The Economic Club of New York

116th Year
705th Meeting

Daniel Huttenlocher
Inaugural Dean
MIT Schwarzman College of Computing

May 3, 2023

Webinar

Moderator: Peter Coy
New York Times Opinion
Economics Writer

Introduction

President Barbara Van Allen

Good afternoon and welcome to the 704th meeting of The Economic Club of New York. I'm Barbara Van Allen, President and CEO of the Club. The Economic Club of New York is recognized as the nation's leading nonpartisan forum for discussions on economic, social, and political issues. And, of course, we're now in our second century, so we've been doing it for a long time, up to and including today.

I'd like to extend a warm welcome to students joining us virtually from Rutgers University, Columbia Business School, and the Gabelli School of Business at Fordham. Also, we have some members joining us from our largest-ever Class of 2023 Fellows – a select group of diverse, rising, next-gen business thought leaders, and we actually have 72 this year. We have a few that are also national virtual fellows from around the country because we have expanded the program to get beyond the borders of the New York-Metro area.

We are really honored to welcome our special guest, Dan Huttenlocher. Dan is the Inaugural Dean of the MIT Schwarzman College of Computing and is the Henry Ellis Warren Professor of Electrical Engineering and Computer Science. Previously, he helped found Cornell Tech, the digital technology-oriented graduate school created by

Cornell University in New York City. He served as its first Dean and Vice Provost. I think probably everyone listening is very familiar with that wonderful accomplishment.

Daniel's research and teaching have been recognized by multiple awards including the ACM Fellow and CASE Professor of the Year. He has a mix of both academic and industry background, having been a Computer Science faculty member at Cornell, researcher and manager at the Xerox Palo Alto Research Center, and Chief Technology Officer of a fintech startup.

He is internationally recognized as a researcher in computer vision and the analysis of social media. His book, *The Age of AI: And Our Human Future*, co-authored with Henry Kissinger and Eric Schmidt, was published by Little, Brown in 2021. It's now in soft-back as of last fall. You can get a copy. He served as a member and chair of the board of the John D. and Catherine T. MacArthur Foundation and currently serves as a member of the boards of Corning, Inc. and Amazon.

He will begin with some opening remarks which will then be followed by a conversation. And we're really delighted to have our Club member and *New York Times* Opinion Economics Writer, Peter Coy, doing the honors of moderating. As a reminder, the conversation is on the record. We do have media on the line and in the room. We'll have the chat box open if folks want to put questions in. We will be able to have Peter

take a look at those and, time permitting, use them. So Dan, if you're ready, the mike is now yours.

Opening Remarks by Daniel Huttenlocher

Absolutely. I'm delighted to be speaking with you all today. Thank you, Barbara, and the Economic Club for having me. And thank you, Peter, for leading our discussion.

I just wanted to open with some comments to help frame our conversation. When you think about AI technology, actually it dates back to the middle of the last century. For instance, the first chatbot – we've been seeing a lot of chatbots lately – the first chatbot, Eliza, which imitated interactions with a psychotherapist, was actually created in the early 1960s. But over the past decade, advances in machine learning have completely changed the nature of AI from what it was before.

With machine learning, AI is now fundamentally different from any previous technology in the history of humanity. Our centuries of experience with technology no longer serve as a guide. Historically, technology was built to be predictable, to be something that could be understood in a reductive manner. You could look at what bolt failed and caused a vehicle to crash or, you know, what line of code was wrong that caused an error in software. But AI today is not like previous technologies because machine

learning produces systems that are different in three essential ways. They're imprecise. They're adaptive and they're emergent. This is very different from previous technology that was very predictable. So you don't think of imprecision and adaptiveness and emergence as predictable.

So technology previously had been somewhat rigid and now technology is very flexible with machine learning, this new revolution in AI. And that kind of flexibility we normally associate with humans or with animals. So these advancements are giving us unimaginable new opportunities, but they also pose substantial new risks. We really have no roadmap for this change and precious little understanding given the fundamentally different nature of this technology. And I want to stress that this is the case despite the definitive sounding statements about the impact of AI that one hears almost every day.

Recent advances in AI demand a new understanding at a practical, conceptual and philosophical level. We need to be wary of purported understanding that draws too close a parallel to previous technologies given these differences with previous technology. But we also need to be wary of purported understanding that draws too close a parallel to science fiction, be that either Dystopian science fiction or Utopian science fiction, and a lot of today's statements really draw close parallels there.

And I just particularly want to emphasize, it's particularly true, the warnings that one hears about sort of the subjugation of humans to AI. This is a distraction from more immediate issues. And I would posit, not always a distraction that's accidental by those who are making it. So we have our work cut out for us, and by us I mean all of us, not just technologists like myself, in building this new understanding and roadmap that's practical, conceptual and philosophical. With that, Peter...

Conversation with Daniel Huttenlocher

PETER COY: Thanks, Dan. I'm Peter Coy, Opinion Writer for the *New York Times*, I cover economics. And I'm coming to you today from Bronx River High School. I spoke to students this morning, and I couldn't get back to my office in time for this. So they put me in a guidance counselor's office. So here I am.

Dan, a great intro and great book, by the way. I read it in preparation for this event, and I recommend it. And although the book was from 2021, you updated it with an afterword last year. So obviously there's no way you can, things are changing so quickly that you didn't get to the most recent iteration. But you got ChatGPT in there so it's a very good tour of the horizon for people who want to understand better what's going on in tech and in AI.

I actually want to ask you, first of all, about your collaboration with Kissinger and Schmidt. What was that like? You mentioned that you had some disagreements, so talk about the process of putting the book together and where you disagreed and how you managed to find a way to put the book together.

DANIEL HUTTENLOCHER: So, especially to The Economic Club of New York, I can proudly say something that I usually say regardless, which is that this collaboration could only have started in New York. Eric Schmidt had long been engaged with the Cornell Tech Project from the time when he was CEO of Google and Google gave us space in their building down at 111 Eighth Avenue.

And I had been working with Eric on the Cornell Tech Project on various things around AI. And then I got introduced to Henry Kissinger by a mutual friend because Dr. Kissinger was trying to understand AI and this friend said, well, you should talk to Huttenlocher. And unbeknownst to me, Eric had also been talking to Kissinger about AI. So all these parallel conversations were going on and, you know, sort of the triangle collapsed.

And we started, not with an objective of writing a book, or even of writing anything, but rather trying to – and this was, you know, sort of, probably 2017-‘18 kind of time frame – just really trying to understand this AI revolution that was happening and the kinds of

things I mentioned in the opening, that AI was fundamentally different than other technologies. And I think, you know, I'd say Dr. Kissinger's curiosity about what this meant for his world is really what drove the collaboration fundamentally. And then Eric and I, of course, come at this from a, both as technologists, but Eric much more from an amazing career developing such technologies and me more from a research perspective.

So I think as to differences, I'd say Eric is more of a techno-optimist, of the three authors. I'd say Dr. Kissinger is more of a – I wouldn't go all the way to techno-pessimist, but a lot of caution. And I would put myself squarely in the middle, kind of a middle-of-the-road guy. And, you know, I think you can see some of that in some of the discussion in the book where I think we all agree on things like the genie is out of the bottle. There's no pushing this back in. And we all agree that these are things that we need to work on broadly around the world and in the United States, and not just the technologists. So those are definitely the points of agreement which are sort of the things I also opened with. I think they're the biggest issues. The disagreements are kind of around specifics, I would say.

PETER COY: Right. So I want to start out with something on the positive side about AI, the story of halicin. Can you walk us through that? It's really, I found that very cool.

DANIEL HUTTENLOCHER: Yes, it is. It's amazing. And actually there's a next chapter to that, that's beyond that. So I think one of the really powerful things about machine learning is the opportunity to really change the way discovery happens in almost every domain – in education, in scientific research, in technological innovation. And we've seen a lot of use of machine learning in the last sort of five years or so, that has really revolutionized areas like certain kinds of diagnosis in the medical domain, particularly radiology-based diagnosis, and also discovery and development of drugs.

And so with halicin what they did, actually they were looking for new antibiotics because of antibiotic resistance. And so the idea was to use machine learning to help identify likely candidates among existing substances that were already being manufactured for other purposes. And they identified something that they called halicin, which was not at that time being used as an antibiotic. And the interesting sort of footnote to that is that there's so many stages in the development of pharmaceuticals after sort of the early-stage discovery, it's proving that halicin in vivo is probably less effective than it looked like early on in the discovery process.

But, of course, this happens with human-discovered substances as well and, in fact, it underscores why it's so important to get more efficient at identifying potential compounds, because really getting, you know, the safety and efficacy was already known because it was an existing compound that was used in humans for other

purposes. But the effectiveness of it is proving to be less than it looked like from the screening process.

PETER COY: The part about the halicin story that I found most fascinating was the way that the machine learning detected a pattern that a human being could not detect. And you liken it to the way a computer plays chess. Can you draw that analogy out a little?

DANIEL HUTTENLOCHER: Absolutely. So machine learning is fundamentally about recognizing patterns and then sort of storing and being able to access those patterns in various ways. And, you know, much of human discovery is about finding patterns in things and so that's why I'm so excited about machine learning for scientific discovery. And, you know, we're limited by our own perceptual abilities as people and so the idea that there's now technology out there that essentially can perceive things about the world that we don't perceive as naturally.

So you could almost think of it as a sort of discovery partner suggesting, hey, these are things you might look at. And I think maybe the most accessible version of that, because, you know, even for me, you know I'm not a chemist. But I think something that's maybe most accessible is if you look at what happened with recent advances in machine learning with respect to chess. So chess is another place where it's a pattern recognition game. And with this chess playing program, AlphaZero, look, it's been since

the 1980s that computers have beaten humans at chess, including Grand Masters.

But computers played in this very sort of rigid and tedious, I would say tedious manner, and new machine learning algorithms, not only can beat all the existing chess programs, but they play in a manner that a human chess expert, instead of calling tedious at wearing them down, would call inventive, exciting, something where the human wants to learn the new strategies and the new tactics that these new computer chess AIs are using. And that fundamental shift from computers, and this is back to the sort of opening comments I made, technology used to be very predictable. AI is moving to being unpredictable. Unpredictability comes with huge advantages in things like scientific discovery and new strategies for almost anything, not just chess, you know, as long as it's done in collaboration with humans. Unpredictability, completely open loop is a different issue. But it also raises new challenges because it's so different from other technologies. It's a great question. Thank you.

PETER COY: You know, Dan, I'm asking you sort of technical questions because some of the more societal questions a lot of people have weighed in on. You have a unique perspective. So I want to ask you another sort of technical question which is the difference that you lay out in the book between supervised learning, unsupervised learning, and reinforcement learning.

DANIEL HUTTENLOCHER: Yes, absolutely. And maybe one comment I should make is that the more technical material in the book, which is in some of the early chapters to try to give people some of the basis for reading about the broader societal impacts, and I encourage all of us who are now bombarded with lots of articles in the popular press about the societal impacts to learn as much as you have the stomach for about the technology because I think it helps you see through some of the aspects of the commentary on the societal impacts.

But with respect to different modes of machine learning, I think there are very good analogies to human learning. So supervised learning sort of conventionally used, and if you think about something like, you know, you want to recognize license plates so you can do automatic tolling, people driving on expressways or you want to have a self-driving vehicle be able to identify if there's a pedestrian or vehicle there or something. So those are done by a lot of data, as with all machine learning, but that data is labeled by experts. Right? So this is the actual license plate, here's a picture of it, or, you know, here are the cars, here are the pedestrians. And that's what's referred to as supervised learning. And you can imagine that doing that at very large scales becomes impossible because just getting humans to label all of this data is not really feasible. But supervised learning has played a very important role in machine learning and continues to, but it doesn't scale to the very largest machine learning problems when there's a lot of data. And, you know, that's not that different from, you know, if you think a lot of sort of

classroom learning. Right? You're being told, you know, this is a this, that is a that. A lot of grade school is a lot of sort of rote learning about how to classify things and so forth.

Unsupervised learning is where a learning system just looks for patterns on its own with no sort of instruction or feedback from any kind of expertise. And so that, you know, in a machine learning context, that generally means finding patterns where things sort of look similar to one another according to some notion of similarity that the system itself derives from looking at a lot of patterns. And, you know, people do a lot of unsupervised learning and I think it's one of the areas we don't understand that well in humans who are translating into __.

And then reinforcement learning is – and unsupervised learning plays a fairly big role in some aspects of training the large language models as does reinforcement learning. So reinforcement learning is where the AI system tries things out and it gets feedback from the world. And that's another thing that humans do a lot of. Right? You tinker with something. It breaks. You throw it out. You get a new one. Or historically, evolutionarily, right, you get eaten because you don't run fast enough from the predator. So reinforcement learning is something that, you know, both in an individual's life but also in an evolutionary sense is very important in human and inanimate development. And reinforcement learning also plays a very big role in these large language models that underlie things like ChatGPT. I think your mute went on, Peter, it looks like.

PETER COY: Sorry. I want to remind the audience that the questions are open. Just go into the chat function and type in your question. I'm going to ask one from Sophie Wong. She says, I'm an MIT alum and excited to be part of this event. Dean Huttenlocher, two questions. One, there are concerns that due to the capital-intensive nature in edge-cutting AI, research will increasingly come from deep pocket private institutions instead of from academia for the benefit of the public. How do you think this could be mitigated? And two, I love how the College of Computing is bringing together interdisciplinary expertise to advance CS research and applications. What do you view as the most promising intersections that the school is looking to foster more of? Thank you very much.

DANIEL HUTTENLOCHER: Great. Two great questions. So I think as we've seen with a lot of technology development over the last, really since modern technology, capital investment is a very important driver of scaling technologies to sort of national and global use. So this is a phenomenon that we've dealt with repeatedly in different ways and it's true here as well. And often society has a hard time dealing with these things. So if you look at the early days of the industrial revolution, it was very capital intensive. And if you look at sort of what evolved in terms of factory automation and industry and labor sort of relations, some of those have been very rocky, especially in the early days of the industrial revolution.

So these are things that go well beyond technological questions. There are things that, I think, require a shared set of understanding and values, require regulatory regimes that reflect and understand those shared values. That's something that worries me a little bit. Right now there's some pressure to regulate in some parts of the world without having articulated clearly what the aims are. But this is a place where I think there's a lot to do that's outside the technology and something that we've been very actively involved in at MIT and working with our colleagues in Congress and the Executive Branch and around the world.

But on the technological side, because I think another aspect of your question that's important is, you know, is the new development in this technological area going to be dominated by the corporate sector and sort of what's the role of more open research that academia tends to do. And I think universities are going to start to play a sort of different kind of role with respect to large language models than just developing the models and the base technology because it often is tens or hundreds of millions of dollars to train one of these models.

But, for example, these models have got to be made much, much, much more efficient, especially at what's called the inference stage. Right? There's a learning stage where we train the model and the inference stage where you use it. And, you know, if you just think about the environmental footprint of people typing all kinds of inane questions at

ChatGPT, it becomes significant. And so I think there's a lot of focus at universities on efficiency in this kind of technology. There's also a lot of focus at universities on sort of more specialized versions of this technology that aren't necessarily the sort of general-purpose things that you're seeing in industry right now.

So I think there will be a lot of driving of those fundamental AI technologies at universities still. But the sort of head-to-head competition which is who can build the biggest large language model is something that I think will be between industry players. But I also don't think that, I don't think that it's, so there I may differ from some of those industry players who seem to be claiming that the biggest models are going to produce magic. I think a lot of advances are going to come from places other than just the biggest models.

PETER COY: And then the second part of her question is about interdisciplinary expertise.

DANIEL HUTTENLOCHER: Yes, so this is phenomenally important in everything that's happening in machine learning these days and in the use of artificial intelligence, as some of Peter's questions touched on and things that we've talked about in the book. But the College of Computing at MIT is attacking this from a much broader perspective even than just AI. When you look at computing broadly construed, including artificial

intelligence and machine learning, it's really changing every aspect of daily life over the last 25 years and now it's accelerating.

And so the College of Computing really has a three-fold mission, one of which is, you know, to continue to push the forefront of computer science and artificial intelligence research. The second is looking at what we call the infusion of computing. And essentially every discipline, certainly every discipline at MIT – now MIT is more technically-oriented than many universities – and look here, it's including our Music Department, you can see our Arts Department and Media Studies and Writing. And it's not just the technology fields at MIT.

And this is something where the computing technology both holds the opportunity to transform these fields, but closer coupling between these fields and core research in computing is also a critical part of informing the development of computing technologies and the broader understanding that one gets from social sciences and humanities engagement in these kinds of technologies in addition to engagement of people in the sciences as well as in engineering.

So the first piece is kind of core computer science and artificial intelligence. The second piece is infusion of computing in the disciplines. And the third piece is the social and ethical responsibilities of computing where we're both doing that through faculty hiring

and through education. Two thousand MIT students a year now are taking computing classes that have material on social and ethical issues integrated in, whether it's in the problem sets, the projects, or the lecture materials. And these are activities we started about two and a half years ago so the fact that they're now up to the scale of 2,000 students a year and growing shows the appetite for this among our student body and our faculty. So a great question. Thank you.

PETER COY: Another question. I'm jumping around here. It came from Emily Spratt. She said, overall the trend in computing trends has been towards open access to everything. Given the recent AI innovation, do you think that there needs to be a change in attitude towards its open-source tech? Does it not threaten our national defense? Do you think that this attitude is now shifting on the level of computer scientists? And then secondly, what is your opinion of the Getty Images lawsuit?

DANIEL HUTTENLOCHER: So open-source computing technology, open-source computing software has been an extremely important advancement in computing technology. I think the AI models do raise some questions there. My belief is that the best path is still open source on these. But let me say what I mean. I think that because of the significant resources required in training these kinds of models, we have to be careful about open release of the actual models that get trained. But the underlying sort of basic, algorithms and software, I think it's very important for scientific advancement

that those things are released in an open fashion.

And I would draw an analogy to another place in the sciences where if you look in the biological sciences now, you know, most published papers that involve synthesis of some kind of molecule, you know, sequencing of something, require open release of the underlying, be it sequence data or molecular synthesis or whatever. I think it's extremely important to scientific advancement that we don't just have published papers, we also have the accompanying artifacts that are important to really advancing the science. And I think in all of those cases it's a tradeoff of looking at risks and rewards to advancement. But I don't think it's something that's just true in the computing domain.

And I think that science being done behind closed doors, I'm much more worried about than science that's being done in the open. You know, I think we understand much more about what's going on if things are being done in an open way. Even if some adversaries are doing some of their work in a closed manner, I think it helps us understand how the technology is really moving forward in ways that you don't get from just papers. So I'm still an advocate of open source.

I saw something pop up there that's very relevant so I'll just – about export controls for LLMs, so it's really the same question, so maybe I'll just...

PETER COY: Yes, jump on that and also if you have a thought about the Getty lawsuit too.

DANIEL HUTTENLOCHER: I don't really. The stuff that I'm really paying attention to right now much more is the policy issues around AI and there's so many.

PETER COY: How about the export controls?

DANIEL HUTTENLOCHER: Yes, so the export controls of LLMs, I would just point out how ineffective and actually sort of problematic to the technological advancement, the export controls on public key encryption work, back, sort of, what, 25, 30 years ago when the U.S. government tried that. I think that the sort of base technologies, it's just very hard to do that in any kind of effective way and that it just tends to get in the way of our actual commercial usage in our country and/or the scientific advances. Right? I mean, you know, people used to walk around with t-shirts that said this is a munition, which, you know, had the code on it for public key encryption.

I mean I think that software is an information asset and information assets, as we've seen in all kinds of places – look at government secrets – information assets are not easy things to keep from spreading if you also want to make them accessible for ready use. And I think that tradeoff comes down on the ready use side in my view.

PETER COY: I'm going to throw in one of my questions and it pertains to some of the later chapters in the book where things get kind of really scary. And I can tell this is where Henry Kissinger was weighing in more. And the idea is that as AI gets better and better it's harder to keep a human being in the loop because human beings, what do they really know? They're just seeing an output from the AI. How can they guess what the AI is saying? So technically there's a human in the loop but if the human being is just rubber-stamping what the AI says anyway, what have you really gained? On the other hand, if you take the human being out of the loop, that's scary too because then you're just surrendering to this black box. So what's the answer?

DANIEL HUTTENLOCHER: Yes, and this is saying something that at least I'm encouraged by the fact that many of our senior defense leaders are thinking very carefully about. And I think, I think this issue is now hopefully much more apparent to the general public because people are believing that, you know, these chatbots like ChatGPT are their friend, are their romantic partner, are all kinds of things. This is AI simulating interactions with humans. I think our tendency to trust technology and to anthropomorphize technology are both extremely important underpinnings of this issue about what does it really mean to have humans in the loop with AI.

And I think one of the things that I both view as an opportunity and as a challenge with AI is that our mode of interaction with these AI systems should always be one of

interrogation, questioning, pushing on the thing to understand what it is that it's conveying back to us. And, you know, I think taking anything that a chatbot says at face value I recommend strongly against.

PETER COY: Dan, can I just stop you there for a second because there are AIs that can be interrogated and you can ask the AI to explain its logic, but that's not all AI. There are some that they don't even understand themselves, how they came to say what they said. If you ask them, they could probably concoct retrospectively some justification but that's not really what happened. So again, how do you deal with that?

DANIEL HUTTENLOCHER: So, yes, and we won't lose the thread of the defense question at the end. So the first thing I want to point out is that when a person explains how they got to something, they often concoct an answer.

PETER COY: That's true.

DANIEL HUTTENLOCHER: Right? And we have all kinds of things that AI doesn't. Right? Like, you know, motivation and embarrassment. So we don't want to look stupid. AI doesn't even know what it means to look stupid. And so sometimes we concoct answers for those kinds of reasons. But I think you made an extremely good point about what's the nature of the interrogatory interaction. So sometimes systems are designed

to really sort of be able to try to show you the sort of underlying logic of how they got to where they got. Certainly for large language models, there is no logic in the ways that we understand logic. These things are about patterns of texts that have been sort of identified and synthesized by looking at huge amounts of text and that's extremely powerful, but it's not logic like we're used to logic.

And so, therefore, but I think what you can get is, you know, you keep sort of pushing on this and you start seeing that the answer shifts to some other answer. Right? Like sometimes you push on these things and they say, oh, okay, yes, some other thing that is either directly contradictory to what they just said or enough of a difference that you start saying, okay, this thing doesn't know anything about what it's talking about here. And sometimes you get actually a pretty consistent set of responses even as you sort of push on it from different directions. And so I think even though that's not an explanation, I think it is an interrogatory form of interaction where you start to build up some kind of sense that, okay, there's something underlying what it's telling me here that's consistent.

So I think what we really care about is in synthesizing this answer to something, that it sort of just picked something because that was sort of the highest probability thing out of a whole bunch of other things that completely contradicted. Right? Is this some sort of stable part of the space that we're in or an unstable one? And so I do think these

interrogatory sort of interactions can be very powerful even with chatbots.

And I think, just as another side point here, one of the things I'm very excited about in the educational domain, you know, if you think about sort of Socratic dialogue as the ability to really have an individualized back and forth discussion and sort of pulling and pushing on ideas, the level at which any society I know of invests in education, we don't have the resources to do that on an individual basis for students. And I do think that, you know, future versions of these large language model technologies will provide that. But students shouldn't interact with those things by believing, they should interact with them by learning how to question them and interrogate them and push on them and really it is the Socratic mode that's going to be extremely important there.

And I think that same mode coming back to how do you come to interact with these sorts of systems in an exigent situation, in a national defense context, it's similar. You need to really learn how to question. And then there's something that's in the book that I don't have an answer to and I don't think there is an answer to, which is that these technologies are also speeding everything up. And so if you need to decide really quickly and you know that this AI can get to a recommendation or a course of action much faster than human reason...

PETER COY: Nuclear attack for example.

DANIEL HUTTENLOCHER: How do you really, and this is where I just think, I do think there's a lot to be done here. You know, what are sort of modes of dialogue, modes of interrogation that give you confidence in the recommended course of action. Sometimes they will need to be done on very short order. They should never be followed, one of these systems, no matter how often it has seemed to be right in the past. And I think this is one of these areas where there needs to be a lot of work still. And, as I said, you know, any senior military leader I've talked to, you know, U.S. or NATO, they are deeply engaged in this.

PETER COY: Yes. Okay, we're coming towards the end, just a few more minutes here. If you have any questions that are urgent, get them in now. I'm going to ask a question from Carla Brite. She asks, do you think there is opportunity to leverage this technology to promote more equity within societies and between societies, e.g. Africa versus more advanced economies?

DANIEL HUTTENLOCHER: Yes, absolutely fantastic question. I'm just writing it down so I don't completely forget it halfway through because it had two parts. So I think equity within societies, I'll start there, I think that is extremely important, a huge opportunity and terribly misunderstood at the moment. I think some sloppiness in the ways that some early AI systems were trained and then rolled out very rapidly, which may be true of what we're seeing right now again also, some of these systems really sort of

exacerbated existing biases and inequities in society because they're trained on a bunch of data about human sort of reported activity.

But I think that there's a huge opportunity actually for these systems to improve, not just to get rid of the bias in these systems but for these systems to be less biased than individual human actors. There's a sort of distinction that one of my colleagues has made that I find particularly useful. Humans are fallible. Algorithms are fragile.

PETER COY: Explain that distinction.

DANIEL HUTTENLOCHER: Yes, so here's, and so if you think about algorithms, part of the fragility comes from the fact that once an algorithm is good, everybody uses it. Look at the amount of use of Google versus other search engines. So you get this sort of dominant effect like if some algorithm is good, people use it. Well, we know from biological systems that diversity is important. So this form of fragility is one that we've become dependent on a single system, and if it does something bad like, you know, exhibits racial bias or other forms of inequity, that's a problem because we're all using it.

But humans, individually they make different decisions but their decisions are very fallible and often very biased. And so the algorithms together with humans actually, I think, give us a way of, and this is back to the sort of human-algorithm neutral

interrogation I was saying before, the algorithms can double-check the people and double-check that the people are not exhibiting certain kinds of bias.

And, for example, you know, the people I know out there now are sort of making effective use of algorithms and things like resume screening and so forth, that's exactly the mode they're using. But they're using it in a way that sort of complements the human and can actually lead us to less biased outcomes so I'm very excited about that. But we still have to watch out for the fragility and the bias in these systems when they're stand-alone.

On across societies, and particularly less-resourced societies, this is another place where I think there's a huge amount of opportunity. In fact, another colleague of mine who I've spent some time talking with in Africa, has been pointing out that, you know, there's been a lot of concern in western countries about use of AI in medicine and very rightly. You know these are heavily regulated domains. Risk-reward is, you know, the risk is very high. But the risk-reward tradeoff is very different in a place where you have no access to a human doctor or very little access to a human doctor, little access to a human expert.

And so these are again places where I think that the ways that say, you know, western governments will approach some of this will be very appropriate to those cultures where

we spend a heck of a lot on healthcare, but not appropriate in countries where there is very low healthcare availability. And I think those are places where there's a lot of opportunity. Again, it needs to be driven, not by western imperialism and wealthy country imperialism but this is, you know, talking to people in these countries saying, you know, we would like access to some of these things to help move our medical care forward. So a great question, both parts. Thank you.

PETER COY: Alright, well, Dan, thank you very much. I'm going to turn it over to Barbara now. I hope you all enjoyed it.

DANIEL HUTTENLOCHER: Thank you all.

PRESIDENT BARBARA VAN ALLEN: Yes, this was just outstanding. I think we could have gone on for another hour frankly. And looking at the chat box, there are many more questions. Peter, thank you for weaving in some of our members' good questions, and just a great conversation.

I just want to give a quick look ahead for the Club. On May 9th, we're going to have our Club Chair and the President of the New York Fed, John Williams, for a Signature Luncheon. And that will be followed by two-member peer exchanges that you can see there on the screen – Current Trends for Family Offices on the 10th. On the 11th, The

Strategic Selective Decoupling of the U.S. Economy from China. On the 17th, we will have the Ukraine Ambassador to the United States, so we're excited to have that as a virtual conversation on May 17th. And, of course, May 23rd, we will be celebrating Henry Kissinger's 100th birthday, and he'll be in a conversation with Marie-Josee Kravis. So I look forward to that event. It's proving to be quite popular. On June 1st, we're going to have Sally Susman of Pfizer talking about her book in a conversation with Reshma Saujani on our board and CEO of Girls Who Code. And we're delighted to be able to have on the schedule Marc Rowan, the CEO of Apollo Global Management, which should be a great event. And then that will be followed with a virtual event June 29th, Karen Karniol-Tambour, the Co-Chief Investment Officer at Bridgewater. And I have to say we have a number of other events that are still being scheduled. So please continue to track our website.

And as we always like to do, I want to just take a moment to recognize those of our 361 members of the Centennial Society, any of those joining us today, as they continue to provide the financial backbone of support for the Club and our many programs. So again, thank you, Dan. Thank you, Peter. And thank you to everyone that attended today, and we look forward to seeing you again soon. Thank you.