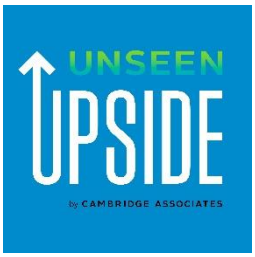

Season Four: Episode Six
AI & Healthcare: The Next Frontier
Launch Date: April 16, 2024

DISCLOSURE: All opinions expressed by host and podcast guests are solely their own opinions. The host, podcast guests, and/or Cambridge Associates clients and employees may maintain positions in the securities discussed in this podcast. This podcast is for informational purposes only and should not be relied upon as a basis for investment decisions. Please listen to the full disclosure at the end of the podcast for additional information.

Meg Lunn-Halbert: I was actually drawn to science since I was really young, probably about six years old. I used to pretend to publish journal articles with my dad. We had a little three ring binder. And so, I would publish articles on different animals in our backyard. We had like a pretty cool creek running through it, so I had like a great horned owl and a fox that lived nearby.

Hillary Ribaudo: Meg Lunn-Halbert is a second year PhD student in Biochemistry at the University of Washington's Baker Lab.

Meg: And so, I wrote little articles and published it in my little designed journal. But when I was 13, my mother passed away from lung cancer. And ever since then, I knew that I wanted all that scientific research I did to end up helping people in some way.

Hillary: The Baker Lab is a hub of researchers from various backgrounds and ages. And what connects them is AI powered protein design software. Meg uses Machine Learning to make vaccines more effective against diseases. Research efforts like Meg's are expected to eventually translate into interactions with real patients.

Daphne Koller: We need to get better at understanding what medical interventions are truly going to provide value to people. And I would say that one of the profound moments that I had in my life was the first time I heard the word evidence-based medicine. And the reason that was profound to me, because I had no idea that was the exception rather than the rule.

Hillary: Computer scientist Daphne Koller, is a MacArthur fellow based in San Francisco.

Daphne: That's unfortunately still true today and partly it's because in order to deploy evidence-based medicine, a clinician needs to stay on top of all of that evidence and there's just way too much of it for any human mind to assimilate and keep track of.

Hillary: I'm Hillary Ribaudo and this is Unseen Upside by Cambridge Associates, where we explore investments beyond their returns. This season we're talking to leaders and investors behind healthcare innovations that could change how long and how well we live. And in this episode, we are exploring how Artificial Intelligence innovations can lead to groundbreaking progress in diagnostics, treatment, and patient care.

Hillary: Artificial intelligence, or AI, is technology that basically is giving computers and machines the ability to simulate how humans think and solve problems.

Hillary: Meet Robo-Hillary she's an AI representation of my voice on the audio editing software we use to create this show. She reads my scripts ahead of time, so the production team can listen and make edits before I come in and record this. Robo-Hillary is just one of the many applications of AI. "You can get to Siri, at any time just by holding down the home button for a couple seconds. And then Siri's listening to you". This is audio from October 4th, 2011, when Apple introduced the world to its intelligent voice assistant. "So let's go ahead and ask Siri about the weather. "What is the weather like today?" [Siri Answers:] "Here's the forecast for today." It is that easy. And more recently, we've seen headlines about AI programs powering the controversial self-driving cars from Tesla and Waymo. "Waymo Robo Taxis are officially in service in L. A. You can now order a ride across 63 square miles from downtown L. A. to Santa Monica without being forced to chit chat with the driver if that's not your thing."

Daphne: For years, you've been interacting with AI and may not have realized this. Even starting with little uninteresting baby things like search that you do on the web has been using AI for well over a decade. There are People who will tell you, "Oh, this AI thing, it just suddenly emerged," and it didn't. I mean, we've been working in this space since the late 1950s.

Hillary: Daphne Koller is the founder and CEO of Insitro, a south San Francisco-based company that uses machine learning and data at scale to develop new medicines for patients. And her work with technology dates back decades.

Daphne: I was actually one of the people who were studying AI long before it was popular. When I graduated from my PhD in 1993, you couldn't say you were doing AI because it was, as my teenage daughter likes to say, "sus." AI was just kind of fringe.

Hillary: Daphne was actually the first machine learning hire at the computer science department in Stanford back in 1995. Since then, she has advanced the field in many ways. A significant one was co-founding Coursera, the online education platform.

Daphne: Many of the ideas that have emerged into prominence today are very similar to some of the, you know, admittedly baby versions of what existed in the '50s and '60s and certainly in the '80s and '90s.

Hillary: The concept of machines that think can be traced back to ancient Greece, but a pivotal moment came with a paper called Computing Machinery and Intelligence published in October 1950 by Alan Turing, the father of modern computer science. He's famous for cracking the German Enigma code in World War II. Many of the AI applications we use now were first built using what we know as traditional machine learning models.

Daphne: Machine learning is the task of getting machines to learn from data to perform tasks that are very challenging.

Hillary: Machine learning uses math to predict and estimate outcomes based on data. Essentially, it's an algorithmic approach to learning from data and making informed decisions.

Daphne: An algorithm is a set of steps that the machine takes in order to perform a task. What machine learning does is, it introduces a set of parameters into that algorithm that subsequently can be adjusted based on data to make the algorithms performance better and better.

The machine learning training phase is adjusting those parameters To maximize performance on something that you're trying to optimize. And then with those parameters having been learned, the algorithm can then go and execute the task again and again. If you want to take a really simple analogy, you can have a really complicated stereo system that has multiple knobs that you can set. So, you can imagine at the time that you're calibrating your stereo system to the room, you can twiddle different knobs and say, "Oh, this sounds better. This sounds worse." But then once you've twiddled the knobs, you can sit there and enjoy the music.

Hillary: Classic Machine Learning models require human intervention to learn, but AI has continued to evolve. And in 2012, there was a breakthrough on something called Artificial Neural Networks, which are computer systems inspired by how our brains work. They're made up of connected parts, kind of like artificial neurons, that can learn from information, recognize patterns, and make decisions with minimal human intervention, allowing for even more advances in the field. Daphne describes this journey as riding an exponential curve for decades, starting with a gradual slope in the early years.

Daphne: What's driven this exponential curve that we see in AI, It's the convergence of multiple forces that have come together, and all of them are essential ingredients. One, which I think is the driver, frankly, behind most of it, is the existence of very large amounts of data. Because data is this thing that feeds the machine learning, and in order to get to the kinds of models that we are all looking at today, you know, the chat GPTs of the world and the performance that they're achieving, that is because they're trained on web scale data that just never existed 10 years ago.

Hillary: Daphne says another essential component is the computer power needed to train the models.

Daphne: Those together, combined with the kind of AI talent that we have in increasing numbers today, is what allows us to rapidly experiment on lots of different models and continually refine them and improve them.

Hillary: AI as a field is constantly advancing, and the line between theory and practice is also in flux. So, definitions are fluid, but we can think of AI in the context of what it can do - its capabilities- and how it does it - its functionality.

There are several types of AI based on its functionality. For example, Netflix recommendations use Reactive Machine AI, a system that has no memory and is task specific. The model analyzes your viewing history data to suggest shows and movies you'll probably like. There is also Limited Memory AI which can remember past events and uses both old and new data to figure out what's best to do. But it can't retain data forever like we do - it only remembers for a little while. As it learns from more data, Limited Memory AI gets better at its job. Examples include self-driving cars, virtual assistants, and something the industry calls Generative AI.

Bob Nelsen: Generative AI, which is the latest things that you hear about Open AI.

Hillary: The makers of ChatGPT.

Bob: Essentially, it's using math to predict the next thing, the next event, usually a word when you're thinking about using it on language, but transformer math can be used on other things.

Hillary: Bob Nelsen is co-founder and managing director of Arch Venture Partners, a bold, early-stage venture firm focused on life science discoveries to prevent, detect, and cure disease. Arch has backed disruptive science companies for over 30 years, including Insitro.

Bob: Think of it all as evolving, right? When it comes to predicting what you might do, for instance, with generative AI in a customer service center, that's easy. It already is better than calling a helpline, and it's going to fix your phone or whatever it is faster. When you think about very disorganized states with less data, like in biology, there's a big question mark as to how it all comes together.

So, it's actually probably the hardest space that you can think of, to think about AI in because of the lack of data. So, you need to create data or find data, whereas in the other applications of AI, you have the internet, right? You have reams and reams of data about everything.

And you can train models and make approximations of things that are radically different than if I tell you, "I want to figure out how Alzheimer's works." That's actually a much, much harder problem.

Hillary: But we might be on the brink of a significant step forward because generative AI models can now create more compelling and more complete outputs.

Daphne: We have these systems that are able to pass the bar exam, that are able to pass the medical exam. You're not going to send your AI to do the bar exam for you, but those systems are going to be able to help physicians create better medical notes and retrieve better insight from the medical notes that they've taken. I mean, the use cases are just going to be so prolific, so profound, so pervasive that it's going to just transform everything.

Hillary: The intersection of AI and healthcare can be divided into two primary areas: diagnostics and pharmaceuticals in the first category, and hands-on patient care in the second category.

Bob: Care itself is already being revolutionized by AI, and it's a much easier case than drug discovery and diagnostics. You've probably seen the articles where ChatGPT, even though it hasn't been trained on healthcare, can already outperform physicians.

Hillary: True story, a study in JAMA Internal Medicine compared ChatGPT against real doctors on Reddit's AskDocs. They tested both on the same medical questions, and independent licensed healthcare professionals, who didn't know who wrote what, actually chose ChatGPT's answers 79% of the time! They judged based on how good the information was and how empathetic the responses felt.

Bob: Right before the pandemic I was speaking at a conference, about half physicians and a lot of healthcare investors, and I suggested that my nanny plus an AI was going to be better than them in five years. And they said, why your nanny? And I said, because she's got a really good attitude and bedside manner, and she's going to be able to get the data out of the AI.

Hillary: But we will still need doctors. Bob says AI can make good doctors a lot better.

Bob: You will need less doctors per patient. And you'll be able to push down the interfaces to much less trained individuals, so that the doctors can focus on more complex cases. On the administrative side, it just makes things more efficient. So, hospitals and practices will be adopting lots and lots of tools that just make things flow better.

Hillary: But to advance on the diagnostics and pharmaceutical front, systems need a lot of specific data.

Daphne: we're at the very beginning of the use of AI in healthcare, I think is probably 3 to 8 years behind where the use of AI is in other applications. Partly because we're dealing with what is rightly a heavily regulated industry, you know, do no harm, is really important. And then this absolute criticality of having lots of data in order to drive really good AI models. And the data in the healthcare space is for multiple reasons less abundant than just general-purpose web scale data.

Hillary: Daphne says there's already lots of books and articles that would help you to pass the medical exam, but that's just a fragment of the data actually needed here.

Daphne: Health records of individual patients with the kind of Language the doctors use, not as abundant because of privacy reasons, of course. And I think that's changing partly because of digitization of these data, partly because of the ability to go and do federated learning within protected environments without exposing patient privacy, all sorts of ways to get to that amount of training data. That is really the barrier that we're about to pass over the next three years.

Bob: You can get data from medical records. You can get it from your blood, and you can get it from other kind of diagnostics.

Hillary: This is Bob Nelsen again**.**

Bob: So, I get every six months a full body MRI scan and I get a grail test in my blood to look for cancer.

Hillary: Bob is talking about Grail's Galleri test, a blood test that aims to detect multiple cancer types from a single blood sample. Grail is a biotechnology company that's using advanced technologies like next-generation sequencing, population-scale studies, and AI to improve early cancer detection. Their results are promising, and they've caught the eye of investors like Arch Ventures, but more research and validation is still needed.

Bob: Grail test is an AI driven diagnostic, the first one actually, and it essentially looks for 50 cancers in your blood. So, you can see where that might be going, which is you have multiple inputs into a system on a healthy person to decide what might happen or what is happening, catch it early. And if you catch cancer early, for instance, it's mostly survivable.

Hillary: And then you have drug discovery.

Bob: That's the hardest one, because you don't have enough data, and you may need to even create that data. you have the laws of physics to define how these molecules work. And it's not clear how it all comes together, but the goal and the potential if you had the right data, maybe you could make a model that essentially figured out how to predict biology.

And that would be the Holy Grail, right? Which is a large language model or some other kind of math that took all of this data that we have about humans and disease and drugs into one model that you could say, "Hey, find me a molecule that looks like a drug that has low toxicity, that doesn't have an immune response that binds to a target that is relevant to these people that we call depressed." so the future is going to be, ask the machine what we should do, and it gives us a suggestion, but it's going to be very concrete, and actionable, and testable, and that will revolutionize drug discovery, I think, but we still don't know the math that's going to unify that. We are in a kind of discovery piece of the AI, more than in other areas.

Hillary: And that's the kind of work companies like Insitro are doing.

Daphne: For drugs that enter clinical development, as in they enter Phase 1 clinical trials, the percent that emerge as a successfully approved drug, depending on how you count, is between 5 and 10%. That's a success rate, not a failure rate. And so, when people ask, why are drugs so expensive. It's not only because the process is expensive on its own. It's because every successful drug has to carry on its back the cost of all of the failures and also the risk.

So, we're feeling our way in the dark a little bit. Biology is really multifaceted, really complicated, multiple biological scales. And the human mind is just not equipped to understand that complexity.

Hillary: Daphne thinks of AI as a framework that will help us navigate that complexity.

Daphne: Over time, when you feed it enough data allows us to create models that make hopefully better and better predictions of if we did this in a human, what would happen? We're nowhere close to that today, but the vision behind Insitro is collect more and more data at more and more biological scales, that allow us to make predictions about that cause effect relationship.

Hillary: In their lab, they collect vast amounts of data by tinkering with cells and tweaking genes to see how they react and change. This includes turning genes off or upping their activity. They also gather data from real patients, looking at how genetic differences affect different traits.

Daphne: All of those data can feed into the maw of this machine learning model that allows us to learn the relationship between genetic interventions and downstream consequences, and hopefully that will allow us to create a more robust set of predictions of what would happen if I intervened in this gene, And if I'm able to get the confidence around those predictions, I can now start to make a molecule that does that intervention and hopefully feed it into a human and maybe intervene in a meaningfully beneficial way.

Hillary: But that is a long journey, because first they have to build the engine to interact with the data, and then they have to turn those insights into actual drugs.

Daphne: Drug discovery is a slow business, and in our case, we actually have to build a platform before we could deploy it towards discovery efforts. So, you kind of have to add that prequel if you will.

Hillary: Currently, Insitro focuses on 3 therapeutic areas: metabolic diseases, oncology, and neuroscience, and their work so far is revealing a lot of potential.

Daphne: The first indication that we picked in neurodevelopmental work was a disease called tuberous sclerosis complex. It's a pediatric epilepsy that is quite severe, these are infants that have 20-30 seizures a day. And when you have that level of seizure frequency, it fries your brain basically. And so, you begin to have a very high risk of intellectual disability.

Hillary: So, Daphne's team developed all the necessary capabilities for studying the disease. This included taking and producing neurons in a dish, mutating them to mimic disease traits, and conducting all the required measurements.

Daphne: That was like three years end to end. And we were extremely excited when we got some of the insights that came out of the machine learning and we're able to take those insights and put them into an In vitro system that basically grows neurons to the point where they're like, 50, 60 days old.

Hillary: And then these neurons started showing seizure activity similar to what we see in human EEGs. They were communicating and seizing together in a feedback loop.

Daphne: And our tool compounds were able to eliminate that seizure activity, which was like really exciting. The second time that we did this, which was in a different genetic epilepsy that leveraged the same platform, but with a different disease-causing mutation, we're not quite at the end of the day, but the thing that took us three years took us three months. Because you have a platform and that's the vision behind what we're doing, that you have to make the upfront investment to build data generation, Data collection and data interpretation capabilities, you spend time building the infrastructure in the engine first, and that takes much longer, and everyone gets annoyed and frustrated. But then once you have it, the engine starts to crank and stuff just starts to come out. So, it's the industrialized engineered approach to discovery that we Placed our bet on. And right now, I think we're optimistic about the bet panning out.

Hillary: But the secret sauce in this process is really the multidisciplinary aspect of the work.

Daphne: The kind of system that we're building draws on the expertise of a tremendous number of disciplines, whether it's life scientists, biologists, chemists, computational scientists, engineers, machine learning scientists, the DNA support functions that enable all of this. Machine learning models come and go. The true sustainable advantage is the culture that you build and the people that build that culture.

Hillary: Another place with a multidisciplinary core that's making strides in this intersection of AI and life sciences is just north of Insitro, in Seattle, Washington.

Sanjay Srivatsan: So, you know, somebody recently visited the lab.

Hillary: Sanjay Srivatsan is an alum of the Baker Lab, and now is a professor at the Fred Hutchinson Cancer Center.

Sanjay: They were a funder and they described it as like Willy Wonka's chocolate factory, but instead of chocolate, it's for proteins, and that's actually a really great descriptor for the lab. It's kind of this magical and whimsical place where literally anything that can be done with proteins is being done with proteins. And it's all being done from scratch using AI tools, but also like this very cool, human, imagination, and creativity that kind of just comes from the people that are there.

Hillary: The Baker Lab is part of the Institute for Protein Design within the Department of Biochemistry at the University of Washington.

At the Baker Lab, a team of researchers -including graduate students and postdoctoral fellows- develop protein design software and use it to create molecules that solve challenges in medicine, technology, and sustainability. If you walk into the lab, you'll see whiteboards, computers, and people milling about exchanging ideas.

David Baker: My sort of metaphor is a communal brain, where every researcher is a neuron and you're connected to everybody else, totally flat hierarchy. And it's a very intense environment where people are collaborating with many different people in many different projects.

Hillary: And the collaboration goes beyond the Baker Lab's walls.

David: When a student wants to design a protein to solve a new problem, I tell them, the first thing you have to do is find the three best people who work on this problem in the world and set up collaborations with them. So, I think it works because everybody knows what else everyone's doing. So, if someone makes a breakthrough, then a week later, everyone's taking advantage of it for their own work.

Hillary: The lab is run by biochemist and computational biologist David Baker. He pioneered methods to predict and design three-dimensional structures of proteins, which play a crucial role in various biological processes. David is also an investigator of the Howard Hughes Medical Institute in Maryland, and a professor at UW. And he has literally climbed mountains with Sanjay.

Sanjay: One way in which David is like Willy Wonka is that like he wants everyone to feel and take pleasure in that same kind of creative aspect. So, he's always bouncing up and down, there are over 100 people in the lab, and he knows exactly what every single person is doing. People across so many disciplines from physics to chemistry, to material science, to cell biology, they all see that vision for protein design and come to this place to do it.

Hillary: Baker Lab started when David came to UW back in the 1990s.

David: At that time, the idea of protein design sounded totally science fiction. It was almost exactly 20 years ago that we showed that it was possible to design a completely new protein from scratch and to do it accurately.

Hillary: Ok, to really understand what they are doing here with AI, we need to define what a protein is.

David: Biology, at least molecular biology, has this aspect of studying the elven runes of like Lord of the Rings, where you have these things from lost generations, and that's the proteins that we have in the world today that are in our bodies.

So, our genomes contain genes. Each gene encodes one protein, and the protein is kind of what does the work. Our genomes are kind of like the blueprints for life, but the proteins actually go and do things. The function of a protein is basically what it does, whether it breaks down your food when you eat it or mediates electric currents going through your neurons when you're thinking. Basically, everything that happens in our bodies and in all living things is being done by proteins.

Hillary: Now, the building blocks of proteins are organic compounds called amino acids, and there are 20 of them in the human body.

David: Each protein has a unique sequence of amino acids. It's like a long word where you have a choice of 1 of 20 letters at each position.

Hillary: And in nature, proteins came through the natural selection evolutionary process.

David: There were proteins hundreds of millions of years ago that mutated and changed and their sequences changed and that let them acquire new functions. The proteins that we're de novo designing are completely outside of this evolutionary process. You can't trace their history back to the origins of life. Our proteins are made, we're designing them completely from scratch to solve whatever problem we want to solve.

Hillary: So, you might be wondering, why create what will eventually evolve?

David: Evolution operates very, very slowly. So, you know, it's basically survival of the fittest. Well, in the cancer case, you don't really want evolution to be the solution. So, protein design lets you design proteins now to solve current real-world problems, rather than having to wait for millennia for new proteins to evolve to solve them.

Sanjay: It's amazing the number of different kinds of things that you can do.

Hillary: Sanjay Srivatsan, again.

Sanjay: You can design cells, you can design anti-venoms, you can design vaccines.

Hillary: Students here are working on all kinds of problems using AI model, things like autoimmune diseases, or trying to combat drug resistant bacteria.

David: For Celiac's disease, we've designed proteins that break down gluten, and those are currently in pretty advanced clinical trials looking very promising. We've designed proteins to attack resistant cancers.

David: we're designing sort of artificial photosynthetic protein systems that you could use for autonomous energy, light harvesting and for a variety of chemical reactions, plastic degradation.

There's really a wide range of different problems that protein design can be applied to. But the underlying methods are the same. And then other students are really trying to develop the sort of the next generation methods.

Hillary: And it turns out that proteins are a perfect fit for AI methods like Machine Learning and Deep Learning because we've got a ton of data about them.

To start, scientists had to figure out how proteins fold, which they did by uncovering their three-dimensional structures. Each protein's unique sequence of amino acids determines how it folds into its 3D shape.

David: Proteins are so complicated and big, to model them you need to use computers.

Hillary: Several projects here leverage Deep learning, a subset of machine learning that requires less human intervention and is particularly effective with large datasets.

David: it's called deep learning because it's a neural network with many connections and hundreds of millions of parameters that take you from the sequence to the structure or in the design case, you specify a problem, and the output is a new protein that solves the problem.

Hillary: Over 2 decades ago, David and his team got started on a software tool for protein science called Rosetta, years later the Baker lab also started Rosetta Rosetta@home, which is a community science project that let many thousands of volunteers run Rosetta calculations on their ideal personal devices like computers and phones. And then in early 2020s a new tool called RoseTTAFold emerged from the lab, harnessing Ai to predict the structure of a Protein from its Amino Acid Sequence. Imagine it like this super smart “how to guide” for the world of molecular puzzles. In 2021, Science magazine named AI-powered protein prediction as its Breakthrough of the Year, featuring AlphaFold which was developed by the Google sister company DeepMind and RoseTTA fold.

David: But neither program was able to predict the structures of proteins interacting with other molecules.

Hillary: So, fast forward to last month and the lab released an extension of RoseTTAfold, called RoseTTAFold All-Atom. It leverages deep learning to model complex molecular assemblies beyond just proteins, including DNA, small molecules, and metals. This approach has allowed for the design and testing of proteins that bind to specific molecules related to cardiac disease, enzyme activity, and light harvesting, and all of it was validated through experiments.

David: it can take in not only the amino acid sequence of a protein, but also, all of the molecules that it interacts with, and it will build a model of the whole biomolecular system, which is important for drug design, for example.

Hillary: In resonance with IPD's focus on accessibility, the code is open source, which means is available to anyone.

Now, this upgrade to RoseTTAFold allowed the team to create a new version of RF-diffusion, which is their powerful deep-learning program that can generate new protein structures. And the inspiration behind it came from an unexpected place.

David: So, you're probably familiar with programs like Dall-E that you can give it a text prompt and it will generate an image. The way Dall-E is trained is a very large number of images are taken from the Internet. Noise is added to them and then a network is trained to remove that noise. Once that is done, you can start with a completely random noise image and progressively denoise it. And what you end up with is a new image that looks like it could be one from the database that you were training on, but actually it's completely new. And you can guide that process with the text prompt. That's why you get images that, you know, look like whatever you want to look like.

Hillary: Baker Lab's RFDiffusion works in a similar way.

They start by taking real protein structures and adding some random noise to them. Then they train their program. After that, they can start with completely random placements of amino acids and use RFDiffusion to clean up that noise, over and over again.

David: And we end up with something that looks like a totally viable normal protein structure, but it's completely new. And we can bias that process by giving it a prompt, like we can give it the protein on the surface of the cancer cell, and carry out this denoising process in the presence of it, and then what results is a protein that binds to that target.

Hillary: But the AI part is just the first part.

David: So there's the computational part of protein design, but then there's the whole, experimental wet lab part where you're actually testing it to see whether it does what you designed it to do.

David: If we're designing a new protein to treat pancreatic cancer, for example, we design an amino acid sequence that's predicted to fold up into a three-dimensional structure that would attack the cancer cell. That's just on the computer though. Then we need to actually make the protein in the real world.

So what we do is we make a synthetic piece of DNA like a synthetic gene that encodes that new protein, we put it into bacteria, the bacteria then produce large amounts of that protein. And then we can take that protein and see whether it actually destroys the tumor.

Hillary: These methods have already been used to design therapeutics.

David: So the highlight so far is my colleague Neil King at the Institute for Protein Design has used the methods I've described to design a COVID vaccine, which has been approved for use in humans in Korea and Great Britain.

It's the first approved de novo designed therapeutic. And other, design medicines that we've made are currently being tested for a variety of different indications. So this is happening now. It's not really the Future.

Hillary: And because the designed proteins are going towards actual drugs, they have to go through preclinical and clinical trials like any other drug to make sure they're safe.

David: Biology is complicated. And so applying AI to say more complex biological protein problems than proteins will require the acquiring really good data sets. But AI is not going away. I think the methods that we're developing will just keep getting more powerful and more broadly used.

Theresa Hajer: So, if you think about the discovery aspects, you're not changing the decisions of how you think through drug development, but you are accelerating some of that early discovery piece. And so, you're not changing the decision maker, but you're really augmenting a scientist in this case early on to do their discovery work. I think that's where we've seen more of the adoption today, and it's made a real impact.

Hillary: Theresa Hajer leads the US venture capital research team at Cambridge Associates. With 20 years of experience, she's seen the field evolve.

Theresa: I started in this position in 2003. And within the first five years the human genome was sequenced, Facebook was founded. The iPhone was launched. So in that period of time, technology markets and venture capital was really just a U. S. Market. And today technology markets are truly global and the access and adoption of technology from a societal standpoint has gone global.

Hillary: The AI-designed medicines these teams are creating could be more potent and stable than what we have today, potentially reducing costs and simplifying global distribution. However, as with AI in other fields, fear is part of the conversation.

Daphne: people look at this technology, and they think back to all of the science fiction movies that they've watched, whether it's 2001 Space Odyssey or the Matrix.

Hillary: Insitro CEO, Daphne Koller.

Daphne: and they're like, "Oh, these things look so smart. Look at the kind of dialogue that GPT can generate. They're almost conscious." And they're not. They're just really powerful pattern recognition engines, but they're not sentient. And it's not like really close to that. But I do think they're very powerful tools, and people have demonstrated in the past, able to take really powerful tools and use them for evil. And I really wish that people would pay more attention to those risks rather than focusing on this sort of mythical the machines are out to get us- Terminator type scenarios.

Hillary: And there is also the concern about jobs that AI will displace. Daphne says every major technological revolution has led to the elimination of entire job categories, like

switchboard operators for example, but it has also sparked the creation of new job categories like computer programmers.

Daphne: The ultimate success here is going to be achieved with a synergy between people and computers, but the intellectual driving force, that is still the human. And I think that will be that way for quite a while.

Bob: Thinking about AI is an opportunity more than a threat.

Hillary: ARCH Venture Partners co-founder, Bob Nelsen.

Bob: I actually think that the threat of overregulation in IA, not just in health care, but everywhere is much greater than the threat of AI is going to kill us all.

Hillary: And from David Baker's perspective, the unseen upside for protein design really outweighs the risk.

David: One of the really important things about protein design in particular is where one goes from a computer model of a protein to the actual protein, and that requires manufacturing a piece of DNA that encodes that protein. We make up this new protein. There's no gene for it. So, you have to make it. So, there's a very clear step where you go from something that's just a computer fiction to an actual thing in the world. And I think that's the step that is important to monitor, and I favor, an approach in which these DNA sequences that are made are tracked.

David: If anything, dangerous ever happens, we can easily go back and see where that sequence came from. And stop anything more coming from that source.

Hillary: And many top scientists agree that AI's advantages in the field far outweigh any potential downsides. Over 90 experts, including David, have signed the Responsible AI x Biodesign agreement, highlighting their shared values and commitments for using AI responsibly in protein design. This agreement stemmed from a workshop held at the White House.

David: the government said look, governments can impose regulations, but they can be very heavy handed. So it's much better If the community who really understand the technology and the problems can basically self-regulate and self-govern, and that's really what those community guidelines came out of.

Hillary: Theresa says that from an investment perspective there are also challenges.

Theresa: just thinking about it at a very high level, it’s still very early. And so if you're an investor thinking about the opportunity set in the role that technology is playing, where are the shorter-term applications versus the longer-term applications. Navigating a lot of uncertainty can be critical.

Hillary: And although these are kind of the early days, the landscape is already very competitive.

Theresa: There's higher barriers to entry if you will, to start a life science or venture capital firm focused on health care overall. And so it's a smaller number of firms. I still think it's very highly competitive and always will be to partner with the most talented founders in the space.

Hillary: But like many, Theresa keeps a positive outlook.

Theresa: I hope it's going to impact more and more patients, both in terms of being able to help identify more therapies and in some cases more cures for different diseases, but as importantly, thinking about the ability to reach more patients, even outside of A I. But when you just think about broader technologies and the ability to bring health care to more people who may not have the opportunity to live nearby some of the excellent institutions that many of us Get to be, I think the access and that impact is something that I do think is possible with some of these advancements. And I think that will be a true measure of success.

Daphne: The world that I think we need to go towards is the next generation is actually data driven medicine, where, as a person goes through their lifetime, there is more and more data that is collected about their health state, from the standard exams as they go into the doctor's office from blood tests and imaging and so on.

And there is an engine on the back end that integrates that information and is able to make Data driven recommendations of what tests to perform, what treatments might work, and of course, it's not going to be an entirely automated procedure. There's going to be a personal loop with a clinician, and really, it's going to be a shared journey between the patient, the clinician and the AI.

But you need to be able to be making better health recommendations. The only way to do that is with data, and the only entity that's able to assimilate that data is going to be basically an AI model.

Hillary: Bob Nelsen says that to really make a difference, we need to rethink the entire care system using AI.

Bob: We need to get more efficient and our single biggest expense, which is health care, and I remind my tech friends who are kind of view healthcare investing as kind of this curiosity that the health care part of the economy is actually bigger than tech.

Which they're surprised to learn. But it turns out we have a big bill, in an inefficient system. So I think about it as just tremendous opportunity to disintermediate an ineffective system, because it's going to benefit society and people, it's going to prevent disease and cure disease and make the system much more effective and save tons and tons of money that we can then apply to the other problems that we have.

Hillary: If you want to learn more, please visit us at cambridgeassociates.com/UnseenUpside or check out the show notes. If you like what you're hearing, leave us a review and tell your friends and colleagues.

At Cambridge Associates, our podcast team includes Michelle Phan, and me, Hillary Ribaudo. Thank you to Megan Morrissey, Robert Scherzer, Krista Matthews, and Deirdre Nectow. And a special thank you to our guests and to Bryan White, Ian Haydon, and the Baker Lab team.

This episode was produced by Sandra Lopez-Monsalve at PRX. Genevieve Sponsler is our editor. Production support by Isabel Hibbard and Emmanuel Desarme. Mastering by Samantha Gattsek. And the executive producer of PRX Productions is Jocelyn Gonzales.

Unseen Upside will return in the fall with more stories about investments beyond their returns. Thanks for listening!

[DISCLOSURE]