

05262021 Build KEY Kevin Scott

Build 2021
Kevin Scott

KEVIN SCOTT: Thank you all for joining us today and welcome to Build. I'm joining you all from my own home maker lab right now. I'm actually working on putting together an AI gadget, this little guy here to help me with remembering to stand up and move around when I'm in long work meetings.

I have a pretty interesting variety of tools here, which I always find super inspiring as I work on creative projects of my own. I'd like to talk a little bit about the idea of tools and how we as human beings use them to empower our own creativity and ingenuity. Not just table saws, CNCs and oscilloscopes, like the things that you see here, but incredibly powerful technological tools like AI and machine learning that are increasingly becoming critical parts of every developer and maker's virtual workshop.

If you think back to the very dawn of history, mankind has always used tools to solve our most vexing problems, turning our constrained zero-sum problems into abundance, opportunity and prosperity. Our tools allow us to expand the frontier of what's possible, from building a safer, more prosperous world to fostering our creativity and building community. Our tools are not just an essential building block of human society but are an essential part of our own humanity.

Over the past few decades, we've experienced a technological revolution that has put tools into our hands that are more powerful than our ancestors could have possibly dreamed of. The potential of technological tools like AI, the intelligent edge and the intelligent cloud to profoundly transform almost everything around us is unprecedented.

But maybe the most amazing thing is that right now these tools can be accessed by practically anybody who wants to use them to help make their positive mark on the world. By combining the democratization of technology, the support of creative communities and our own innate human curiosity, all the incredible power of these tools can be at your fingertips as programmers and creators, and the scope of what you can do with them is limited only by the boundaries of your imagination.

So today I want to encourage you to fundamentally reprogram how you think about applying technology as a creative tool. We're going to explore creative labs and spaces where people are using cutting-edge technology to create Grammy Award-winning music, explore the possibilities of programming for machine learning-powered edge and IoT devices to start new businesses and build new applications with supercomputing scale AI, and solve some of the biggest healthcare challenges on the planet with biotechnology.

I recently had the great pleasure of chatting with our first guest on my podcast Behind the Tech. Ben Bloomberg is a creative technologist who imagines, designs and builds

everything from electroacoustic musical instruments to AI-driven performances and tours. He's a recent graduate of MIT, earning his Ph.D. Ben and his creative partner, Jacob Collier, have an incredibly broad curiosity and ambitious creative vision. They've used existing technology and innovative ways to make art and to build community around it.

And when their creative vision runs into obstacles, like how to turn Jacob's layered, harmonically sophisticated YouTube performances into live shows in venues around the world, they've invented a new technology to make the impossible possible. In their hands, technology, even AI, is an instrument. Let's have Ben tell us more about how technology has been able to fuel his life-long creative journey.

BEN BLOOMBERG: I am a creative technologist and I experiment with everything from sound design to software and hardware design. And I work with artists to realize crazy interesting ideas. I grew up in a very musical family.

My grandmother was actually a music teacher, and my grandfather was an electrical engineer. So, I got really interested and excited about putting those things together. I had about 35 computers in my room, and I liked to take them apart and tinker with them. I love to imagine something and then build it.

I found one of Jacob's videos in – at the end of – 2014, actually, and just on a whim, sent him a message. And then he replied, "I've got a performance and I'd like to be able to make the videos that I make for YouTube live in front of an audience. So, you know, do you have any ideas about technology we could build that would achieve that?"

We really approached his project from the perspective of giving him more limbs. If he had 17 arms, you know, he could play 17 instruments. We were thinking of the technology sort of more as like a bionic prostheses. We want to give him superpowers. The first question we had was, you know, how can we do this without making the audience sit and wait for him to build up all of these layers, over and over, and over again?

The first part of that, we thought, should be what's called a harmonizer. Jacob specifically wanted to be able to use his to layer up these really complex harmonies, really quickly. We only had, actually, about 16 days to build it. There is a Windows computer that is completely sealed and passive, running in the box, and then there is some custom analog hardware, so it was really the combination of the harmonizer and this looping system together, that allowed us to pull off his first world tour.

I think it's really important to have the freedom to mess up and to experiment and to take risks. One idea that my adviser, Tod Machover, pioneered at the MIT Media Lab was this concept where you have a performer expressing themselves through more media than just the instruments. We talk about sort of wielding an entire performance venue, lighting and sound and robotics and set automation, and Scenix, it will happen that like everything in our lives will sort of become instrument-like.

You know, I think IoT plays a huge role there. You know, in the morning, you might try to find a way to choreograph your lights, and your alarm clock, and your coffeemaker, and the window shades, to create these special, you know, experiences.

Technology presents an incredible new language, with a huge potential for expression in new ways. The most important technologies going forward are going to be technologies that encourage us to be more ourselves. I would encourage everybody to build systems that help people celebrate who they are and how wonderful it is that we're all different.

KEVIN SCOTT: Thanks, Ben. I'm amazed by what you and Jacob have created so far and really can't wait to see what you'll do next. Ben mentioned how excited he was to use a variety of home IoT devices to create immersive new experiences. The potential for using the combined power of the intelligent cloud and edge to program for a wide variety of IoT and edge devices is going to increase exponentially in the coming years.

The opportunity for developers to create on this new ecosystem is going to be enormous. It's incredible to see what makers are already doing with technology and how powerful the hardware and software for innovating on the edge is becoming.

I suspect that any of you who are part of that group will already be pretty familiar with our next guest and their work. MIT engineer Limor Fried founded Adafruit in 2005 with a goal to create the best place online for learning electronics and creating products for makers of all ages and skill levels.

Today, Adafruit is one of the fastest growing U.S.-based manufacturing companies, a certified minority- and women-owned business enterprise and one of the biggest names in the maker world, creating phenomenally popular community-driven products and code.

Let's join Limor to learn more about how you can get started programming machine learning at the edge with Adafruit and Microsoft technology today.

LIMOR FRIED: Hello, everybody. It's me, Ladyada, and I'm here at the Adafruit factory in downtown Manhattan, where we manufacture all of our electronic goodies, from accessories for the Arduino, or Raspberry Pi boards, to our very own Feather and Circuit Playground Express. And we have a huge community of makers and engineers as well, with almost 30,000 Discord members, 1,500 GitHub repos and weekly live shows, almost every single day.

So, we have a lot of insight into what makers and engineers are doing with our products, and we're always so impressed with how creative they can be. I'm going to show you a demo using a Adafruit Braincraft Hat for Raspberry Pi and Microsoft Lobe. The Braincraft Hat fits on top of the Raspberry Pi 4 and makes it really easy to connect hardware and debug your machine learning projects.

There's a 240x240 color display, so you can see what the camera sees, which is great for vision projects. There are two microphones for audio input. And then, of course, you can connect things like relays, Servos, LEDs or other mechanical devices that you want to control from your Raspberry Pi. Microsoft Lobe is a free tool that you can use to create and train machine learning models that you can then deploy almost anywhere. This takes care of the hardest part of machine learning, which is creating and training a new model.

I've been playing with the idea of using a Raspberry Pi, the Braincraft Hat and a Raspberry Pi camera to recognize these delicious baked goods for my local deli. Now, these baked goods don't have QR code or bar codes. And so that's why using a camera with vision recognition would be an excellent way to identify and price each individual baked good.

Let's show how easy it is to use Microsoft Lobe to train a new machine learning model. We're going to create a new project called the Hello Bakery. Now it's time to import image data. Click on import and select your webcam. Now it's time to train the model on different images of baked goods. Let's start with the cinnamon roll. Place the object you want to train in front of the camera, and select the label that you want for these images. Now take about 20 images of your baked good. You'll want different poses, angles, let's see – if you flip it upside down.

OK, I've got about 20 images of this cinnamon roll. I'm going to continue doing this for each one of my baked goods. After I verify that my model is correctly identifying all my delicious baked goods, it's time to export that model. Go to the export tab and you can see all the options. You can create a web app. You can export it to TensorFlow, to JavaScript. You can even make your very own web server, web app or, when on a mobile phone, using iOS or Android. In our case, we want to export it to a TensorFlow light model file that will then copy over to our Raspberry Pi.

Let's start by SSHing into the Raspberry Pi, then CD into the directory where I've downloaded the software. Now you can run the basic prediction project that's written in Python. Because I've already deployed the model, it's going to immediately start doing predictions based on one the Raspberry Pi camera sees. To start, it sees nothing. But when I put a cinnamon roll in front, the text will update telling me the confidence and the label of the object detected.

I can also try a cross-bun or a bagel. Now, because we've deployed our machine learning model onto a Raspberry Pi, it's made it really easy for us to kick off all sorts of other cool hardware. So, I've got, in addition to the display, three LEDs we're going to glow green when something is detected. I've got a speaker and some text-to-speech code that's going to speak out what's detected. And I've even got a receipt printer hooked up that'll print out the product and the price.

So, let's try it with the cross-bun. One of my favorites, the cinnamon roll. Or the bagel. There you go. Thanks, Lobe. We sell the Braincraft Hat and Machine Learning Kit for Lobe at the Adafruit shop, and we've got tons of tutorials to get you started on your

machine learning journey. Now you've seen how easy it is. I can't wait to see what kind of creative projects you come up with. Thank you to Microsoft and Lobe for making machine learning training so easy. Now back to you, Kevin.

KEVIN SCOTT: Thank you, Limor. I actually took a shot at building an Adafruit Lobe project of my own. The device that I was working on when you all first joined me was a little computer vision system built on an Adafruit Raspberry Pi Lobe kit, running a vision ML model train with Microsoft Lobe. It will keep an eye on me when my monitor is engaged and remind me to get up and move around a little bit, every 15 minutes or so.

I started my engineering career over 30 years ago, designing electronic control systems and writing embedded software. I'm just stunned by how capable these systems are today and how much I'm able to do just tinkering around in my spare time. Like I seriously just put this little device together in a couple of hours over the weekend.

Now, let's change gears for a few minutes to talk about what we can do with truly immense computing power. Last year at Build, I announced our partnership with OpenAI to build one of the biggest supercomputers in the world. Since then, we've seen a ton of progress in large AI models and how we use them. We've gotten really smart about how to harness a certain set of algorithms and just massive amounts of compute in ways that have really allowed us to advance at a jaw-dropping rate.

It's a mind-boggling amount of computing power and we are speeding forward at an accelerating rate. But it's really important that we recognize that behind these huge supercomputers and behind the hype, AI is just another tool. It might be the most powerful and useful tool that we as human beings have ever seen, but it's only as useful as what we choose to do with it.

Our modern AI systems are being used to solve a broader and broader range of problems every day. Big ML systems have gone from being useful to predict ad clicks and ordering information on results pages, and in feeds, to helping us explore space, solve problems in aerodynamics, design molecules, and to tackle NP-hard combinatorial optimization problems like those at the heart of the logistics businesses on which we all depend.

When reading science and nature every week, I'm beginning to see folks using ML techniques in places where numerical optimization and differential equations might have once been used. In this sense, AI is more like a new form of mathematics, a tool to help us better explore and understand the world, than it is a substitute for human intelligence and ingenuity.

Much of the world's progress on AI, since the '50s and even since the beginning of the deep learning revolution that started about 10 years ago, has been about AI systems getting better and better at narrow tasks. And in some cases, the systems have become superhumanly better at those narrow tasks.

But the most interesting thing, and what we've clearly seen over the past year, is that we're seeing an acceleration and broadening of the cognitive tasks and of the complexity of problems that AI systems can tackle. I'm particularly excited about how AI can help us to be more creative and how it can give us more space and time to do so.

Last year, OpenAI launched its language model GPT-3, and we were blown away by some of the early outputs we saw, from writing poetry to auto-completing lines of code. There have already been hundreds of businesses that have used GPT-3 to create new applications, including Fable Studio, which is creating a new genre of interactive stories and using GPT-3 to help power their story-driven virtual beings.

Just a few months ago, OpenAI launched two new models, CLIP and DALL-E. These two models work together, hand in hand, and are intended to better equip machines to understand the world in the same way as we as humans do. We take in information, not only with what we read, but also in what we see and hear. So, if we can expose AI models to data in the same ways that it is absorbed by people, they should learn concepts in a way that's more similar to humans.

CLIP, which stands for Contrastive Language Image Pre-training, offers a new way to make computer vision significantly more flexible in general, because the models learn directly from natural language, as in regular English, like I'm speaking right now. Computer vision is important because that's how these models actually see the world.

Next, we have DALL-E. The DALL-E model can generate an image based on a natural language description. Again, a plain English description, such as a baby penguin in a red coat playing the piano, or a pig with a scarf flying an airplane. We've even used DALL-E to harness the massive power of AI supercomputing to solve one of the most vexing problems facing creators and developers today, creating infinite variations of Clippy.

OpenAI is still uncovering the full extent of these models' capabilities and evaluating whether they can be deployed in a safe, responsible way, but we're extremely excited about what's to come here and the possibilities that all of this will unlock.

As I mentioned earlier, we know it's you, developers, who can use powerful tools like GPT-3 to create ambitious applications that will leave a positive mark on the world. So I'm thrilled to introduce Sam Altman, CEO of OpenAI, who has some exciting news to share.

SAM ALTMAN: Hi, I'm Sam Altman here at OpenAI. We're a research and deployment company dedicated to making sure that powerful general purpose AI is safe and benefits all of humanity. Last year, we made our 175 billion parameter language model, GPT-3, available to developers through an API and private beta. We've been amazed by what people have done with it so far.

Developers are using GPT-3 to create realistic dialog, summarize complex documents, answer customer service questions and make search better than ever before. And those

are just a few of the hundreds of applications in production today. We want to help push the boundaries of what powerful AI models can do and support really ambitious projects aimed at solving complex problems of the highest order.

So, today, we're delighted to announce the OpenAI Startup Fund, a \$100 million fund, managed by Open AI, investing in startups with big ideas about how to use AI to transform the world. We're very happy to have Microsoft as an investor in the fund as well.

This is not a typical corporate venture fund. We plan to make big early bets on a relatively small number of companies, probably not more than 10. And we're looking for startups in fields where I can have the most profound positive impact, like healthcare, climate change and education.

We're also excited about markets where I can drive big leaps in productivity, like personal assistants and semantic search. We think that helping people be more productive with new tools is a big deal, and we can imagine brand new interfaces that weren't possible a year ago.

These aren't the only applications we'll consider, but they're at the top of our target list, and we're especially excited to hear from startups led by founders from underrepresented groups. We hope to treat funded companies as close partners and work alongside the founders here. They'll get early access to future OpenAI systems, discounts on Azure and support from our team. We're really excited about the opportunity for startups, for the industry and for people everywhere, who can put AI to work, improving their lives.

So if you're a developer or an entrepreneur looking to build something transformational with AI, we'd love to hear from you. Applications are being accepted now at openai.com/fund. Thanks very much and enjoy the rest of Build.

KEVIN SCOTT: Thanks, Sam. Microsoft is thrilled to be able to support this fund. We think this will be a fantastic opportunity for folks to really make a difference with AI. The range of what we can do with these kinds of large AI models, when applied to programming, is awe inspiring. But what can really change the world in ways we've only previously dreamt of is when we stretch the boundaries of our imagination and apply the benefits of massive ML models and AI supercomputing to solve problems in fields outside of traditional computer science.

For those of us who are engineers or developers, we've spent our lives dealing with engineered systems. Those systems are complex, but nowhere near as complex as evolved systems like human biology. With today's massive computing power, with high-volume experimentation and laboratory automation, and with innovations like large-scale AI, we're at an inflection point of what's possible at the intersection of technology and biology.

Using these technologies together, alongside the many advances in biosciences over the past two decades, increasingly allows us to treat these evolved biological systems the same way that we treat engineered systems, affording us greater understanding of how these complex systems behave, and in many cases, the ability to engineer them to do new and useful things.

This means that solving problems that previously would have taken years of painstaking experimentation can be solved much more quickly, sometimes within days or even hours, by reusing engineered biological components or even creating brand new ones.

We can now also split the hard work of experimentation, efficiently, between digital and wet lab environments. It's basically recasting biology itself as computation, giving us an immensely powerful computational bio lab. And it's a revolution that's accelerating incredibly quickly right before our eyes.

Let me share one very timely example of how this is happening today. Microsoft is excited to be working closely with Dr. David Baker, a professor of biochemistry and the director of the Institute of Protein Design at the University of Washington. His research group is focused on the design of macro molecular structures and functions.

When Dr. Baker began his work at the University of Washington, he decided to focus on the problem of something called protein folding. Each gene in our genome encodes a unique protein that carries out a unique function. And it does so because the DNA sequence in the gene encodes a unique sequence of amino acids, which then in turn fold up into a unique three-dimensional protein structure. These protein functions encompass every important process in our bodies and in the entirety of the biological life on the planet.

So if you want to understand how these processes work, how to understand to prevent disease or better understand the world around us, you first have to understand the interactions between proteins. You can think of these interactions almost like the relationship between a lock and a key, where they have to fit together very precisely to work. It's incredibly important to understand the geometry of the structures and how they interact, the same way you need to understand how the parts of a machine fit together to make it work correctly.

Until now, figuring out the chemistry and the physics of each individual three-dimensional structure has been incredibly challenging and expensive. What Dr. Baker and his team have been able to do is harness powerful computing technology to develop methods to go directly from the amino acid sequence of a protein to its 3D structure, bypassing incredible amounts of time and effort, and enabling breakthroughs never before possible.

One of these breakthroughs took place over the past year as we battled a global pandemic, creating new drugs to combat COVID-19 at unprecedented speed. As soon as the genome sequence of the COVID-19 virus was made available, Dr. Baker and his team

used the computational methods they've been developing to predict the three-dimensional structure of the protein on the surface of the virus, the so-called spike protein.

When that spike protein is searching for an entry point for the coronavirus to attach to and infect human cells, it looks for something called an angiotensin converting enzyme 2, or an ACE2 receptor for short. The ACE2 receptor is the key to the spike protein's lock, that when engaged, allows the virus to slip into cells and wreak havoc in the organism it is infecting.

Dr. Baker's team has designed small proteins that fold in such a way that they have a better match to the lock, becoming a more attractive target for the virus to latch onto. In doing so, they were able to create entirely new biological compounds that bound to the virus almost a thousand times more tightly than ACE2.

These were completely novel proteins unrelated to anything seen before in the natural world. The team and their collaborators were able to determine in the lab that the small proteins behaved exactly as expected, blocking the virus from getting into cells. And today, their discovery is now headed for clinical trials.

Just imagine having the benefit of having processes like this ready for the next pandemic, or to design completely personalized drugs to fight debilitating chronic diseases or cancer, or practically anything else that could be imagined in the biological world, all using software as a tool to custom build the precise biological molecules you need. It's something that would have been science fiction a decade or so ago, but it's being manifested as reality today because of the creativity and imagination of folks like Dr. Baker and his team.

That innate human curiosity and imagination is exactly what I want to encourage all of you to give yourselves permission to use, to paint on an almost infinite canvas with the amazing technological tools we have at our disposal. We've showed you a bunch of these mind-bogglingly great technological tools today.

Now, we at Microsoft, and a whole lot of other people have been talking about this kind of technology for a pretty long time now. And it's understandable at this point that you might think, well, it looks and sounds cool, but so what, does all of this AI stuff actually matter?

Here's the thing. The challenges that we have in front of us as a species are as big as we've ever dealt with in the course of human history. Just think about the scope of what we have to do to fight climate change, provide high-quality healthcare and eldercare for a rapidly aging global population, combat the next global pandemic, and deal with a host of other crises that are going to demand an extraordinary amount of time and resources, and impact your life no matter who or where you are on the planet.

So, all of this technology only actually means anything so far as it can empower the people who have access to use it to create a better future for all of us. We're building this

technology in service of people and the problems they have. We're already seeing examples of how folks are creating some acute and impressive technological solutions for problems that we're going to face in the future.

Using machine learning to identify plastic pollution in rivers and simulate how it moves in the ocean, these insights power passive clean up systems to help remove plastic that impacts our ecosystems, or developing opensource platforms for identifying and tracking wildlife, combining the strengths of AI and citizen scientists to help fight extinction. Our AI for Health team have developed a dashboard to track vaccine administration in the U.S. and worldwide.

To come to fruition at a larger scale, all of these kinds of things are going to require the innovation and ingenuity that only you can contribute. Part of the joy of using technology to create is that it should exist in a state of play. If it's not fun, we're doing it wrong.

So, I want everyone to know that Microsoft is committed to empowering whatever your creative dream may be, big or small, whether writing a song, creating software or IoT devices that help solve the world's hardest problems, or potentially saving millions of lives through cutting-edge biomedicine.

But to do so, we all have to expand our imaginations to unlock new and creative uses of technology that can be applied in the most powerful ways imaginable. If we do, what we can accomplish together can be truly amazing. Thank you so much for joining us. Until next time, be safe, be well and never stop creating.

END