users.

A One-Stop Shop for Farmers

Kuaishou owes its growth to users like Li Bing, a 21-year-old who lives in a village in Gansu, a poor province in China's northwest. For Li Bing, Kuaishou has displaced WeChat, Tencent's social network and messaging app, which nearly everyone in China uses.

He logs onto Kuaishou to share short videos with friends, relatives and thousands of strangers. Most of his videos are very simple recordings of his daily life. In one, he showed his dad scooping up freshly harvested corn and laying it in a cart. Li also spends hours every day watching videos other Kuaishou users in his village or nearby regions post.

In some parts of China, especially in the northwestern region where Li lives, Kuaishou is much more than just a video app for fun. For users in rural areas, the app is their everyday tool for keeping in touch with friends and acquaintances. They watch each other's videos and send messages through Kuaishou's chat function. Some users say Kuaishou, for them, is a source of information on their towns and villages, from neighborhood gossip to new restaurant openings. Li and his friends like Kuaishou because the app is a social platform where they see other people who are more like themselves, not online celebrities from glamorous distant cities.

Kuaishou's core users aren't just farmers in the countryside. The app is also widespread in lower-tier cities in northern China—obscure cities whose names aren't even familiar to those who live in Beijing and Shanghai.

Users like Li are also starting to buy things through Kuaishou. He recently purchased toothbrushes and hand cream after watching a live-streaming video on Kuaishou promoting the products.

"When I'm bored, I like to watch what other people in the countryside are doing," Li said. While Li has downloaded Douyin on his phone, he prefers Kuaishou because many of his friends and relatives are on it. "Even my grandpa uses Kuaishou," he said.

Embedded vision startup Prophesee teams with Sony to shrink its pixel size to less than 5 micrometers

By Samuel K. Moore



Images: Prophesee

There's something inherently inefficient about the way video captures motion today. Cameras capture frame after frame at regular intervals, but most of the pixels in those frames don't change from one to the other, and whatever is moving in those frames is only captured episodically.

Event-based cameras work differently; their pixels only react if they detect a change in the amount of light falling on them. They capture motion better than any other camera, while generating only a small amount of data and burning little power.

Paris-based startup Prophesee has been developing and selling event-based vision sensors since 2016, but the applications for their chips were limited to systems that needed only VGA resolution. The sensor resolution was limited because the circuitry surrounding the light-sensing elements took up so much space. In a partnership announced this week at the IEEE International Solid-State Circuits Conference in San Francisco, Prophesee worked with Sony to put that circuitry on a separate chip that sits behind the pixels.

"Using the Sony process, which is probably the most advanced process, we managed to shrink the pixel pitch down to 4.86 micrometers" from their previous 15 micrometers, says Luca Verre, the company's cofounder and CEO.

The resulting 1280 x 720 HD event-based imager is suitable for a much wider range of applications, including surveillance and monitoring, augmented reality, virtual reality, and drones.

The company is also looking to enter the automotive market, where imagers need a high dynamic range to deal with the big differences between day and night driving. “This is where our technology excels,” he says.

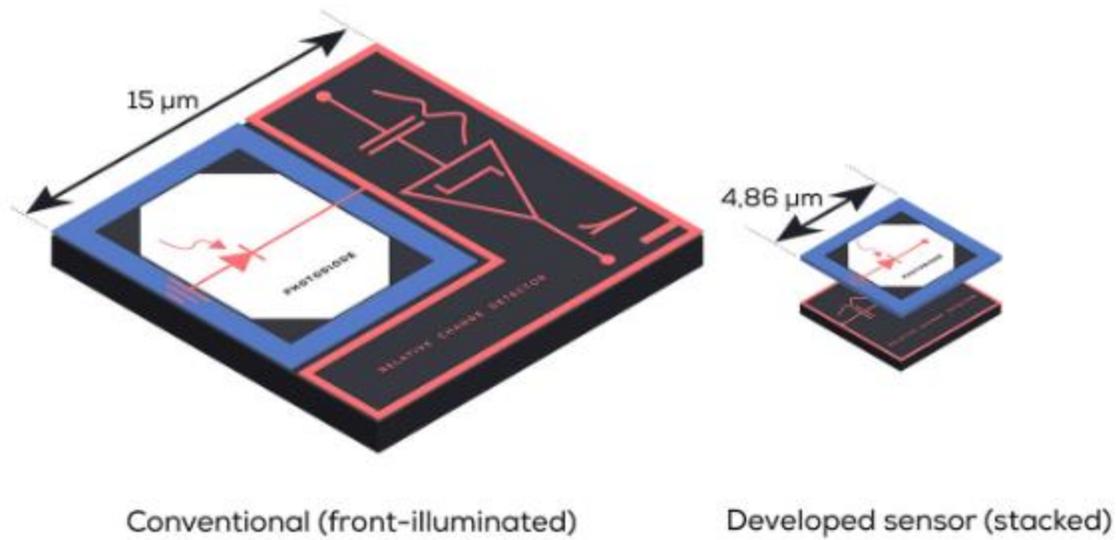


Illustration: Prophesee

Conventional versus stacked pixel design

Besides the photodiode, each pixel requires circuits to convert the diode’s current into a logarithmic voltage and determine if there’s been an increase or decrease in luminosity. It’s that circuitry that, in the new sensors, is put on a separate chip that sits behind the pixels and is linked to them by a dense array of copper connections. Previously, the light-sensing area made up only 25 percent of the area of the pixel, now it’s 77 percent.

When a pixel detects a change (an event), all that is output is the coordinates of the pixel, the polarity of the change, and a 1-microsecond-resolution time stamp. The imager consumes 32 milliwatts to register 100,000 events per second and ramps up to just 73 milliwatts at 300 million events per second. A system that dynamically compresses the event data allows the chip to sustain a rate of more than 1 billion events per second.

Uber's Old Rivals in Southeast Asia Hold Merger Talks

By Juro Osawa

Grab and Gojek, Southeast Asia's two largest ride-hailing and food-delivery apps, valued at \$14 billion and \$9 billion respectively, are discussing a merger that would create one of the world's most highly valued startups, people familiar with the talks said.

A deal would be the latest step in a consolidation of the broader food-delivery and ride-hailing markets, as companies try to stem losses caused by costly fights for market share. But the two sides are far from an agreement on terms, including how to value the two businesses. SoftBank-backed Grab, which operates all over Southeast Asia, competes most fiercely with Tencent-backed Gojek in Indonesia, where Gojek is based.

The two companies' management teams have had occasional meetings over the past few years, but conversations about a potential merger became serious in the past few months, the people familiar with the matter said. Earlier this month, Grab President Ming Maa and Gojek CEO Andre Soelistyo met for the two companies' latest discussion, according to two of the people.

The talks are believed to be still ongoing. A major stumbling block to a deal is coming to an agreement on control of the combined firm, say the people familiar with the discussions. Grab recently told its major investors that Gojek wants a 50-50 deal but Grab wants to control a significant majority in any merger scenario, according to one of the people. In addition to those differences, any attempt to combine two of the region's biggest tech companies would also face regulatory hurdles—a challenge that Grab's earlier deal with Uber faced.

Both Grab and Gojek are unprofitable. Each has scaled back some of its subsidies and promotions to cut costs. But their rivalry could still impede their efforts to become profitable. Both companies used to compete with Uber but the U.S. ride hailing giant sold its Southeast Asian operations to Grab in 2018.

Some shareholders on both sides are open to the idea of a deal that could boost the value of the combined business. Aside from SoftBank, Grab's major investors include Uber, Didi Chuxing and Toyota Motor, while Gojek's backers include Google, Visa and KKR, as well as Tencent.

The discussions between Grab and Gojek come as each company is trying to reassure its investors that its business can eventually make money. After Uber's disappointing IPO last year and WeWork's near-collapse, investors looking at future IPOs are getting more skeptical about money-losing startups with inflated valuations. In an interview with The Information in December, Grab's Maa said he was "extremely confident" that the company has a clear path to profitability.

It is possible that Grab and Gojek will keep fighting their battle to become profitable separately. The two companies are continuing with their efforts to improve their finances and raise more money. Last week, the Nikkei reported that Mitsubishi UFJ Financial Group, Japan's largest bank, had agreed to invest \$700 million in Grab as part of the two firms' strategic partnership. Gojek, which has struggled to raise money over the past several months, is in advanced discussions with multiple potential strategic investors, according to people familiar with the matter.

The stakes are high for Grab and Gojek. Southeast Asia is home to more than 600 million people, and Indonesia, where Gojek is based, is the region's most populous country with 250 million people. A report published in October by Google, Temasek and Bain estimated the size of Southeast Asia's overall internet economy at \$100 billion and predicted the region's ride-hailing market would triple to \$40 billion by 2025.

Attackers could use cheap consumer drones to project fleeting images onto roadways, fooling Tesla and Mobileye navigation systems and causing havoc, says Ben-Gurion University team

By Luke Tress



In a Ben-Gurion University study, a Tesla perceives a projected image as a real person, left, and Mobileye's 630 PRO autonomous vehicle system considers an image projected on a tree as a real road sign, right.

Autonomous vehicles can be fooled by “phantom” images displayed on a road, wall or sign, causing them to unexpectedly brake or veer off course and making them vulnerable to attackers, Israeli researchers said.

Semi- and fully-autonomous cars perceive and respond to two-dimensional projections as real objects, according to researchers from the Ben-Gurion University of the Negev.

Attackers could exploit the vulnerability to put vehicle passengers in danger, for example by projecting an image of a person in front of a car, causing it to brake suddenly, or by projecting fake lane markers onto the surface of a road, directing the car into oncoming traffic or onto a sidewalk, endangering pedestrians.

The team from the university’s Cyber Security Research Center used the Tesla Model X and the Mobileye 630 PRO system in the research, which was published earlier this month through the nonprofit International Association for Cryptologic Research science organization.

“This is not a bug. This is not the result of poor code implementation. This is a fundamental flaw in object detectors that essentially use feature matching for detecting visual objects and were not trained to distinguish between real and fake objects,” researcher Ben Nassi said in a statement. “This type of attack is currently not taken into consideration by the automobile industry.”

The team demonstrated how attackers could carry out an attack remotely by using a drone to project an image onto a road, or by hacking a digital billboard to insert a fake road sign into an advertisement. Mobileye's system, a leader in the field, could be fooled with a fake projection lasting just 125 milliseconds, the researchers said.



Cars in the Mobileye fleet of autonomous vehicles leave the Mobileye garage for test drives November 5, 2019, as part of the 2019 Mobileye Investor Summit. (Walden Kirsch/Intel Corporation)

For only a few hundred dollars, a bad actor could purchase a drone equipped with a portable projector and use it to carry out a terror attack by tricking cars into hitting pedestrians, for example. Criminals could create traffic jams by falsely projecting slow speed signs onto roadways. Fraudsters could have someone cause their self-driving Tesla to crash, then sue the company.

Previous studies have shown that autonomous driving systems are vulnerable to attack, but only by skilled attackers carrying out lengthy preparations, and in ways that would require the attackers to approach the scene, putting themselves at risk. Manipulating vehicle sensors would leave little to no evidence behind.

The sensors in autonomous vehicles do not verify what they perceive with other systems, and therefore react to stimuli independently, in what the researchers call a “validation gap” that could be exploited by hackers.

Sensors that perceive depth exist, but the researchers said that the systems likely employ a “better safe than sorry” policy that regards the projections as real objects.

Computer vision algorithms identify their surroundings based on geometry and patterns (the researchers call them “feature matchers”), such as the shape of a stop sign, and do not consider the context of the image, its texture, or how realistic it looks. So, a partially transparent stop sign floating in a tree, or a pixelated image of a man lying in a street, could cause a car to suddenly brake.

The systems do not have any concept of what a fake object would look like, and do not take into account the possibility of phantom image attacks.



Mobileye's system relies solely on computer vision algorithms, which assumes all objects are real. The system registers any image of a stop sign within a certain size range — if the sign is too small, the system “thinks” it is far away, and if it is too big, it thinks the driver is too close, and unable to stop. The researchers found the system does not even register the colors of the sign.

In its experiment on Mobileye's driver assistance system, the team used a DJI Matrice 600 drone with a projector disguised as a delivery box, and a Renault Captur car.

The drone projected a 90 km/h (56 mph) speed sign of onto a wall for 125 milliseconds in an urban environment, tricking the vehicle into telling its operator to drive dangerously fast. They were able to carry out the same attack using an 8.8 ounce projector on a smaller DJI Mavic drone, which retails for \$400.

The study notes that billboards could be hacked and signs hidden in an advertisement video, making them difficult for humans to notice.

The researchers inserted a road sign into a Coca-Cola advertisement video on a billboard for only three video frames, again, tricking the car into trying to drive dangerously fast.

To test attacks on semi-autonomous cars the researchers used a Tesla Model X with its “Hardware 2.5” autopilot capabilities. Although it is called autopilot, the system is only meant to help a human driver, and the company intends to release fully autonomous cars in the future.

Tesla's obstacle detection system employs an array of cameras, sensors and radar.

The research team projected the image of a person onto the road ahead of the vehicle, which it perceived as an actual human. (The “phantom person” the team projected onto the road was an image of Tesla founder Elon Musk.) The car's radar and sensors, which are meant to monitor other vehicles, are unable to effectively detect people.

Next, the team projected an image of a car in front of the Tesla. To its surprise, the Tesla registered the phantom image as an actual vehicle, suggesting that the car's obstacle detection system does not validate objects it sees with

the vehicle's sensors. The researchers said they contacted Tesla for an explanation, but the company declined to comment.

Attackers could exploit the flaw by projecting images onto a highway, for example, causing cars to brake suddenly and putting the occupants of the car and other vehicles in danger.

The researchers also tricked the Tesla into driving into an oncoming traffic lane by projecting false road markings onto the street ahead of it.

To fix the problem, the researchers recommended employing add-on software that could validate objects by having the cars rely on more than the cameras to navigate. The cameras should take into account information beyond shapes and patterns, such as an image's size, angle, context, surface, and lighting, they said.

The paper said that, in response to the findings, Mobileye said: "There was no exploit, no vulnerability, no flaw, and nothing of interest: the road sign recognition system saw an image of a street sign, and this is good enough, so Mobileye 630 PRO should accept it and move on."

Tesla dismissed the findings, saying that its autopilot system "is intended for use only with a fully attentive driver who has their hands on the wheel and is prepared to take over at any time. While Autopilot is designed to become more capable over time, in its current form, it is not a self-driving system, it does not turn a Tesla into an autonomous vehicle, and it does not allow the driver to abdicate responsibility."

While current automated driving systems are meant to help, but not replace, human users, many of today's drivers overestimate vehicle capabilities. A study last year found that 48 percent of drivers thought it would be safe to take their hands off the wheel while using Tesla's autopilot system.

Mobileye is an Israel-based subsidiary of Intel, which took over the startup for \$15.3 billion in 2017 in the largest-ever acquisition of an Israeli company.

Keithley I-V Tracer Software uses the touchscreen interface of modern Source Measure Unit (SMU) instruments to recreate the familiar user experience of a curve tracer for two-terminal devices.

By Joseph Gorley

Curve tracers are electronic test devices that are similar to oscilloscopes in many ways. Fundamentally, these instruments operate by varying a parameter, then measuring a separate one to produce data for analysis referencing the characteristics of semiconductors such as diodes, transistors, and thyristors. They're particularly useful for semiconductor failure analysis and parametric characterization.

By 1955, when Tektronix introduced the industry's first curve tracer, the company had already been producing oscilloscopes for nearly 10 years. Curve tracers were a logical extension of the oscilloscope business, approximating a power supply and an oscilloscope packaged in the same box. They work by applying a swept voltage to two terminals of the device under test and measuring the amount of current that the device permits to flow at each voltage level. The resulting I-V graph is then displayed on a scope display.

The first curve tracer, the 570, was introduced to display characteristic curves for vacuum tubes, followed over the years by models for testing transistors, diodes, and other solid-state devices. Tektronix curve tracers such as the 576 and 370A/370B went on to broad industry acceptance until production eventually stopped in the mid '80s.

Continuing with the Classic

Over time, curve tracers have gone to become more sophisticated, complex, and expensive for applications such as comprehensive semiconductor-device-level characterization. Meanwhile, classic curve tracers never stopped being useful, and have continued to remain in demand to this day, primarily for failure analysis and education applications. Their popularity is due, in part, to a simple interaction model that modern curve tracers have failed to replicate. Since they're no longer being manufactured, the sustained demand has led to a robust market for used traditional curve tracers, with refurbished 370Bs for example fetching \$20,000 or more in online auctions.

For labs that prefer to continue using traditional curve tracers, the reliance on vintage instrumentation presents its share of obstacles. Given the cost of usable old instruments, labs typically will share one unit across all of their failure analysis engineers. Keeping existing instruments functional requires sourcing and stocking old replacement components, and the rather large footprint consumes limited lab space. And since the storage device is a floppy disk (remember them?), capturing and sharing data can present a challenge.

A Software Version

Recognizing the ongoing interest in traditional curve tracers, the Keithley division of Tektronix has introduced I-V Tracer software to bring many of the features loved in classic curve tracers to a modern instrument, namely Keithley SourceMeter Source Measure Units (SMUs). Since an SMU can source voltage or current while measuring voltage and current, it has similar hardware qualities to a curve tracer.

The new software leverages the touchscreen interface of 2400 Series Graphical SMUs (Fig. 1) to re-create the familiar user experience of a curve tracer for low-power two-terminal devices. I-V Tracer uses the full capabilities of supported SMUs, including the dual high-speed digitizers of the 2461 for example, to perform tracing with AC polarity and pulsed DC, in addition to standard DC polarity. This maps to the 576, for example, that had +DC, -DC, and AC polarities, which means the output is either +voltage, -voltage, or both + and - voltages.



1. Keithley I-V Tracer software leverages the touchscreen interface of 2400 Series Graphical SMUs to recreate the familiar user experience of a curve tracer for low-power two-terminal devices.

Rather than pooling resources for shared, outdated equipment, labs can now equip individual engineers with their own I-V curve tracer. Keithley SMUs are portable and fit easily on testbenches, as shown in the comparison with a 370B (Fig. 2).

Moreover, users still have all of the source, sink, and measure capabilities of Keithley SMUs with curve-tracer functionality a click or tap away. It's also easy to export curves or screenshots, and there's no need to scour online listings for spare parts.



2. An SMU with Keithley I-V Tracer software is a fraction the size of a classic curve tracer.

For failure analysis of semiconductor devices, curve tracers are so popular, in part, because they provide precise control and immediate results. If you source too much power through a device, it's possible to destroy the sensitive evidence that points to the root cause of failure. I-V Tracer simulates this functionality by providing direct control over the output level, letting failure-analysis (FA) engineers slowly ramp up to an I-V curve anomaly then seamlessly creep into the behavior, with a minimum of 500 nV (or 500 fA) resolution on the sourced output.

In addition to failure analysis, curve tracers have long been considered a must-have in engineering classrooms for their simplicity in allowing students to directly apply their learning to electrical devices. I-V Tracer with an SMU offers the same benefit, giving real-time, direct control to students, allowing them to experiment themselves and solidify understanding on a wide range of electronics.

Priced at \$1,499 plus the cost of the SMU, the I-V Tracer offers a cost-effective way to gain the usability benefits of a classic curve tracer without the burden that comes with trying to maintain a vintage instrument.

For more information, go to: tek.com/keithley-i-v_tracer.

By Li-Ann Dias

The tech industry is built on the venture capital model where hockey stick growth and selling to a larger company or going public are markers of success. But the traditional VC model does not leave much room for startups that might not be the next unicorn but still generate revenue — just not the type of returns investors are looking for.

This is where exiting to the community comes in.

“A lot of times, selling to the public doesn’t necessarily make the company or its service a better experience for the user or the workers,” Start.coop founder Greg Brodsky previously told TechCrunch. “Often it gets worse. It’s only really better for the investor.”

Brodsky, who helps cooperative startups through the Start.coop accelerator, pointed to this exit to community idea as an option for startups looking to transition out of the more traditional Silicon Valley model. In this framework, some portion of the company is sold back to the workers or end users, he said. This idea is being spearheaded by Nathan Schneider, a Start.coop board member and professor of media studies at the University of Colorado, Boulder.

“The idea with exit to community is how can you create a model where the whole point is to create a vibrant community that will become its eventual stewards,” Schneider tells TechCrunch. “It seems like a natural fit, especially in a moment where we’re looking for increased accountability and the wealth distribution problems in the startup economy.”

Through the Exit to Community project, Schneider is exploring ways to help startups transition from investor-owned to community ownership, which could include users, customers, workers or some combination of all stakeholders. Schneider is holding a series of meetings with people interested in this challenge to try to chart a clear pathway.

“I don’t think it’s going to be one playbook,” Schneider says. “We’re looking at a variety.”

The variety of approaches, however, does not include cooperatives, which we’ve explored here. In short, co-ops shift the ownership model from investor-owned to worker or user-owned. Schneider has explored additional possibilities due to barriers startups face with access to venture funding and other types of structural support. So far, Schneider says he has identified three different ways to go about doing this.

The first is a stockholding trust, inspired by more traditional employee stock-ownership plans in the U.S. The goal of the trust would be to acquire 100% ownership in the company over a period of time by buying out pre-existing outside investors. We’ll save the nitty-gritty for another time, but this approach would offer investors a way to recoup their initial investment and then some while giving the ownership to founders, employees and users.

Schneider explored this user-buyback option with Meetup co-founder and chairperson Scott Heiferman. WeWork, which bought Meetup in 2017, later decided to put it up for sale last year. When that happened, Heiferman started thinking about what it would take to buy it back.

“The question I am fascinated by — why I’m here today — is just that another world is possible,” Heiferman said during a webinar in December. “I don’t think we have seen the future yet of what it means to reimagine platforms that the world is depending on as sort of the lifeblood of software eating the world. The internet still remains a

network of people, not a network of computers, and hopefully not a network of corporations. For those of you involved at venture at earlier stages, I hope you'll be bold in imagining different possibilities."

For Meetup, it's looking like it's not going to work out for Heiferman to buy it back. He said there's another deal in the works that he's not involved in.

"It doesn't look like we'll be able to make a competitive offer and that's what poses the challenge for a business where its revenue comes directly from users," Schneider says. "It's so aligned with its users in that way, and it's all very member-centric and has that spirit to it. So why can't we do that? Why is that not an obvious model for it to be owned by users? We don't have the maturity as an ecosystem to really enable this at this point but this is an example of why we need to make this possible."

The second potential conversion strategy is the creation of a federation. In this scenario, a company could reorient itself toward the community by distributing decision-making powers to moderators and users of the platform. In order to seamlessly do this, the startup would need to first convert to a benefit corporation. The third option involves tokenization via blockchain to reduce the number of intermediaries involved and make shared governance more transparent and direct.

The plan for this year, Schneider says, is to keep gathering people around these ideas and to further flesh out the possible pathways to converting into a business that exists to the community. Beyond empowering users, Schneider sees exiting to the community as a way to increase inclusivity in the tech industry and better serve diverse communities of people.

"The thing with inclusivity is it's not just about representation, it's about the model," Schneider says. "It's the fact that our business models are not meeting the needs of diverse communities and we need more diverse business models for technology and startups if they're really going to be an inclusive community."

By Kurt Schlosser



The produce section, and the swarm of cameras overhead, inside the new Amazon Go Grocery store on Capitol Hill in Seattle.

Amazon Go is going full grocery.

Two years after launching a chain of convenience stores without cashiers or checkout lines, Amazon is opening its first “Amazon Go Grocery” store in Seattle on Tuesday morning, enlarging the footprint for surveillance-style shopping and signaling a larger challenge to the broader world of brick-and-mortar retail.

The debut is also the answer to a longstanding mystery about the 7,700-square-foot space, at 610 E. Pike Street in Seattle’s Capitol Hill neighborhood. Amazon’s plans for the property have long been under wraps. Last fall the company confirmed that its Amazon Go team was “running internal tests” at the location, but declined to say more until now.

GeekWire got a sneak peek at the store during a recent media preview, entering by scanning a smartphone app and strolling the aisles of the completely stocked store. The banks of cameras and sensors overhead track everything put into a shopping cart, with the help of artificial intelligence — rendering unnecessary the old-fashioned ritual of scanning and paying at a checkout stand. Items are charged to a shopper’s Amazon account shortly after they walk through the exit.

Apart from the larger size, the concept is very similar to the Amazon Go convenience stores that first opened to the public in Seattle in January 2018. Amazon Go has expanded to 25 locations across cities including San Francisco, Chicago and New York. That smaller concept, sized between 450 and 2,700 square feet, ushered in an era of grab-and-go shopping.

“What Amazon Go did for central business districts — like locating it very close to where people work so you can get breakfast, lunch, snacks — Amazon Go Grocery does the same thing, but closer to home,” said Dilip Kumar, vice president of Physical Retail & Technology for Amazon. “It’s a new format, it’s not just a bigger Amazon Go. It’s a much more expanded selection that caters to what people are looking for shopping for groceries.”

What Amazon is looking for is yet another answer to traditional retail, where it’s leveraging convenience and technology in a \$675 billion U.S. grocery industry. The tech giant scooped up Whole Foods in 2017 in a bid to take on the sizeable brick-and-mortar footprints of Walmart, Target, Kroger and others. Those companies have consistently responded to Amazon’s digital pushes around online grocery ordering and delivery.

Amazon posted \$4.4 billion in revenue last quarter in its physical stores category, which includes Whole Foods and Amazon Go stores.



Shoppers scan a QR code in the Amazon Go app on their smartphones upon entering. (GeekWire Photo / Kurt Schlosser)

Kumar declined to say how many Amazon Go Grocery stores are coming, where the next one might be, or whether they will all be the same size. Plans for an even larger grocery concept in Los Angeles and elsewhere are “something else” entirely, he said, but he likes what they built first in Seattle.

“I think it fits the neighborhood, it’s the right size,” Kumar said. “It’s an interesting challenge to be able to fit all the selection that people would care about, in a store.”

No matter the size, the continued push toward tech and automation is sure to fuel the ongoing debate around human workers being replaced by machines. Amazon Go Grocery will staff just a handful of associates.

After entering the new store through the kiosks which scan a smartphone QR code, a familiar sight greets traditional grocery store shoppers: a line of shopping carts at the ready. Free, green shopping bags are also offered.

Hundreds of cameras in the ceiling overhead make up the key technological component of the just-walk-out concept, and they’re put to the biggest test in the produce section, where a variety of individually priced fruits and vegetables are available.

“Most of the things at Amazon Go are packaged, or they’re single items like a can of Coke,” Kumar said. “But here, people are shopping for potatoes or they’re shopping for onions — there’s a lot more browsing and rummaging that tends to happen. That’s what makes this problem a lot more complicated.”

Amazon’s goal is to generate accurate receipts, no matter how long you stand over the avocados or apples, shifting them around and picking them up before settling on three and then changing your mind to two.

The cameras are keeping track of those “interactions” with the product and know exactly what is being taken off shelves and put back. Allowing people to do this type of “considered shopping” plays into the Go Grocery concept of making sure that customers don’t have to do anything unnatural when it comes to how they shop.

“They’re used to seeing produce laid out in [a traditional] way,” Kumar said, joking about how it’s almost necessary, as a shopper, to get spritzed by the misters in the lettuce section.

Kumar called a robust produce section the hallmark of any good grocery store, and Amazon Go Grocery sources its organic produce from the same farms that supply Whole Foods. Its 365 organic label is on prominent display.

Up and down aisles throughout the store — there are 5,000 unique items — national brands are mixed with local favorites that Amazon believes its neighborhood customers would expect the store to stock.

There is no meat or seafood counter and no food preparation on the premises. Fish, chicken and beef products are brought in several times a week, individually wrapped. Signage near cases advises customers on the differences between cuts of meat or wild caught seafood vs. farmed fish. There is also an artisan cheese area where people can get the same sort of quick education via signage rather than from a human cheesemonger.

In an aisle featuring other grocery essentials and household supplies, such as cleaning products and deodorant, there is also pet food, treats, toys and cat litter.

“A store wouldn’t be complete if you didn’t really have an ode to pets,” Kumar said. “In Seattle it’s a must have.”

And it’s another indication that Amazon Go Grocery goes beyond Amazon Go.

Back near the front of the store, the quicker grab-and-go nature of what Amazon likes about its Go concept is more readily on display. It’s here where the fresh baked goods — donuts, bagels, fritters and more — and self-serve coffee and espresso stations are located. There’s a sizable alcohol section — where you’ll run into a human who has to check your ID. And around the corner is a large section called “Meals Made Easy” that caters to the what’s-for-dinner shopper with entrées including pasta, salad, pizza, sushi and more.

What to grab at the end of the day was a big driver in Amazon’s decision to extend Go into grocery, closer to where people live.

And people certainly live in this part of Seattle, where the Capitol Hill neighborhood, like many across the city, has seen a boom in apartment construction as tech workers and others seek housing close to the urban core. The store itself sits below five stories of apartments and the blocks around it feature similar new buildings.

The entire footprint for the location, including space for back stock and more, is 10,400 square feet. But the store will not serve as a hub for grocery delivery, the company said.

And it won’t replace Whole Foods or other methods that shoppers appreciate because Amazon said it has come to realize that customers want to shop in a variety of different ways for a variety of different needs.

“Some people want their food delivered, some people want to go shopping at Whole Foods, some people want to shop at a different kind of store,” Kumar said. “The single biggest thing that people say is that they don’t have enough time to do all the things that they need to do. One of the key things that we always index on is how we can provide the convenience that customers expect in places where they are.”



Self-serve coffee stations near the front of the store.

By Northwestern University



One hundred small robots line up in the laboratory. Credit: Northwestern University

For self-driving vehicles to become an everyday reality, they need to safely and flawlessly navigate one another without crashing or causing unnecessary traffic jams.

To help make this possible, Northwestern University researchers have developed the first decentralized algorithm with a collision-free, deadlock-free guarantee.

The researchers tested the algorithm in a simulation of 1,024 robots and on a swarm of 100 real robots in the laboratory. The robots reliably, safely and efficiently converged to form a pre-determined shape in less than a minute.

"If you have many autonomous vehicles on the road, you don't want them to collide with one another or get stuck in a deadlock," said Northwestern's Michael Rubenstein, who led the study. "By understanding how to control our swarm robots to form shapes, we can understand how to control fleets of autonomous vehicles as they interact with each other."

The paper will be published later this month in the journal *IEEE Transactions on Robotics*. Rubenstein is the Lisa Wissner-Slivka and Benjamin Slivka Professor in Computer Science in Northwestern's McCormick School of Engineering.

The advantage of a swarm of small robots—versus one large robot or a swarm with one lead robot—is the lack of a centralized control, which can quickly become a central point of failure. Rubenstein's decentralized algorithm acts as a fail-safe.

"If the system is centralized and a robot stops working, then the entire system fails," Rubenstein said. "In a decentralized system, there is no leader telling all the other robots what to do. Each robot makes its own decisions. If one robot fails in a swarm, the swarm can still accomplish the task."



Still, the robots need to coordinate in order to avoid collisions and deadlock. To do this, the algorithm views the ground beneath the robots as a grid. By using technology similar to GPS, each robot is aware of where it sits on the grid.

Before making a decision about where to move, each robot uses sensors to communicate with its neighbors, determining whether or not nearby spaces within the grid are vacant or occupied.

"The robots refuse to move to a spot until that spot is free and until they know that no other robots are moving to that same spot," Rubenstein said. "They are careful and reserve a space ahead of time."

Even with all this careful coordination, the robots are still able to communicate and move swiftly to form a shape. Rubenstein accomplishes this by keeping the robots near-sighted.

"Each robot can only sense three or four of its closest neighbors," Rubenstein explained. "They can't see across the whole swarm, which makes it easier to scale the system. The robots interact locally to make decisions without global information."

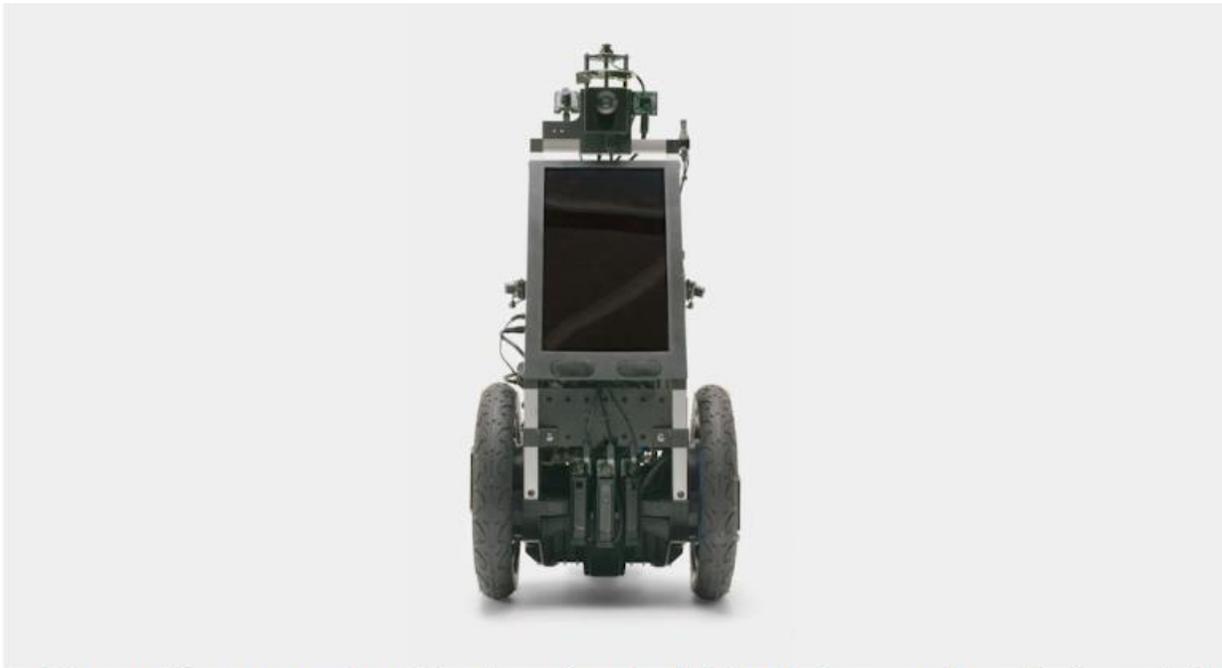
In Rubenstein's swarm, for example, 100 robots can coordinate to form a shape within a minute. In some previous approaches, it could take a full hour. Rubenstein imagines that his algorithm could be used in fleets of driverless cars and in automated warehouses.

"Large companies have warehouses with hundreds of robots doing tasks similar to what our robots do in the lab," he said. "They need to make sure their robots don't collide but do move as quickly as possible to reach the spot where they eventually give an object to a human."

More information: Hanlin Wang et al, *Shape Formation in Homogeneous Swarms Using Local Task Swapping*, *IEEE Transactions on Robotics* (2020). DOI: [10.1109/TRO.2020.2967656](https://doi.org/10.1109/TRO.2020.2967656)

Journal information: [IEEE Transactions on Robotics](#)

By Steve Crowe



Anki was working on two versions of the Bingo robot: the Mini Bingo for home security applications and this larger Bingo robot for commercial security applications.

We reported last June Anki already had a prototype of its next consumer robot prior to shutting down. The Robot Report has learned there were multiple products in the works, and neither one is likely what you'd expect.

First, meet Bingo, a security robot with a German Shepherd-inspired design. Anki was working on two versions of Bingo: a small consumer version called Mini Bingo for home security applications and a more robust, rugged version for commercial and military security applications.

The consumer version was designed to roam around the house, go to recharge its battery, and then continue its patrol. If something was out of place, it would send you an alert. It would even bark like a real dog. The larger model was about four-feet tall and had capabilities including thermal imaging, room scanning, mapping, facial recognition and more.

The Robot Report broke the news about edtech startup Digital Dream Labs acquiring Anki's IP assets in December 2019. H. Jacob Hanchar, Founder of Digital Dreams Labs, told us home security-type functions are some of the most requested features customers want added to Vector, the second-generation consumer robot from Anki.

"People just connect with [Vector] and begin to trust it," he said. "It's an amazing level of trust and emotional connection they give this robot. It's unlike any other emotional connection I've ever seen in robotics."

However, Hanchar said Digital Dreams Labs likely won't pursue Bingo as new products. The larger version seems like "hardcore DARPA stuff" and the functionality of Mini Bingo could be embedded into Vector. If the latter happens, Hanchar said it won't be until 2021.

“We’ll think about the security features inside of Bingo, what people are asking for and enhance Vector with these features,” Hanchar said. “Vector’s internal workings are sophisticated enough to handle these features, so I don’t see a need to build a new robot.”



Anki designed Mini Bingo to be about twice the size of a soda can.

Stripped down Anki Cozmo 2.0

Anki was also working on Project Whiskey, which was a stripped down Cozmo 2.0. It was designed to be sold as a toy rather than a robot.

“It is less sophisticated, there isn’t as much memory, the plastic is thinner, the components are lighter, and it can’t do as much enhanced computing as the current version of Cozmo can,” Hanchar said. “We will always consider how to cut costs, but the demand is fever pitched for Cozmo as it was. So we probably won’t consider this.”

Anki raised more than \$200 million and had \$325-plus million in revenue since it was founded in 2010. However, as this teardown of Vector points out, the manufacturing process for Vector was quite complex, unsustainable, and likely contributed to Anki’s shut down.

Hanchar said Anki was gearing up for a \$1 billion IPO and caused the company seemingly lose its focus.

“Toward the end, they were throwing everything they could at the wall,” he said. They were hurrying to show they were a full suite robotics company. Things were very hurried, very rushed. A lot of it doesn’t make a lot of sense to me.”

Vector Kickstarter

Digital Dream Labs is planning to revive and manufacture more units for each Anki product line: Overdrive, Cozmo, Vector. The company has also discussed a subscription-based model that could reduce the upfront cost of the robots and offer different levels of functionality to customers.

Digital Dream Labs also launched a kickstarter to enable users to take Vector off of his external servers while still remaining fully functional. The campaign surpassed its \$75,000 goal in less than one day and had raised more than \$120,000 at press time.

“The Kickstarter goal was the bare minimum to fund the project,” Hanchar said. “The more we raise, the faster we can work. Vector’s updates will be released in the Fall.”

Hanchar said manufacturing of Cozmo units is on track to begin in June with the product “hopefully on the shelves for Christmas.” To meet at these promises and deadlines, Hanchar said Digital Dream Labs plans to hire 30-plus employees this year.



The collaboration is expected to produce the highly integrated, modular architectures required for high-volume vehicle applications.

By Murray Slovick

Electric and hybrid vehicles must become more affordable for consumers and more profitable for producers. For that to happen, suppliers have to find new and better ways to make electric propulsion more economically viable.

Standardizing electric drive systems with integrated inverters and motors—combining the electric motor, gearbox, and power electronics into a single package—could reduce cost, weight, and packaging for car manufacturers. In turn, it will make electrification more affordable.

To that end, GKN Automotive and Delta Electronics of Taiwan, a global provider of switching power supplies and thermal-management products, announced their collaboration aimed at enabling more rapid acceleration of next-generation integrated 3-in-1 eDrive systems. An electric motor and inverter supplied by Delta Electronics is integrated with GKN Automotive's eMotor and single-speed transmission module in a 3-in-1 solution. GKN's Automotive SPICE Level 2- and 3-certified software engineering processes enable it to manage the system's complete integration.

According to the two companies, within three years, the new 3-in-1 eDrive units will be available for the start of production in a multitude of vehicle types, from A-segment city cars to D-segment executive cars and SUVs. The units will support a range of torque outputs from 2,000 to 5,800 Nm, with nominal power outputs of between 80 and 155 kW.



GKN's Family 3 eDrive system covers small family "compact" cars and "midsize" vehicles with peak axle torque of 2,700 to 4,100 Nm. (Source: GKN)

Using a quiet electric engine unmasks high-frequency noise caused by the transmission. GKN says reducing noise, vibration, and harshness (NVH) is possible with its new Powder Metallurgy (PM) gears tailored for EV or HEV modular drivetrain platforms. The compact eDrive setup with electric motor, gearbox, and inverter all in one housing also can be optimized to suppress noise and vibrations, delivering greater refinement and efficiency than powertrains incorporating separately packaged elements (for maximum serviceability, the inverter will retain its own separate housing in the 3-in-1 setup).

Hybrid Topology

GKN Automotive's eDrive systems are said to be capable of providing either full-electric power, or hybrid-electric power, to support an existing internal combustion engine (ICE). The hybrid-electric topology can be characterized by the positioning of the main components of the hybrid-electric system on the vehicle:

- P0: The electric motor is connected to the ICE through a belt on the front-end accessory drive.
- P1: The electric motor is connected directly with the crankshaft of the internal combustion engine.
- P2: The electric motor is side-attached (through a belt) or integrated between the ICE and the transmission; it's decoupled from the ICE and it has the same speed of the ICE (or a multiple of it).
- P3: The electric motor is connected through a gear mesh with the transmission; the electric machine is decoupled from the ICE and its speed is a multiple of the wheel speed.
- P4: The electric motor is connected through a gear mesh on the rear axle of the vehicle; it's decoupled from the ICE and located in the rear axle drive or in the wheel hub.

P0 and P1 architectures don't allow for the mechanical disconnection of the electric motor from the engine. On the other hand, P2, P3, or P4 configurations disconnect the electric motor from the engine through a clutch.

The high efficiency of P4 electric axle drive architectures will make them central to most pure electric and plug-in hybrid strategies going forward. In the P4 hybrid-electric drive systems, the electric motor is decoupled from the ICE—the inverter and motor are coupled into a compact transmission module that applies torque directly to the rear wheels and recovers energy directly from the wheels when braking.

To provide the best balance of cost and performance, single-speed systems are expected to account for 94% of P4 eDrive volumes. By 2026, annual production of P4 electric-drive systems is forecast to reach around 13 million units, according to GKN.

In a statement, GKN Automotive said it's the world's leading supplier of P4 electric drive systems. This next phase of its electrification strategy will produce the highly integrated, modular, scalable architectures required for high-volume vehicle applications. Standardizing these systems will make electrification more affordable and will enable OEMs to offer consumers optimized e-mobility "at the right price."

Meeting Emission Rules

Regulation remains a main driver for electrification in Europe and the U.S., with CO₂ penalties providing a strong financial incentive for automakers. Under EU rules, carmakers must lower average CO₂ emissions from the vehicles across their fleet to 95g per km, or face penalties that could run into hundreds of millions, even billions, of euros. Industry forecasts indicate that by 2026, one in every 10 vehicles sold worldwide will be a battery electric vehicle.

Together, GKN and Delta will invest about £100 million a year (approximately \$130 million) in developing standardized eDrive solutions for the next generation of battery-powered vehicles. They say their new solution will cover a large portion of the future eDrive market, which is projected by IHS to be worth more than £12 billion (approximately \$15.6 billion) by 2030.

Sourced by EnterpriseInsights

Huawei said today (Feb 25, 2020) at a London launch event for new 5G gear that it would prefer to go via the operator community to serve the Industry 4.0 movement, using slices of their licensed public networks to create dedicated networks for industrial automation and intelligence.

It follows the line from Swedish rival Ericsson, to go through its traditional networking customers as it seeks to thread the industry with wireless connectivity. Nokia, the other of the big three kit vendors in the traditional telecoms market, has been more aggressive, stating it will go direct as required – although it has also toed the operator line more recently.

Speaking during a roundtable session at a product launch in London, where the Chinese vendor proclaimed its tech leadership and unveiled new innovations for 5G radio, core, and transport networks, the company said operators are best to serve the enterprise set using new 5G slicing and edge computing mechanisms.

Ryan Ding, executive director of the board and president of Huawei's carrier business group, said: "The most efficient approach is to allocate spectrum to operators. That has been our view, and it has not changed. 3GPP already has two key technologies for private networks. The first is network slicing. The second is multi-access edge computing. With these two, we will be able to slice a part out from the public 5G network to serve enterprises. This is the most efficient approach to serve enterprises."

At the same time, the firm said it will go direct if enterprises are adamant they want to keep operators out of their private network operations. "As telecom provider, our global strategy is to advocate spectrum allocation to operators... [But] as long as enterprises want to buy equipment from us, we won't say no," said Ding.

Among a flurry of innovations revealed in London (check-out the write-up here), Huawei trailed its new FlexE-based 5G slicing solution, afford industry-beating slicing precision of 1 Gbps, according to its author.

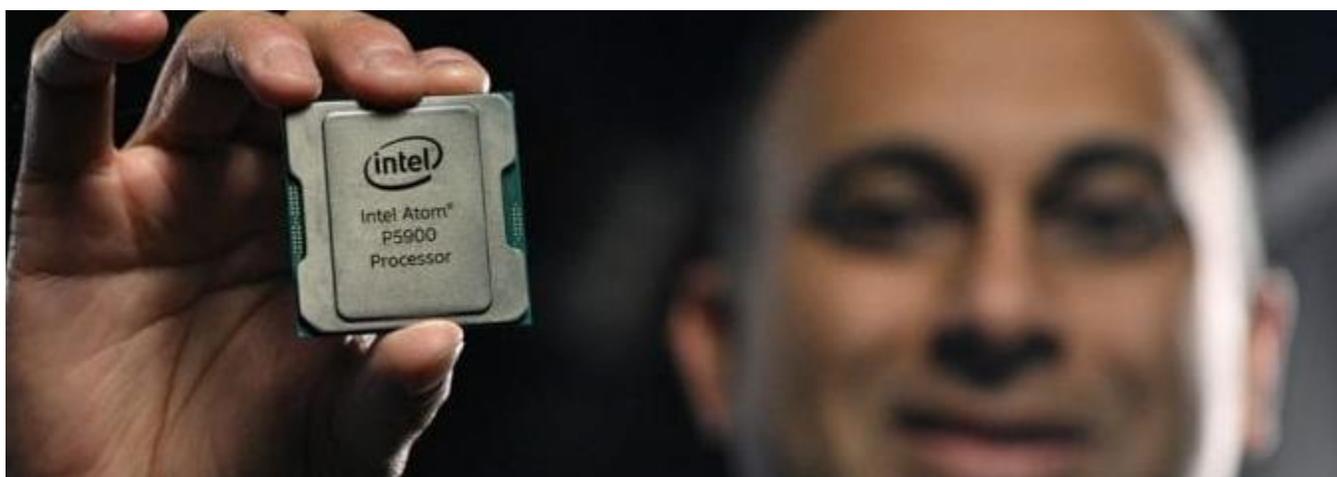
"This is five times the industry average," the company stated. It means best-effort IP networks can be transformed into "SLA-committable IP networks", which cut can independent slices on the IP network for high value services – "like a dedicated lane on the road so emergency vehicles can pass traffic jams", explained Peng Song, president of marketing and solution sales in Huawei's carrier group, on stage in London.

In effect, critical communications and high-value industry can, now, rely on operators to underwrite slices of their public networks. "Because today's IP networks cannot meet the requirements of vertical industries. And vertical industries cannot rely just on best-effort [networks]. They need SLAs to be committed," he said.

Vodafone, notably, has said mobile operators must take more risk to make more revenue from new customers in the 5G era. The operator community must be prepared to put in place service-level agreements (SLAs) around 5G services offered on network slices to industrial operatives, if they are to win business from the digital transformation of the broad enterprise sector, it has said in these pages.

The new Huawei 5G network solution also utilizes in-situ Flow Information Telemetry (iFIT), meaning faults can be located in minutes and the network is highly available. Song described the mechanism as "like Google Maps", where network data, like networked devices, is tagged to report latency, traffic status, and any blockages on the road.

By Sean Kinney



Ericsson and ZTE will use Intel 5G chips

Intel today (Feb 25, 2020) announced the Atom P5900, a 10-nanometer system on chip designed for use in 5G wireless base stations, and right on time as the silicon giant projects the need for six million 5G base stations deployed through 2024.

Intel's Navin Shenoy, EVP and GM of the Data Platforms Group, said in a statement the company views 5G network infrastructure as “a \$25 billion silicon opportunity by 2023. By offering customers the fastest and most effective path to design, deliver and deploy 5G solutions across core, edge and access, we are poised to expand our leading silicon position in this growing market.”

Intel is targeting roughly 40% share of this market and projects it will be “the leading silicon provider in base stations by 2021,” down from previous forecasts of 2022. This would hinge on adoption by major players like Ericsson and Nokia but also represents an opportunity for the rapidly-developing open RAN space, which is focused on driving multi-vendor interoperability to decrease network costs and, ideally, foster innovation.

In addition to the P5900, Intel detailed its second generation Xeon Scalable processor which the company says will provide a 36% performance boost. Also debuted today is Intel eASIC, a project code-named Diamond Mesa, that is meant to essentially split the difference between more costly field programmable gate arrays and cellular application-specific integrated circuits for cellular equipment. In conjunction with the product launch, Intel announced a win with Chinese network infrastructure vendor ZTE.

Intel's Dan McNamara, SVP and GM of the Programmable Solutions Group, called out how 5G use cases are marked by a massive increase in overall data volume and said the company's “structured and standard ASICs enable ZTE to achieve critical cost goals and cement their position in the exploding 5G market.” ZTE used FPGAs to prototype various 5G implementations but said it needs to lower costs for scaled deployments.

Last year Intel pivoted its 5G strategy after pulling the plug on development of modems for devices and honed its focus on supporting the infrastructure side of 5G networks.

By Catherine Sbeglia

5G requires more high-powered, RRUs and, in some regions, a shift to massive MIMO architecture, says CommScope

The implementation of 5G requires more radio units, more cell sites, more power. But at the same time, operators are looking for solutions that have minimal impact on existing network infrastructure. This can be a difficult balance to strike. According to CommScope Director of technical sales Colin Bryce, 2020 is 5G's year, and for its part, the network infrastructure provider has announced a round of [recent products designed to help operators achieve smarter, faster 5G implementation](#).

“Operators want to deploy 5G antennas on sites that already support multiple technologies and frequency bands—with minimal impact to existing site infrastructure,” Bryce told RCR Wireless News. “CommScope is meeting this challenging demand with a comprehensive antennae lineup that integrates multiple technologies and features such as high port count, multiple arrays, multiple frequency beamformers, as well as support for Time Division Duplex (TDD) and Frequency Division Duplex (FDD).”

The narrow-width antennas, which come in multiple lengths and bands, add capacity to sites with zoning restrictions or where structural loading is limited can be particularly challenging. Further, the new CommScope antenna are the only ones in the industry to offer 4 mid-band arrays for 43 cm and 3 mid-band arrays in 39 cm, allowing operators to attach more radios to a single antenna.

Bryce also explained that rolling out 5G requires more high-powered, remote radio units (RRUs) and, in some regions, a shift to massive MIMO architecture. In addition, the higher-wattage remote radio units require thicker, more expensive copper cables.

He went on to explain that CommScope's plug-and-play dc power supply PowerShift eliminates the need for a tower-top converter.

“Put simply, PowerShift optimizes the transmission of electrical power by regulating voltage at the RRU or remote. PowerShift also increases the amount of time a battery backup runs, allowing RRUs to stay active up to 35% longer in the event of a power outage,” he said.

The PowerShift is a single-rack unit, supporting up to 2,000 wattage per circuit, with power connectors in the front.

Bryce also discussed the state of standalone 5G, offering that pre-commercial technology demos have already been taking place, and operators, particularly in the U.S., are eyeing initial launches in 2020.

“However,” he added, “real-world mass rollouts are probably two or three years away in most markets. Power and backhaul are limiting factors, and CommScope is working on solutions to mitigate both of these issues.”

2020, though, will still be a big year for connectivity, according to Bryce. On the Wi-Fi side of things, he said that Wi-Fi 6 access point shipments will increase across multiple and diverse verticals such as healthcare, education and hospitality to support high-bandwidth applications including 4K video, eSports, AR/VR, facial recognition and public safety.

“Wi-Fi 6 APs, which support up to a four-fold capacity increase over preceding Wi-Fi 5 Wave 2 APs, are expected to represent the majority of access points shipping by the end of 2020,” he stated. “Multiple Wi-Fi 6 APs deployed in dense device environments can collectively deliver required quality-of-service to more clients with more diverse usage profiles.”

And when it comes to 5G, Bryce explained that a successful, cost-effect implementation requires both passive and active radio solutions in sub-6 GHz spectrum bands that are capable of utilizing FDD and Time Division Duplex TDD band plans, and that operators must consider a number of tradeoffs and factors including performance versus cost, EMF emissions and deployment constraints.

He added that while network operators still have a lot of work to do to “make 5G technology live up to its full potential,” CommScope truly believes that “2020 will be the year that network operators put the pieces into place to deliver the 5G promise.”

By Jonathna Shieber

One of the biggest roadblocks to reducing costs in the American healthcare system is the system's inherent lack of transparency. Most healthcare networks and hospital systems can't even accurately account for the doctors they manage and which insurance plans those doctors accept — let alone how good those doctors actually are at providing care, according to Ribbon Health chief executive Nate Maslak.

The former healthcare consultant founded Ribbon Health to address just that issue, and the company has raised \$10.25 million in new financing to roll out its software services to a broader network of payers, providers and digital health companies. The new financing was led by Andreessen Horowitz, and included Y Combinator and the New York-based investment firm BoxGroup. Individual healthcare executives like Nat Turner, the chief executive of Flatiron Health; Vivek Garipalli, chief executive and co-founder of Clover Health; and Eric Roza, the former chief executive of DataLogix, also participated in the financing.

It's the first deal for Andreessen's newest healthcare-focused partner, Julie Yoo, and is in an area with which Yoo is quite familiar. The former serial healthcare entrepreneur developed a similar business to tackle better data collection and delivery for hospitals at Kyruus.

Taking an API-based approach, Ribbon Health is building on the Kyruus approach, Yoo said, with the potential to expand across the entire breadth of the American healthcare system.

Simply, Ribbon Health is trying to create an accurate database of what doctors and health plans have, which specializations offer their services to which insurance providers, and produce the best outcomes for patients.

“\$700 billion wasted because of poor decisions,” said Maslak. “The information not flowing to the right place at the right time. Over a third of healthcare spending is wasted and we think that over half is data-addressable.”

“The majority of decisions in health care rely on data about a provider or health plan, yet our industry lacks the systematic infrastructure to centralize this information and contextualize it for those who need it. There is a clear need for a single platform that can provide comprehensive, up-to-date data to enable informed decision making across health care, and we believe Ribbon is poised to lead in this space,” said Yoo, in a statement.

Along with the new financing, Ribbon also unveiled a tool that provides cost and quality information for patients to understand their potential out-of-pocket cost estimates based on their deductible, plan design and provider prices.

“So much of the innovation in health care relies on accurate data. Our goal is to provide these companies the critical data infrastructure needed to improve quality of care, health outcomes, and control costs,” said Nate Fox, co-founder and chief technology officer at Ribbon Health, in a statement. “Our platform and seamless API make it easy for customers to trust us to deliver the most comprehensive, accurate data, allowing them to focus on what they do best on the front lines of health care.”

The company is already working with Oak Street Health and Well (Well Dot, Inc.), and will use the additional funding to expand its sales and marketing efforts and increase adoption.

“Provider data is a basic building block of every healthcare transaction,” said Yoo. “Whether it's you or I trying to enroll... or referral claim processing... there are tens of billions of transactions, all of which require information about a provider.”

By Alex Wilhelm

On the heels of her [conversation-driving Twitter thread](#) on 2020's venture fundraising climate, [Hustle Fund's Elizabeth Yin](#) converted her thoughts into an op-ed for TechCrunch. In keeping with her expansive thread, we asked her to adapt her thread for a TechCrunch column and join us for an extended conversation.

What follows is an interview between Yin and myself that came after I read her piece, digging into venture capitalist fear, the ability of established founders to raise outsized rounds, her advice on growth and how some Series A and Series B-stage companies posting impressive revenue expansion might be nigh-unfundable in this, the new fundraising reality.

What follows is an edited, occasionally condensed transcript of our chat. Let's go!

TechCrunch: Okay, question one. You said, "VCs have gotten scared, almost to a fault." Aside from the WeWork IPO implosion, what are the leading drivers of this recent increase in fear?

Elizabeth Yin: Taking a step back, I think we have to ask ourselves, what is even the place of venture capital in the first place? And when you think about the original venture capital industry, you know, decades ago, those VCs were taking big, big bets, like at those moments in time during the 90s, or even before that, for some of the chip companies or even Apple Computer, there were many bets happening there.

If you had to segment it: there were markets, there were technical bets. And basically a combination of all of the above. Like, could you technically build the Apple computer? And was there a big enough market for it, etc. So VCs, when you think about it, were taking a lot more risk back then. Because back then, you actually needed a lot of technical know-how to even be able to set up servers or, or build, you know, huge undertakings of devices, etc. These days, frankly speaking, a lot of the risk is already off the table for most software VCs. Very rarely am I looking at a startup and thinking, "I wonder if the founders can actually build this website," because in many cases, a lot of the building blocks to be able to do so are pretty streamlined and straightforward.

You still obviously need to be a pretty good engineer or product person, but I think it's just a lot easier now to remove some of the risks, including technical risk — and in other cases, market risk is also removed in cases where a company may be trying to just replace another company that is already a big behemoth, but may not have a great user interface or whatnot.

So I think if we kind of look at how VC has evolved, it's evolved over the years where VC today, I would argue, in software, at least, don't take a lot of risks.

Most B2B SaaS companies are even less risky than some of the other bets, such as some of the scooter companies, where there have been a bajillion marketing software companies that have been built and have turned out to be big, but scooters are relatively new, who knows if that will end up being a big company. But I think you're seeing even more of a shift to take even less risk. And I think that's interesting about this time.

So there's been a decrease in risk tolerance, which I call Increase in fear. I think the WeWork implosion, when the S-1 dropped through the kind of the IPO cycle that ends up being pulled, is generally called the catalyst for this. But it doesn't seem to be enough.

Are there other things that have happened in the last maybe like six or nine months that have led to this either increase in fear or decrease in tolerance?

I would point to Uber as well. And I think if we're going to kind of blanket statement it, a lot of people have been looking at SoftBank's Vision fund, just in general and looking at some of their investments that may not have gone very well in the later-stage markets. And if you had to further blanket statement here, it seems that there are companies in certain categories that people now feel were not really meant for VCs to get involved with.

And those tend to be in areas like marketplaces, or they tend to be in areas that are very capital-intensive, such as WeWork, etc. And so I think people are kind of looking to all of that as an overall change in how they *think* they should think about things.

You said in the piece that you're seeing "Series A and Series B companies with 30% month-over-month growth that were popular before are currently struggling to raise rounds, because they're not profitable."

Do the companies that you just mentioned — the marketplaces and maybe e-commerce players — fit into that box? And if so, are they uniquely unprofitable to the point today and just become essentially unfundable?

No. And so this is where I think the pendulum of risk tolerance kind of swings back and forth. At least in my view, there are many hot companies in, let's just call it marketplaces in the e-commerce category a year ago, a lot of VC guys thought, "these are good investments, let's fund them, let's invest in them," where they've continued to grow on the trajectory they were on even a year ago.

But now, because people's mentality is shifted a little bit just having seen the late-stage market reactions to some of the IPOs — the Uber IPO, or, or the reaction to WeWork trying to IPO — I think there's now this new fear of, "oh, gosh, maybe the public doesn't like these kinds of companies, or maybe later-stage investors to me won't like these companies."

Since I need their capital to continue funding my companies, maybe these are not investments I should be in — and that has a trickle-down effect all the way to early-stage companies. So there's nothing really fundamentally different those companies everybody thought were pretty interesting last year, to this year, they continue on their trajectory. But e-commerce and marketplaces are not hot or interesting or as interesting as before, and so there are fewer people who want to be throwing money at those, and therefore it's just harder to raise money if you're one of those kinds of companies.

Going back to your point about SaaS being very fundable both in the piece and what you said a few minutes ago, because the late-stage and IPO market for SaaS companies is relatively good. I presume that's driving the extra interest of the earlier stages for SaaS companies today. Is that fair?

Yeah. So I think there are a couple of drivers for that. I mean, I think if you look at some of the late-stage SaaS companies like some of the ones that went IPO last year, or even if they didn't go IPO, but have raised later-stage rounds, I would just say as a blanket statement, many of those have done phenomenally well. Those ideas generally went quite well.

If you look at Zoom, that's great. The second component is looking at their characteristics of "oh, well, with SaaS, the margins are high and the upsells can be good and the retention as well." So therefore, just from a business mechanics perspective, even if other funders are not interested in pouring money into my SaaS companies in my portfolio, they can still make their business survive and thrive and really grow fast, even without a lot of capital.

So I think combining those two things together, then VC just think, "aha, well, maybe I should go into SaaS." And I think just, you know, I'm not talking about people who were strictly in marketplaces before now going strictly into

SaaS. But I think, you know, a lot of software investors are generalists. And so, you know, people may start to favor doing more SaaS investments than you know, marketplaces or e-commerce.

One of the things that you mentioned is that raising in San Francisco is very hard. People think it's easy when it's really, really not. And you said there's more funding there, but there are also a lot more startups. And so I'm kind of curious, for early-stage founders looking at this kind of more profit-focused world, would you ever recommend an early-stage company, maybe around the seed stage, move into San Francisco today? Or would you generally recommend that they stay in their home market and build from there?

That's a great question. And I think, like everything else, it really depends. There are a couple of things happening — one is that you see the rise of other ecosystems outside of San Francisco, and there are particular ones. So for example, one of my favorite ecosystems to scout in for startups is actually Toronto.

And Toronto, you know, has had a number of companies come up, I think Shopify being the most notable. There are a lot of alums from Shopify, but also from Google and Facebook spinning out and starting their own companies there. And so they have an ecosystem that is really growing and thriving.

I think they're a great example, but there are, of course, plenty of others in the lower 48 states as well. We all know them: Boston, Seattle, Austin, Denver, Boulder, New York, of course, LA, etc. So I think you see a lot more cities just really growing their startup ecosystem.

Now, it's unclear that you necessarily need to move to San Francisco, at least at the earliest stages, because you can get a lot of advice within your city, as well as online in most cases. A good exception is that they're obviously a lot of really, really small cities in the U.S., and I think if you're trying to build in a small town, it may be worthwhile to move to a larger startup ecosystem that's nearby. And so I wouldn't necessarily say move to San Francisco, but I would move to a place that has a thriving startup ecosystem, at least at the earliest stages.

In the later stages, it is worthwhile to move to San Francisco because as you're growing your company, there are a lot more people in San Francisco who have built high-growth companies before, there's a lot of knowledge that I think is still insider knowledge in San Francisco itself. But at the earlier stages, I don't think that that's necessary.

And then the flip side is San Francisco itself has a lot of problems with housing and cost of living and all this other stuff, that make it pretty cost-prohibitive for a startup founder who has no money to be able to have enough runway to really make an early-stage company work out. And so I think those are kind of all the tensions and why San Francisco is less interesting than perhaps 5-10 years ago.

One thing that you talked about at the start of the piece was the discussion about how for certain founders, it's very easy to raise. And for other founders, it's very, very hard. And you talked about how for some people with very established records or resumes, they can kind of command outside-the-scene face valuations and outside seed-stage raises. You also noted that founders who raise a lot of money often struggle to raise later on; why are people so willing to put large amounts of money to work in these well-pedigreed founders if they tend to struggle later on? Haven't investors caught on to that trend in those deals?

I think there are a couple of tensions here. One is if you're trying to look for patterns as to who does well, it is still true that a lot of the unicorns that you see coming out all have a certain kind of resume, like maybe they went through YC, or maybe they graduated from Stanford, or maybe they're ex-Facebook or Google, or maybe some permutation.

I think a lot of that is still true. But... it's hard to say whether it is actually helpful to have that on your resume. I personally don't believe that actually helps you build a successful company. But I think a lot of people just sort of erroneously look at the results and not who doesn't make it. That then drives who ends up getting funding, because

people then see all the people who built the companies all tended to come from like a FANG company, or they went to a certain school or whatever.

That's just sort of not looking at the data properly. But that happens a lot in this industry, ironically. So that's number one. And number two is, to a certain extent, people who come out of some of these institutions have already been filtered, in the sense that if you were an engineer at, let's say, one of these FANG companies, hopefully that is a signal that actually you are pretty good at engineering, because you've gone through all the rigorous engineering interviews. In this day and age, where what we talked about before in core businesses these days, unlike, say, 10 years ago, you actually don't need to be the most brilliant, technical founder for most generic software ideas. I think that you need to have good user experiences and things that customers like.

But that doesn't mean that you need to have a certain level of knowledge in the way that you did 10 years ago. And I think even further back, like if you need to build your own servers, and there was no information on the web about how to even build a website, the differences between building a company back then and building a company now actually doesn't mean that you need the same level of skills or knowledge now as you did back then.

So the market has evolved in terms of what it rewards more quickly than venture capitalists have amended their pattern recognition or models that they use to invest. And there was a mismatch between how to put capital to work efficiently in certain deals with certain countries and certain countries. I mean, that just sounds right to me, frankly.

Yeah, I think those are the two biggest drivers. But the third one is just networks. Networks, I think are starting to be pulled apart a little bit. But certainly, you know, if you worked at one of these FANG companies and you know, like the co-founder of your company and they've done well, they will likely put money into your startup because they know you and they like you. And they're investing because, you know, you're their friend or whatever.

In these networks, if your friends have money, then it's easier to raise money... so that's that third piece of more insular networks, you're starting to break down that by the rise of accelerators. Like if you get into YC and you were a nobody before, then you automatically get elevated into a network that you didn't have access to previously. And once you get momentum on your round, because your friends invested, then that makes it a lot easier to close a round, even with people who don't know you. So I think those are kind of three drivers for that.

Summarizing your advice to founders, it felt like you were saying “grow quickly, don't lose too much money, and don't raise too much money too quickly.” I know that's way too broad of a stroke to kind of summarize your whole piece, but it felt like you were almost advocating for a more historically normal venture cycle with people taking risk, but not insane amounts of risk and not trying to raise all the capital in the world when they're still like, you know, pre-revenue.

Are we just getting back to like a more normal venture capital market after maybe a five-year period a lack of discipline?

As a very early-stage VC, of course, you know, this is self-serving advice. But I also think that with my founder hat on, it is what founders need to do. Because as a founder, when you don't know exactly what you're building for which exact audience and how to get these exact people, there's a lot of experimentation that needs to happen. And throwing \$3 million into a problem at that stage is not going to help you solve that problem any faster.

I think to a certain extent, putting money into faster experiments helps, but it only goes so far. You can't drive speed of execution with money; it happens through the speed of learning. And I think that's something that there's an upper bound as to how much and resources you can put into increasing that speed of learning in the earlier stages.

And so I think, I'm very fearful that and, you know, we've got a number of founders in our portfolio also who have raised like, say \$5 million at the seed and my hope for them is that they will be very cautious with their money during this experimentation phase and not burn it too quickly.

The companies that we talked about that were growing very quickly, and the companies that are out of favor in the wrong categories — do you think they're going to be able to reduce their growth rate and burn sufficiently to survive until they either come back into favor or they're improving profitability and making them attractive again?

So one, I think there are investors in the Valley who are truly contrarian and will pick up on this arbitrage opportunity. And so I think that very good, let's say marketplaces, are e-commerce businesses, even if they're not profitable, but they're growing quickly.

They may take longer to raise to find those investors. We actually don't do a whole lot of marketplace investments and we don't do any e-commerce investments, either. If we did, that would be probably the first category I could try to put my money into, because it's always better to go where other investors are not looking and they're still great companies.

We will also see startups in those categories die where if they had access to just a tinge more capital, they might have made it. But this is where I think you see separation of the savviest entrepreneurs from kind of everyone else. And I think part of being savvy is also being frugal or effective with your money, too. So we're not going to see all marketplace or all e-commerce companies die, for sure. I mean, every economic state, like they're always winners, and so I think it's just even more important for them to be cash efficient.

But I do think it's more challenging for a marketplace or an e-commerce company to have that growth and be profitable. And that's going to be a challenge just in general for the category.

Your line about how VCs want to invest in both companies that are high-growth and profitable and like you know, who wouldn't? That takes the venture out of venture capital.

Yeah, well, I mean, we've seen this ship, right. Like, I think if you just look far enough back, like if you look at Apple Computer, and if somebody came to you with that as a napkin idea, I think there are very few VCs who would want to back that business.

With the rise of SaaS, mechanical knowledge of different metrics, ratios, expectations and growth rates becoming kind of standardized, it seems like people want to de-risk venture capital entirely. They want companies to just fit a certain model or curve, and then they know they're guaranteed success. It's almost done remarkably less adventurous you know, less cowboy, ask less.

I think on that last point, if you play this out further, like what happened, let's just say to SaaS startups. Well, it's interesting because I think you're gonna see a couple of things happen: one, as long as VCs are going after the same SaaS startup. And granted like these days, I would say that the market for SaaS startups has expanded but, I think it's going to be more competitive for VCs to get in SaaS companies and you see valuations go up even further even though the multiples are already quite high. Yeah, so that's number one.

But then number two, I think the other thing is for the savviest SaaS founders — especially the ones that with a fair bit of traction, where there's a level of predictability — they can see that their retention is strong, their upsells are good. And you're just talking about I just need some working capital to basically pay for my customer acquisition, and that's going to pay back in like two months or whatever, the savviest SaaS founders are not going to go to VC because they'll realize that venture capital money is incredibly expensive when you think nothing is free, when you have that big exit, VCs will make hundreds of multiples over on their money in a great business.

And those founders if they had seen that predictability from the beginning, they wouldn't need to take that extensive capital. They would then go to get a debt line or revenue-based financing sort of debt. I think we're now starting to also see the rise in those funds, as well, competing for the same business. So you're just going to see a bloodbath, not just amongst VCs there, but, but you know, other types of funders.

And then I think even arguably, some of the companies like, you know, brax, has a credit card, and you can get a 60-day, interest free loan, essentially. And so if you have customers who pay you back, like fully within 30 days, that's a no-brainer, you just borrow capital constantly for free, and just fund that customer acquisition. So I think we're going to see a lot happening in alternative financing as well.

I've been tracking the rise in revenue-based financing, venture debt and other things like term loans, various kinds of repayment functions. And to me if you're a SaaS business post-Series B and half your next raise isn't some form of debt facility, why would you take on extra dilution when effectively that's when you would have spent the money anyways? And this cuts out some VCs from the rounds, which probably makes the competition higher, which will lead even higher prices, ironically, and then we're going to end up with what 30 – 40x ARR multiples I mean, I am perplexed at how we've gotten to a roughly 13x multiple on public cloud/SaaS companies as a group, and some are at a 20-25. Even as a SaaS optimist, I do not understand those prices at the early stage, they can be even more stretched. So I mean, how much more room is there to go? I would have said "no more" two years ago, but I would have been flat wrong. So I don't want to put a cap on it. But I see what you're talking about and the tensions that you've described. I just don't know what's going to happen.

Yeah, exactly. So to your point, I mean basically like valuation supply and demand right. So if, let's say debt providers take up now half the round or third of the round. That's like a third less supply for investors like VCs to be investing in. So that will create, you know, an even greater rise in valuations to your point. And I agree with you. I think right now, it's quite pricey. I don't love that as a VC investor, but I think that's what's going to happen.

Newcomers face challenges, but there's never been an easier time to raise for some founders

By Elizabeth Yin

When I was a founder many years ago, I felt like I heard constantly conflicting advice and opinions on raising money for my startup.

It's easy to raise. It's hard to raise. If it's easy to raise, you should raise a LOT of money. You should raise a little money. You should try to go for a high valuation. You should raise at a "normal valuation" so it doesn't bite you later.

It was hard to understand what was going on and what I should actually do.

Many years later, now as a VC, it turned out that most of the things you hear people say about fundraising are generally true and generally good pieces of advice. All at the same time. Even when these ideas conflict. How is that possible?

Because, like anything else, different pieces of advice are apt for different types of companies and founders. Today's fundraising landscape is particularly an interesting time of bifurcation that's worth laying out in detail.

For some founders, it's never been an easier time to raise

In the San Francisco Bay Area, if you're a founder who has a "well-branded" resume, it's a fantastic time to raise money at the earliest stages. It almost doesn't even matter what company you're building. You will get funding. You could be leaving Pinterest to start a company. Maybe you went to MIT and then did a 10-year stint at Google. Or maybe you were a former YC founder who is taking a second crack at a company. Or maybe you sold your last business for \$10 million. If you did any of these things, it's a great time.

For these founders, I'm seeing massive party rounds here in San Francisco — \$3 million – \$5 million seed rounds. Sometimes \$10 million rounds right out of the gates! My friend, a fantastic serial entrepreneur with an exit, raised \$8 million recently at \$30 million+ post-money valuation with only a very early version of a product. Investors literally threw money at her and her round was oversubscribed.

SaaS is hot

And then, even if you don't fit this profile, you can still generate a lot of heat on your fundraise. In the last few months, VCs have become very concerned about profitability. It's not enough to be working on a fast-growth startup anymore. In part, we've all seen big-name startups that were once the darlings of Silicon Valley flounder in the late-stage markets because of high burn rates and being nowhere close to profitability.

And VCs have gotten quite scared. Almost to a fault.

So, I'm seeing companies at the Series A and Series B stages with 30% MoM growth that were popular before now struggle to raise their next rounds because they are not profitable. The feedback they receive is to "come back when you're profitable or really close to it." This mentality change has had a huge impact on marketplaces and e-commerce companies — companies that don't have predictable repeat customers or high margins.

On the flip side, SaaS companies have become the new darlings VCs have gone gaga for. SaaS businesses have repeat customers, strong lifetime values and upsell potentials. They are capital-efficient, high-margin businesses.

And if you are growing well as a differentiated (differentiated being a key word) SaaS company, you probably have many VCs knocking on your door — at all stages early and late even if you are not on the coasts.

For most founders, it's still challenging (as always) to raise money

For everyone else, after reading news stories about such large fundraises, it can be confusing to understand why their own fundraiser is so challenging. Why is it so hard for me to raise money?

It turns out that fundraising is still hard for everyone else. Even in the Bay Area, if you don't fall into the categories above, it's hard. People often erroneously think that just being in San Francisco will miraculously make fundraising easier. That's far from true. There are certainly many people who get funded there, but there are also just many more startups in San Francisco than elsewhere. Outside the SF Bay Area, it's even harder to raise. So we have a weird Goldilocks and the Three Bears situation. Some companies are really hot. Most are really cold.

The press mostly writes about the hot deals, like companies that raise \$5 million seed rounds and went through YC. After all, no one wants to read about how someone's fundraising process is going horribly — that's just not a news story that sells. So now, everyone thinks Silicon Valley is littered with gold just by reading the news. The reality is that San Francisco mostly has poop on the ground and a small number of people will find a Benjamin once in a while.

Valuations are all over the board

I'm seeing valuations well above \$10 million post — even \$20 million post for hot seed-stage companies. And then for companies that are cold, the valuations are where they've always been — largely anchored based on geography. As low as \$1 million post within U.S. and Canada. And it can even be lower elsewhere globally.

So when people ask me what a fair valuation is, it's a really hard question. It depends on where you are, what you're working on and what your background is. Many people think valuations are based on a company's progress. That's just not how it works. Valuations are based on supply and demand. Supply of your fundraising round. And investor demand for your fundraising round. Valuations go up when more investors are interested in investing. There's no such thing as a "typical" valuation.

Everyone's mental model will be shifting

Friends outside of Silicon Valley often ask me if I think this time VCs will favor profitable companies over fast growth.

I think the answer is VCs would love to back profitable companies with fast growth.

(That, of course, begs the question in this day and age with other debt or revenue-based financing options why such a company would raise a lot of VC money, but that's besides the point.)

That said, I do think that in this new era we are entering in 2020, companies that focus on profitability will separate the winners from the losers in the next few years. Thriftier founders will win.

Now, here's the irony. As we go into this new age where frugality is a strength, I think that the startup journey will actually be harder for the founders who are able to raise their large seed rounds so quickly at high valuations. From past experience, I've found that founders who can raise easily in a first raise really struggle later on subsequent raises because they don't know just how hard a fundraiser can be. Moreover, founders who can raise large amounts in the beginning tend to be less frugal and often burn through too much cash before their progress really kicks in. In contrast, overlooked founders who have often found it challenging to raise know that they need to be frugal by default, because it's unclear how hard the next fundraiser will be. These founders know they need to make the business work with or without investors.

The ironic twist is that investors throw money at founders with particular resumes because they believe those founders will be the most likely to succeed with big exits. A strength can quickly turn into a weakness in this market.

My hope for all founders in 2020

My hope for all founders is that they focus on staying thrifty, watch cashflow and chip away at getting to profitability so they can own their own destiny. By focusing on customers, instead of investors, you can sell more and sell quicker. Ultimately, the end goal for a company is to be able to serve customers sustainably and effect change in our larger society.

And that's what I wish all startups find in 2020, so they don't have to care about the whims and fancies of investors as they change with the times.

By Kevin Dugan

Casper Sleep's nightmare of a public market debut earlier this month has demonstrated that investor skepticism toward money-losing tech companies persisted into 2020. And that raises questions for several high-profile tech names contemplating going public this year, including two of the most highly valued, Airbnb and DoorDash.

Unless they can demonstrate a reversal of 2019 trends and turn a profit, both companies risk taking a haircut on their valuations when they hit the public market, bankers say. DoorDash was most recently valued at \$12.6 billion, while Airbnb was valued at \$31 billion.

These are among a slew of big-ticket tech initial public offerings possible this year. Postmates, Instacart, Sweetgreen, GitLab, Snowflake, JFrog, Asana and Jamf are all companies valued at or above \$1 billion that could go public, according to people familiar with their plans.

All of them have to wrestle with investor wariness of tech firms that lack a clear path to profitability. Another issue they have to navigate is a shortened IPO calendar in this election year, which means companies are likely to want to go public before the summer. Few executives want to replicate the experience of last year's poorly performing IPOs, most notably Uber, Lyft and Peloton. Those stocks are still trading below their IPO price. Concern about such an outcome could prompt some companies to delay their public debut to beyond this year.

The total number of tech IPOs is likely to be down a little from last year's 49, which in turn was down from 54 in 2018. But last year tech IPOs raised more money—\$25.7 billion, versus \$20.8 billion in 2018. The amount raised in 2020 will depend on whether companies go public through a conventional IPO or a direct listing, where they don't sell new shares—which is still a question for many of them.

Here are a few of the biggest issues hanging over the IPO market:

- **At least until this week, market conditions have been good**, according to investment bankers who are advising companies say. The economy doesn't appear to be about to fall into a recession. The Federal Reserve is likely to keep borrowing costs at or below their current levels for the rest of the year.

But as demonstrated by Monday's big sell-off—continuing so far Tuesday—conditions can change quickly, particularly as the coronavirus continues to spread. “Signs that certain segments of tech are trading at unsustainable valuations are supported by the record level of speculative call option activity and outsized gains in certain tech stocks,” Marko Kolanovic, analyst at JPMorgan Chase, wrote in a note this week. “We caution investors that this bubble will likely collapse—i.e., this time is not ‘different,’ with valuations reverting closer to 2010–2020 average.”

- **Timing is key.** Most companies will want to go public by the summer. Closer to the November election—and the possibility that voters could remove President Trump, under whom the stock market has soared 63%—the market is likely to get more volatile. That shortens the calendar for some companies to the next few months.

“There's the first half of the year; then we'll have the election,” John Chirico, Citigroup's co-head of banking, capital markets and advisory for North America, told The Information. “There are a lot of folks who, if they have something to do, are trying to get it done sooner rather than later.”

But that window might be too small for companies hoping to shore up their financial results. For Airbnb, say, whose business has taken a hit from the coronavirus, a summer public debut may not be practical. A debut in the fall would allow Airbnb to show growth through the popular late-summer travel season. It is also sponsoring the Summer Olympics in Tokyo, which shareholders hope will be a boon.

- **Investors still value profits over cash-burning growth.** One of the first tech companies to hit Wall Street this year was Casper Sleep, the money-losing direct-to-consumer mattress seller. Things have not gone well for the company. Casper slashed its offering price to \$12 from \$18 before the offering, but once shares started trading, they quickly slipped to around \$10. That puts Casper's market capitalization at about half its last private market valuation. Its experience should be a lesson for other companies contemplating going public, says one banker.

“A company that is highly capital intensive—meaning they have to spend a lot of money to generate revenue, [and] that burns cash—going public should be the last thing on their mind. The first thing on their mind should be how to stop losing money,” Joe Voboril, managing partner at merchant bank Farvahar Partners, told The Information. “More so than ever before, profit matters.”

That's an issue for DoorDash, which was projected to lose \$450 million in 2019, and potentially for Airbnb. The travel accommodations firm lost money at least through the first nine months of last year, a reversal from its performance the previous two years. In cash flow terms, however, Airbnb remained positive, say people familiar with the situation, which is what the company is likely to emphasize to investors.

Airbnb also has told investors and employees that increased investments in engineering and marketing ahead of a public listing raised expenses last year. This year the company expects to spend more on verifying listings and improving safety. It has discussed ways to offset some of those rising expenses. Late last year, Chief Financial Officer Dave Stephenson asked department heads to come up with plans that would freeze headcount growth and cut contractor positions, a person familiar with the matter said.

The question investors will have is whether Airbnb, and other companies, are showing improving trends in terms of profitability.

- **Staying private is an option.** DoorDash may opt to stay private to avoid a disappointing public debut. It could file confidentially to go public but not proceed if the market doesn't recover. While the company doesn't have oodles of cash—we reported in December it likely had between \$500 million and \$1 billion—it was considering going public via a direct listing where it wouldn't raise cash. That implied it had other means of raising cash if necessary. One possibility is that DoorDash could raise money privately before making any public debut. In that case, it could easily delay the public move.

Another company that may delay is Postmates, according to a person with direct knowledge of Postmates CEO Bastian Lehmann's thinking. The delivery company, which shelved an IPO last year, is waiting for a period when markets are less choppy, and it's still possible that the company could wait until 2021, the person said.

Airbnb, however, is unlikely to push back its public listing past this year. It has been consistently stating over the last two years that it would go public by 2020. Executives told employees two years ago that the company needed to meet that timeline before batches of restricted stock units expire after seven years. Airbnb started issuing these units in early 2014. A company spokesman declined to comment.

- **Subscription software firms remain in favor.** Zoom Video Communications, one such example, had a spectacular IPO last year. Its stock continues to trade at nearly three times its IPO price. Investors like the

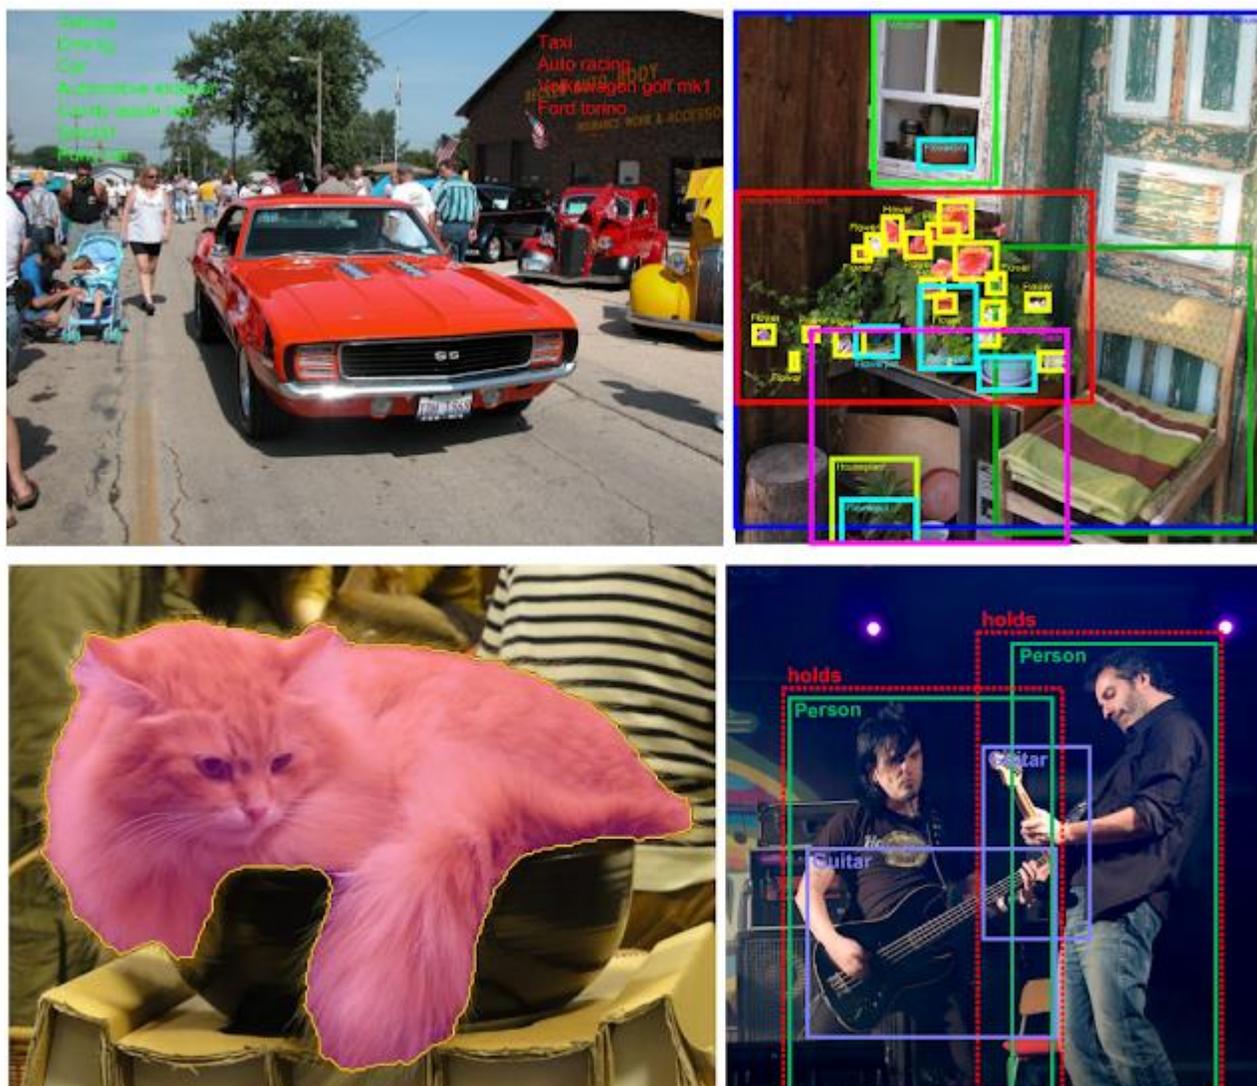
reliability of revenue from these companies. This year, companies including collaboration software firm Asana and cloud database firm Snowflake are going public.

“Zoom-ness’ is really good and ‘WeWork-ness’ is really bad,” said Rett Wallace, CEO of Triton Research.

- **Direct listings remain in vogue.** There will be about five or six direct listings this year, according to bankers who are in talks with these companies. Among those thinking about that route are Airbnb, DoorDash, and Gitlab, though those plans could still change depending on market conditions, those bankers said.

Posted by Jordi Pont-Tuset, Research Scientist, Google Research

Open Images is the largest annotated image dataset in many regards, for use in training the latest deep convolutional neural networks for computer vision tasks. With the introduction of version 5 last May, the Open Images dataset includes 9M images annotated with 36M image-level labels, 15.8M bounding boxes, 2.8M instance segmentations, and 391k visual relationships. Along with the dataset itself, the associated Open Images Challenges have spurred the latest advances in object detection, instance segmentation, and visual relationship detection.



Annotation modalities in [Open Images V5](#): image-level labels, bounding boxes, instance segmentations, and visual relationships. Image sources: [1969 Camaro RS/SS](#) by D. Miller, [the house](#) by anita kluska, [Cat Cafe Shinjuku calico](#) by Ari Helminen, and [Radiofiera - Villa Cordellina Lombardi, Montecchio Maggiore \(VI\) - agosto 2010](#) by Andrea Sartorati. All images used under [CC BY 2.0 license](#).

Today, we are happy to announce the release of Open Images V6, which greatly expands the annotation of the Open Images dataset with a large set of new visual relationships (e.g., “dog catching a flying disk”), human action

annotations (e.g., “woman jumping”), and image-level labels (e.g., “paisley”). Notably, this release also adds localized narratives, a completely new form of multimodal annotations that consist of synchronized voice, text, and mouse traces over the objects being described. In Open Images V6, these localized narratives are available for 500k of its images. Additionally, in order to facilitate comparison to previous works, we also release localized narratives annotations for the full 123k images of the COCO dataset.



0:26 / 0:46

Localized Narratives Example

Watch later

Share



In the front portion of the picture we can see a dried grass area with dried twigs. There is a woman standing wearing a light blue jeans and ash colour long sleeve length shirt. This woman is holding a black jacket in her hand. On the **other** hand she is holding a balloon which is peach in colour. on the top of the picture we can see a clear blue sky with clouds. The hair colour of the woman is brownish.

Image and Trace:

Caption:



In the front portion of the picture we can see a dried grass area with dried twigs. There is a woman standing wearing light blue jeans and ash colour long sleeve length shirt. This woman is holding a black jacket in her hand. On the other hand she is holding a balloon which is peach in colour. On the top of the picture we see a clear blue sky with clouds. The hair colour of the woman is brownish.



Sample of localized narratives. Image source: [Spring is here:-\)](#) by Kasia.

Localized Narratives

One of the motivations behind localized narratives is to study and leverage the connection between vision and language, typically done via image captioning — images paired with human-authored textual descriptions of their content. One of the limitations of image captioning, however, is the lack of visual grounding, that is, localization on the image of the words in the textual description. To mitigate that, some previous works have a-posteriori drawn the bounding boxes for the nouns present in the description. In contrast, in localized narratives, every word in the textual description is grounded.



The man at bat readies to swing at the pitch while the umpire looks on



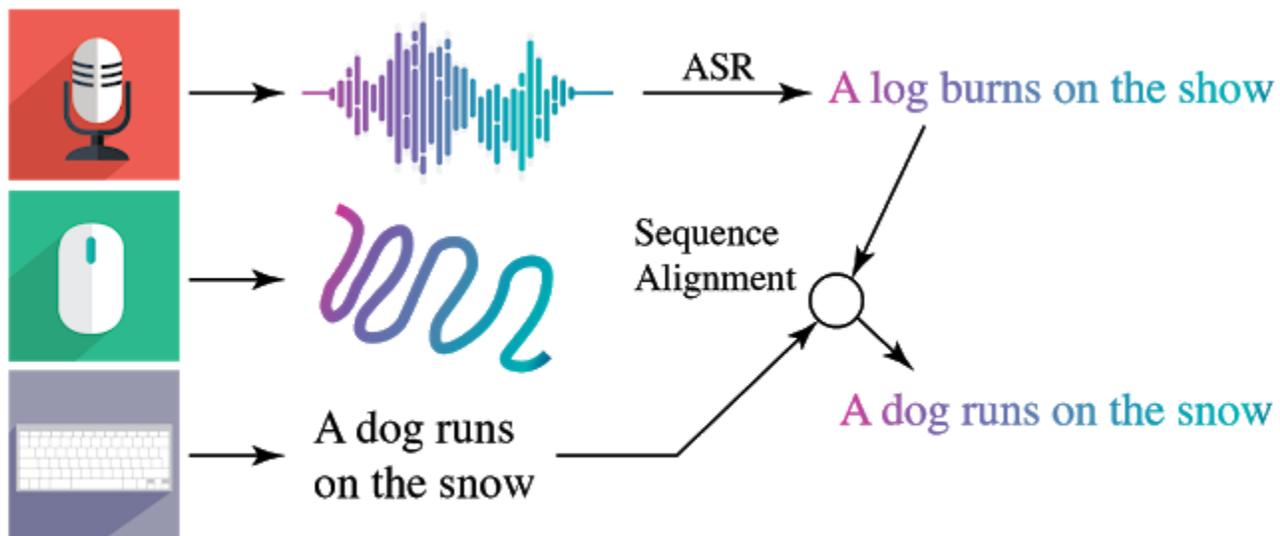
A man with pierced ears is wearing glasses and an orange hat.



A person sitting on the grass besides a plant with the basket. She wears a cap. On the background we can see many trees. And this is the sky with heavy clouds.

Different levels of grounding between image content and captioning. Left to Right: Caption to whole image (COCO); nouns to boxes (Flickr30k Entities); each word to a mouse trace segment (localized narratives). Image sources: COCO, Flickr30k Entities, and Sapa, Vietnam by Rama.

Localized narratives are generated by annotators who provide spoken descriptions of an image while they simultaneously move their mouse to hover over the regions they are describing. Voice annotation is at the core of our approach since it directly connects the description with the regions of the image it is referencing. To make the descriptions more accessible, the annotators manually transcribed their description, which was then aligned with the automatic speech transcription result. This recovers the timestamps for the description, ensuring that the three modalities (speech, text, and mouse trace) are correct and synchronized.



Alignment of manual and automatic transcriptions. Icons based on an original design from Freepik.

Speaking and pointing simultaneously are very intuitive, which allowed us to give the annotators very vague instructions about the task. This creates potential avenues of research for studying how people describe images. For

example, we observed different styles when indicating the spatial extent of an object — circling, scratching, underlining, etc. — the study of which could bring valuable insights for the design of new user interfaces.



Mouse trace segments corresponding to the words below the images. Image sources: [Via Guglielmo Marconi, Positano - Hotel Le Agavi - boat by Elliott Brown](#), [air frame by vivek jena](#), and [CL P1050512 by Virginia State Parks](#).

To get a sense of the amount of additional data these localized narratives represent, the total length of mouse traces is ~6400 km long, and if read aloud without stopping, all the narratives would take ~1.5 years to listen to!

New Visual Relationships, Human Actions, and Image-Level Annotations

In addition to the localized narratives, in Open Images V6 we increased the types of visual relationship annotations by an order of magnitude (up to 1.4k), adding for example “man riding a skateboard”, “man and woman holding hands”, and “dog catching a flying disk”.



Image sources: [IMG_5678.jpg by James Buck](#), [DSC_0494 by Quentin Meulepas](#), and [DSC06464 by sally9258](#).

People in images have been at the core of computer vision interests since its inception and understanding what those people are doing is of utmost importance for many applications. That is why Open Images V6 also includes 2.5M annotations of humans performing standalone actions, such as “jumping”, “smiling”, or “laying down”.

Finally, we also added 23.5M new human-verified image-level labels, reaching a total of 59.9M over nearly 20,000 categories.

Conclusion

Open Images V6 is a significant qualitative and quantitative step towards improving the unified annotations for image classification, object detection, visual relationship detection, and instance segmentation, and takes a novel approach in connecting vision and language with localized narratives. We hope that Open Images V6 will further stimulate progress towards genuine scene understanding.

Inside the mind of an autonomous delivery robot

By Luke Dormehl

In the summer of 2014, Ahti Heinla, one of the software engineers who helped develop Skype, started taking photos of his house.

There is nothing particularly unusual about this, of course. Only he kept on doing it. Month after month, as summer turned to fall and fall gave way to winter, Heinla went out to the same exact spot on the sidewalk and snapped new, seemingly identical pictures of his home. Was the man who had played a crucial role in building a multibillion dollar telecommunications app losing his mind? As it turned out, there was an entirely logical reason for Heinla's actions — although it might have nonetheless sounded a bit crazy to anyone who asked what he was doing. Ahti Heinla was helping future autonomous robots learn how to see.

More than half a decade later, the world (or, at least, a few select parts of it) are reaping the rewards of Heinla's seemingly oddball experiment. As the co-founder of a startup called Starship Technologies alongside Skype co-founder Janus Friis, Heinla has helped build a fleet of self-driving delivery robots. These robots, which resemble six-wheeled coolers, have traveled tens of thousands of miles all around the world, making some 100,000-plus deliveries in the process. They are particularly prevalent on a growing number of university campuses, although they have also traversed streets in cities ranging from San Francisco to Milton Keynes in the U.K.

To order something from one of Starship's delivery robots, a customer simply selects the item they want from one of Starship's delivery partners. For a small delivery fee, the robot will then pick up the item and autonomously deliver it to the spot of your choosing. All the customer has to do is unlock the robot using the app and retrieve the order. Simple, right?



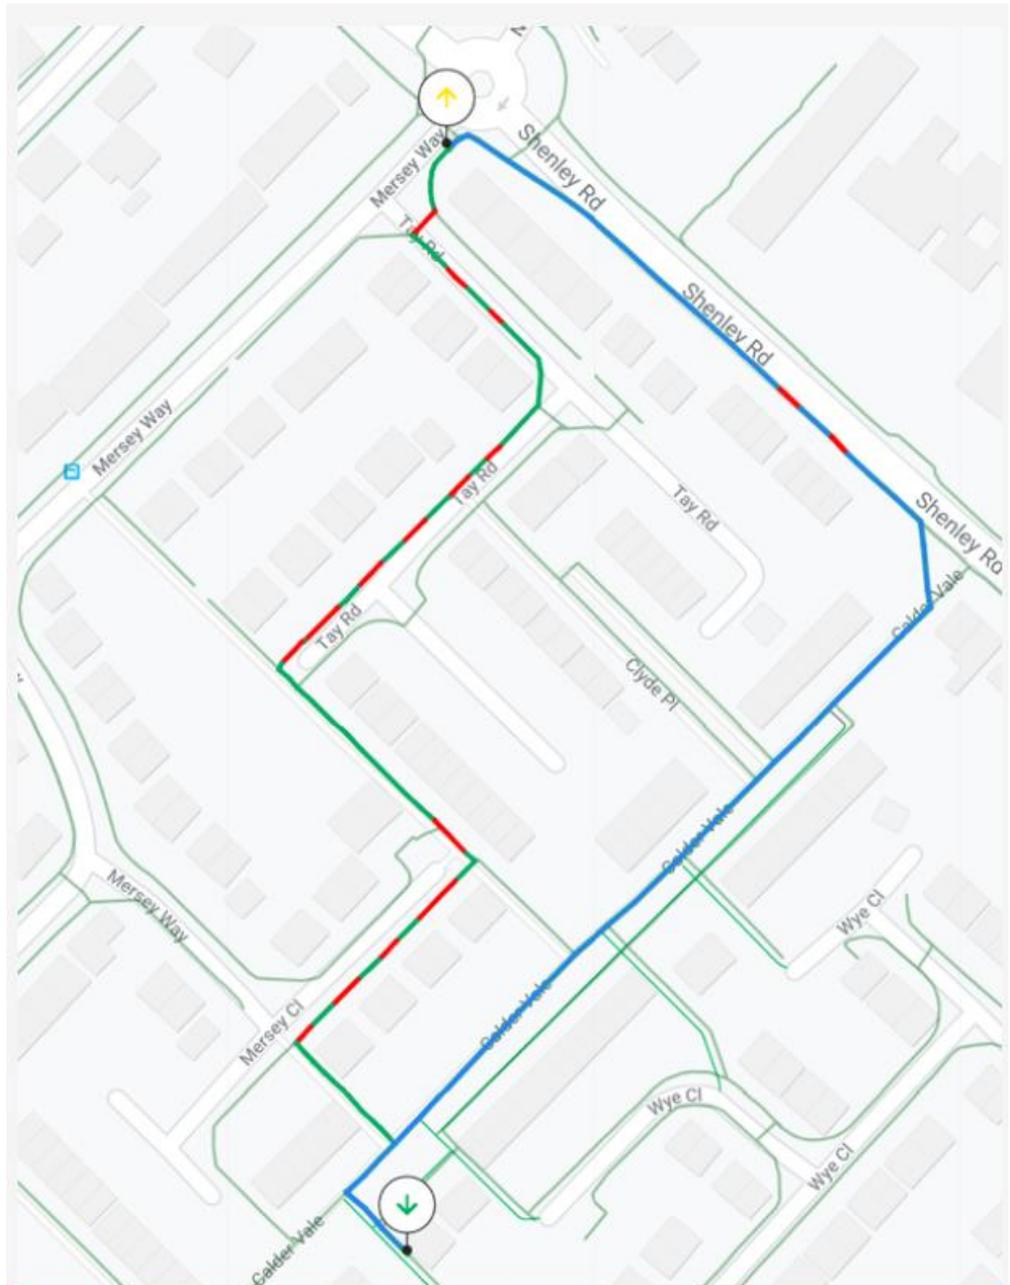
As with any such solution, however, the simpler things seem from the user perspective, the more complex they are technologically. Here in 2020, we're used to hearing about self-driving cars that are able to navigate through the world with impressive levels of ease. As one of the first companies to roll out self-driving vehicles without human safety drivers, Starship Technologies has helped play a key role in making autonomous technologies such as this a part of everyday life.

We shouldn't take these tools for granted, though. Not only are they amazing feats of engineering and computer science, but the choices currently being made surrounding these technologies will help determine the future of human and robot interactions.

Maps aren't built for robots

Do you remember the overwhelming feeling of starting a new school as a kid and having to navigate your way around? Perhaps, if you lived close by, you even walked from home to school on your own or with friends. Normally, these trips were preceded by ones on which we were accompanied by a parent or guardian who's able to give us tips about how to navigate the world around us. They might walk with us the first few times to ensure that we are familiar with a certain path. They will probably point out certain landmarks, such as signs or particularly memorable buildings. Before long, we form a mental map of where we are going and how to navigate there.

This ability, which most of us take for granted, is what Starship Technologies has worked hard to develop for its robots. In some ways, it is a surprisingly complicated one. Take maps, for instance. When Starship's robots set out to navigate from point A to point B, they start by using satellite imagery to help them plan out the journey. A routing algorithm is then used to figure out the shortest and safest path for the robot to take. So far, so simple, right? Except that it isn't.



The shortest route (green) is not always the fastest and safest. The robot will prefer the route that is longer in distance, but faster and safer

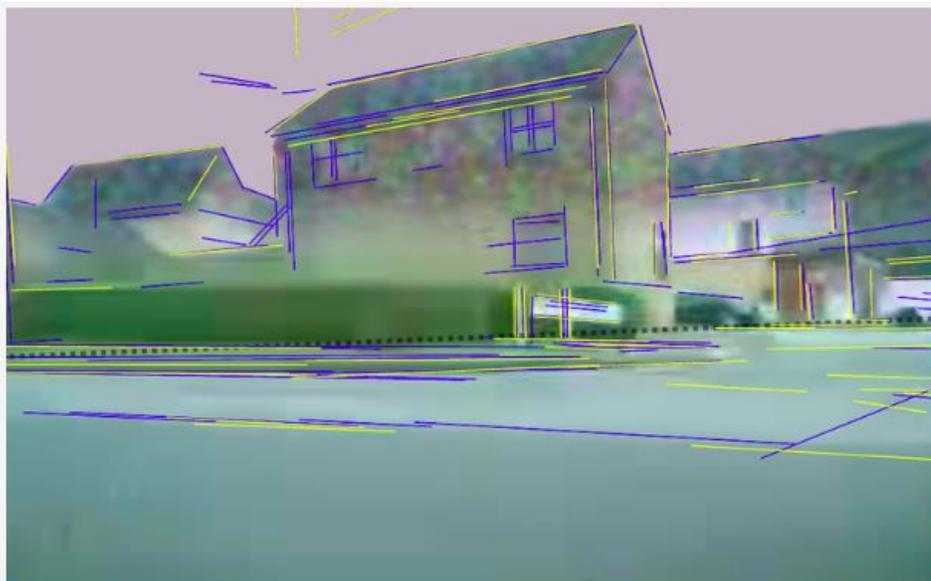
As Heinla says: “We can’t use a lot of existing maps because they are not really made for robots; they are made for humans.” Existing mapping systems assume a level of human knowledge, such as an understanding of which part of the road we should walk on, and how we should maneuver on a busy sidewalk. These are all things a robot doesn’t necessarily understand. There are plenty of added complexities.

For example, think about how your behavior while walking across a driveway differs from a regular sidewalk. We might not think of them as being especially different, but they are. If one of Starship’s robots encounters an obstacle on the sidewalk, its response is to stop in its tracks. That’s because stopping is the safest thing to do. But stopping on a driveway, or while crossing a street, blocks access for vehicles. It requires learning a totally different type of behavior.

To help understand the kind of behavior its robots should use, Starship has developed machine learning tools that can segment maps into a series of interconnected colored lines representing sidewalks (in green), crossings (in red), and driveways (purple). Rather than simply selecting the shortest route in terms of distance, the robot determines the quickest route by attaching a cost to every scenario that the robot will encounter over the course of a journey.

Recognizing the world around them

After this, Starship’s robots head out into the real world, using a bevy of 10 cameras to identify the 360-degree world around them through observation. Special image-recognition systems divide the world up into thousands of lines, giving it a simplified wireframe view of the world to use as guideposts. Over time, as the company’s robots spend longer in one area, they can build up collaborative three-dimensional wireframe maps of entire areas, making it far easier for future robots to understand the scenery around them.



“It’s just like the way you might direct a person: continue until you hit a yellow building, then turn right and continue until the church,” Heinla said. “The robot also has landmarks, but they’re not yellow buildings or churches; they’re abstract shapes.”

The last stage of the robots’ mapping process is to work out exactly how wide and where the sidewalk is. This is done by using both its onboard cameras and its 2D map taken from satellite imagery.

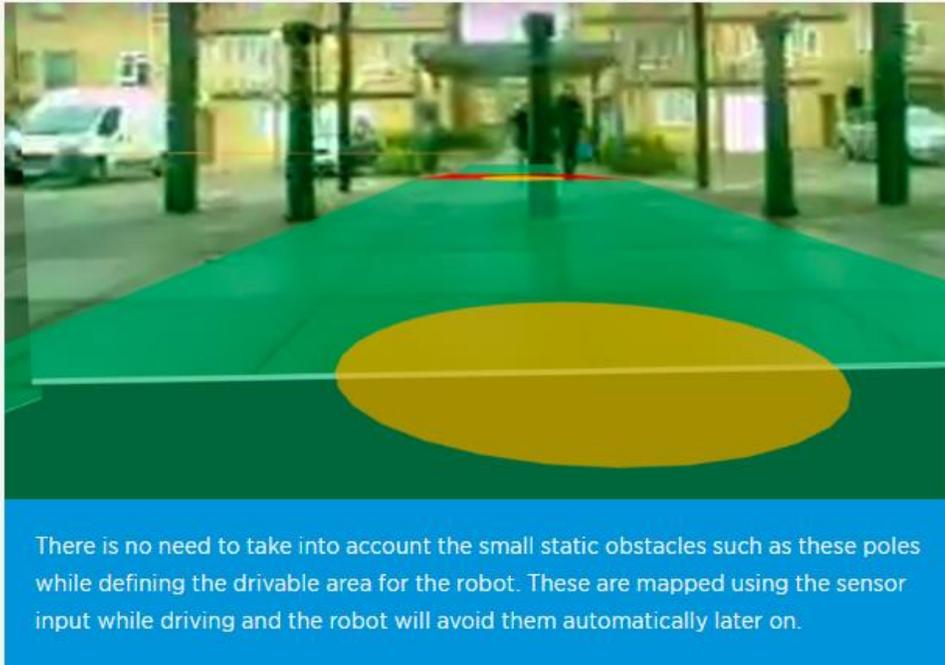
“Even something as simple as walking down the sidewalk is something we’ve learned from the time that we were very young,” Heinla said. “We take it for granted. But for machines, it’s something that needs to be taught. There are

things like whether you pass an approaching person on the left or the right. If someone slower than you is walking ahead, do you slow down or pass them? If you slow down, how close should you get to the other person? If you get too close, the other person will get uncomfortable. All of these we have to teach to the machine.”

Should all go to plan (and, to date, it has), Starship's robots will be able to navigate to the destination users select on the map.

How do we want robots to interact with people?

This isn't a challenge that's unique to Starship Technologies. A number of other companies, ranging from Nuro to BoxBot, are exploring their own self-driving robot-delivery services. But it goes far beyond robots that can bring us takeout or groceries when we're too busy (or lazy) to go to the shops. As robots play a bigger role in our lives, the question of how to integrate them within our world is becoming more pressing.



Robots have traditionally performed very well in lab conditions where every variable can be perfectly controlled. They have also been largely separated from people for safety reasons. Now they are moving into the real world in a big way. If we're not used to the sight of robots on our streets now, we sure will by the time the 2020s comes to an end.

“Every week in our autonomous driving team, we have a meeting where, for one hour, our safety team shows the autonomous driving engineers some of the most interesting things that have happened during the last [seven days],” Heinla said. “These

interesting things are either places where there has been some discomfort, the robot has done exceptionally well driving, or [where there have been] some unusual weather conditions or objects.”

Some of these t problems involve robots being able to comprehend our world. That is what Heinla was testing when he took photos outside his house in the early days of Starship Technologies. He wanted to know whether a robot would be able to recognize his house as, well, his house, regardless of whether it was a sunny summer's day or a rainy winter's evening. It turns out that it could — and that insight helped spawn a whole company (or maybe even an entire delivery industry).

Research such as this — part engineering, part sociology — is all about finding answers to how humans and machines can better coexist. Is it worse for a robot to be overly cautious or too reckless? What happens when delivery robots encounter guide dogs? Data from this new field of research is being gathered and used to tweak the algorithms that power robots made by companies such as Starship Technologies.

One day, we'll thank them for it. For now, though, it's just important that we understand the decisions they make — and the reasons they make them.

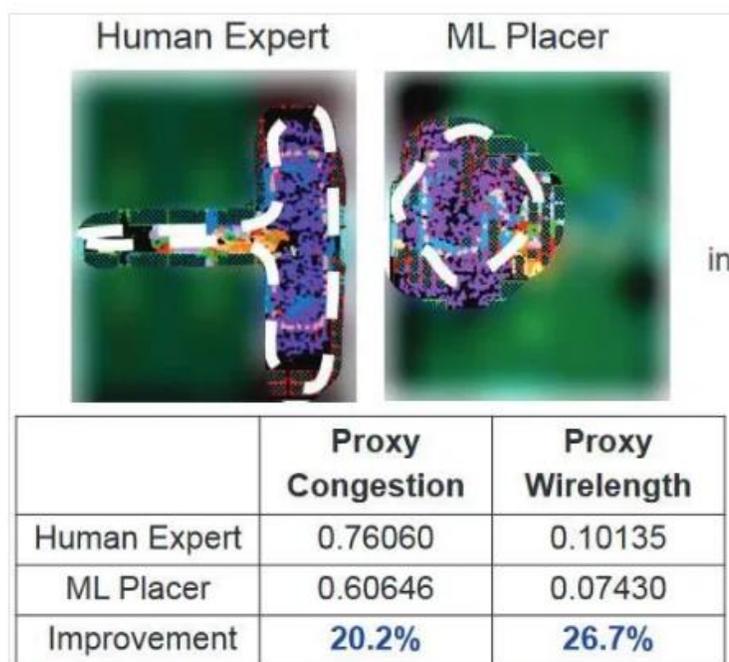
By Brian Santo

Google is experimenting with machine learning (ML) to perform place-and-route in IC design and is getting great results. The revelation, announced last week at the ISSCC conference held here, is as important for artificial intelligence (AI) as it is for circuit design.

AI has been the most massive thing in the electronics sector for years, pulling an extraordinary amount of semiconductor research in its direction (along with venture capital and headlines). Acknowledging the obvious, the theme of this year's Integrated Solid-State Circuits Conference (ISSCC) was "Integrated Circuits Powering the AI Era," and the opening plenary session was constructed to map the extent to which AI has warped semiconductor space.

The four plenary speakers explained how the requirements of AI are, for example, driving a new category of processors architected specifically for AI applications (alongside CPUs and GPUs); spurring innovations in structure (e.g., chiplets, multichip packages, interposers); and are even influencing the development of quantum computing.

The plenary session's first speaker was Jeff Dean, the lead at Google AI. Dean delivered an update of the overview of machine learning (ML) that he's been presenting in one form or another for more than a year to lead into the discussion of the ML place-and-route tool.



Results of a human expert at placing and routing an ASIC design versus the results from a low power ML accelerator chip. Google deliberately obscured parts of the images. (Source Google Research / ISSCC)

He started with a quick overview of the history of AI and ML, starting with machines that learned how to play backgammon in 1995, and running through machines that learned to excel at chess, and then at go, and can now negotiate complex video games such as StarCraft "with remarkable success." ML is also being used in medical imaging, robotics, computer vision, self-driving vehicles, neuroscience (analyzing microscopy of brain scans), agriculture, weather forecasting, and more.

The basic idea that drove computing for decades is that the bigger the problem, the more processing power you throw at it, and the more processing power you have, the bigger the problems you can solve. For a while, that applied to problem-solving with AI.

Where that broke down was when problem spaces got so mind-bogglingly vast it was simply not possible to amass enough CPUs (and/or GPUs) to solve them.

It turned out that AI/ML doesn't need typical CPU/GPU power, however. The math required can be simpler and requires much less precision. That realization had practical ramifications: processors dedicated to AI/ML don't have to be as complex as CPUs/GPUs.

That’s one of the fundamental insights that led to specialized processors designed for inference, such as Google’s own TensorFlow processors, now in their third generation. Google, by the way, is commonly expected to come out with a fourth generation TensorFlow one of these days, but if anyone had hoped Google would reveal anything about it at ISSCC, those hopes were dashed.

The realization that less precision is necessary for inference was followed by the realization that less precision is needed for training as well — that is relatively new. EE Times editor Sally Ward-Foxton explained the concept in her recent blog *Artificial Intelligence Gets Its Own System of Numbers*.

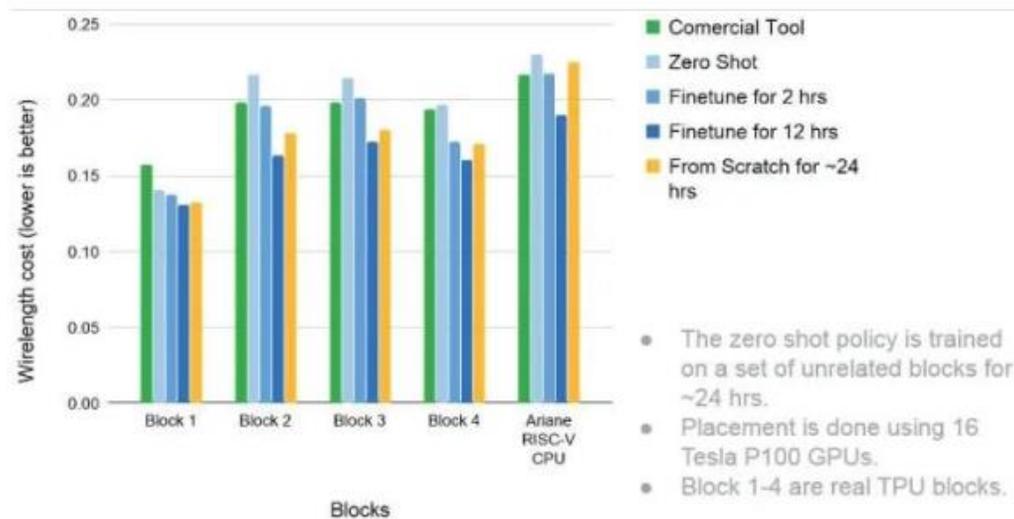
AI/ML processors can be relatively simple, and therefore relatively cheaper, and we now have AI/ML processors that are powerful enough to train pretty rapidly, even on enormous data sets. All of that is making it easier to push machine learning farther out into the network edge, Dean explained. A specific example is speech recognition; Dean said that as of 2019 Google has had a pretty compact model that works on smartphones.

Each AI application — autonomous driving, medical imaging, playing go — results from tweaking a dedicated AI/ML system to learn each. We have basically one AI per application. The next question was: is it possible to take an AI that learned one thing, and then see if it can apply what it’s learned to some other task that is similar?

“I brought this up because we began thinking about using this for place-and-route in ASIC design,” Dean said. “The game of place-and-route is far bigger than the game of go. The problem size is larger, though there isn’t as clear goal as there is with go.”

Google created a learning model for place-and-route, and then set out to find if the tool could generalize. Could it take what it learned on one design and apply it to a new design it had never seen before? The answer was an unambiguous “yes.”

Furthermore, Dean said, “We’ve gotten super-human results on all the blocks we’ve tried so far. It does a little bit better, and sometimes significantly better than humans.”



Google compared the results in the performance of an AI that used machine learning (ML) to teach itself to place and route components of an ASIC. The test circuits were several different blocks, including an Ariane RISC-V CPU. Google compared the performance of the same ML after progressive intervals of additional tuning, all against the performance of a commercial tool. (Source: Google Research / ISSCC)

“Better” includes performing place-and-route in extraordinarily less time. It might take a human expert weeks and weeks to accomplish the task. An ML placer typically does the same job in 24 hours, and its layouts typically have shorter wirelengths, Dean reported. The ML placer also did well against automated place-and-route tools. (Read more about ML and place-and-route in “Machine learning in EDA accelerates the design cycle,” written by Cadence’s Rod Metcalfe, in EE Times’ sister publication EDN.)

ML might also be extended to other parts of the IC design process, Dean said, including using ML to help generate test cases that more fully exercise state space in ASIC design verification, and perhaps also using ML to improve high-level synthesis to get to more optimized designs from high-level descriptions.

What all this means for ML, however, is as important as what it means for accelerating IC design schedules. If an ML can generalize within a category, can it generalize to perform tasks in other categories?

“What might future ML models look like?” Dean asked. “Can we train one model to generalize to similar tasks? Ideally we’d want one model that can learn to do thousands or millions of tasks.”

The artificial intelligence Internet of things (AIoT)

Kou-Hung “Lawrence” Loh, the senior vice president and chief strategy officer at MediaTek spoke of how AI is transforming just about everything connected to the Internet, and that the AI Internet of things (or AIoT) will rapidly expand from the tens of billions of devices today to encompass an estimated 350 billion devices worldwide by 2030.

AI is moving toward the edge in part because it can (as Dean mentioned earlier in the session) and because in many cases it has to, for several reasons including alleviating the growing processing burden on data centers, minimizing the traffic on networks, and because some applications require, or will work best, with local processing.

Local processing will have to be fast, it will have to be designed specifically for AI computation, and it will have to be extremely energy-efficient.

These are by nature a new category of processor. Loh called them AI processor units (APU). Others have referred to them variously as neural processing units (NPU), brain processing units (BPU), and other names. An APU might be less flexible than a CPU, for example, but by virtue of being purpose-built, APUs can be as much as 20 times faster at 55 times less power, he said.

Loh said that APU developers are working on devices that will reach 1 TOPS at 3 TOPS/Watt. He said he believes 10 TOPS at 10 TOPS/W is achievable. It might eventually be possible to get to 100 TOPS at 30 TOPS/W, he said.

Not coincidentally, MediaTek researchers presented at ISSCC a separate paper proposing a “3.4 to 13.3TOPS/W 3.6 TOPS Dual-Core Deep Learning Accelerator for Versatile AI Applications in a 7nm 5G Smartphone SoC.”

That’s at 7nm. Performance improvements will be gained by racing along the curve of Moore’s Law to smaller process nodes for at least one more step, from the present 7nm to 5nm. Moore’s Law still applies, Loh said.

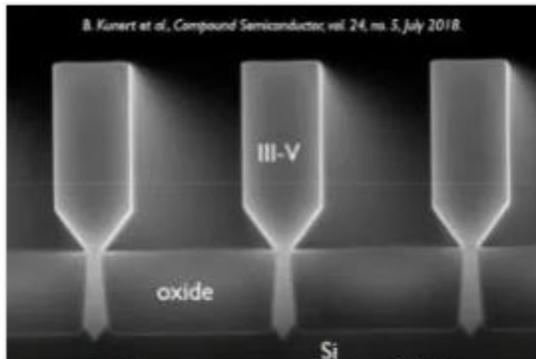
Not without caveats, however. Transistor counts are increasing with integration, continuing to follow the classic Moore’s Law curve, “but the cost per transistor is not following,” Loh said. Furthermore, due to the complexity of chip design, and because process steps are getting more complicated, costs for leading-edge devices are soaring, prohibiting smaller companies from using the technology. There are also yield issues.

A common solution to many of these problems is splitting the die, Loh said. As a practical matter, that might mean using approaches such as chiplet technology. “It can lead to doing better than Moore’s Law,” he said. Whether it’s chiplets or some other architectural approach, it all means more challenges in interconnect.

System technology “co-optimization”

Nadine Collaert, program director at Imec, brought the plenary theme forward the next step, going over the need to separate die and figure out alternative structures and architectures for integrated circuits in the future. She called it system technology co-optimization, or STOC.

Moore’s Law is likely to pertain for years to come, but scaling CMOS is getting more challenging, she said. She illustrated the point with a series of examples of ever-more-complicated device structures, including (but hardly limited to) FinFETs, nanosheets, and forksheets, that can indeed be used to achieve further CMOS scaling at the chip level.



Imec demonstrated the ability to grow an unspecified III-V material on a silicon-on-insulator (SOI) substrate in a nano-ridge formation. (Source: Imec / ISSCC)

Something eventually has to give, however, she explained. A new approach is needed and “we believe 3D technologies are the best way. That includes multi-die packages, using bonding or, even at the device-level, fine-grade connections with other standard cells.”

Figuring out which technology to use will require matching system requirements against the properties of the options available. “That’s going to be a complicated exercise,” Collaert said. That is going to put pressure on EDA vendors to provide tools that will enable designers to weigh their options.

Front-end modules for wireless communications systems are going to be a particular challenge. “Generally these are the most diverse systems — they have many different components with different technologies, and that complexity will increase with more antennas, more PAs, more filters...”

The industry is moving to higher frequencies and higher efficiency. One option is combining III-V materials (e.g. GaN and SiC) with CMOS to get the benefits of both materials. That can be done with 3D integration, she said, showing several examples including an image of a 3D nano-ridge with a III-V material grown on a silicon-on-insulator (SOI) substrate, “but there’s a lot of work that needs to go into enabling this.”

As for memories? “New apps like AI and ML are driving the roadmap,” Collaert said. They need fast-access memories. “There’s a push to look at compute in memory, and as you bring logic and memory closer, 3D packaging is of course very important.”

Moving forward, using flash in advanced applications will be about stacking more tiers, she said. There’s also a desire to improve channel current in these memories. “To do that, we have to look at channel mobility, and again, that means looking at III-V materials.” And by extension looking at 3D architectures that stack a layer of silicon with a layer of a III-V material.

Meanwhile, in DRAMs, capacitors are growing from squat cylinders to pillars — yet another shift in the third dimension. Other memory options include magnetic memory for cache replacement, and 3D storage class memory — Collaert noted that Imec has demonstrated a vertical FeFET (ferroelectric field effect transistor) that still needs more research.

The development of all of these memories, she said, “is all in the context of machine learning. AI is booming. A lot of this is in the cloud, but for various reasons we want to move it to the edge, where there will be constraints on energy.”

Imec is more optimistic than MediaTek, in that it believes it might be possible to get to 10,000 TOPS/W. “Scaling continues. The party is not over!” she concluded. “New memories might not make it into the roadmap, but they may have applications in machine learning.”

Quantum computing

Dario Gil, director of IBM Research, wrapped the plenary by addressing “what’s next,” which he said is generalized AI, which will be almost certainly achieved on quantum computers. That said, the key thrust of his talk is that the greatest benefits will probably be derived from the complementary use of bits (digital processing), neurons (AI) and qubits (quantum computing).

He noted that IBM opened access to its first quantum computer through the cloud in 2016, and that it now has access to 15 quantum computers available, including its latest 53-qubit model.

It's not the who-is-winning-the-race narrative that many might expect

By Kirsten Korosec

The California Department of Motor Vehicles released its annual cache of autonomous vehicle testing and disengagements data that, depending how one chooses to interpret the data, shows either stunning progress or stagnation.

The data, which every company testing autonomous vehicles on public roads in California must submit, tells a winding and sometimes contradictory tale of growth, consolidation and priorities. The total number of autonomous miles driven in 2019 rose 40%, to more than 2.87 million, thanks largely to a notable uptick in public on-road testing by Baidu, Cruise, Pony.ai, Waymo and Zoox, as well as newcomer Lyft.

And yet, the rise in total autonomous miles and permitted companies don't tell the whole story. While the number of companies with testing permits grew to 60 in 2019, the percentage of companies actually testing on public roads fell to about 58%. In 2018, about 62% of the 48 companies that held permits tested on public roads.

Some companies scaled back public testing in California, either to move operations out of state or prioritize simulation. Aurora, for instance, saw its total on-road autonomous testing drop 59%, to 13,429 miles. Meanwhile, Aurora ramped up its simulation efforts, conducting more than 735,000 tests per day, an increase of over 100 times from 2018.

“While on-road testing is useful for collecting targeted data and performing late stage validation of self driving systems, we find that large-scale, on-road autonomous testing is a slow, and inefficient approach to development relative to more sophisticated, virtual techniques,” Aurora co-founder and CEO Chris Urmson wrote to the DMV.

Others, like Drive.ai, no longer exist. Two companies, Roadstar.ai and Ximotors.ai, failed to submit a disengagement report and have had their testing permits revoked.

The upshot: It's not the who-is-winning-the-race narrative that many might expect or try to tell. Those kinds of rankings and comparisons are nearly impossible, for a number of reasons, including the fact that testing on public roads is conducted in areas with varying degrees of complexity. Additionally, companies aren't required to report testing on private roads or tracks, out of state or in simulation, all of which provides a better assessment of an AV developer's technology.

But the biggest issue is how companies interpret “disengagements,” a term that describes each time a self-driving vehicle disengages out of autonomous mode either because its technology failed or a human safety driver took manual control for safety reasons. Companies not only have different views of what qualifies as a disengagement, but that interpretation can change over time.

The DMV contends that these reports are not intended to compare one company with another or reach broad conclusions on technological capabilities. Instead, the agency told TechCrunch that it uses the reports for public awareness.

“From the reports we can see that as a whole, autonomous miles driven continue to increase annually, as do the number of permit holders, test vehicles and safety drivers,” a DMV spokesman wrote in an email.

Now, industry grumbling over these disengagement reports is moving from behind-closed-doors lobbying to public commentary on social media and other forums.

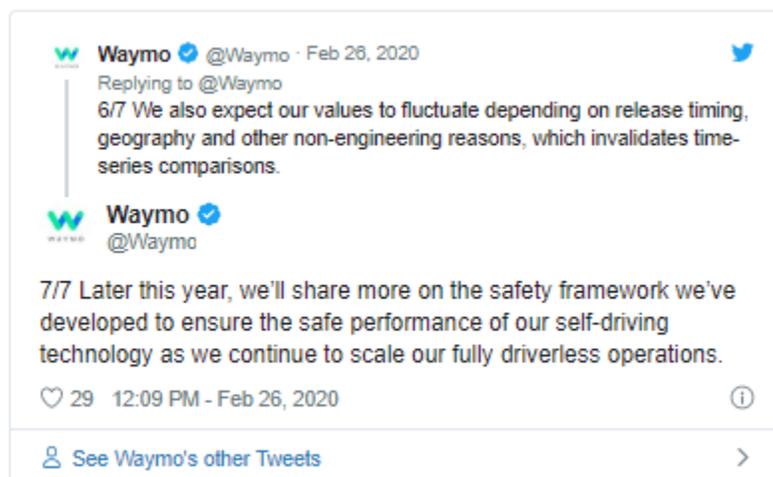
This year, a growing number of companies, including Aurora, Cruise and Waymo, issued public statements that DMV disengagement reports don't provide relevant insights into performance and are a poor way of measuring progress or competency.

Moments after the DMV released the disengagement reports, Waymo took to Twitter to log its concerns, noting that the report doesn't "provide relevant insights into the capabilities of the Waymo Driver or distinguish its performance from others in the self-driving space." Waymo also noted that most of its public road testing is outside of California, in markets like Detroit and Phoenix. The "real-world driving" that Waymo does conduct in California is "predominately engineering development, and not production releases."

Waymo's public criticism marks a shift within the company. In previous years, Waymo has celebrated its progress in glossy reports. This year, the company has become a vocal critic, even as this latest report shows a year-over-year improvement in its disengagement rate as it increased its total number of miles. Waymo drove 1.45 million miles in autonomous mode in 2019, a 200,000-mile increase from the previous year, while its disengagement rate dropped from 0.09 to 0.076 per 1,000 self-driven miles.

Other companies, as well as analysts and industry watchers, echoed Waymo's sentiments. Several weeks ago, Cruise co-founder and CTO Kyle Vogt published a blog post that argued these disengagement reports should not be a proxy for the commercial readiness or safety of self-driving cars.

This airing of grievances did not produce an alternative metric that would accurately measure competency, readiness and progress. Waymo did say in its series of tweets that it is preparing to share more on a safety framework it has developed. Vogt's post also suggests that Cruise is working on a more comprehensive metric.



The reports have their shortfalls. However, they're often the only window into a company's autonomous vehicle program. Comparisons between companies might be ineffective, but examining multiple years of data from one AV developer can be helpful in connecting the dots on a business strategy or an imminent demise.

Take Cruise as an example. The company has amassed a \$7.25 billion war chest, and a chunk of that capital is being poured into putting more vehicles on the road for longer periods of time. Cruise reported 228 registered autonomous vehicles in 2019, a 40% increase since the previous year. Over that same time period, Cruise's total mileage has increased by more than 85%.

Or take a look at Pony.ai. The company, which announced earlier this week that it has raised \$400 million from Toyota Motor Corporation, reported 22 registered autonomous vehicles in 2019, three times more than the previous year. The startup reported 16,356 total AV miles in 2018. That figure skyrocketed to nearly 175,000 miles in 2019.

Despite all of the data — flaws and all — that these reports provide, they get no closer to revealing what metric companies use internally to determine progress and competency and answer the critical question of how safe is safe enough?

By Priya Anand and Kevin McLaughlin

Amazon is angling to become a toll collector on physical retail just as it is in e-commerce, with its vast marketplace that outside merchants access for a fee. But brick-and-mortar retailers may want to approach Amazon's latest idea with caution.

As Amazon on Tuesday unveiled in Seattle its biggest Amazon Go outlet—grocery stores equipped with technology that eliminates the need for cashiers—The Information confirmed a Wall Street Journal report that the e-commerce giant was considering licensing the Go technology to other retailers. That could extend its tentacles throughout the brick-and-mortar landscape. Such a move could elevate the Go technology from its current experimental status into a money-making business.

For retailers, though, striking a deal with Amazon would align them with a company that many of them view as an existential threat. If history is any guide, they should run for the hills. In the early days of online shopping, some retailers, such as Toys“R”Us, Borders and Target, cut deals with Amazon to run their e-commerce operations—and it didn't end well for them.

And that presents retailers with a conundrum. Turning down Amazon's offer may make it harder for them to launch stores without cashiers. That could put them at a competitive disadvantage. Eliminating cashiers could both reduce their costs and improve the experience for shoppers who don't like waiting in line to check out.

And there are alternative sources of the technology that retailers could turn to. For example, San Francisco-based startup Standard Cognition has been demonstrating its cashier-less tech, using cameras and artificial intelligence, in a convenience store in downtown San Francisco for the past year and a half. It is also testing its technology in three U.S. stores, though a company spokeswoman would not say which retailer it is working with or where those stores are located.

Other contenders include Berkeley, California-based Grabango, which has at least three pilot stores—one in Northern California, one in Pittsburgh, and another in a location the company would not disclose. Grabango's technology is much simpler than Amazon's, said Andy Radlow, the company's chief business officer. It deploys cameras, shaped like fluorescent light fixtures, which use computer vision to identify what shoppers are placing in their baskets. He says that makes it much easier to retrofit an existing store with the Grabango technology, compared to the more extensive hardware Amazon relies on.

“The vast majority of retailers consider Amazon a fierce competitor,” Radlow said. “I haven't had to exert a lot of salespersonship to convince retailers to work with us as they prepare for higher levels of competition from Amazon and others.”

Radlow said Grabango has four clients, including Giant Eagle, a regional supermarket chain.

There's also San Francisco-based Zippin, which has partnered with Brazilian retail chain Lojas Americanas to put its checkout-free tech in stores, opening its first such outlet in Rio last year.

Meanwhile, Microsoft reportedly is developing similar technology, although it has yet to announce anything. Microsoft has openly courted partnerships with retailers by positioning itself as a safe alternative to Amazon. It has struck deals to sell its cloud services and other software applications to Walmart, Albertsons and Kroger. Microsoft has sought to demonstrate that major retailers are flocking to its offerings to run their retail stores.

Far Ahead

Amazon, though, looks to be way ahead of its rivals—at least in having its technology being used by consumers. It now has about two dozen Go stores in four U.S. cities: San Francisco, Seattle, Chicago and New York. Those stores prove that its custom software and hardware works. Amazon has been working on the Go tech licensing plan for at least six months, according to one person familiar with the matter. (The Wall Street Journal earlier reported on the plan.)

With its long lead time, Amazon believes it can price its tech low enough to overcome competitive concerns, say two people familiar with the matter. Its Go technology includes hardware such as custom-built video cameras that identify which items shoppers remove from shelves in Go stores and in what quantity. It also includes roughly two dozen different types of sensors, including ones that detect the movement of shoppers within stores and others that track the weight of store shelves to determine when it's time to restock.

Key to Amazon's technology is an Amazon-designed device called a networking switch, which coordinates the flow of data collected by shelves and sensors in Amazon Go stores, according to one of the people with direct knowledge. Some of the data, such as sensor data that lets Amazon Go staff know when they should restock shelves, is analyzed in the stores. Other data, such as that used to improve AI algorithms, goes to the cloud for analysis, said the person.

Amazon has also developed AI algorithms to more accurately account for unusual situations, such as when a shopper picks up five items from a shelf and then puts back four, for example. The algorithms ensure that the customer, in that scenario, is billed for the one product they held onto, according to the person with direct knowledge.

Sorry History

But retailers should be wary of giving the reins for their stores to Amazon. Consider the fate of Toys“R”Us, Target and Borders, which gave Amazon management of their online stores in the early 2000s, before online commerce took off. Borders ended up taking back its website in 2008, three years before it filed for bankruptcy. By giving Amazon control of its website, it had allowed the then-small firm to lock in its stranglehold on e-commerce.

Toys“R”Us at one point gave up its own website, directing users to Amazon instead, in exchange for being the exclusive retailer of toys on Amazon. Later, though, Toys“R”Us noticed other merchants selling toys on Amazon. The dispute led to a court battle between the two companies. And Target eventually ended its partnership with Amazon as well.

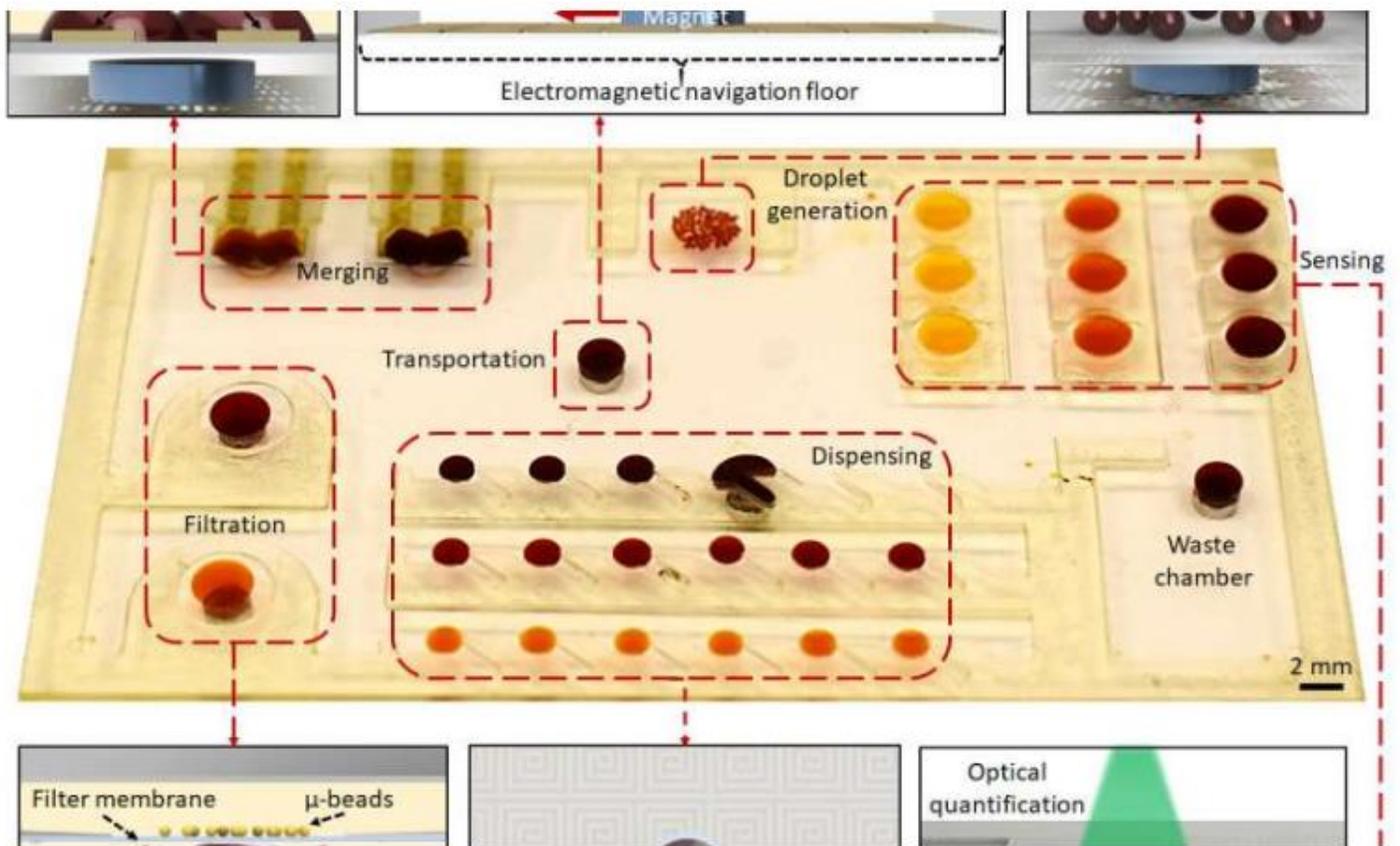
Retailers who license Amazon Go technology could find themselves forking over data on foot traffic, sales and other shopper behavior. One could make the case that giving Amazon access to such data would improve the technology for all the retailers licensing it. That's because the data helps train the AI algorithms to more accurately measure consumer behavior. But it's not clear whether retailers want to inadvertently help any rivals also using the Go technology by sharing their data.

An Amazon spokesperson said in a statement: “We have not commented on the rumors about selling Amazon Go's technology, but I can tell you that all Amazon third-party technology services follow long-standing company policies to protect customer data and its use, and employ a multi-faceted approach to enforce those policies.”

Companies that partner with Amazon often find that Amazon goes into competition with them. Antitrust regulators in Europe have examined whether Amazon uses data about sellers on its platform to undermine them, given that the company sells on its own platform as well. Cybersecurity companies that have initially benefited from promotion by Amazon Web Services have later found themselves competing with AWS's own cybersecurity products.

For retailers willing to take the Amazon Go plunge, no one can say they weren't warned.

By University of California, Los Angeles



Overview of microfluidic system concept and mechanism. Credit: Yu et al., *Sci. Robot.* 5, eaba4411 (2020)

UCLA engineers have developed minuscule warehouse logistics robots that could help expedite and automate medical diagnostic technologies and other applications that move and manipulate tiny drops of fluid. The study was published in *Science Robotics*.

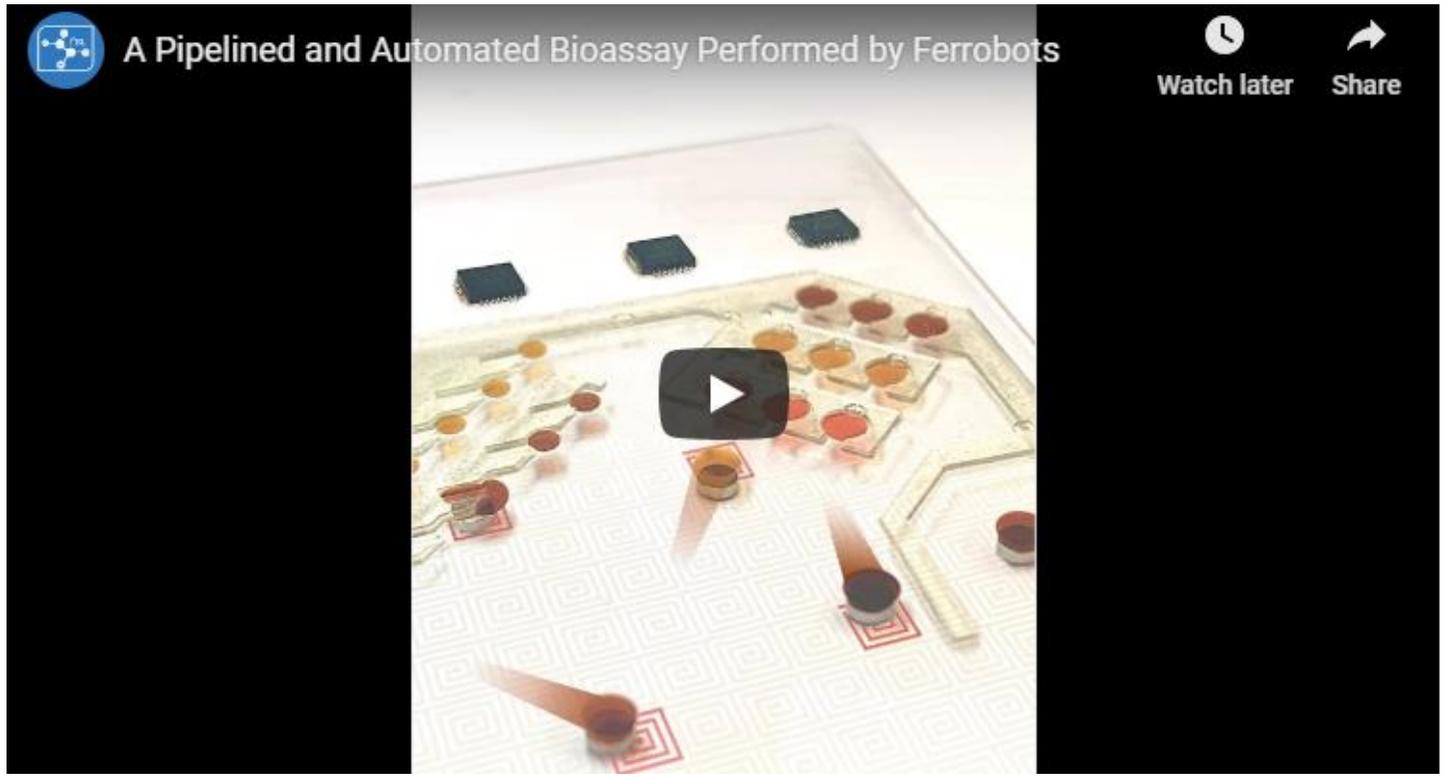
The robots are disc-shaped magnets about 2 millimeters in diameter, designed to work together to move and manipulate droplets of blood or other fluids, with precision. For example, the robots can cleave one large droplet of fluid into smaller drops that are equal in volume for consistent testing. They can also move droplets into preloaded testing trays to check for signs of disease. The research team calls these robots "ferrobots" because they are powered by magnetism.

The ferrobots can be programmed to perform massively parallelized and sequential fluidic operations at small-length scales in a collaborative manner. To control the robots' motion, electromagnetic tiles in the chip pull the ferrobots along desired paths, much like using magnets to move metal chess pieces from underneath a chess board.

"We were inspired by the transformational impact of networked mobile robot systems on manufacturing, storage and distribution industries, such as those used to efficiently sort and transport packages at Amazon warehouses," said Sam Emaminejad, an assistant professor of electrical and computer engineering and the study's corresponding senior author. "So, we set out to implement the same level of automation and mobility in a microfluidic setting. But

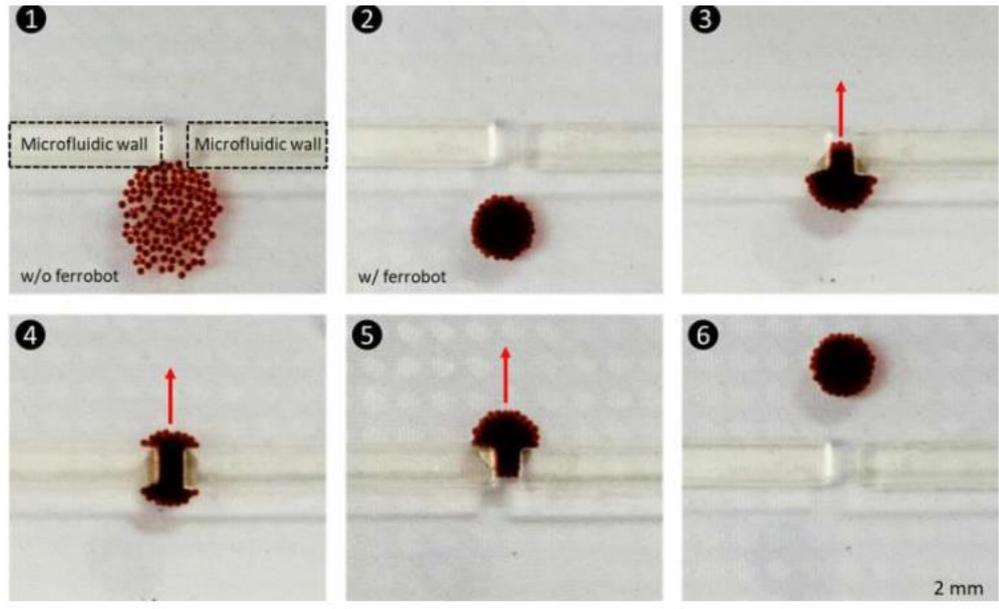
our 'factory floor' is much smaller, about the size of your palm, and our goods, the fluid droplets, are as small as a few tenths of a millimeter."

The "factory floor" is an index card-sized chip, designed by the researchers, with internal structures that help manipulate fluid droplets transported by the robots, as demonstrated in this video:



"In the same way that mobile and cross-collaborative Amazon robots transformed the logistics-based industries, our technology could transform various biotech-related industries, including medical diagnostics, drug development, genomics, and the synthesis of chemicals and materials," said study co-corresponding and senior author Dino Di Carlo, UCLA's Armond and Elena Hairapetian Professor in Engineering and Medicine. "These fields have traditionally used refrigerator-sized 'liquid-handling' robots. Using our much smaller ferrobots, we have the potential to do a lot

more experiments—and generate significantly more data—with the same starting materials and in the same amount of time."



Collective transportation of nanoliter droplets by a bigger droplet. Sequential imaging indicates the collection and transportation of generated nano-droplets by a millimeter-scale droplet through a tiny pore. Credit: Yu et al., *Sci. Robot.* 5, eaba4411 (2020)

"We programmed when and where the tiles were switched on and off to guide ferrobots through their designated routes," said Wenzhuo Yu, a UCLA electrical and computer engineering graduate student and a co-lead author on the paper. "This allows us to have several robots working in the same space, and at a relatively fast pace to accomplish tasks efficiently."

The robots moved at 10 centimeters per second and performed more than 10,000 cyclic motions during a 24-hour period in the experiments. In addition to transportation, other functions such as dispensing, merging and filtering of fluid samples were demonstrated as ferrobots interacted with structures on the on the chip.

More information: Wenzhuo Yu et al, *A ferrobotic system for automated microfluidic logistics*, *Science Robotics* (2020). DOI: [10.1126/scirobotics.aba4411](https://doi.org/10.1126/scirobotics.aba4411)

Journal information: [Science Robotics](#)

Einride ready to recruit first remote operators for autonomous trucks

Sourced by ROBOTREPORT



In the U.S. and other countries, aging populations and growing logistics demand have resulted in shortages of truck drivers. Autonomous trucks could help relieve those shortages. Einride AB today announced that it plans to hire what it called “the first autonomous and remote truck operator in the freight mobility space.” The Stockholm-based company said it will hire drivers in Sweden next month, followed by the U.S. in the third quarter. The remote operators would begin commercial services in Sweden in Q3 2020 and in the U.S. in Q4 2020.

In 2019, the U.S. had a shortfall of nearly 60,000 truck drivers, **according to** the American Trucking Associations. An aging workforce, a lack of female and millennial drivers, and concerns about emissions have increased interest in **autonomous** and electric vehicles, said **Einride** (*Chambiz DF 26 Jan 2019)

Remote operators promise savings

Einride designs, develops, and sells electric and autonomous “pods,” and its Einride shipper platform is a cloud-based transport-execution system for logistics. The company is working toward **SAE Level 4** self-driving vehicles. In addition, Einride said the use of remote operators would provide the following benefits:

- Reduce fuel/energy costs by 70%, from 60 cents to 18 cents (U.S.)
- Lower transport costs by 30%, bringing the ratio to one human driver per 10 vehicles
- Cut operating costs by 60% and increase productivity by 200%
- Reduce CO2 emissions by 90%

Hiring, training process

Einride said it will hire a former truck driver as the first dedicated autonomous truck operator. The remote operator will go through an extensive safety and technology training program and provide feedback to help develop Einride's remote driver station. The company also plans for a nine-month trial period and to recruit additional truck drivers to be remote operators.

“Today, our autonomous pods are operated by developers — robot engineers trained to drive trucks. A commercially scalable solution must rely on truck drivers, trained to remote-operate robots,” stated Robert Falck, founder and CEO of Einride. “We are excited to open up an entirely new category of jobs that will not only benefit the industries currently employees with improved hours, working conditions, and knowledge, but [also] reinvigorate a dying employment sector for the next wave job seekers.”

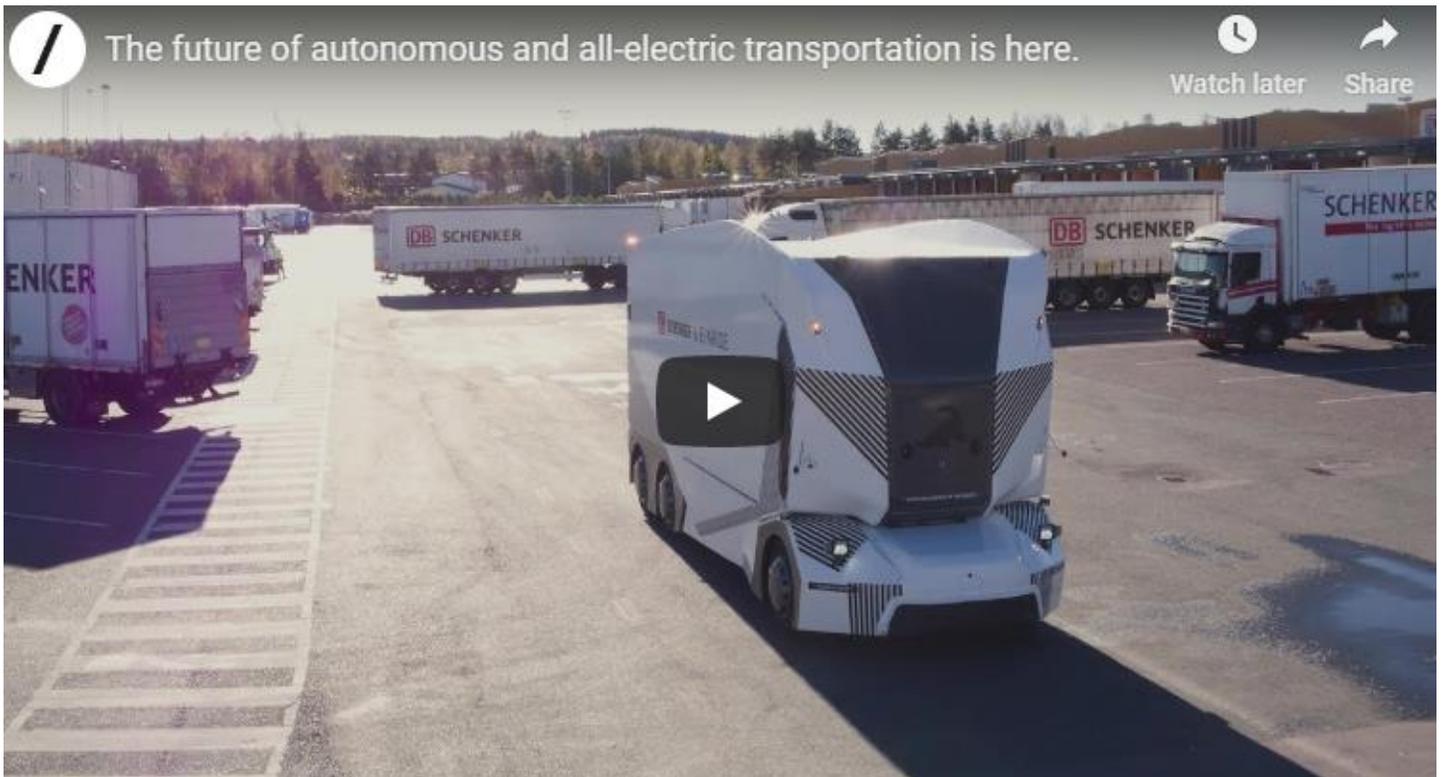
Falck responded to the following questions from **The Robot Report**:

What sorts of sensors are Einride's vehicles using? Are you creating the artificial intelligence for self-driving trucks yourself, or are you working with partners?

Falck: The Einride pod is equipped with lidars, radar, and cameras. We develop parts of the AI ourselves — such as basic safety systems — and have partners for others, like object recognition.

Can you describe the remote-control interface?

Falck: The remote station is a dome-like structure offering a view similar to that of a driver's cab, but with better side and backward visibility, and the option to switch between different views. But exactly how this working environment will look in the future is something we're still researching, and we are now involving truck drivers into that process.



During the testing period, will there be a safety driver inside the truck in addition to the remote operator? Once the vehicles are more autonomous, will there still be an option for remote or manual override?

Falck: The Einride pod doesn't have a driver's cab, so no. We start out with high levels of autonomy, and expand the operational design domain as things progress, rather than the other way around. For the foreseeable future, the vehicles will be supervised remotely, with an option for manual override. We believe in a human-to-truck ratio of approximately 1:10.

Among the challenges for semi-autonomous driving is the handoff between a human driver and the vehicle. How will Einride manage that with the remote driver?

Falck: The challenge is different, as Level 4 autonomous vehicles have fewer **disengagements**. Although it is possible to remote-control our pods, a typical disengagement doesn't require the operator to take over — i.e., to drive — only to give instructions.

When does Einride envision fleets of fully autonomous pods being on the roads?

Falck: By 2022 to 2023, we believe we'll have autonomous fleets operating in Europe and the U.S.

Are the biggest remaining technical challenges training the vehicles, scaling up to meet demand, or something else?

Falck: Scaling is our biggest concern — we are currently laying the groundwork for industrialization. But there are legal challenges, too, and some technical challenges, mainly related to the **lidars**.

How are autonomous vehicle regulations in Sweden similar to or different from those in the rest of Europe, or say, in the U.S.?

Falck: Well, there are vehicle regulations, and there are traffic regulations. As for traffic regulations related to autonomous vehicles, the U.S. is a patchwork, much like Europe, as the legal situation varies from state to state. Vehicle requirements, on the other hand, are regulated by federal authorities, so you have to deal with two levels. Sweden has a more streamlined process.

On the other hand, we don't have the progressive self-driving legislation of some U.S. states. Effectively, removing the driver is easier in the U.S.; removing the driver's cab is easier in Sweden. On the whole, the U.S. is moving faster.

Are you already serving logistics customers?

Falck: Small scale, yes. We have a commercial pilot for DB Schenker in Sweden, with a pod serving a continuous flow between two hubs in an industrial area.

How cool would it be to be able to fly your drone during an entire day instead of losing most of the time charging batteries or replacing them, like nowadays? This dream is now more reachable thanks to hybrid drones.

Six years ago, Quaternium startup, made a revolutionary breakthrough developing the first hybrid fuel-electric drone ever. This drone was inspired by hybrid cars from the automotive industry. In 2015 the company made a disruptive flight of 3 hours and 10 minutes with their multirotor HYBRiX, crushing all previous flight-time records. Since then, the achievements of hybrid drones have not stopped growing.

Quaternium is the leading European brand of long-endurance multirotors. Pioneers in the drone industry since 2008, the company has created very popular multirotors and gimbals, such as the Spidex drones. The company is located in Valencia, Spain and has thousands of customers in the five continents.

An entire new category of drones have emerged with the promise to unlock demanding commercial applications. Hybrid multirotors multiply flight time more than 10 times compared to electric solutions. Amazing, right?

The leader of hybrid drones, Quaternium, just broke the rules again. Pulverizing any previous endurance record with an astonishing flight of 8 hours and 10 minutes non-stop. All this, with an experimental version of their new drone HYBRiX 2.1, of 25 kg MTOW.

HYBRiX UAV took off at 08:52 am in Valencia, Spain, and landed at 17:02 pm. Performing a stationary flight of 490 minutes. With this breakthrough, Quaternium set the New World Record of Endurance with the longest hybrid-drone flight on Earth. HYBRiX's flight in Spain lasted nearly one hour more than the previous flight record, from Richen Power.

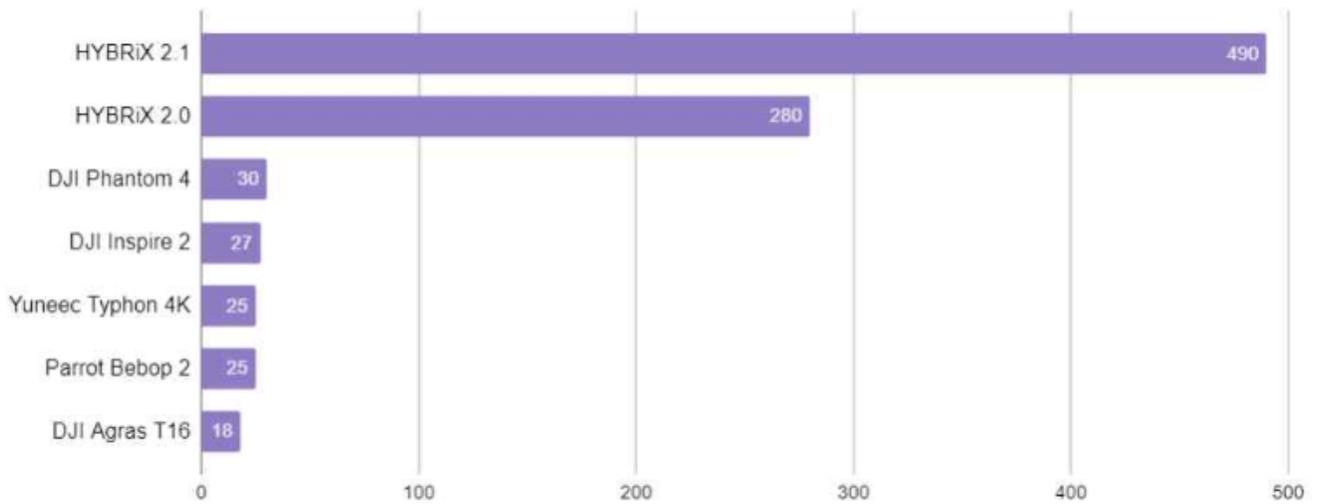
Below is a flight video with the demonstration of the Record.



Most multicopter drones in the market offer flight times ranging from 25 to 40 minutes, forcing the pilot to carry multiple batteries to the field while losing most of their time charging them. Compare this tedious work with just spending the whole time flying and getting all the information you need. With this game-changing innovation, it is possible.

Most multicopter drones in the market offer flight times ranging

from 25 to 40 minutes, forcing the pilot to carry multiple batteries to the field while losing most of their time charging them. Compare this tedious work with just spending the whole time flying and getting all the information you need. With this game-changing innovation, it is possible.



Flight time of most popular multirotor drones compared to HYBRiX.

HYBRiX project is supported by the European Union through their innovation program Horizon 2020, which fosters disruptive projects in innovation, excellent science and industrial leadership. The project has also received multiple Awards by institutions such as AUVSI, CDTI, or the Spanish Government.

Thanks to this huge advantage in flight time, HYBRiX is the ideal aircraft for aerial missions such as Surveillance or First Response. In every emergency situation, when time runs against the clock and lives are on the line, a long-endurance drone is a strong ally.

Today, the world is suffering from recurrent catastrophes caused by climate change. From floods to forest fires, generating huge damages to critical infrastructures, communications systems access and basic services. For over 10 years, quadcopters have proven their value in evaluating the aftermath of earthquakes such as the case of the Sichuan earthquake or Hurricane Florence’s landfall. In both disasters, multirotors were used to search the land and identify as many people as possible. Although they were useful, they all had the same limitation, flight endurance. Being only able to cover small areas.

During these situations, every second counts. This is what motivated Quaternium to develop a drone that could have a much longer endurance. With HYBRiX, rescuers and firefighters will be able to identify hot spots and access property damage while searching for survivors. Thanks to the unique features of HYBRiX, these situations can be managed safely while achieving faster results.

Quaternium’s HYBRiX long endurance solution is already available worldwide. More info is available on their website quaternium.com.