MIT CSAIL Alliances | Andrew_Lo_Project_4

Welcome to MIT'S Computer Science and Artificial Intelligence Labs Alliances podcast. I'm Kara Miller.

[MUSIC PLAYING]

On today's show, sometimes underneath tough technical problems, there are even tougher human problems.

When markets go down by more than 25%, something like 4% to 5% of households will completely liquidate their equity holdings and put it into cash.

Andrew Lo, a professor of finance at the MIT Sloan School of Management and a principal investigator at CSAIL says, regular investors need access to low cost, personalized financial advice. And he's trying to figure out if artificial intelligence can help.

Can large language models dispense financial advice that is accurate enough and relevant enough to a given individual that makes the risk worthwhile compared to similar risks in humans?

And Lo says, AI financial advisors could not only help consumers, they could also help human financial advisors.

Imagine an AI that can answer questions, deal with relatively simple transactions for that 80% of the financial advisors book of business, allowing the human financial advisor to focus much more on the 20% that really require his or her particular human talents.

That's all coming up in just a minute. But first, in so many of our conversations here, we talk about generative AI. It's one of the game changing technologies of the last few years. So how can you understand it better? How can it help your business?

MIT, CSAIL, and MIT xPRO are offering an online course this fall designed to answer those questions. If you're interested in more info, just email us podcast@csail.mit.edu. Listeners to the podcast get 10% off the course. So again, the email is podcast@csail.mit.edu.

And if you're looking for another great podcast, our friends at Harvard Business Review have one you might like. Their Tentpole Show HBR IdeaCast, every Tuesday, editors at HBR bring you insights from the world's best business and leadership experts to help you manage up, manage your business, and manage yourself. Listen to HBR IdeaCast for free wherever you get your podcasts. HBR IdeaCast, it's all you need to lead.

Back in 2013, Andrew Lo was asked by a mutual fund to give a talk to potential investors. The talk was about rebalancing your portfolio. And at the end of it, a former student approached him.

Who came up to me and said, Professor Lo, I just want to let you know I really appreciated your comments. In the face of 2008, I liquidated all of my holdings and moved them into cash, and therefore I missed out on a lot of the losses that happened in 2009.

Pretty quickly, Lo realized that what seemed like quite positive news getting out of the stock market way before it bottomed out wasn't quite as positive as it had initially seemed.

And I said, that sounds great. So, why are you here at this seminar? And he said, well, I just want to know, do you think it's time for me to get back into the market? This is five years after the financial crisis.

So a few numbers for you here. In the fall of 2007, the Dow closed above 14,000. About a year and a half later, so now we're at the spring of 2009, the Dow had fallen to under 7,000 from over 14,000 to under 7,000. It was a really, really rough 18 months.

But after the spring of 2009, the stock market started to come roaring back. And in 2013, which is when Lo gave that speech and talked with his former student, the Dow hit fresh highs, ultimately closing above 16,000. That is a pattern that has been repeated again and again throughout the history of the stock market.

And he basically missed out on the rebound that occurred a few months after the big losses that happened in 2008. That's the problem that we're dealing with and you have in some cases, some very sophisticated investors who nonetheless are reacting emotionally, not rationally to these losses.

Lo has seen that emotional reaction on a broad scale, and it worries him a lot.

We've done a study where using data from a large brokerage firm in the US, about 600,000 accounts across various different household demographics. When markets go down by more than 25%, something like 4% to 5% of households will completely liquidate their equity holdings and put it into cash.

5% he says, may not sound like a lot of people, but these are people's retirement accounts, money that they need to invest prudently to make sure it grows over time. And it's important to note, Lo says, that those who liquidated often selfidentified as having sophisticated financial knowledge.

So I think maybe knowledge is a dangerous thing, particularly when you're not using that knowledge, but really reacting emotionally to financial losses. And we all react emotionally at one point or another, question is, are we prepared to be able to deal with those overwhelming emotions?

Lo says that liquidation is not a terrible thing if there are underlying structural problems in the market. But not getting back in, which is often what happens is terrible for wealth building. So how can you address this problem and help people more effectively manage their money? And could AI offer a solution? Lo says, he first started to consider that possibility when he studied the data of the people being impacted.

It's generally not the ultra high net worth investors that are engaging in these freak out moments. It's typically investors that have a small amount of money that they're trying to manage to get to their retirement. And often times it tends to be older individuals that may be a bit more sensitive about their wealth, maybe more concerned about the news, maybe not quite as informed about macroeconomic events, and are therefore frightened about the headlines that they're seeing when stock markets do come down.

And so once I saw that, I realized this is a really serious problem because if you're an ultra high net worth investor and you miss out on a couple of years of a bull market, that's not going to affect your retirement one way or the other. But if you're a librarian or a pharmacist that ends up missing out on two or three years of a reversal from a serious market decline, that's going to have a major impact on the quality of life of your retirement.

And so at that point, I started thinking about ways of using technology to support the decision making that all of us are challenged with and how to deal with the emotions that can sometimes overwhelm the rational process by which we all would like to make our investment decisions.

It sounds also like an upside of AI here is that it's both inexpensive to consult versus a personal financial advisor and that it's available all the time because people get worried at 2:00 in the morning.

Exactly right. I mean, imagine being able to consult with a financial advisor anytime, day or night, whenever the mood strikes you, whenever the fear starts building up and you want to have some reassurance. Imagine having a financial advisor that has read every possible article about finance as well as all of the various different news sources.

And imagine an advisor that can track financial markets and is better at numbers than any one of us can be and understands exactly what the implications are for any decision that we're likely to make on what happens to our retirement 10, 20, 30 years from now. That's the power of AI and the potential for doing a lot of good with some really interesting technology.

Now, I mean, as you know AI makes mistakes and so I can imagine people thinking really AI that's maybe the last place I would go to get financial advice. So how do you instill confidence and in fact, how much confidence should you instill?

Well, absolutely. I think that that's a very legitimate concern. And five years ago, I would not be suggesting that we could possibly use AI for any type of financial advice, maybe for some financial calculations, maybe for supporting some human financial advisor. But really, a lot has changed in the last five years.

In fact, a lot has changed in the last 18 months with generative AI making incredible strides, both because of software but also because of hardware and data. All three areas have had tremendous progress within the last few years, and that has now changed the landscape for what is possible.

So I think that we are now in a period of time where with the right use of AI technology, we might be able to solve this age old problem of being able to dispense reliable, accurate, customizable, and trusted advice in a number of different domains, including financial investing.

I know you've asked ChatGPT before, what you should do if you lose a whole bunch of money. Let's say, you lose a quarter of your life savings and what does it say back to you?

When ChatGPT first came out, I gave it a try like everybody else I knew. And one of the questions that I fed it was, what should I do if I just lost 25% of my retirement portfolio? And it gave me predictable answers like you should reevaluate your risk tolerance, you should consider all of your different investments and make sure that you're comfortable with it, and so on.

But then it gave some advice that I didn't think was so good, like you should engage in dollar cost averaging. Now dollar cost averaging is a perfectly legitimate investment process that a number of people can benefit from, but not everybody can benefit from it, and certainly not everybody in the circumstance of having just lost 25% of their retirement assets.

In fact, a number of investors, even if they could benefit from it, the fact that you would be asking them to invest more money as the stock market goes down and they lost a bunch just simply ignores their emotional state. So clearly, this is not appropriate advice for the given individual that I thought would need it. So I left it at that.

And a few months later, an updated version of ChatGPT came along. And when I tried the new and improved version, the results were eerie because now the advice that it generated was advice that I thought was spot on in every respect. And it's advice that I think a reasonably competent financial advisor would be giving to their clients.

And at that moment, I began to see that this might actually be possible to create a large language model targeting individuals who need financial advice and either can't afford it or aren't the proper client type that investment houses are going to want to serve because they simply don't have enough assets. But those are the ones that need it most in my view. So this really has opened up a whole new vista of possibilities.

And in your view, the issue of hallucinations or is this mistakes, is not a deal breaker, it sounds like.

Well, it certainly has to be addressed. And let me add that hallucinations are not just a phenomenon of large language models, humans hallucinate as well. And so financial advice is one of these areas where there's a minimum level of competency that you, of course, want every financial advisor to have. But there are many other qualifications that are necessary, and not all financial advisors have those qualifications.

Even though they are perfectly licensed to practice, they may not be able to provide the customization that we would expect of the very best financial advisors. So the question is now a more nuanced one which is, can large language models dispense financial advice that is accurate enough and relevant enough to a given individual that makes the risk worthwhile compared to similar risks in humans?

And I think that this is a question that needs to be answered over the course of the next few years, that's what my coauthors and I are focusing on right now.

Now, if you have a financial advisor who's a person, generally that person knows a lot or a good amount about you. They have a sense of how much money you have, how that money is allocated, where you are in your work. Are you 45 or are you 65?

How much road is ahead of you in terms of making money and what maybe you've probably answered questions about, what are my expenses month to month, year to year. Is it important, do you think, if AI is to step into this space for the AI to know something about the 100,000aire versus the 5 millionaire?

Absolutely, I think that's critical. It's one of the three key aspects of any advice that we humans would benefit from. The three areas that we've identified for these challenges are, number one, domain specific knowledge. And I would argue that right now large language models like ChatGPT have a fair bit of that, if not all of it.

Second, they need to be able to tailor and customize the advice to fit a particular client profile. And third and most importantly and probably most difficult of all, they need to develop a relationship of trust with human clients. And that sounds really complicated and challenging, but in the context of financial advice, we think we have an approach to address it.

But to answer your question, certainly that second piece of customization is critical, and I don't think we're as far away from that as many other people think. In other words, I think that with the given demographic information that a typical client provides to their human financial advisors, I think using that data and using the domain specific knowledge that large language models can easily be endowed with, I think we're almost there at being able to provide the appropriate tailored financial advice for a given individual.

Now, when companies think about integrating AI, depends on what business that company is in, but there's often this hesitancy to release their data or let's say, a large company built the model that's not them, where's this data going to live? Who's going to be able to access it, whatever? How concerned should an individual be about that same thing because to tell people a lot about you're my financial life, you've got to tell them a lot about private things?

I couldn't agree with you more. I think that that's a huge concern, not just with financial data, but with all of our data. And I feel that policymakers need to step ahead of this and really think about legislation that protects our data and essentially gives consumers' rights that they currently don't have to say no to vendors that are going to provide us a service and use our data along the way.

Because right now, if we think about medical data, we do have laws like HIPAA that protect the use of that data and require certain security provisions. We need HIPAA for all of our data, including financial data.

When you think about where this is all headed, do you imagine that in a few years that people will be able to turn to an AI financial advisor and that will be potentially commonplace?

Well, in my vision, there are two things that will happen. One, we will be able to allow human financial advisors to greatly leverage AI so that they can manage a much larger book of business than they're doing right now, simply because 80% of financial advice is relatively simple and standard and can be dispensed without a lot of human talent.

Obviously, humans have to be properly trained and they are and they've been certified financial planners, financial advisors. There are all sorts of licenses that one has to get in order to provide that financial advice. But overall, 80% of it is relatively straightforward.

It's the 20% that is very, very difficult to automate because there you need to have not only the knowledge and trust of the client, but you need to have the human creativity to think about potential solutions that are appropriate for the emotional state of a client, not just their financial state.

So imagine an AI that allows a human financial advisor to quickly deal with the 80% by having the most incredible automated assistant that can answer questions, deal with relatively simple transactions for that 80% of the financial advisors book of business, allowing the human financial advisor to focus much more on the 20% that really require his or her particular human talents.

So that's an example of humans and AI collaborating and leveraging their respective abilities to improve outcomes for those clients. That's one vision. The second vision is that there are a number of underserved potential clients that are simply not wealthy enough to capture the attention and interest of financial institutions.

We talk about the mass affluent as the typical target of the largest brokerage firms, and of course, the ultra high net worth are already being catered to by a number of different service providers. But below the mass affluent level, you've got individuals who really need financial advice, and it's not very complicated in their situations.

They, of course, may have some personal challenges like medical expenses or child care and so on that have to be addressed. But those are individuals that I think could be dramatically helped by generative AI tools. And so I think that being able to provide free or very, very low cost automated financial advice for those individuals who can't afford financial advisors, that's one portion that could really be effective using these tools.

And then the second is to be able to leverage financial advisors that want to expand their business but don't have the resources. This automated assistant could be dramatically transformative for those situations.

So can you give me an example? I mean, it can be a real one, I don't know if you've spoken with financial advisors, but when you talk about that 80-20 breakdown of many things are typical, they happen over and over, and an AI assistant could potentially help you out with them. And then 20%, it sounds like are a little more this is an Andrew issue like let's get a human in here to talk about it. What things are in that 20% where you'd want a human in there talking to you about it?

Well, for now and obviously all of this could change as AI becomes more and more sophisticated, but for now, imagine an individual who is an entrepreneur and his company or her company is about to engage in an IPO, and the individual needs some advice on how to dispose of their assets in a way that allows them the freedom to deal with some of their life expenses while at the same time still allowing them to capture the upside of the shares that they own from their company.

That's a relatively complicated problem that depends on many different factors, including the health of the individual, the size of their household, the bequest motives for the next generation, and tax laws, and the likelihood that those tax laws will change, not to mention the macroeconomic conditions that they're making this decision in.

I believe that that's a situation where I think you'd want humans involved because the complexity of the various different possibilities are far too great for typical generative AI models to be able to sort through, at least at this point.

So when you think about the other 80% of issues that people encounter, is it possible that those issues could be addressed by following typical generic general investing principles? Or do you really need to turn to AI here and have it play a role?

Well, those questions can be answered typically by reading a number of investment textbooks.

Not something everybody does.

Exactly. How likely is it that you're going to get somebody who's working 40, 50, 60 hours a week and doing overtime and thinking about paying rent and dealing with various other stresses in life, how are they going to be positioned to read those textbooks and be able to come up with the answers? That's one of the reasons why we need financial advisors.

And in some cases, even though they know the answers, that's still not easy for them to implement because they need a coach, they need some support, they need a bit of a nudge. And that's where AI can provide that type of benefit if it is situated to understand the psychology of a given investor, the emotional state of that investor and their particular objectives, constraints, and ultimate life goals.

We've certainly seen this in other domains, but I wonder how you think about the question of liability. I mean, I don't know how it works right now with financial advisors who I'm sure are sometimes they've got to be wrong occasionally. But how does it work if AI says here's what I do and it doesn't really pan out?

That's a great question, and I think that question applies not just to financial advice, but any advice and any automation. For example, driverless cars, what happens when a driverless car kills a pedestrian? Who's responsible? So I think we need to work out those issues as a society.

And until and unless we do, we're not going to see tremendous progress in the application of AI to these very sensitive contexts, including financial advice, legal advice, medical advice, and so on. But we've been focusing on financial advice for a while now and we think we have an answer on how to address the issue. First of all, there are already laws in place that address this issue for human financial advisors, and it has to do with something called fiduciary duty.

So this is a legal term. And what fiduciary duty means is that there are certain special relationships in the eyes of the law that require individuals to put their own interests second to their client's interests, and that's what a fiduciary is. When you hire a financial advisor, that's a fiduciary relationship where the advisor is a fiduciary, meaning that she or he is a trusted provider of advice and therefore must put your interests ahead of his or hers.

And if that relationship is violated, if it turns out that financial advisor ends up engaging in business practices, that puts their interests ahead of yours, then they can be sued for civil and criminal penalties. So the same applies to an AI. And the question that we've been working on at MIT is, is it possible to create a piece of software that will satisfy the legal definition of a fiduciary in the eyes of financial regulators?

We think the answer is yes, but we're not there yet. We're actually currently engaging in research to figure out just how to structure it. The good news is that at least in financial markets, there are all guidelines as to how one needs to discharge their fiduciary duties. There are also lots of case law of situations where fiduciary duties are violated.

So using the combination of those guidelines, the codes of ethics as well as the voluminous body of knowledge that we've developed over many, many decades of various inappropriate financial behavior, we believe that we can come up with a large language model that will actually pass that fiduciary duty test.

That's fascinating to me because when you're talking about a person, obviously it's not OK if they don't tell you that on the side, they're getting a kickback whenever you buy this type of mutual fund. That would not be OK. But it's so funny to think of making sure that the AI knows that on the side they're not getting anything. I mean, is that what you're trying to do, trying to make sure that the AI is not recommending anything that is in its own best interest, which is hard for me to even wrap my head around, what is in its best interest?

Well, let me give you a few examples, because I think examples really help to highlight this issue. So imagine if a piece of AI is being employed by a brokerage firm.

And as part of the instructions of the AI by the firm's IT department, we program it so that it will give financial advice subject to the constraint that the advice has to benefit the brokerage firm, either by generating sales commissions or by other benefits that are directly against the benefits to the consumer. We can certainly do that, it's just a simple piece of software programming.

But that's an example where we will have violated the fiduciary duty of that relationship. Now imagine a different scenario where we program the AI to focus on the client's objectives first and foremost and put our interests, the interests of the brokerage firm second. We can certainly do that too.

With one single line of code, we can basically instruct the AI to behave in a manner that is consistent with its fiduciary duty versus something that would typically violate that duty. And in a way that's just so much cleaner, isn't it, than instructing a human financial advisor to behave him or herself and make sure that they don't focus on their end of year bonus or when they're making a recommendation or a decision.

So I think that this is really where we're excited by the whole vista of different possibilities. We think that using AI in these contexts could actually produce a better solution for consumers while at the same time generating more business for the users of these technologies and not destroying jobs, but actually creating more jobs and more value for all parties concerned.

What do you hear from financial advisors? What do they say?

So we've actually met with a number of financial advisors at the very outset of this project because obviously they have to be part of the solution. They have to understand what we're trying to do and they have to contribute to this process, otherwise, it's going to be a solution that they will not embrace. And what they started out saying was, we don't want to have any part of this if it's going to take away our jobs.

And I get that, make sense. But as we described to them what our goals were and what the power of AI can be for their business and for expanding what they're doing and doing it better, they started getting really excited about this possibility.

So what we're doing now is engaging a number of financial advisors to work with us to do some A/B testing, actual clinical trials, where we have subjects that will be using human financial advisors as well as others that will be using AI financial advisors and comparing and contrasting those experiences and trying to design the software so that it actually develops this trusted financial relationship that we hope will exist with every financial advisor and their clients.

So the financial advisors that we're working with right now are tremendously excited and they're also incredibly helpful in telling us where some of the pitfalls as well as the opportunities are in using these technologies to help them do a better job.

If you had to take a shot at what the horizon is, the time frame for when you could imagine that financial advisors could start using this tool, what do you think?

Well, I've been wrong before in the optimistic side of things. I thought that we would have had this AI financial advisor a decade ago. But what's happened over the last couple of years makes me wonder whether or not my forecasts are going to be wrong the other direction, and I'll be too conservative.

I would have said that maybe within the next three to five years, we're going to see AI financial advisors being used routinely, both by human financial advisors looking for their assistance as well as human clients that aren't being served by the typical financial institutions. So three to five years was my initial guess, but what I've seen over the last six months tells me that the speed of innovation in AI has just been picking up month after month.

And I wonder now whether or not it might be possible within even a year or so before some of these tools can be ready, simply because the power of both the hardware and the software is growing, it seems exponentially month by month. So I don't know, I'm very excited about this set of opportunities. And I think we're going to see some really remarkable announcements of new hardware and software platforms in the next few months and beyond it, I think, is anybody's guess as to what technologies will come out of that really exciting milieu.

And regulators are plugged in enough to make it happen potentially in that time frame.

I think so, certainly in the financial realm. The SEC has made some significant progress in upgrading its technological capabilities, hiring very sophisticated people and really paying attention to what's going on in the technological side. So I'm optimistic that although regulations in general will lag innovation as they should, you don't want to be regulating prospectively. You want to regulate only after an innovation is being adopted.

But I think rather than being three or four steps behind, I think the regulators are now only one or two steps behind, which is exactly where they want to be. So I'm hopeful that they will be there to be able to address some of these issues. The real challenge, as we talked about before, is to get policy makers to think about data privacy laws much more proactively than they have. I think that's a major issue that I'm hoping the new president, whoever he or she may be, will take on these challenges directly.

Andrew Lo's professor of finance, Director of the Laboratory for Financial Engineering at the Sloan School of Management, he's also a principal investigator at CSAIL. Andrew, thank you so much. I really appreciate it.

Thank you Kara. It's been a pleasure being with you.

[MUSIC PLAYING]

And before we go here, a reminder that CSAILs Gen AI course is starting soon. It's a technology that is reinventing all jobs. So learn how to use, apply, and strategize with this new course, driving innovation with generative AI created by MIT CSAIL and MIT xPRO. If you're interested in more info, just email us podcast@csail.mit.edu. Listeners to the podcast get 10% off the course.

Again, the email is podcast@csail.mit.edu. And if you're looking for another great podcast, check out Me, Myself, and AI. In each episode, expert hosts and researchers talk with AI leaders from organizations like NASA, Upwork, GitHub, and meta to explore how organizations achieve success with generative AI and what challenges and ethical considerations they face along the way.

Listen to Me, Myself, and AI wherever you stream podcasts. I'm Kara Miller. Our podcast is produced by Matt Purdy with help from Audrey Woods. Join us again next time and stay ahead of the curve.

[MUSIC PLAYING]