[SQUEAKING]

[RUSTLING]

[CLICKING]

GARY GENSLER: On Monday, we talked about financial technology as a whole, and what we're going to do in these next four classes is really talk about three major technologies. The broad subject of artificial intelligence, machine learning, and deep learning, we'll talk about today, and next Monday, we'll move on and talk about the marketing channels, what I'll broadly call OpenAPI and some of those conversational agents and the relationship that is happening there. And then next Wednesday's class, I might be off. It's the following Monday's class, we'll talk about blockchain technology and cryptocurrency. So we're taking three big slices of technology, and then we're going to go into the sectors.

And so what I thought of about in structuring this class, as we've already discussed, is I take a broad view of what the subject of fintech is, that it's technologies of our time that are potentially materially affecting finance. So in any decade, any five year, it might be different technologies. It's technologies of a certain time that are materially affecting finance, and it's the broad-- any competitor can use fintech.

Now, I recognize that a lot of people use the subject or fintech, and they narrow it down, and they just use those terms for the disruptors, the startups. And there were really good questions of why I think of it more broadly. I just think that the incumbents are very much involved in this fintech wave, and to ignore that would be at your peril if you're starting a new business, if you're an entrepreneur.

And of course, if you're Jamie Dimon running JPMorgan Chase, or you're running a big bank in China or in Europe, you would be at your peril if you didn't think about it. And big tech, of course, has found their way in here. The Alibaba example's in China, but all around the globe, whether it's Brazil, whether it's India, whether it's-- your big tech has found their way in here.

And so I've also decided to sort of structure it around these three big thematic technologies, and then we'll start to dive into the sectors themselves and take a look at four or five sectors before we call it a day at the end of the half semester. So I'm about to pull up some slides, but, Romain, are there any broad questions?

ROMAIN: Nothing so far, Gary.

GARY GENSLER: OK, thank you. So just give me a second to make sure I pull up the right set of slides. I think that-- all right, so two readings today. I don't know if you were able to take a look at them. I hope you had the time to take a look at them.

And these were really just a chance to sort of grab hold of a broad discussion of what's going on. Again, I went to the Financial Stability Board. It's a little dated because it's 2017, but I thought the executive summaries and these various sections were helpful. And then a shorter medium post really on six examples, and today, we're going to talk about what is artificial intelligence, what is machine learning and deep learning, and then what are the eight or 10 major areas in finance that it's being used. And then next Monday, we're going to talk about a lot of the challenges and go a little deeper with regard to this.

Now, I said that I had study questions. We're going to see if we can get this to work, and I'm going to ask for volunteers to speak up. If not, Romain might just unmute everybody, and then we'll see who I can cold call on or something.

But who would like to take a crack? And remember, we're all pass emergency and no credit emergency, so this is just about trying to get the conversation going. But who would want to answer the question, what is artificial intelligence, machine learning, deep learning?

And you don't need to be a computer scientist, but these terms are really important if you're going to be an entrepreneur and do a startup, or if you're going to be in a big incumbent or big tech company, just to have this conceptual framework and understanding of artificial intelligence, machine learning, and deep learning. So I'm going to wait for Romain to either find somebody who's raised their blue hand, or, Romain, you get to help me cold call if you wish. But hopefully, somebody wants to just dive into this.

ROMAIN: I'm still waiting let's see who will be-- ah, we have our first volunteer of the day.

Thank you, Michael.

GARY GENSLER: Michael.

STUDENT: Sorry, I forgot to unmute. So my understanding-- artificial intelligence just is more of an over encompassing term, being-- just computer is kind of mimicking human behavior and thought, so that's more--

GARY GENSLER: Let's pause there. And a very good answer, so computers mimicking human behavior. And, Michael, just a sense-- do you have a sense of how-- when did this come about? Was this in the last five years, or was it a longer time ago that somebody came up with this term "artificial intelligence?"

STUDENT: More like the early to mid 1900s. It's been a while.

GARY GENSLER: So it's been a while. It's actually way back in the 1950s, "artificial intelligence--" the concept of a computer mimicking humans. In fact, at MIT, we have the Computer Science and Artificial Intelligence Lab. It's not a creature that was just invented in the 20-teens or 20-naughts. It goes back decades.

It was a merger of two earlier labs, but we've had an artificial intelligence lab at MIT for multiple decades. So who wants to say what machine learning is? Romain? I hand it off you to find another--

ROMAIN: Who will be the next volunteer for today?

STUDENT: I think the machine [INAUDIBLE], there's very limited or even no intervention with human, and learning are you solve problems step by step in a sequential manner. So that is like the machine solve problems step by step without much intervention by humans.

GARY GENSLER: So I think you're raising two points, is the machine solving a problem without intervention, which is good. And then you said briefly, learning. Do you want to say something more, or anybody want to say, what does it mean that the machine is learning?

ROMAIN: We have Rakan who raised his hand.

STUDENT: Yes, well, essentially, it means that you're feeding the machine data, and as you feed the machine data, the machine learns not to do specific tasks. And you get better results as you feed it more data and more and more.

GARY GENSLER: All right, so the concept of machine learning, again, is not that new. It was first written about in the 1980s and 1990s. It had a little bit of a wave-- you might call it a boomlet, a little bit of a hype cycle-- and then it sort of tamped down. But the conceptual framework is that the computer, machines, with data, are actually learning, that they actually-- whatever algorithms or decision making or pattern recognition that they have gets better.

So machine learning is a subset of artificial intelligence. Artificial intelligence is this concept of computers, a form of machines, mimicking human intelligence, but it could mimic it without learning. It could just replicate or automate what we are doing. Machine learning was a new concept-- it didn't take off. It didn't dramatically change our lives at first-- where, actually, the computers could adapt and change based upon analysis of the data, and that their algorithms and their pattern recognition could shift. Does anybody want to take a crack at deep learning?

ROMAIN: Anyone? Albert raised his hand.

STUDENT: So deep learning involves using large neural networks to perform machine learning, so it's sort of a subset of machine learning. But it can be very powerful, and most of the time, it just gets better and better as you feed it more data. Other machine learning algorithms tend to have sort of a plateau and can't improve no matter how much data you feed them.

GARY GENSLER: So important concept there in a phrase was "neural networks." Think of our brain's-- neural, neurology. It's about our brains. Early computer scientists started to think can we learn something from how the human brain works, which is in essence, a biological computer that takes electrical pulses and stores data, analyzes data, recognizes patterns.

Even all of us right now are recognizing voice and visual patterns. That's in our human brain. So when looking at the brain, the conceptual framework is could we build a network similar to the brain, and thus, using these words, neural networks.

In deep learning and machine learning, there's pattern recognition, but deep learning is a subset of machine learning that has multiple layers of connections. And in machine learning, you can take a base layer of data and try to find patterns, but deep learning finds patterns in the patterns. And you can think of it as putting it through layers.

Now, if you're in Sloan, and you're deeply involved in computer science, and you also enjoy the topic, you can go much further. But in this class, we're not trying to go there. The importance of deep learning is that it can find and extract patterns even better than machine learning, but it takes more computational power and often more data.

Deep learning is more an innovation of the 20-teens, and by 2011 and 2012, it really started to change things. And in the last eight years, we've seen dramatic advancements even in deep learning. It's a conceptual framework of taking a pool of data, looking at its patterns, and going a little higher. I'm going to talk about an example a little bit later in this discussion, but please bring me back to talk about deep learning in facial recognition, deep learning in autonomous vehicles and just thinking about it there, and then we'll pull it back to finance.

ROMAIN: Gary?

GARY GENSLER: Yes, questions?

ROMAIN: We have two questions-- one from Rosana, who's asking, what is the difference between representation learning and deep learning? And then we'll give the floor to Pablo.

GARY GENSLER: So very good question. Representation learning, you can think almost as in between machine learning and deep learning. If we said, this is almost like those Russian dolls.

The AI is the big vessel. Machine learning is a subset, and deep learning was a subset of machine learning. Representational learning is a subset, in this context, of machine learning, and it's basically extracting features.

So think of machine learning that's looking for patterns. It's extracting features out of data. So photo recognition, our standard Facebook photo recognition that might recognize Kelly versus Romain is extracting certain features of Kelly's face. I'm sorry if I'm picking on you Kelly. You just happen to be on my screen.

But it's extracting features, so some people call it feature learning or representational learning, but extracting features as opposed to specific data. And then deep learning is generally thought of as a subset even though there is some debate as to how you would categorize these. And you said there was another question?

STUDENT: Yeah, hi, Gary. This is Pablo. So I just wanted to verify, because my understanding is that, also, one of the big differences between machine learning and deep learning, or artificial intelligence in general besides deep learning, is that deep learning used unstructured data, whereas typically, for machine learning models, you need basically tables of data that you have organized and labeled and everything. And just the check whether that's the difference besides the additional complexity and multiple layers.

GARY GENSLER: So let me step back and share the important vocabulary, and again this is important vocabulary well beyond being a computer scientist. It's important vocabulary if you're running a business, and you're trying to get the most out of data and your data analytics team. So I'm going to assume we're talking as if you are now going to be in a C-suite, and you want to get the most out of your data analytics teams and so forth-- is this concept of structured data versus unstructured data, and then I'll go to that specific question. Data that we see all the time with our eyes, that we read, that comes into us, you can think of as unstructured sometimes because it doesn't have a label on it. But if it has a label on it, all of a sudden, people call it structured data.

So machine learning, conceptually, is that the machines are getting better at recognizing patterns. It's primarily about pattern recognition, that the machines are getting better recognizing patterns off of the data. And the question here is, is machine learning always structured learning? And structured learning means that the data is labeled, that you have a whole data set, and it's labeled.

An example of labeling I will give you, that we all live with in our daily lives-- how many people-- just use the blue hands if you wish. How many people have ever been asked in a computer, will you please look at this picture and tell us whether there are any traffic lights in the picture? We want to make sure you're not a robot. All right, so here, I'm just going to-- Romain, just cold call somebody that put their hand up. I want to ask them a question.

ROMAIN: Sophia, please?

GARY GENSLER: All right, Sophia, are you are unmuted?

STUDENT: Yes I'm on.

GARY GENSLER: All right, so, Sophia, when you go in, and you tell the computer that you're not a machine, why do you think they ask you that question?

STUDENT: To make sure that there aren't any bots who are trying to take advantage of any service that the machine is trying to offer.

GARY GENSLER: Really good answer. You're correct, but you're not completely correct. What's the other reason that they're asking you whether that's a traffic light or not?

STUDENT: And also to collect data so that they can use that data for labeling for future purposes as well.

GARY GENSLER: So they are labeling data. They're using Sophia, if I might say, as free labor. Sophia is training the data that Google or whomever is putting together there.

And thank you, Sofia. You're labeling that data so that autonomous vehicles will work better in the future. You're also frankly labeling data to put millions of people out of jobs as truckers, but I don't want you to get sort of wrapped up into those sort of social and public policy debates, but that's what's happening.

So back to the earlier question. Data can be labeled by us. An earlier form of labeling was labeling what is an A, what is a B, what is a C, what are all the letters of the alphabet so that our postal services now can use a form of machine learning to read all of our written scratch on envelopes. If we address an envelope, it can all be read by computers rather than humans.

To your earlier question-- I went a long way around this, but machine learning can be both unstructured and structured. The question was, is machine learning always labeled? Is machine learning always structured?

And the answer's no. Machine learning can also be unstructured and unlabeled. Deep learning can be both labeled, which is structured, or unlabeled.

Some of the economics and some of the computer science are worthwhile to understand. Labeled data can be trained faster. Label data can, in many regards, lower your error rates faster and have a certain-- extract correlations better, but it comes with a cost. You need to label the data, and so there's some tradeoff of getting Sophia to label data or other humans to label data versus unlabeled data.

Think about radiology in the practice of medicine, and looking at body scans or mammograms or any form of radiology to identify whether there's an anomaly. There's something that needs to have further investigation to see whether it's a tumor or not. Radiology is dramatically changing in the last three or five years based upon machine learning and deep learning. Remember, deep learning just means there's multiple layers of pattern recognition in these neural networks.

Labeled radiology, labeled mammograms, or labeled MRIs will train the machines faster, but unstructured, unlabeled data can also be used. You need bigger data sets. So I went I went off a little bit, but I hope that that's helpful. Other questions, Romain?

ROMAIN: Yes, we have one from Victor.

GARY GENSLER: Please.

STUDENT: Hi, professor. I just wanted to double click on the-- take a step back in the initial definitions between machine learning and deep learning, because we were discussing that deep learning had the feature that it keeps learning despite the data growing exponentially. It doesn't plateau. But I don't fully understood the differentiation between two concepts beyond that.

GARY So I didn't disagree or agree with that comment, and I apologize, I can't remember

GENSLER: who said that deep learning keeps to grow exponentially. I think both machine learning and deep learning-- both machine learning and deep learning learn from data, and this word "learn" should be explored a little bit more. What machine learning and deep learning can do is extract correlations.

I hope nearly everybody in this class has taken some form of statistics at some point in time. You might have hated statistics, but we all took it, and some of us took more advanced statistics where you use linear algebra and the like. But just thinking about a standard regression analysis-- a standard regression analysis finds a pattern, generally a linear pattern, or a quadratic pattern if you move on.

Machine learning and deep learning find pattern, and they're really remarkable tools to extract correlations. And one of the features of both machine learning and deep learning is they look at error rates, particularly versus data sets that have been labeled. And traditionally, what you do is you have a big data set-- maybe it's millions of pieces of data that you're training on, and you take a random sample of it and put it to the side, a random sample on the side that you label. And then you compare what comes out of the machine learning with the test data on the side and see what's the error rate. And this labeled set on this side might say, these are men, these are women, this is a stoplight, this is a traffic light, whatever the labeled data is, and you see the predictive model, what's the error rate.

Both machine learning and deep learning continue to do quite well. And I apologize I cannot remember who said it earlier, which student said, deep learning continues to grow further than machine learning. It can, but I wouldn't accept that machine learning can't get better and lower error rates. Now, once you get down to very low error rates, that's another circumstance altogether.

The difference between deep learning and machine learning is that-- I'm going to use photo recognition software. If you put a photograph into a computer, where does it start? Does anybody-- what does it see at the very beginning, it's base level of data?

STUDENT: The pixel?

GARY GENSLER: What did I hear?

STUDENT: I said, the pixels.

GARY GENSLER: Pixels. So the only thing a computer can read is pixels. It has to start with the pixels and build up, and the next layer at most-- and again, I once pretended to know something about computer science. But I programmed in Fortran and APL years ago, and that was before many of you were born. But I guess I used to know how to program something.

But if you read the pixels, the next layer up to find a pattern in the pixels is just small changes of shade, and then you can think of the next layer up from those little-- you can see edges. So the computer has to sort of go through layers from the pixels up to this is a traffic light versus a stop sign, and so deep learning interposes multiple layers of pattern recognition. Some would say that you need many layers, and then other research shows that, no, once you get to about three layers, there's less and less return on this. So let me sort of-- unless, Romain, is there other questions, or can I--

ROMAIN: No, we're all good.

GARY GENSLER: All right, so again, this is just a broad thing, and we're not going to spend as much time. But what's natural language processing? Does anybody want to-- just what are these words broadly mean?

ROMAIN: Any volunteers?

GARY GENSLER: I'll volunteer, then. So natural language processing is just simply taking human language, natural language, and processing it down to computer language, all the way down to machine readable code, or going the other direction. So you can almost think of it as input and output to the computer, or we have many, many languages represented on this call right here with 99 participants. You can think about it as translating French to German and German to French, but instead it's natural language, what we do, down to the computer.

And this is really important in terms of user interface and user experiences, and we'll get to that a little bit more. And then we're going to talk a lot about which sectors in financial services are being most affected at this point in time. So I won't call on the class right now because I want to keep moving to go forward.

So we're in to talk about the financial world and fintech, and then-- oops, I didn't change this. This slide will-- I'll shift, because that's from the other day. So we looked at this slide the other day, and it just helps us-- what is AI machine learning?

Extracting useful patterns from the data, using neural networks that we talked about, optimizing to lower error rates. Optimizing so that you actually say with 99% or 99 1/2% of the time, this is a traffic light, this is a stop sign. Actually optimizing.

There's lots of programs that you can use. Google has TensorFlow, and-- I don't know. Have many of you ever used TensorFlow? I don't know all of your backgrounds, but any show blue hands? I'm not going to call on you to describe it. I'm just kind of curious if there's many people that have used TensorFlow or not.

ROMAIN: How about we go with Devin?

GARY GENSLER: So Devin's actually used it. I wasn't going to pick on them up, but if you wanted to say anything about it, Devin.

STUDENT: Yeah, I can very briefly. So it essentially gives like a plug-and-play method to do machine learning and build neural nets in Python. You don't necessarily have to have a full understanding of how under the hood works. You can just add bits as you want, take bits away as you want, and it speeds up the whole process.

GARY GENSLER: And so what's important about that is, just as somebody that came of age in the 1970s and '80s didn't have to learn how to computer code all the way down to machine readable code, they could learn how to use-- by the 1990s, C++, or C And C++, and later, Python. And many people in this class know how to use Python. In the machine learning area, there's been plug and play programs like TensorFlow, where you don't need to actually know how to build the data and things like that.

Most importantly, and this is if you're thinking about being a data analyst or actually building a business around it, it's the data and the questions you train on the data. And most studies have shown, as of 2019, that 90%, 95% of the cost of data analytics is what some people might call cleaning up the data, making sure the data's well labeled. We talked about structured versus unstructured data earlier.

Really important about that labeling, the cost. If you have 1,000 people working in a machine learning shop, or 500 people or five, it is quite likely that a big bulk of their time is standardizing the data, so to speak cleaning up the data, making sure that the fields are filled, and ensuring that you can then train on this data, meaning it's labeled, or enough of it's labeled. And then thinking about what the questions you're really trying to achieve, what you're trying to extract.

Why is it happening now? This is off of Lex Freeman's slide again, but why? Because the hardware, the tools, the analytics-- a lot has shifted in the last five or eight years.

To give you a sense of what it's being used for, all of these, we know. We know already. It's dramatically changing our lives. When we're sitting at home, sheltering at a home, and you're thinking about the next movie, and you're on Netflix, Netflix is telling us what they think the next thing we should watch. That's training off of not just the knowledge of what each of us has been watching, but it's about what others are watching.

It's the Postal Service that no longer has to have a human reading the text, our scrawl on the envelope. It's Facebook with the facial recognition programs and so forth, and autonomous vehicles that are now being tested on the roads but are very likely part of our future. Now, will they be rolled out in a dramatic way in five years, or will it be 15 years? But I would feel comfortable that we will have autonomous vehicles on the road sometime at least by the 2030, but maybe others would be more optimistic.

So it's changing a lot in many, many fields. The question is now, how is it shifting this field of finance? Why do I put it at the center of what we're doing here?

So we talked about this. This was just my little attempt, and so now you have a slide that does what we chatted about before. One important thing also is happening is, in finance-- [CLEARING HIS THROAT] excuse me-- people are grabbing alternative data, using alternative data. I said earlier that the most important questions are what's the good data? So then we think about, in finance, what type of data do we want to grab?

Data analytics in finance goes back centuries. That is not new. The Medicis, when they had to figure out to whom to land in Renaissance era had to figure out who is a good credit or not. And two data scientists from Stanford started a company called Fair Isaac-- those were their last names, Fair and Isaac-- and that became the FICO company, the Fair Isaac Company.

So the data analytics in finance and the consumer side has certainly been around since the 1950s and 1960s, but where we are now is to say what is the additional types of data that we might take, and not just banking and checking and so forth? But Alibaba can look at a company, look at a company very closely, and do a full cash flow underwriting. Alibaba, because they have AliPay, can see what that small business is spending and what that small business is receiving.

Amazon Prime can't quite do it as much, but they can do a bit of it as well. And even Toast, which is a fintech company in the restaurant business until this Corona shut down, Toast could see a lot about what's happening restaurant by restaurant in their cash flows. They had the revenue side more than the expenditure side. Alibaba, much better data sets than Toast, but I wouldn't put at rest a Toast started earlier in the last year to do credit extension to the restaurant business.

Now, do you need deep learning and machine learning to do that? Not necessarily. You can still use plain old regression and linear statistics, linear regression analysis, but machine learning and deep learning help you go further.

And then there's, of course, everything about our usage, our browser history, our email receipts, and so forth. If we look at China, they've stood up a broader social credit system, and in that system, they can tap into data about users in many different platforms. Romain, do I see-- are you waving at me? Is there a question?

ROMAIN: No, I'm not. Sorry for that.

GARY GENSLER: That's all right. So natural language processing, I mentioned. I just want to say a few more words about it.

Think of it as computers input and output interpretation. This sort of going from German to French, or going from computer language to human language and back again. That's this important back and forth, and so it's natural language understanding, meaning a computer understands something, and also natural language generation.

So it can be audio, image, video, any form of communication-- even a gesture. This hand wave can be interpreted-- if not now in 2020, within a few years will be interpreted. A movement of your face will be interpreted as well.

And so how it's being used is really quite interesting, but we all know about chat bots and voice assistance already. That's shifting our worlds. Yeah?

ROMAIN: We have a question from Nadia.

GARY GENSLER: Nadia, please.

STUDENT: I have one question related to the chat bots. What kind of factors do you think will encourage people to use chat bots? Because now, I do think people prefer to talk to a person rather than chat bots.

GARY GENSLER: Well, I think, Nadia, we might still prefer to talk to a person, but there's a certain efficiency in-- and I'm just going to stay in finance for a minute. But there's a certain efficiency that financial service firms find that they can use chat bots instead of putting a human on the phone. So even when you and I call up to a Bank of America, and we want to check in on something on our credit cards, we're put through a various series of push one if you want this, push two if you want that, and we're pushing buttons. That's not high technology, by the way, but that's an efficiency that Bank of America has interposed into the system instead of having a call center of humans. And so if they can move from a cost center of humans to an automated call center of chat bots, they can provide services at a lower cost and to more people.

Now, you and I might still want a human on the other side, but business is interposing an automation, and that automation means, often, quicker response time. So many of us, you go into a website today, and there's a little bot window that comes up. And the first thing that comes up on so many websites-- and this is true if it's a financial site. It's true if it's a commercial website where you're buying something online. It's probably true of dating websites, that there's some little bot window that comes up and says, can we help you. That's not a human, but it does give us greater service. It gives us an immediate recognition somebody's answering a question. Nadia, do I sense you'd prefer not to have the chat bots interposed? Is Nadia still there?

STUDENT: Oh, yeah. Yeah, because I do think sometimes, chat bots, we ask a question, but their answer is not really related to our questions.

GARY GENSLER: So you would prefer a human because you think the human will interpret your question and be able to answer it better?

STUDENT: Yes.

GARY GENSLER: So if the chat bot could answer as well as Kelly could answer-- again, I'm sorry, Kelly. You're on my screen. But if the chat bot could answer as well as Kelly or Camillo or others in this class could answer, you'd be all right with that?

STUDENT: Yeah, it's faster.

GARY GENSLER: See, if it's faster, and it can answer as well. So those are some of the commercial challenges.

ROMAIN: I think Ivy would like to contribute as well.

GARY GENSLER: Sure.

STUDENT: Yeah, I just wanted to offer a little bit of some of the consumer studies that we did when I was working for a startup where we were building chat bot. And interesting enough, I think most people are actually pretty-- and we did a pretty large survey, and most people were pretty open to the idea of working with a chat bot because I think that's become so pervasive. But then there's this idea, they want to know that it is a chat bot, that the company is very transparent about that, because people change their behavior when they are speaking to the chat bot or some kind of virtual assistant, as long as they know, just to build that trust. And also, I guess we try to be more explicit in our wording, both verbally as well as written texts.

And then secondly, I think because-- I guess I wanted to pose this as a question, too. When I think about chat bots and things like that, the technology is not necessarily there, or it takes a long time. And so I'm just curious what your thoughts are in terms of-- for me, I see it as like AI gets us 80% of the way there, but we need the human touch 20% of the way there. And so I actually see a lot of companies either having that human at the end of the-- you ultimately still need a human at the end of the day. So I mean, I just wanted to explore that a little.

GARY GENSLER: I think what Nadia and Ivy are raising, and I'm sure we're all grappling with this. We are living in a very exciting time, and I'm not talking about this corona crisis. That's a different type of challenge.

But we're living in an exciting time where we can automate a lot of things that humans have done. We've automated so many things that humans have done for centuries, but we're now automating this interface, through chat bots and conversational interfaces, voice assistance. I would dare say that, of the 90-plus people in this class now, that most of us, if not all of us, at some point in time, have used Siri. I mean, if we're driving along a highway, and we're supposed to be hands free, we might talk and start up an app or something legally, legally.

And there's a lot of automation that's going on. How many of us have called to arrange a reservation at a restaurant, and we're not quite sure if we're talking to a human or a conversational agent? But I think that what Ivy's saying is that there might always need to be a human somewhere there.

I don't know, Ivy, if that's correct. That's where we are in 2020. Let's think about autonomous vehicles.

Right now, we're not comfortable enough. The manufacturers aren't comfortable enough. The computer scientist aren't comfortable enough. The regulators aren't comfortable enough. The public's not comfortable enough to have autonomous vehicles on the road with no humans whatsoever, but that's not really necessarily where we'll be in 2030.

Or take radiology. Right now, at least in advanced economies like in Europe and the US and elsewhere, in advanced economies, we say we still want a doctor's eyes on a radiologist report. So the mammogram might be read by some artificial intelligence machine learning trained data, but we still have a human. But is that really the tradeoff we'll make in a few years? And is it the right tradeoff to be made in less developed countries, where they don't have the resources to have the doctors?

And now, we even look in the middle of this crisis, the corona crisis, if-- this is sort of God willing. If the Baidus of China and the Googles of the US and others sharp analytic AI shops come up with a way to extract patterns and develop some recognition as to who's most vulnerable, are we going to rely on that, or are we going to say a human has to also interpret it and be involved in it? And I don't know.

So I think we're at an exciting time where we're automating more and more. I do agree with you, Ivy, there's always going to be a role for humans. I'm not terribly worried that we'll all be put out of a job.

200 years ago, our ancestors, all of our ancestors, were, by and large, working on farms. That's the economies. And we have found other things to fill those roles and those needs. I think we'll still have the humans, but not in every task.

So let me go through a little bit-- so the Financial Stability Board, this is their definitions I'm going to pass on, but this was in that paper that you all read about what big data and machine learning was. When I show this page to computer scientists, when I show it to colleagues of mine at MIT from the College of Computing, they look at it, and they say, jeez, that's funny that a bunch of financial treasury secretaries and central bankers and their staffs define big data machine learning this way. So I partly put it up because this is kind of what the regulators define as to what it is. Machine learning may be defined as a method of designing sequence of actions to solve a problem, known as algorithms, and so forth. Computer scientists would name it a little bit differently.

So I said to you the other day that I think of financial technology as history is building on these things, but machine learning and deep learning is at this top level. In the customer interface, it's the chat box we were just talking about, and on the risk management side, it's extracting patterns to make better risk decisions. So it's in these two broad fields, I sort of think of it is the customer interface and then lowering risk and extracting patterns. And sometimes, it's not just lowering risk. It's enhancing returns.

And so I was going to go through and chat about each of these eight areas, and not all of the areas we're going to talk about are as robust. Asset management, right now-- asset management from hedge funds all the way to the BlackRock and Fidelity are exploring the use of machine learning and AI. By and large, most high frequency trading shops, most hedge funds today, most asset managers today, are not using much machine learning and AI.

I view that as an opportunity in the 2020s. I view that as a real possibility of a shift, a very significant shift. But where is it being used so far?

So BlackRock has been announcing and saying that they're already using BlackRock as one of the world's, if not the world's, largest asset manager. Before this corona crisis, probably $6 or $7 trillion of assets. It's, of course, a lower number now. And BlackRock and others have been saying we're using machine learning to actually listen to all of the audio files, all of the audio files of the major companies when they announce their quarterly earnings. And they're also putting in news articles, digital news or articles about those announcements, and also feeding in some of the actual financial statements that are released.

And that takes a little bit of natural language processing. You need some form of taking the audio files, taking the written files, and interpreting that. But with that data they're looking for sentiment. They're trying to interpret the sentiments and see if the stock price are moving based on all of that data.

Now, that's BlackRock with $6 or $7 trillion of assets. They're deeply resourced, fidelity, and so forth. But if you go down the value chain, if you go down to smaller asset managers, they're not doing a lot, I would say, yet.

But there are hedge funds that are specifically saying, we are data analytic hedge funds. We want to move a little further. We want to try to use this machine learning because it's a better way to extract correlations, and that's what it's looking for. It's pattern recognition.

I think that you're going to see more and more high frequency trading shops and hedge funds exploring this. But one conversation I had in the last couple of months with the high frequency trading shop-- and it was a shop that had about 100 employees. It was not big, but it was big enough. It was certainly making money in those days. I don't know how it's doing now.

But they said, look, we feel pretty good about what we do, and when we look for algorithms, when we look-- our algorithmic trading doesn't need all of that expenditure and all that resource intensive thing of machine learning and cleaning up the data and finding the patterns. And in fact, we think that is not flexible enough yet for us. We're looking for short term opportunities, and we think that regression analysis and our classic linear and correlation analyses are enough at this moment. What's going to come in 2023 is a different thing, but what they're doing now, [INAUDIBLE] not really needed. Questions about asset management just before I go to a couple other fields?

ROMAIN: Now is the time to raise your hand if you have any. I don't see any, Gary.

GARY GENSLER: So cost--

STUDENT: Sorry, Gary, I raised it in the last second, so Romain didn't see it.

GARY GENSLER: [INAUDIBLE].

STUDENT: My question is-- so I fully understand their claims that this is not the claims of the asset manager, you were talking of high speed trading. I don't see why machine learning or any algorithm besides linear regression doesn't have enough flexibility to actually replicate what they're doing, but with a bit more of accuracy or computing performance, et cetera, because in the end, my understanding is that you need the same-- except when you go to deep learning. But you need the same data sources that you currently have, and the only thing that you're going to do is, instead of just having linear models predicting the different traits that you want to do, you have other models that can find alternative basically trends and et cetera. So I don't see the limitation there of using it.

GARY GENSLER: So I think you raise a really good point. We're in a period of transition. I personally don't think that machine learning and deep learning is the answer to all these pattern recognition challenges in finance. But you will find I'm more to the sort of center maximalist than center minimalist, and those of you that know me, that when we talk about blockchain technology, you'll find I'm more to the center minimalist side.

But what I mean by that is I think that the pattern recognition out of deep learning, machine learning-- and by pattern recognition I mean the ability to extract, with remarkable ability, correlations and then create certain decision sets based on those correlations-- is better than classic linear algebra, classic regression analysis. But it comes with a cost, and that's the tradeoff in 2020. Maybe in a handful of years, it'll be less cost, but I'm going to use an example, which is not about asset management, but it's about lending.

I had this conversation just a few weeks ago with the CEO of a major peer to peer lending company, and as this is being recorded, I'm just going to maybe not say his name. And I said, do you use machine learning and deep learning to do your credit decisions, all your credit decisions? And he said, yes, we use a lot of alternative data. We run it into the decision sets, and we see what patterns emerge.

And I said, so you extend credit then based on that? He said, well, not exactly. He said, what we do is, we look for patterns, and then when we find them, we then just use classic algebra and linear flagging when we actually extend the credit.

So we use it to look for patterns, but then we sort of use the traditional way, and I ask why. He said, well, there's two reasons. It's less costly than running the whole thing all the time, and two, they can explain it better to regulators, and they can explain it better to the public.

And at least in consumer finance, there were laws passed about 50 years ago-- in the US, it's called the Fair Credit Reporting Act. These laws were passed and said if you deny somebody credit you have to be able to explain why you're denying them credit. And of course, there are other laws in many other countries similar to that, and there are also laws about avoiding biases. In our country, we call it the Equal Credit Opportunity Act or ECOA, and in other countries, similar things about avoiding biases for gender and race and background and the like.

So I'm just saying that, whether it's asset management, whether it's consumer credit, there's some tradeoffs of these new data analytic tools. And those tradeoffs, I think, in the next handful of years, will keep tipping in the way towards using deep learning. But they don't come-- they're not cost free is what I'm saying.

ROMAIN: José has his hand up.

GARY GENSLER: Please.

STUDENT: So you talked a bit about how this is impacting the high frequency trading shops. Is there something similar happening on more like value investing long term or in the hedge funds? So I heard some of them are using satellite image to predict the number of cars-- to see the number of cars [INAUDIBLE] and predict the sales for that year, things like that. But I don't know, do you see a lot of headroom in this area?

GARY GENSLER: I think there is headroom where I'm saying, I've mentioned two areas, but I think you're right. There's a third area. The two areas are sentiment analysis, just sentiment analysis around either an individual company or the overall markets and seeing what the sentiment, the mood, the sense of the crowd is off of words, images, or the like. And then also the high frequency traders and so forth just looking for the patterns in the short term trading.

You're talking about more on the broader value orientation, and I absolutely share your view that we'll see more of that develop. I haven't heard a lot of it, but you're right about sector after sector that you might be able to analyze. But again, it has to have some ability to extract a pattern better than classic linear regressions in analysis.

So let me just try to hit a couple more of these because we're going to run out of time, but we are going to talk Monday more about these. We talked about call centers, chat bots, robo-advising, and so forth, so I think you've got that. A lot of that is not just efficiency, but it's also inclusion. You can cover many, many more people by automating some of these tasks.

It's just reality, it's the tradeoff of efficiency and inclusion, that you cover far more people. You can also be more targeted. You can be more targeted with advice and so forth as these are automated. It comes with the tradeoff that Ivy and Nadia were talking about earlier.

Credit and insurance-- this is the concept of basically how you allocate or extend or price, either alone or price insurance. To date, the insurance companies and insurance underwriting are starting to grapple with this, starting to move on, and we'll talk about some fintech companies in this space. I think that insurance companies have been a little bit slower to do it than, let's say, the credit card companies, but the allocation-- I think this will be dramatically shifting.

I think if you look at what's going on, again, in consumer credit and small business credit in China, through WeChat Pay and AliPay, they're much further along than we are, frankly, here in the US. We're still largely reliant on a 30 or 40-year-old architecture around the Fair Isaac Company, the FICO scores that are used in about 30 countries around the globe. These are still quite limited, but FICO itself, FICO itself rolls out new versions of FICO every few years.

I think they're rolling out FICO 10.0 this summer, or were before the corona crisis. I think, if you look at the end of the 2020s, either FICO will not exist at all, or FICO 14.0 or 15.0 will look a lot more like a machine learning, deep learning type of model. What it is right now is pretty rudimentary compared to what it could be in 5 or 10 years.

This is an area that's being used a lot-- fraud detection and prevention. The credit card companies, if you look at Cap 1 and Bank of America, Discover, American Express, they're deeply now using machine learning tools for fraud detection. Many of us probably remember just a year or two ago, you would still call up your credit card company, and you would say I'm traveling to France, I'm traveling to Italy, wherever I'm traveling. I want to put a flag on there that I might be using my cell phone-- using my credit card in one of these countries.

Well, most companies now don't ask you to put a travel alert on it. Now, part of that is because they know where we're traveling because we're walking around with these location devices. The banks no longer need us to call them to tell them we're going to be in Paris because they know we're in Paris-- this location tracking company device. But in addition to that location device, they are also using machine learning to do fraud detection and prevention in the credit card space.

Similarly, they're using it to track and try to comply with laws called anti money laundering. These two areas, fraud detection and any money laundering, which I might call compliance broadly, are two of the areas most developed right now in 2020. That doesn't mean they'll be the most developed later. I think a lot more will happen in the underwriting space. A lot more will happen in the asset management space.

Robotic process automation-- I want to just pause for a minute. Does anybody have a sense of what these three words together mean-- robotic process automation? Romain, you get to see if there's any blue hands up.

ROMAIN: Andrea.

STUDENT: Hi. So robotic process automation is very simple. You have a lot of manual processes or manual work done in, for example, back office of the banks. And the idea here is, instead of people doing that, low skilled work or workforce, you can actually teach robots or the algorithms with the PC to do it instead of you. So for example, it can be anything as simple as just going through the forms and copying or overwriting and rewriting some of the words or parts of the forms to some other place.

GARY GENSLER: Right. So robotic process automation can be as simple as you're giving your permission to a startup company, a fintech company, to access your bank account. And one of us gives a fintech company-- maybe Credit Karma.

We give Credit Karma the right to go in and look at a bank account of ours. Credit Karma might not have permission from Bank of America go in, but they have my-- I'm permissioned them. I've given them my password and my user ID, and they go in, and they want to automate.

Credit Karma wants to automate. They don't want to have a human actually have to type that all in. It can be as simple as just automating inserting the user name and the password and the like, but then you go further than it can navigate the web page. It can navigate and click the right buttons and get the right data and so forth.

So robotic process automation can be helping a startup company say Gary Gensler, we want you as a client of Credit Karma. We, Credit Karma, will figure out how to interface with Bank of America and with Chase and the others. But also, the banks are using robotic process automation to automate so much of their both back office and their data entry. And one form, just to say, is many of us now feel very comfortable to deposit a check from our cell phone.

So you can take your cell phone, take a picture of a check, and somehow, that gives an instruction, a digital instruction, to move money. Well, part of that's natural language processing, that the cell phone could take a picture, read all that data, put it into computer language, and actually move a digital. Part of that is robotic process automation. Romain, was there a question? I think I saw some flashing chat rooms.

ROMAIN: Yes, sorry. You have 15 minutes left.

GARY GENSLER: Oh, OK. Then was there a question or no?

ROMAIN: No question at this point.

GARY GENSLER: So now, trading-- trading is an area I spent a lot of time with Goldman Sachs. And in that trading of the day, we were automating everything we could automate, and this is in the 1990s. And ever since, anything is-- a trading floor in 2020 looks very different than a trading floor in the 1990s in terms of the day to day trading, and this is trading at the center of the markets, the platforms themselves, and of course, the high frequency traders on the other end of the market is basically just like asset management.

What patterns can you see? Now, this is less about value investing. This is the patterns right in the nanoseconds and milliseconds and so forth.

I'll make one note in terms of trading, which is not related to machine learning, but just related to the corona crisis that we are all living in. I have an overall belief that this coronavirus crisis will accentuate trends that we've already seen. In industry after industry, if we're locked down for two, three, or four months, or God forbid, for 18 or 24 months-- if we're locked down for that long, we're going to find new ways to engage with each other in economic activity and social activity and the like.

And we've already had some trends, deep trends-- we talked about them Monday-- that we're unlikely to be using many paper money and coinage money. Three months from now, nearly 70% or 80% the world will have forgotten how to use paper money and coinage. In fact, it will even be viewed as a disease delivery device. It might be dirty. It might be something we don't want to use because it could be a problem.

Well, let me talk about trading for a second. The New York Stock Exchange and the world's largest stock exchanges are now electronic. They could have done that two years ago.

When the Intercontinental Exchange, which is a big public company, bought the New York Stock Exchange a handful of years ago, five or so years ago, Jeff Sprecher, the entrepreneur who started the Intercontinental Exchange in 1998 or '99, he's always been an entrepreneur, an innovator, to do electronic trading. They could have taken the New York Stock Exchange fully electronic, but guess what? That's what happened in the last three weeks.

So after we get out of this lockdown period, will we bring back the floor of these London and New York and Shanghai and Mumbai and so forth? Will we bring back the floors? I think quite possibly not. Not sure, but there's a lot that's shifting on, I think.

Natural language processing-- we talked about in customer service, process automation, and sentiment analysis. These are sort of the slices that I think about in these fields. So I was curious how many people have ever used Siri? Probably almost every hand would go up. But how many people have ever used Erica?

ROMAIN: So perhaps we have Shaheryar who'd like to share his experience. Sorry for your name-- mispronouncing it.

STUDENT: Yeah, so it's essentially like Siri, but you actually-- it's a product of Bank of America, and you can use it to check your spending habits. You can also use it to, if you need things with regards to check depositing, or if you want to know something, it can do that. But currently, I believe it's not as refined as Siri, and I still think there is a lot room over there for improvement.

GARY GENSLER: Why do you think it is that Erica-- and JP Morgan has one as well. I can't remember what her name is. They all do seem to be mostly female voices, if I'm not mistaken. But why is it that Erica and the like, a virtual assistant, as they're called in finance, aren't as developed as the Siris and Alexas, do you think?

STUDENT: Sorry, can you repeat the question?

GARY GENSLER: Anybody can answer. Why is it that the finance virtual assistants like Erica are not yet as fully developed as the Home and other ones like Siri and Alexa?

STUDENT: I believe some of-- I think it's got something to do with the fact that the number of users for financial assistance is way lower as compared to Alexa, Siri, or Alexa or whatever. So I believe that is something which may explain it, the disparity between these two kinds of assistants.

GARY GENSLER: Yeah. Any other thoughts on that, one why--

ROMAIN: Nikhil has a different take, and that we have Laira.

STUDENT: Probably along similar lines, banks so far, the interactions probably have been in person or over phones, and they weren't used to processing data and processing requests. I think they have a smaller data set to go through and understand what problem customer questions are. And that's probably a limiting factor versus, say, Siri, they have much more data on everyday users. I think that's probably the biggest differentiator.

GARY GENSLER: Yeah, and was there another comment, Romain, you said? Sure.

ROMAIN: So I think Laira had her hand up, but I think she withdrew it. So perhaps we can hear from Brian.

GARY GENSLER: OK.

STUDENT: So in addition to the data, I think there's also a human capital element. It's possible that Apple has better human capital capabilities than do these financial institutions, so it's harnessing that data as well.

GARY GENSLER: Yeah. So what we've just talked about was data, human capital, experience-- all true. Also, Erica and the financial firms only started more recently, and so forth.

But the voice recognition programs and then taking that data that an Apple has or their competitors in big tech around the globe-- because it's not just here in the US-- is really remarkable now. Even to the extent that I don't know how many people use earbuds, but if you look closely at the user agreement on the airbuds that you use, it says-- and if it's an Apple, if I'm mistaken, the Apple lawyers will chase me. But the user agreement says that they can listen to that to help you to make sure that there's not a drop between your earbuds and your cell phone.

They are picking up vast amounts of data, vast amounts of data from our text messaging. If you look at something like Google, they're picking it up from Google Chrome, Google Maps, Gmail. Multiple places that they can pick up, and we talked about this conceptually, big tech versus big finance versus startups in this triangle of competitive landscape. And why I wanted to sort of close on Bank of America Erica and this discussion of Erica versus Siri and Alexa is big tech using Google, just as an example, has this remarkable network that they're layering activities.

Remember, we said data network activities-- that's the Bank of International Settlement way to put it, and what a perfect example to show. Google has Gmail. It has Maps. It has Google Chrome. It has Android, the operating system.

So all of these different ways to build their network, and they layer activities on top of it, and then vast amounts of data come in and the human capital that was mentioned at the end there. And they have more experience to move it forward. Apple, similarly. Baidu and Alibaba in China and so forth, similarly. If I were a CEO of a big incumbent, yes, I would be very focused on the fintech startups, but I tell you-- I'd be looking at big tech in a way that their advantages are really significant, very significant advantages. Romain, I see some hands up maybe.

ROMAIN: We now have Laira who has her hand up.

STUDENT: Yeah, I'm just curious to know what do you think? So currently, we know that tech has kind of-- or AI has kind of emerged to the financial and payment space in the form of virtual assistants, but what do you anticipate the next step would be in terms of this integration of technology and AI into the payment space? What next after the virtual assistants?

GARY GENSLER: We're going to have a whole class on payments specifically, but I think that what's happening in the payment space is we've seen specialized payment service providers. Of course, we've seen a lot of the competition starting with PayPal in 1998. This is not a new space for disruption. But what we've seen more recently is, in the retail space, whether it's companies I've mentioned earlier, like Toast, that got into one vertical, one slice within payments, which was restaurants-- they can provide a better product for that slice. And then can collect back to AI enough robust data within that slice-- this is using Toast as an example-- that they can provide better software, better hardware, and also less risk loans.

Basically, as Toast started to provide lending to restaurants within that space, built upon the payment. So it's the marriage between the user experience providing the users-- in that case, the restaurants in the payment space-- but providing the users in that space something that the generalized platform-- a bank payment system is generalized. It's multi sector. It's a general product, and Toast was able to say, no, we can provide something that just restaurants-- there might be something a little unique about the restaurant business that we can provide software and hardware.

In their case, it was tablets. They were providing tablets for the servers to walk around and take the orders. They could integrate the menu right into the payment app. So there was something a little bit unique about that.

But then based on that, they get a bunch of data, and that data helps them with, I would say, underwriting decisions based on-- it doesn't have to be machine learning. But it's enhanced data analytics because of machine learning. So I hope that helps. I think on the conversational agents and the virtual assistants, what we're seeing in the payment space, because that was your question, is we're moving from card authorized payments to mobile app QR codes. Then the QR codes is not based upon virtual assistants, but it's an interesting question whether we'll get to some voice authenticated payments.

There are a lot of uses of voice authentication already. Vanguard and many other asset managers, where you can have your brokerage accounts, you can call in and get voice authenticated before you can do a trade. And that voice authentication is just like other forms of authentication, but it's not perfect as of this moment.