Numis

www.numis.com

Numis/MMC AI Conference Transcript

24/10/2017

Table of Contents

Part 1

Garri Jones: We have Al-Noor from the Pru, Clare, fresh off the Today program from Essentia Analytics. I didn't mess up that name, did I?

Clare FlynnLevy: No.

Garri Jones: Thomas from Man Group AHL. I'll ask them to introduce themselves very quickly and their early experiences with AI and then we'll just go into two or three questions from me and maybe one or two questions from the audience. Great.

Al-Noor Ramji: I've mostly been in software and serving customers, so AI is an obvious thing for us [00:00:30] because the organization doesn't have infinite memory and to prevent our customers running out of money before they want to, ideal.

Clare FlynnLevy: I run a company called Essentia Analytics, which uses technology to help human fund managers make measurably better decisions. I used to be a fund manager. In fact, Garri was one of my brokers back in the day. Got into this because [00:01:00] I was very frustrated with the extent to which technology was not helping me do my job. It was helping me execute trades and for a really long time, that's all it really was doing. What I wanted was a data-driven feedback loop that could just help me focus my energy on doing more of what I was good at and less of what I wasn't good at. That's what we do at Essentia today.

Thomas Flury: Man AHL is a systematic investment manager, so that means we build statistical models [00:01:30] for financial markets. We've been around for 30 years, so it's in our DNA to try and use statistical tools and statistical models to predict financial markets from that point of view, it's been a very gradual transition into adopting more and more of those AI tools. For me personally, a few years ago I was trying to predict or build a model and thought, "Oh, this is really hard to do this idea." I read about this algorithm that would solve my problem and, "Let's give it a try." I tried it, and it worked really well. Since then, [00:02:00] I want to do more of that.

Garri Jones: Thomas, you said as part of your intro, we're going to transition here, how far through, from Man AHL's perspective, are we on this journey? Are we right at the beginning, or have we been at it for a few years now?

Thomas Flury: From our point of view, we have been at it for a few years. We have been doing machine learning research for a few years already. We have a three-year plus live trading track record of trading AI-based signals for our clients. How much further can we [00:02:30] go? Well, I hope we can go a lot further. To me, it still feels a bit we're at the pioneering front, where it's not fully clear how far can we go. I'm very excited. There's still a lot of ideas and new tools that we can bring in from the academic world or from the non-financial world into the modeling of financial markets. I think it's also in connection with the other big buzzword, which is big data. I think these two AI, machine learning and big data, interact and play quite a lot. There, I think, is also [00:03:00] a lot of open questions in how far can we push our models or the insight into markets by tapping those data sources, by using machine learning tools.

Part 1

	how far can we push our models or the insight into markets by tapping those data sources, by using machine learning tools.
Garri Jones:	Clare, the answer from the fund management, perhaps more the active management side, is probably a little bit earlier, right? We've only just started. What's your experiences?
Clare FlynnLevy:	Well, so everyone's aware of the topic in active fund management and really, what most people have done is hire a bunch of PhD's to sit in a [00:03:30] room and experiment, which is cool, but not particularly helping the business today. It's firms like mine, which are using AI in that augmented intelligence definition of the word, so not necessarily building an autonomous thing that's going to replace a fund manager, but rather creating tools that can be useful today to help a human do the thing that he's already doing better. [00:04:00] That is starting to catch on. We've seen just this year alone, since the start of the year, Black Rock made an announcement about diverting energy from active into systematic and that freaked some people out and so all of the sudden, we started getting a lot of calls and we're having a great old time right now. It is practical application.
Garri Jones:	Can you give us an example? If you're a fund manager, one of your clients, what are they doing now that they perhaps weren't doing [00:04:30] a year or two years ago?
Clare FlynnLevy:	Well, so historically ... Everybody has a history of all the trades that they've ever done, right, and what their holdings have been every day going back in history, but they haven't really done much in the way of analysis of that data and they haven't been able to detect patterns and behavior around anything beyond maybe performance attribution style stuff: stop picking versus asset allocation decision making. Whereas what we're [00:05:00] doing is saying, "Okay, let's look at the actual decisions that you're making: a picking decision, an entry timing decision, a decision about how big to go and how quick to get there and so on and so forth. Let's look at that as a skill and say across each fund manager, 'Which ones are you good at and in what circumstances,' not so that we can embarrass you or make you feel bad or fire you, but so that we can help you improve like an athlete who's detected that he's [00:05:30] really good at this part of the process but not that one." That is something that the industry can grasp at the moment. They've already got the data. The need to improve is there. The pressure is on and there's a generation of fund managers coming up who are open to the idea of data-driven feedback beyond just, "Here's what your risk exposures are."
Garri Jones:	Got it. Al-Noor, from that wider financial services perspective, are we still at the starting gate, [00:06:00] or what's your thought?

what's the very big picture look like, as far out as you can see? Where are we going?

Al-Noor Ramji: Well, it's hard to say, but I think it's very important that boards don't treat this lightly. I would say, "Test and learn." What we're finding and what we think would happen in 7-10 years, firstly the cost [00:09:00] structure will change dramatically, but also the ability to do many more things than we do today. For example, our job is really to protect the health and wealth of a customer but really try and optimize those two, so not just the Fitbit things, but can we make sure you are healthier for longer? What we call ... Can we prevent first, which would be the greatest thing? If not, can we postpone it? If we can't, then [00:09:30] protection at the end as opposed to today, the insurance industry is, "Let's protect," and frankly, my kids and maybe you look like the age of my kids out here, probably don't understand what insurance is. People don't understand what savings is, really, or mortgage products. How do we use AI to make these things very seamless and almost don't have to make the decision? In 7-10 years, I think we'll automate a lot of this so you're saving without knowing you're saving.

[00:10:00] You're being protected without knowing that, so you get routed to your doctor. There's a company later on who's going to speak. Can we direct you to, and if needed, or just tell you to take an aspirin and stop winge-ing. Those things will appear completely normal, I think. It's really giving us a complete cover but without losing sight of individuals. I think it's going to be very specific to you. 7-10 [00:10:30] years, there's a company in China, for example, already can do your, simulate your whole person in a digital form and conduct pharmaceutical experiments in digital form. Anyway, I'll stop here.

Garri Jones: Got it. Clare, five years from now, active fund manager?

Clare FlynnLevy: Well, I think that there will be fewer active fund managers, for sure, but there will be better ones. I think that it's not just in the fund management industry that my optimist view of what happens [00:11:00] in five years is that people who are in the service business will be able to do a much better job of serving because of AI. It's not that the computer will do everything for us. There might be instances where that's appropriate, like somebody who doesn't know they need insurance, and that sort of thing. The picture that I'm focused on is about ... It speaks to the point that was in the report, the AI [00:11:30] report that these guys have brought up about the convergence between service and product and the fact that there won't be services that aren't technology-driven and there won't be products that don't have services attached to them.

There's a human interface that will become a lot more powerful because it's backed up by technology that's teeing up the important facts, the important decisions that need making, and so on. That, to me as a consumer, is exciting because I look at things like Google. One day, I'm using Google [00:12:00] Maps. I park my car and Google Maps says, "Here's where you parked." Oh, my god. Thank god, right? I used to take a picture. This is embarrassing, but I used to take a picture of where I parked because I can never remember. Somebody figured out that I'm not the only one that has this problem and decided that that's going to be part of Google Maps. It knows you stopped your car. It surmises that you've parked and it pins a little thing there. That's just made my life so much easier in a small way and I think those little [00:12:30] incremental changes will make it possible for all of us to do our jobs much better.

Garri Jones: Thomas, specifically on machine learning and the big data piece, which is exponential, exponential over five years gets very big. Where are we heading?

Thomas Flury: First of all, I think there will be a lot more uptake, and I think if you do not change your way and adopt and try out those new algorithms, I think you will be left behind. I think there are these new data sources that [00:13:00] potentially can add a lot of value and if do not use those data sources, you just will underperform. I think potentially in these big data, data sources and also algorithms, I would expect a lot of them can become a bit more commoditized. The way you sign up to your Bloomberg feed today, you will sign up to this alternative data feeds, potentially.

From that point of view, but I think in the financial industry, the point of view of making predictions for financial markets, [00:13:30] I think it's a bit of a different world from where you see all the huge progress that's been made with AI. The key element for us is that the signal to noise ratio is terribly low in the financial market. It's extremely hard to extract a tiny bit of signal to be able to predict whether the price of crude oil is going to go up or down. Where you have seen these huge advances in AI, it's in image recognition, speech recognition, and all those tools or hand written digit recognition, where [00:14:00] today the error rate in the best models is pretty much zero.

We are not living in a black and white world. It's a lot more grayscale, a lot more noisy. I would also expect there will be more setbacks, in a sense, where people get excited about algorithms and they don't live up to the promise because we live in a changing world. Financial markets are a complex system. It's not a stationary world where it's just, once you've figured out how to read a number three, it sorts it out. I don't think we live in that world in the financial market. From that [00:14:30] point of view, I don't think I'm worried that I'll be out of a job in 5-10 years time because the computer builds its own models.

I think there's a lot more progress required from the academic side in terms of developing AI algorithms that can do actual reasoning, not just learning or inference from data. A lot of these deep learning algorithms, when you look at them and how they learn, how they make their decisions, you often see it's not the same way a human would make a decision. I think that first needs [00:15:00] to be solved to make the real game changer, the next level where I will start being worried about losing my job.

Garri Jones: Thomas, is the value in the data or is the value in the algorithm?

Part 1

Thomas Flury: I think it's in both. To me, with a machine learning algorithm, I can do one thing, which is I can extract more value from the data that I already have. I can model relationships in a better or different way, or I can use the algorithm to represent the world in different ways. [00:15:30] For example, deep learning is a lot about feature extraction. Whereas before I write down my momentum trend following indicator and put that on a computer and it trades, versus now potentially a deep learning model can work out how it's the best way to represent a trend. That's one end, and the other end, I really see it as enabling big data. An example, again, coming back to what you mentioned before, if it evolves the purpose of text, we can't read that and come up with a conclusion [00:16:00] whether that's bullish or bearish for a company, but a natural language processing algorithm, something as simple as sentiment modeling, but it can go a lot more sophisticated and model the topics that people talk about.

That, I expect, really allows you to tap this information source, which before, you needed a human to read the report and after five days, he or she will tell you this is a good company. Whereas for us, we can do that in a second now.

Garri Jones: Sure.

Clare FlynnLevy: Can I ask Thomas a question? As more and [00:16:30] more firms start employing this technology in predicting price changes, that change is the way that prices change. How do you deal with that, sitting where you sit, the fact that you can model on the past, but let's face it, the past is irrelevant? Maybe for right now, it's relevant, but extend a few more years and that's not even what the world looks like. I just sat on a panel with a bunch of robo-advisors talking about ETFs [00:17:00] and we give you this quiz and then you allocate money and duh, duh, duh. I asked them, which maybe I shouldn't have done, "Aren't those models all based on historical price behavior? If we know anything, it's that the people who were playing then are not the people who are playing today and won't be the people playing tomorrow. What do you do about that?" They all just looked at me like, "I can't believe you just asked that." For you, I imagine you're building that in, right?

Thomas Flury: Yeah, so I think the key element here, I think this is not even about [00:17:30] using AI or machine learning. This is about building models for statistical models, doing statistical modeling. Fundamentally, once you build your model, you upload it on the computer, let the computer trade. It has to be able to deal with different [inaudible 00:17:44], right? You have to, in your research process, you have to be aware of that. Man AHL, we've been around for 30 years developing systematic trading strategies, so from that point of view, it's no different to what we had been doing before. You always face that challenge, that the world [00:18:00] might change, where the model that I build today is no longer valid tomorrow.

Maybe the difference is, whereas in the past, if you build an economic model, you know exactly the hypotheses that go into it and you can verify that hypothesis is no longer true. That model is invalid. With AI, it's a lot harder

	because you really have to make ... We do that, a lot of effort shining a light into the black box to understand, what has it learned and the monitor, "Is that still going on in the market?" We constantly monitor our models [00:18:30] and we de-commission models when we feel they're no longer appropriate for the world we live in. In the short term, actually, I think there's a lot of divergence going on because as I say, it's like a pioneering front at the moment. Everyone is trying out new things, so behavior will diverge. You try out one algorithm. I try a different algorithm, so we might not necessarily end up capturing the same effect. In the shorter term, I'm a bit less worried about that.
Garri Jones:	Okay, good.
Al-Noor Ramji:	If you actually take a look at a [00:19:00] company like [inaudible 00:19:01], we can trade in algorithms. If you think we are heading towards fewer, you'll be pleasantly surprised or shocked. If you look at robo-advisors, they're not all created equal. Almost everyone that I hear about invests in about 20 or so portfolios of some combination of ETFs. If you start from scratch and say, "What are the correlations between returns of just any security in the world, broadly defined, including [00:19:30] property or whatever you want across assay class," and then try and put that together, now it becomes feasible. Mathematically, it's still impossible, right? It becomes feasible. We've got a case where we've got about 300,000 portfolios now, which more importantly, we can maintain daily. Any idiot can create them. Not an idiot, but you know, reasonable person. Maintaining them on behalf of the individual and making sure that their needs are constantly being monitored, [00:20:00] that's where we're heading. I think so. You're quite right. Eventually, we'll have this holistic missile, anti-ballistic missile thing going on, but right now it's still early days. There's money to be made here, still. I'm only pitching a couple of companies, but I don't have any invested in them. I'm just saying there's lots to do and there's people who will sit behind them, will give you money if you write an algorithm, and you can trade with each other. Someone like you has plenty to do, yes. You [00:20:30] guys are ahead of the curve, clearly.
Garri Jones:	You mentioned a word earlier, which we need be a bit careful here. I'm interested in the experiences that the financial services industry, which is arguably ahead, can give to the wider companies and particularly to the senior management level who aren't in the big data area. What are the challenges about explaining this more generally? You talk about customer focus to management teams and [00:21:00] to other industries, would you say?
Al-Noor Ramji:	I'm not sure what the question is.
Garri Jones:	What don't people understand?
Al-Noor Ramji:	Ah. Well, I don't understand lots of things myself, so hard to, but I think fundamentally, they fall into two camps right now. They're either ignoring it, or

Part 1

They think it's magic. They don't really understand the class of problems you ought to be tackling, for example, face recognition the Chinese get completely. We call it spying, but they call it security. [00:21:30] You can trace a human walking through a crowd into a city, and they'll still pick him up. The class of problems to which you apply a particular algorithm, not understood. They'll just have words like deep learning and they go, "Should solve everything." I think all I'm advising anybody who wants my advice is to start, but doing live experiments, not just poking, not just using [00:22:00] start-ups, not just doing digital garage-type nonsense, which a lot of people do.

If you want to get excited, go to Silicon Valley for a day and come back and do nothing. Lots of boards go there. They come back doing nothing. Innovation tourism, which people like you might want to sponsor. When they come back, they don't do anything. I think that's ... If you just do something, you'll learn from your own mistakes. Everybody at board level and so on, presumably, is clever enough to figure [00:22:30] it out, but there are not enough experiments, especially at scale, Garri. Lots of people doing tiny POCs and claiming victory, but have you actually done it for the whole firm? Does the whole firm adopt AI? That is yet to be the case.

Garri Jones: Experiment to scale.

Al-Noor Ramji: Also, sorry, one more thing. Data, because there's so much regulation around it, people are frightened by it. They'll actually go speak to the FCA, whatever. They'll let you play with the sandbox and it's much more contained, all the mess in Singapore.

Garri Jones: Okay. Before [00:23:00] we open up to questions, Clare, biggest frustration and biggest concern, or one of the above?

Clare FlynnLevy: You mean in running Essentia, or you mean-

Garri Jones: No, when you talk to your customers.

Clare FlynnLevy: This speaks to your point, actually, that there ... We deal with people who are interested in having a go and they'll start in a small way. There's often a huge disconnect between the business and the tech function or the data security [00:23:30] function. The business will commit to working with start-ups and smaller companies, but the IT security function is not committed to that, and so you get quite far down the line before somebody says, "Hang on, you're not ISO27001," and you're like, "Yeah, I know, because you're not paying me anywhere near enough to get that accreditation, plus I have 18 people. That's not going to happen. Now, I can tell you what we are doing and the steps we're taking to get to that [00:24:00] point, but I need you to take a flyer at this stage. It's just a trial. It's just a POC. Let's be practical here." It's tough for the business because the info set guy's job is not to say, "Oh, yeah, okay, fine." His job is to cover every base. I find that incredibly frustrating.

Part 1

Garri Jones: If you can't roll out slack in your organization, what chance have you got in rolling out AI?

Clare FlynnLevy: If you can't, as a vendor, serve an organization because you're using slack, it's [00:24:30] like, "Wait, what?" There's a lot of tools that are making it possible for us to do really innovative things that will help your business, but if you don't let us use those tools because it somehow threatens your business theoretically, it's very complex situation to get into and it wastes a lot of resource and time and money. I mean, we are seeing companies start to become a little bit more flexible along those lines and the way that it ... You're not talking anymore?

Thomas Flury: No, that's okay.

Clare FlynnLevy: [00:25:00] The way that that happens is by saying to them, "You know this company, that company, your peers, basically, large companies whose IT security you're not going to doubt, have said yes to this." You say that and suddenly their minds open a little bit more, which is helpful. There's a critical mass that's starting to be a little bit more practical around this stuff, which is ...

Garri Jones: You might want to talk a little bit louder here, Thomas, but you've been at this for 30 years, your organization, but I'm sure there's still frustrations internally about what maybe you're [00:25:30] not getting or what you could move quicker at. What would be top of your list?

Thomas Flury: To be honest, generally, I'm reasonably happy because I think after initial skepticism, mainly from senior management about this more black box type of algorithms, I think that we've seen the success we've had trading those algorithms and so the uptake is very real and I think senior management is fully behind it. We can extend our collaboration with the Oxford Man Institute that we continue to sponsor. It's been refocusing to a [00:26:00] machine learning institute where we can collaborate with academics there. I think in that sense that the senior management, I think is fully behind it. Of course, I always would like things to go faster, so it's maybe in the side of how fast do I have to turn over in terms of developing new models.

Garri Jones: There you go.

Thomas Flury: Even things as simple as, it's extremely difficult to recruit talented people, an extreme competition for skill, and at the moment, these people with [00:26:30] PhDs in machine learning or post-docs, financial industry is not their top choice. It may be you're second or third, third option for them. They all want to go work for the Googles, Facebook, Ubers. Next is a tech start-up, smaller companies, and then maybe financial industry.

Garri Jones: The battle for talent. We have time for one or two questions from the audience, if anyone's got anything burning. Yeah, right in the back there, Nick. You might just want to wait for the mic. [00:27:00] Don't worry. Just shout. Go ahead.

Part 1

Audience Member: I'm interested, what are the issues in the automation industries as far as [inaudible 00:27:07] that actually, at the system level, [inaudible 00:27:10]?

Garri Jones: How do you guys-

Audience Member: Part of that is that, [00:27:30] if you ... To eliminate some of this value destruction, what you can try to do is actually lower the cost of manufacturing the product in the first place. Candidly, listening to you, it doesn't seem like it's completely clear that you guys help in that.

Garri Jones: Because you're to some extent, offering additional services. You yourself said we'll still be having people making AI programs, which then we'll sell to active fund managers. That sounds just like more costs, [00:28:00] which effectively, ultimately the individual investor pays for. I'd be interested how you see all this at a system level.

Clare FlynnLevy: The way I look at it, there are two levers there. One is the fund manager can perform better. You may or may not believe that that's possible, but I can show you the data that proves that it is. The fund manager's been flying blind all this time with just his or her own performance as a measure of skill, which, it's very noisy. It's not a good measure of skill. [00:28:30] With a better feedback loop, you can actually perform better, and by "better," we're talking about, if you can even do 25 basis points of better performance, that will make a huge difference to the net performance you're delivering to your customer. The other pieces, as you point out, is about lowering the cost.

I think there's an opportunity in the sector to have more funds run by fewer people and by [00:29:00] potentially more junior people or less expensive people. Maybe they're not more junior, but the cost/income ratio of a fund management firm has been out of whack for forever and if you think about what a junior person with a feedback loop, a data-driven feedback loop, could do to get up to speed much faster, there's a lot about fund management that time isn't the best teacher of. Some things, yes. [00:29:30] How to manage your emotions as a fund manager, I would say time is what you need. Living through markets in real time isn't necessarily the only answer and technology can speed that process up so you can get to a point where you have tailored portfolios, which is what the market wants, where one person is running thousands of them and they're not necessarily having to re-balance every single one every day and do all of this tweaking [00:30:00] that they would have to do today because the technology does it.

Garri Jones: Al-Noor, final comment?

Al-Noor Ramji: I think you're misunderstanding. I think we're being blindsided by looking at people in this room. Fund management, et cetera, et cetera, there are tech companies who are going to attack this space. They've already attacked this space. They're on a completely different cost structure, so even if in our case, we'll get it down to eight basis points, cost not pricing, but even that won't be enough. In terms of both returns ... [00:30:30] There are three things, right? What does a customer want? Some of our customers just want downside protection, whatever. There's a portfolio manager. There's also the cost structure of the whole thing. If you look at some of the companies we look at, all three are being solved in a very, very different way. I think we, people in this room, we'll be threatened by people we don't know as opposed to people we do know. I've seen some of them already and frankly, it's to your point earlier about the boards and so on, at least in our case, board is worried about it. We [00:31:00] are visiting places which are not Silicon Valley, is all I'll say.

Yes, they do frighten us, but on the other hand, the experience is in this room, and if we did what they're doing, you'll get a better answer. That doesn't seem to be the case. People begin to slowly get their fees down and try to defend their jobs, which is perfectly human. I'm sure I've got a shot at some of this myself. That's the way I think the big changes will happen. Ten cents into insurance, [00:31:30] who knew that? Ali Baba is into things that you do know about: mutual funds, investment banking, all sorts of stuff, and they will discover new paradigms.

Garri Jones: Thomas, final comment?

Thomas Flury: To me, I think this is a very exciting point in time. I think these new tools enable us to do a lot more and maybe coming back to your question, with the same amount of resources. You can trade a lot more assets. We haven't grown tenfold, but we trade a lot more different markets. [00:32:00] We have a lot more different models, so I think from that point of view, it's a great enhancer of productivity, I think. I think we're just at the forefront right now and am very excited to find out how far can we push this.

Garri Jones: Cool. We'll have you back next year to tell us. Thanks very much to our panel. [inaudible 00:32:18] Thanks so much.

Al-Noor Ramji: Thank you.

Part 2

David Kelnar: Very good morning to you all. My name's David Kelnar. I'm an investment director and head of research at MMC Ventures. We're a venture capital firm based in London that invests in early stage technology companies, particularly those that use software or data science to unlock value for information workers or for consumers alike. I am fortunate to lead a team at MMC called the [00:00:30] insights team. Our goal is to identify emerging areas of value creation, to understand those areas deeply, and then to identify and invest in the very best early stage companies in those spaces.

It's a pleasure to talk to you today about the important and I think exciting subject of artificial intelligence, as we publish our [00:01:00] inaugural State of AI report for 2017. My goal for the report has to be accessible and jargon-free, but while making clear to investors, to executives, and to entrepreneurs, the state of AI today, what is to come, and most importantly how you can take advantage.

[00:01:30] AI has been described as the ultimate breakthrough technology by none other than the CEO of Microsoft. Indeed, five of the world's ten most valuable companies today are repositioning to become AI first organizations. Now while hype around AI is clearly frothy, and while sometimes results [00:02:00] in the short-term may fall a little bit short of the loftiest expectations, we believe strongly that AI represents a paradigm shift in technology, and that this is a technology that warrants the attention that it's receiving.

In 2017, we believe AI reached an inflection point driven by new milestones around capability, [00:02:30] around applications, around adoption, and around entrepreneurship. It's some of these I hope we've elucidated a bit in the report. Today though I'd like to do five things. Obviously our time is a little bit limited. While awareness of AI is high, I'm conscious that most folks have yet to have the opportunity to actually look at AI and understand what it is. So I thought it would be helpful if I spend a few minutes at the beginning just explaining for the non-specialist, [00:03:00] what is AI?

I'll then describe the proliferation of applications of AI that we've seen. The profound implications that I think AI will have over the decade ahead. A tipping point in adoption that we're seeing as AI moves from innovators and early adopters to the early mainstream. And finally, a second wave of investment and activity among entrepreneurs that offers, [00:03:30] I think, unprecedented opportunities both for entrepreneurs and for investors.

So let's begin. What is AI? AI is a general term, coined in 1965, it just refers to hardware or software that exhibits behavior which appears intelligent. Now of course, intelligence comes by degree. We've [00:04:00] actually had basic AI for decades via simple rules-based programs that exhibit rudimentary displays of intelligence in specific contexts. Meet Prospector. Prospector is an example of

Part 2

an early form of AI called an expert system. Built in 1977, the idea of this expert system and of expert systems in general was to take a body of knowledge, [00:04:30] and to build an inference engine, to combine the two, and to replicate the way that a human specialist would tackle a problem. In the case of Prospector, the goal was to replicate how a geologist would assess the viability of a drilling site. Now, Prospector actually did pretty well. It discovered some new molybdenum deposits in Washington State.

But [00:05:00] generally the efficacy of these expert systems of this early form of AI was pretty limited. Why? Quite simply, most real world problems, from how to optimize an industrial asset through to making a medical diagnosis are just too complex to be solved by programs [00:05:30] following sets of rules written by people. The data sets are too vast, the variables are too complex, the relationships between those variables are too difficult for humans to optimize.

But what if? What if we could transfer some of the burden of making intelligent programs from a programmer [00:06:00] to her program? That's the promise of modern artificial intelligence, where progress has exciting and rapid. Excitement around AI today relates to a set of techniques called machine learning. Machine learning is a type of AI. All machine learning is AI, but not all AI [00:06:30] is machine learning. The key to machine learning is that for the first time we enable machines to learn instead of being programmed with explicit sets of rules.

How do we do this? Well, with machine learning we give computers the ability to learn through training. [00:07:00] So in a typical and slightly simplified case of a machine learning algorithm, we take our machine learning algorithm, we provide inputs, which are usually training data, whose correct outputs we actually already know the answer to. We then specify to the system the variables and features that we think matter. The algorithm will then perform a variety [00:07:30] of statistical mathematical techniques to process those inputs and will make a prediction. We can then compare that prediction to what we know is right. If it's wrong, the function in the algorithm can change and tune to try and improve its next prediction.

Now, you can imagine that the first few times you run a machine learning algorithm you get pretty lousy results. But as we input vast volumes of training data, [00:08:00] and as we operate numerous cycles, the predictions made by a machine learning algorithm become better and better tuned, to the point where they can deliver results that are far better than traditional rules-based software.

Fraud detection is a great example of a real life problem that's very difficult to solve by following sets of rules. You'll [00:08:30] probably know this if you've ever traveled overseas and you get your credit card blocked when it shouldn't have been, or so on. It's very difficult to catch instances of credit card fraud in all and only the correct cases by writing complex sets of rules. With machine learning though, we can input all of these big data sets, information about buyer, the vendor, the location, bunch of other data, specify those different

Part 2

variables, and let the algorithm figure out the best way to optimize those to identify fraud in all and only [00:09:00] the correct cases.

Now, even with the power of general machine learning, there are a range of tasks from recognizing objects in images through to recognizing and translating languages that are still just too difficult for us to solve. [00:09:30] Why? Because in these cases we can't even specify the features for the system to look for in the first place. Dalmatians or ice cream? Image recognition is a really difficult problem to solve. I can't write a set of rules that specifies the features of a dog that will work [00:10:00] in all and only the correct cases. Of course dogs come in all different shapes and sizes, all different colors. When represented in a picture they can be partially obscured, and so on. Even if I could somehow write an exhaustive set of rules for what an image of a dog always looks like, it wouldn't be a scalable solution. I'd need to write an exhaustive set of rules for every type of object we wish to identify.

Enter [00:10:30] deep learning. Deep learning has electrified the field of artificial intelligence over the last few years and is delivering breakthrough results in areas including computer vision, language, and much beyond. Deep learning is itself a subset, a type of machine learning. All deep learning is machine learning. But not all machine learning is deep learning.

How does it work? You'll see. We [00:11:00] go a stage further. We offload even more of the burden from the programmer to her program. We don't just offload the task of optimizing the variables. We also offload the task of extracting features from the data to the software as well. In terms of how it works, you'll need to see details in the report, [00:11:30] little bit pressed for time today. What I wanted to emphasize today is just the breakthrough applications that deep learning has had, particularly in these areas of language and computer vision.

This is how an autonomous vehicle sees the world, and this is made possible by deep learning. We can see it identifies different types of objects around it. We can see how it identifies [00:12:00] traffic lights, barriers. For the first time, deep learning gives computers the ability to see and to understand the world around them. So you can see why this kind of technology has pretty game-changing applications.

Now, over the last 12 to 24 months in particular we've seen new milestones in capability around [00:12:30] machine learning and deep learning. I describe the seven key enablers of AI in the report. Today I'll just touch on three. The first is new algorithms. We've got new machine learning algorithms, techniques like convolutional neural networks for image recognition, that are again delivering much better results than these systems could in the past. In fact, just earlier this year, Microsoft used a [00:13:00] relatively new technique called a recurrent neural network and a particular form of that recurrent neural network to deliver better than human language recognition for the first time in history.

medical diagnosis, and more personalized retail experiences. If we get better at reasoning, we can do better in legal services and start to automate some of the features [00:17:30] in our set management that we do. If we're better at planning we can optimize a lot of the physical and digital networks that we have. By giving computers the ability to see, we usher in the era of autonomous vehicles, and also implications for surveillance. Finally, if we can teach computers to recognize human language, we can enable a new class of voice [00:18:00] activated devices, virtual assistants, and more.

But there is, I think, a simpler and kind of broader truth here, which is just that we live in an information economy today. The reality is virtually every sector and virtually every business process draws on some of those five fields. As a result, AI has an incredibly broad range of use [00:18:30] cases. I'll highlight just a few use cases in four sectors. First is in retail. AI will help usher in the era of personalization in retail. The ability of AI to interpret language, to draw on broader data sets, to find more subtle patterns in data, will unlock a range of possibilities, because [00:19:00] virtually every stage of somebody's interaction with a company can be improved through AI. So in terms of content personalization, why is conversion of online stores and indeed real life stores generally so lousy and averages about 1% online? Well because most of the content we present to people is just not customized according to what they're looking for. AI can help with that by looking at social and other data about a customer, finding more complex patterns in that data, [00:19:30] and presenting more relevant goods.

Also around price optimization, a 1% change in price can impact a company's profitability by up to 10%. The reality is, most rules-based systems of software are just not good enough to balance all the various different multi-variables involved from an end of life of a product to a new introduction, to a customer's propensity to pay, to optimize that problem, whereas machine learning can.

[00:20:00] Manufacturing, to me, is one of the most interesting sectors for benefit for AI. Overall equipment effectiveness, OEE, is a measure of the overall productivity of a factory. OEE varies enormously by industry. It can be as low as 75% in industries in a hyper-competitive world of manufacturing. AI [00:20:30] can help, predictive maintenance is a core use case for AI in manufacturing. For every hour that an automotive assembly line goes down because of an equipment failure, the cost to an automotive manufacturer can be £1.5 million. With AI though, we can extract subtle patterns, there was a nice phrase earlier, the signal from the noise from sensor data embedded in all of these pieces of equipment and predict ahead of time when machines are failing, [00:21:00] swapping out parts ahead of time and preventing that expensive, unplanned downtime.

Finally, utility optimization. There are a lot of really difficult problems in manufacturing. If you want to optimize the output of a gas turbine, you need to moderate the flow of fuel to the microsecond in such a way that it will maximize the output of the asset while minimizing wear and tear and minimizing environmental emissions. [00:21:30] That's very difficult to do with a rules based system. AI can do that much more effectively.

Asset management, I thought it would be helpful to talk a bit about today given the nature of the audience we have in the room. AI's ability to synthesize data from within and beyond firms and to enable human to machine interaction will have application in a range of areas within asset management. Investment strategy is a core use [00:22:00] case. AI can synthesize research and data from within and beyond the firm, and then better optimize the different types and range of objectives that a firm is looking to manage to give better returns.

By leveraging chatbots both within and outside a firm, client service can be dramatically improved. If we deploy chatbots within a firm, it can enable [00:22:30] an asset manager to query much more quickly and get information back about a client's portfolio to enable them to deliver the kind of service to the mass affluent that's only been available to the high net worths previously. When deployed outside the firm it can be a valuable tool for client self-service.

Finally, legal and compliance. AI's ability to understand written spoken language, to extract concepts, [00:23:00] to reason, to understand, obviously has great value in the legal industry. Document review is a core use case. In a typical M&A environment, the average due diligence data room has 75,000 pages of information. AI can parse that information in seconds and look for patterns or clauses that look abnormal. Of course, much bigger than beyond just legal and compliance, but cyber security, it's one [00:23:30] of the defining problems of our age, in my view. Again, AI's ability to establish baseline levels of behavior in a network or system and to detect deviations can help catch threats in realtime much better than human operators can.

So we've seen so far, let's just sum up, machine learning for the first time is a way for computer programs to learn through training instead [00:24:00] of being programmed with sets of rules. We've seen over the last 24 months that the use cases of AI have been proliferating. But to me, the most important part of all this, and also the most interesting actually, are the profound implications that I think AI will have for us as consumers, for companies, and for society, in the decade ahead.

I highlight eight core implications of AI [00:24:30] in a report. Today I'll discuss just four or five. When I started to think a bit about the implications of AI, I thought a useful place to start would be to try and abstract really what the benefits of AI are at heart. I think it's these four. I think it's innovation, the fact [00:25:00] that AI makes possible new types of products and services, things like autonomous vehicles or voice control devices in the home. The second is efficacy, the ability to do tasks more effectively than we've been able to do them in the past, from fraud detection through to industrial asset optimization. The third benefit is velocity, the ability to perform tasks more quickly because for the first time we can automate a lot of previously human capabilities, [00:25:30] things like legal document review. Finally, scalability, because we are enabling machines to do things that people did in the past, AI offers the kind of scalability that humans could never provide. That's enabling things like automated medical diagnosis.

When we understand that those are the real benefits at the heart of AI, the implications I think become a little bit clearer. The first is new market [00:26:00] participants. Access to a range of markets today from healthcare to financial services are limited to subsets of the global population. Take healthcare, diagnosis, primary care. In the west, access to primary care is expensive or an inconvenience. [00:26:30] But in developing economies, access to primary care is often non-existent. Why? Because diagnosis is delivered by experienced professionals, their training takes time and money, and their scalability is of course limited.

AI though will enable the automated diagnosis of [00:27:00] a growing range of medical conditions. The cost of a marginal diagnosis from AI will be nil. Because AI is of course inherently digital, access won't be a problem either. By 2020, 57% of people in developing economies will have access to a smartphone. As costs and barriers to entry fall [00:27:30] for these kinds of technologies due to AI, you'll see an influx of new market participants, people who are able to use these services for the first time in history. That will obviously have profound implications for both providers of those technologies directly, but also the companies around them.

The second thing we'll see is profound shifts in sector value chains. By that I mean where and the extent to which profits are available [00:28:00] within a sector. Take the insurance industry, 42% of all insurance premium revenue come from car insurance today. Yet AI has enabled the age of autonomous vehicles. As adoption of autonomous vehicles grows over the years ahead the frequency and severity of car [00:28:30] accidents will reduce, but with it the price and revenue of car insurance. In fact, the profitability from car insurance in the UK is forecast to fall by up to 80% in the decades ahead as autonomous vehicles gain adoption. How will the insurance industry adjust if it stands to lose 80% of the profitability from 40% of [00:29:00] its revenue. I think you'll see comparable shifts in value chains beyond just that.

Thirdly, shifting companies' competitive positioning. Every paradigm shift in technology destabilizes incumbents and creates new sets of success factors for the rest. I think you'll see the emergence of four categories of company in the age of AI. [00:29:30] You'll see platforms and disruptors and leaders and laggards. Among suppliers of AI you'll see platforms and disruptors. Platforms are companies like Google, Amazon, IBM, Microsoft. Companies that provide both the infrastructure and machine learning services that are at the heart of this AI revolution, given those picks and shovels pieces, [00:30:00] vast resources, world class talent, vast non-public data sets. I think those companies are very well positioned to attract value in the years ahead.

Part 2

A second group though will be what we call the disruptors. These are typically early stage machine learning led software companies that have AI at the core of their value proposition. These are the kinds of companies that we invest in as a venture capital firm. [00:30:30] For incumbents the disruptors are a double-edged sword. For buyers with the vision and ability to embrace the disruptors, then I think they'll be a powerful enabler. But for incumbents that for whatever reason don't have the vision or organizational ability to embrace the disruptors, they are likely to be a highly disruptive influence to them.

Among then buyers [00:31:00] of AI, big companies today, I think you'll see a bifurcation into leaders and laggards. Leaders again are the companies with the organizational vision to embrace AI, the skills to attract, retain, and integrate data scientists within the organization. They possess large, non-public data sets to train their machine learning algorithms. In the age of AI they will, I think, extend their competitive advantage, for [00:31:30] a couple of particular reasons. The first is, we, as I mentioned, live in this information economy, data is power. Companies that can extract information from data more intelligently will disproportionately accrue economic returns.

There are also data network effects to be enjoyed by leaders in AI. The more data you have, the better you can train your algorithms. The better your algorithms, the better results you can deliver. The better results [00:32:00] you deliver, the more clients you get, who bring with them more data. That creates a virtuous circle that's very difficult for other companies to break. So because of disproportionate returns to those who can process information better, and because of the data network effects that they will enjoy, I think leaders will quite dramatically strengthen their competitive advantage through this paradigm shift.

Laggards on the other hand will squander the data sets they have. They [00:32:30] will not be able to attract machine learning talent. They will fail to embrace AI and the disruptors more broadly. I think you'll see them losing competitive positioning and ultimately relevance surprisingly quickly.

We should talk a bit about new business models. To me the most powerful thing about new technologies is not actually their new technical capabilities, it's the [00:33:00] new business models they enable. I think we'll see a lot of the same here. The combination of AI, X as a service, and subscription business models, will obviate a range of existing business models in a number of sectors and offer great new opportunities.

AI for example will alter completely [00:33:30] the economic fabric of ownership of cars and associated insurance. Cars spend 96% of their lives today sitting parked unused. The reason though of course that that crazy model has been necessary is because cars have offered point to point convenience, spontaneity, privacy, and security. But in the age of autonomous vehicles we'll be able to summon on demand an autonomous vehicle [00:34:00] from an optimally distributed local fleet to take us where we want to go.

Part 2

With the cost of the driver removed and the cost of the vehicle and the fuel amortized over a much greater number of journeys in a given period, the cost of that service and the cost of a marginal journey will be vastly lower. I think you'll see the emergence of transport as a service business models where I pay a really pretty low flat monthly fee for on demand access to a fleet of autonomous vehicles [00:34:30] whenever I need it. That will clearly disrupt enormously the business models of today's automotive manufacturers as private ownership of vehicles declines. There'll be secondary consequences down the supply chain. Local repair centers, local fuel, and so on.

Finally, I talk at some length in the report about the fact that AI will offer profound benefits to society but also pretty serious risks that we [00:35:00] need to have, I think, a mature debate around. I describe a number in the report from job displacement through to risks of bias, privacy, and conflict. Just one I touch on today which is around the risk of job displacement. There is a debate today about whether AI will create more jobs than it destroys. I don't yet know the answer to that. But what worries me is that there [00:35:30] are a relatively small number of jobs that AI will be able to fully automate in the short to medium term. There are 3.5 million truck drivers in the US, there's quite a few telemarketers. My worry is that whether or not AI creates more jobs than it destroys, the speed with which quite large groups, quite large numbers of people could be displaced, will prevent them from being reabsorbed into the workforce [00:36:00] in time. The social implications and the political fallout could be quite material.

I'll cover two final areas briefly. One is a tipping point in adoption that I believe we've seen in the last 12 months. Interest in AI has increased significantly over the last 12 months. In fact, it's pretty striking. In 2016, [00:36:30] the term AI didn't even feature in the top 100 terms searched on Gartner.com by Gartner clients. These are among the largest or substantially sized most innovative companies in the world. By May 2017 though, AI was the seventh most searched term overall.

Now, more importantly, that interest is translating into increasing activity. Adoption of AI [00:37:00] is clearly still [inaudible 00:37:03]. But it is changing. According to a survey that McKinsey has done of more than 3,000 AI aware executives across more than 14 countries, one in five companies have now deployed at least one AI technology either in a core part of their business or at scale within the business. We believe AI [00:37:30] is moving, AI is crossing the chasm from innovators and early adopters, towards the early majority.

I think the pace of adoption around AI is also going to accelerate materially. 75% of executives plan to deploy AI to some degree within the next three years. So expect that shift in adoption to accelerate [00:38:00] quite strongly. Adoption is going to be quite uneven I think. This is an excerpt from that same McKinsey study, which shows current AI adoption by sector, and indeed propensity to increase spend in the years ahead. Now please note, this is an excerpt from that chart that only shows sectors that have the highest adoption today and the highest propensity to spend, this is the top right of the chart. What we see is that high tech, financial [00:38:30] services, automotive, energy, media, and transportation are leading the adoption of AI today. Particularly significantly these companies that are leading AI adoption today are also likely to be those that continue to increase their spend most robustly in the future. So I think we'll see a bit of a bifurcation by sector.

Just finally, a point for executives in the audience, you'll see a lot of media attention and a lot of pilot projects about chatbots. Chatbots [00:39:00] are great. But be aware, most of your peers are actually using AI more typically for their core functions, to improve their decision-making, to improve the recommendations on which they act, for process automation.

I'll finish very briefly by talking about a second wave of entrepreneurship that we're seeing. There are about 400 early stage AI [00:39:30] software vendors in the UK. We've met so far with about 250 of them. There's just two themes I'll touch on. The first is that activity is flourishing. We've seen enormous change in this over the last couple of years. The number of new AI startups founded in the UK annually have doubled since 2014 compared with the prior period. On average since 2014 a new AI company has been founded in the UK every five days. [00:40:00] But alongside that there's been a profound shift in focus both for investment activity and also for entrepreneurship.

The first wave of AI investment and entrepreneurship really focused on companies building research in AI, companies like DeepMind, or companies providing core AI technologies in areas like computer vision and language. Now we're moving into the age of applications. [00:40:30] 9 out of 10 of the companies in the UK of these early stage innovators are using AI to address specific problems in given sectors or given business processes. There's a big focus in marketing and advertising. One in seven of all of these AI startups are focused on that function. By sector, finance is the sector where more are focused than any other. [00:41:00] I think this is going to give unprecedented opportunities for both the entrepreneurs themselves but also investors.

So let's sum up. We've seen that AI is a paradigm shift in software development because it enables machines to learn through training instead of being programmed with sets of rules. We've seen that the applications for AI are proliferating. The implications will be profound, from shifts in sectors [00:41:30] value chains, the obviation of old business models, the creation of new ones. Then in the last 12 months AI is crossing the chasm of adoption. Finally, we're entering now a second wave of AI investment and entrepreneurship.

I'll finish with a final thought. Industrial revolutions of the past have focused on the creation [00:42:00] or transmission of power or goods. So the first industrial revolution, steam power in the 1780s was really about the mechanization of production. The second, 100 years later, electricity enabled mass production. Of course the electronics revolution of the 1970s enabled production and communication on an unprecedented scale.

Part 2

Today though we're entering the fourth industrial revolution. [00:42:30] Today our primary source of value creation is in the processing of information. Ultimately, it's AI's ability to process information more intelligently that will yield benefits both humble and I think truly historic. Thank you.

Part 3

William Tunstal: It's been around for decades and again it's been steadily improving. There's been a little bit of a step change with deep learning in the last few years that would reduce the error rates. Never the less it was on a path to steady improvement over a number of year. That's obviously important if the device can't actually understand what you're saying or can't hear the words you're saying, you can't get the experience.

Another part of the technology is natural language understanding. So once you have that text, to the computer that is [00:00:30] a sequence of letters. What is the meaning, what is the intent of the customer, what is the customer trying to ask Alexa, ask the assistance. That is also extremely difficult technology that's also been improving. If it's a questions, how do you actually answer that question? How do you reason to produce an answer? How do you action the intent once you understand it? Then finally there's also the piece about turning the result, you know what Alexa wants to say, back [00:01:00] into text, back into sound. How do actually create that experience of speaking to a device, it understanding you, doing what you want, answering your question, and then producing something straight back, that's exactly what you're looking for.

All of these technologies are really difficult, they've been around for a while, they've been steadily improving and they kind of crossed this magical threshold point and I think they'll continue to improve. The commercial success that Echo and Alexa have had is resulting in massive investment in the space ... [00:01:30] massive institute parameters in hiring thousands of people other companies like Google are producing competing products also fighting for that space. So there will be continuous incremental improvements in all of the technologies, which will also result in a better overall experience for the product.

Nick James: Great, Mark at Babylon you are working in the medical space and producing diagnostic A.I. How do you think about what type of A.I. to use, [00:02:00] whether it be deep learning or probabilistic algorithms?

Mark Tsimelzon: So, it's a great question and I'll echo what William just said. It's never just one technology unless you're doing something very narrow. It's a combination of technologies. So we try to think about a human doctrine, their brain and then what are the parts of that brain that we need to implement in A.I. So it starts with the knowledge. How do you capture the knowledge of medicine, which is a fascinating subject, because any doctor goes through years and years of training? How do you train [00:02:30] your A.I. machine to do the same? And then, of course, you need to be able to communicate to people whether through text or speech, so that's a big part of what to do.

Part 3

The biggest part is what you mentioned, is probabilistic. It's all about diagnosing because diagnostics is not an exact science. It's never the case that if you have Pneumonia you're guaranteed to have a high fever. There's always a probability there associated between a disease and a symptom or a risk factor and disease. [00:03:00] You have to think in a very probabilistic way and deep learning a part of what you do, but it's not the most important part at the moment. I think deep learning has been at the forefront of news cycle. A.I. is a much bigger field than just deep learning.

Nick James: How important is the audit trail of what the A.I. is doing to you?

Mark Tsimelzon: It's extremely important alright because we are in a heavily regulated field so you can't just say, The machine diagnosed the patient and the patient thinks they have cancer [00:03:30] and we don't know why, but the machine said they have cancer. So that's actually where probabilistic are much stronger than deep learning in terms of having this trail and then saying, well if you look at all the risk factors and all the symptoms, and the tests, then that's how the machine came to this conclusion. So, yes for us, it's hugely important.

Nick James: Great. Martin, at Peak, you're working with companies that probably don't have much knowledge of A.I. so how do you educate them on the capabilities, [00:04:00] how do you educate them that this is worth doing and the benefits they can get out of this technology?

Martin Sutton: It's a good question actually. I think that one of the challenges of our job to help people understand what is possible. First of all, what is the problem they're actually trying to fix? A lot of people don't know what that problem is. They kind of hear all this buzz about A.I. machine learning and it's: I want some, but I'm not sure why. We tend to start at the real simples of what are trying to achieve as a business and it can be as broad as; we just want to improve our [00:04:30] EB daft from 6 million to 10 million by next year. Then we're getting some hunches of where that improvement could come from. Then we would start to drill a little bit deeper around this data that this business actually collects and sits and going to be useful for us to be able to work using our A.I. capability and our data science team to drive that performance that they are looking to achieve.

So I think we got to be kind of careful ready because we focus a lot on technology as people. The media is all about the next acceleration of tech, but I think we need to simplify a little bit around [00:05:00] what is it where trying to achieve. Let's find the right tech for that, not let's build a tech and hopes it fixes something we've not thought about.

Nick James: Great, I want to move on to think about differentiation for an A.I. business because we hear a lot about these massive compute platforms from Google or Amazon or IBM. It seems many of the algorithms for how you do A.I. are pretty much open source. So how does a business differentiate using the Alexa [00:05:30] platform as an example? Some people consider that the leading platform for doing this human voice recognition and understanding of natural language. Do you think William, is that a permanent advantage or other people kind of catching up quite quickly on the technology side?

William Tunstal: Specifically, in the voice assistant space, yeah Alexa is regarded as the leader, but I don't think there's a massive [00:06:00] difference in terms of A.I. and machine learning across the really big technology companies. I think they're all investing heavily. They all have very large amounts of capital.

There's a big community of machine learning scientists and specialists. Take speech-recognition for example, there was a community of people who had been working in turning sound into text for a very long period of time, for a very long period of time. They all know each other. They all go to the same conferences. They all collaborate and write papers. They all move from company to company [00:06:30] and they all try to recruit each other from within each of these companies. [crosstalk 00:06:34]There are large numbers working for Amazon. There are large numbers working for Google. There are large numbers of them working for Apple. So yes, you can certainly point to particular products and say Google is better at this, or Amazon is better at this. It's not clear that one company is going to massively dominate it. Sometimes you get advantages from the data [00:07:00] if you have a product out. [inaudible 00:07:01] data that can give you an advantage that's difficult to catch up. There are ways around that. There are ways of bootstrapping. The big companies, the really big companies, I don't think there's a huge difference between them is my take.

Nick James: Janet, the story stream business, I understand, has been repositioning all around A.I. and we just talked about these big companies like IBM or Google who are also trying to sell to companies to help them [00:07:30] maximize their marketing using Watson for example. How in story streaming do you build a competitive advantage in an area where you have these giant Goliaths possibly trying to do the same thing?

Janet Bastiman: Well, it's all about specialization and really understanding the data and the problem that you're trying to solve. If you're looking broadly and you just want a generic object recognition or generic text understanding, then yes, the big players will give [00:08:00] you that. That won't give you that much more insight, you have to go a lot deeper and really understand your problem and really understand your data, otherwise you'll end up with a solution that doesn't do what you need it to do. It'll give you the wrong results because it's only as good as what you put into it. So understanding what's the signal to noise ratio, are you putting that right data in, is it biased and then being able to calculate the outcome of that. That specialization is where companies like my own can really add value.

Nick James: [00:08:30] Mark, a similar question to you. We hear that IBM and Watson have digitized all of the health records in the world and see they can do wonderful

Part 3

things with that. Babylon, how do you go a step beyond and how do you build a sustainable advantage?

Mark Tsimelzon: I now have to echo Janet this time. Specialization, that's what startups do. Babylon is a fairly large company, but we're still a startup compared to IBM. You absolutely right that open source algorithms are available, [00:09:00] but these are very generic algorithms. How to adapt them to a particular domain that's a very hard subject. And that's what we have been doing. Also, again healthcare is a very deep space. So there is a lot of knowledge there that already available. You can't start training systems from scratch just saying give us all the data we'll start training systems. So one of our secrets is we have a lot of doctors working with us and we are building systems that capture their knowledge. When you marry [00:09:30] these two approaches, learning from data and learning from experts, then you build the system that can win. I can not comment too much on Watson, but that's our approach.

Nick James: Great, okay I wanted to move on to building a team in an A.I. led company. We hear a lot of reports about talent being scarce in the data scientist area. Martin, do you find access to talent of data scientists to be a limiter in the peak business at the moment?

Martin Sutton: I'm pleased to say no. [00:10:00] It's quite unusual really. We're fairly early stage business, we coming on to our third anniversary. We seem to have no problem at all attracting data scientists and data scientists are obviously a key component to our business. I think it's because we offer a variety of use cases, industry expertise, working cross-industry, they're working in a fast-paced environment. They're learning a lot of new skills, business acumen skills we bring to our data science team so they can converse [00:10:30] with board level execs around what the data means in board level exec speak, versus an algorithm speak. Because of that, we seem to attract really good talent from a mixture of backgrounds coming out of big industry, people who may be about to join the likes of KPMG or people coming out of academia. We haven't struggled at all, but we do see our clients struggle, I think for the reasons I mentioned earlier. If you're a top of your game data scientist, you want to go and work in a construction business, [00:11:00] for example, just giving one of our clients as an example, that would be much more difficult hire for them, but we're okay at the minute. Thank you.

Nick James: Great, Janet what are your perspectives on building the right team?

Janet Bastiman: It's pretty similar. You got to have a really good value proposition to attract the talent, that's not necessarily the highest salaries, it's about giving a really interesting problem. A lot of the good talent, they're really excited about having variety and [00:11:30] having something that's intellectually stimulating. If you can offer that, then you can attract the people. I've never had a problem hiring in big and small companies for A.I. talent. It's also about getting the right mix of people in the team. It's no good having five to six people who are all thinking the same way and working the same way. You've got to build a complementary

Part 3

team that can challenge each other and push things forward. So having that difference of opinion to help, having [00:12:00] constructive arguments in the team as the right way of doing things. Much like you'd build any development team. You don't want people all saying the same thing and doing the same thing because you'll never innovate.

Nick James: Great. I wanted to move on to ethical considerations of A.I. and David touched on this in his presentation. We've got people like Elon Musk saying that this could be a threat to humanity itself, there are some reports of machine learning models actually reinforcing biases in judicial outcomes, [00:12:30] parol judgments. William, do we need to be worried about the consequences of all of this and how do we deal with biases in A.I.?

William Tunstal: I think there's kind of two tracks here around this sort of ethics thing. One of which is science fiction and one of which is kind of true now. The Elon Musk story and some of the things you hear about the end of humanity is about an A.I. that doesn't exist yet. It's a scenario [00:13:00] where A.I. is intelligent enough to improve itself and it also relies on assumptions like intelligence being some kind of linear thing that can exponentially improve. So you can imagine a system that's clever enough to improve itself. You get this sort of runaway A.I. scenarios. This is a science fiction scenario right now. Nobody can just prove that it can never happen, but the technology that we have right now is nowhere near that kind of capability. So I think that's really important to understand when you hear about the end of human race type stories around A. [00:13:30] I. It is worth worrying about because consequences are very severe, but it's not something that's close.

The other side the ethical side of A.I. is about automated systems. Buses are a very good example. Automated weapons I think will be a great example. The advances in vision, for example, could create automated weapons. So is it acceptable to have a sniper that will identify the enemy and kill him without any human being in the loop or an automated bomber that will select targets and destroy [00:14:00] them?

There's a system in the U.S. that gives advice about whether somebody should be let out on parole. It doesn't take race into account, it isn't one of the factors, but it takes into account all sorts of things that collate all sorts of things that collate very strongly with race and it's been shown to be unintentionally biased. There are people who are in jail in the U.S. as a result of an A.I. system that might not otherwise be.

A.I. systems for loans and fraud. I think fraud was mentioned earlier. I think the current fraud systems [00:14:30] that decline credit cards do use A.I. and they do seem to arbitrarily decide they're going to stop a transaction. People that are declined loans that don't get an explanation are just told the computer system has told them that they can't get their loan. So explainability is important. We have rules about decision making, about what is reasonable to include in a

Part 3

decision. If you have a black box that is statistically determining a result using data and [00:15:00] it's completely opaque, that does raise ethical concerns.

Nick James: Mark, in the medical world, how do you handle the risk of a missed diagnosis?

Mark Tsimelzon: So, that's obviously a very big concern for us. The good news is that medicine has been thinking about these problems for years. Human doctors also make mistakes. In fact, the scary part about doing what we do is, we actually compare doctors to each other and doctors to machines, [00:15:30] we see how often doctors disagree on a diagnosis or how often they make mistakes. So we want to make sure that you have the right standards, your machine doesn't perform worse than the doctor, but then you also want to be able to learn, right. If you make a mistake, how then do you examine your model and try to understand why did this mistake happen and then how do you correct it next time? So nobody expects perfection from doctors, fortunately. We are getting very close in [00:16:00] some areas, our company, and others to matching human performance. In some areas, machines are beating doctors, if you're looking at X-rays and say: "Is this cancer or not cancer". There have been already numerous competitions and studies that show machines doing better than doctors and I think over time, the trend will continue.

Nick James: Excellent, Martin how do you work with your customers on this type of stuff? How do you educate, there's not going to be a risk of bias be introduced their business process?

Martin Sutton: Yeah, there's a couple of things [00:16:30] really I guess on any project we would work on. We would drip feed in our insights to start changing their processes gradually. We wouldn't just advise big bang, we cracked the nirvana, your biggest challenge, believe us and let's see what happens. So we can experiment without actually using live data in some cases. For us, it's more about the A.I. plus human. So I think what's been echoed here particularly in the medical world, research I've done, humans are much better at detecting things like false-positives or false-negatives in [00:17:00] large amounts of data and we see the same. So if we're working with a particular client of ours, and we suddenly see something change by a larger magnitude than makes sense to a human, it gets stopped and it gets flagged immediately and it gets picked up. A human from the data science team will look at it, iron out the anomaly and then we will re-train the model and go again. So I guess to answer that question, do this at a slow pace, don't let the machines make every decision and always have smart people around who can spot anything that may be an anomaly.

Nick James: [00:17:30] Just before we go to the audience and get some questions, I just wanted to ask one high level one. We hear about Google, Facebook, Amazon all having these wonderful A.I. technologies, but we've got this huge problem at the moment of fake news, which is apparently influencing election results. Given A.I. is where it's at, how are we having fake news become the top of Google and Facebook, and not trying to answer [00:18:00] on their behalf. Janet, if you have any thoughts on this area. What's going wrong?

Part 3

Janet Bastiman: Part of it is just the self-reinforcement of people's own cognitive biases. They want to read things that support their own views. So you end up with polarizing viewpoints floating to the top. Where you A.I. that says: "Oh you're interested in this, I will show you more of the same." You end up with that perpetual reinforcement and the people in the middle tend to [00:18:30] either get flipped to one side or the other on very, very minor things, which is a huge problem. What we need to apply the A.I. to do is to strip away sensationalism from what's the underlying fact. We as humans are astonishingly bad at that. Until we can be better at it, it's going to very hard to create an A.I. that's going to good at it too. But we need to be more aware of what's going on and how we're influenced.

Nick James: William, any views on that?

William Tunstal: Yes. I think there are two different [00:19:00] things here. I think what you're talking about largely is the polarization of the newsfeed. The echo chamber, the fake news as in stories that are untrue. The next explanation for that is simply is that A.I. can't do that yet. To work out that a story is not true requires reasoning and understanding of the world that automated systems simply don't have. The scale of stories we are seeing being published each day for Google or Facebook to automatically determine that [00:19:30] a story is false when there probably are signals from the domain [inaudible 00:19:34] something like that. At the moment that requires probably human intervention. It's not something that can be automated with current technology.

Janet Bastiman: Absolutely, you add on to that the variation, so the misinterpretation of what could be a factual story suddenly becoming ever so slightly a variation of it that's necessarily true is very hard to pick up.

Nick James: Great, let's take a question or two from the audience.

Adrian: [00:20:00] Hello, Adrian's my name. Please to meet you. I was wondering if you make a comment on the economic model of some of this A.I. and where I'm specifically speaking is around the models of decision support where ultimately a human might still be responsible [00:20:30] for taking the rap if the decision is wrong versus actually the A.I. actually making the decision. I'm interested by the regularity implications that ... my background is deeply in healthcare so I'm very familiar with sort of Mark's story. But I was wondering if there's any experience from some of the other industries in terms of this because clearly there's a profound impact if it's only decision support because you still need an expensive human there taking the rap.

Nick James: Martin.

Martin Sutton: [00:21:00] So I'm just trying to think in a scenario where we have ... I'll boil this down to a real-life use case. It's a non-sexy industry or use case, so apologies. We're doing some work with a [inaudible 00:21:14] listed construction business. And what we're trying to do is driven optimization in their organization. If they want to improve the return on capital investment, that's their number one metric. So we've identified using A.I. and all ... clever people, that we can save them 27 million pounds this year by getting [00:21:30] them to sell off a whole bunch of assets in their business and return that to the bottom line. Now you can imagine the nervousness in that boardroom when we said we're smarter than you. It's quite antidotal, the comment that came straight back to us is: "You don't know our industry, you can't mix sand and cement with a spreadsheet." So we sort of said: "Okay what we'll do with this we will for the first three to six months we will bet against you that we can predict certain variables in your business that are going to happen. We're not going to ask you to take those actions, [00:22:00] you carry on as normal and you see if you can be smarter than us over the first three to six months."

What we've proven is that we're way more accurate using the models we've got and developed for them. So they now have complete faith to invoke those insights into their business then they're actually going to see those sort of returns that were predicted at the beginning. So it's a question of don't jump in with both feet. Maybe model things out to see how things would have panned out versus just doing it and hoping and continue to let the human do the stuff that their doing and see if you can outsmart [00:22:30] them would be our way of attacking that.

Nick James: Cool, any more questions?

Speaker 7: Just following up on that point, just taking it a step further. If you've gone through all the effort for your customers and clients to get them to believe that your model is better than they are, do you think that trust is something that will continue to stay or do you believe on the first time that it goes wrong, [00:23:00] that will be faith lost. How do you maintain that?

Nick James: Yes, sorry.[crosstalk 00:23:06]

William Tunstal: We touched on explainability a couple times today. I think that is part of the trust issue. If the black box just spits out a decision or a result, it may be [inaudible 00:23:23] trust over time when you see it's working, but that's nothing like as good as the A.I. system being [00:23:30] able to give you an explanation as to how it came to that decision, why it works. Also, if it's wrong, that explanation often can flag that to the people that need to trust it as well. So, there's a lot of research ... the modern sort of machine learning statistical techniques, you know deep learning are notoriously unexplainable at basically massively complicated calculations that are totally opaque to humans. That is a serious [00:24:00] problem in a lot of applications. Some applications its fine, other applications it isn't.

Nick James: Mark, in the medical world I guess in drugs they go through phase one phase two trials and do we need to do phase one phase two of A.I. to get this confidence?

Part 3

Mark Tsimelzon: We are doing very extensive trials of A.I. It's regulated by the same people who regulate human doctors. It's a very new technology, but the way you have to think about it is in many ways you can leverage a lot of the ideas that have been around regulating [00:24:30] human doctors, human hospitals, the drugs that are designed by humans. People talk about regulating A.I. as if it was something completely brand new and yes, A.I. is brand new, but the problems that it's trying to solve is not. So we don't think that's a massive problem.

Nick James: Janet, if you got[crosstalk 00:24:49]?

Janet Bastiman: Yeah, I think one of the biggest problems is as humans, we don't like devolving control particularly well and we accept devolving that to other humans but we're holding A.I. to a higher standard [00:25:00] than we'd hold ... especially in a medical profession. If you just said: "Oh it's just as inaccurate as doctors" then people would be a bit nervous and would prefer a human doctor every time. They want it to be more accurate, and yet if you told them, that the result was by a human, not an A.I. there would be much more comfortable accepting it, and that's just how we are. Until we start seeing a more level playing field understand that A.I. is giving us is just as good as a well-trained human, [00:25:30] then we're going to get the first problem and people are going to want to switch to something else. Whether it's a different A.I. solution or go back to humans. And that's something the industry is going to need to overcome.

Mark Tsimelzon: Just to follow up. We are in actually not just in the U.K., we are also in Rwanda in Africa. We're working with many other countries. There is just not enough doctors there. So we can argue for ages what's better, human doctor or A.I. doctor are they like this or not. If you are in Rwanda and you don't have access to a doctor period and here you get access to an A.I. [00:26:00] doctor, well they're signing up in huge quantities to use our technology there. So we're not making a ton of money there, but we are bringing A.I. there because it's the right thing to do.

Nick James: Excellent, well thank oh, there's one more question.

Speaker 8: What's the panel's view on basic universal taxing robots? How do you think that's going to play out?

William Tunstal: I think we have like 16 seconds left [00:26:30] you really want to open that topic up. I think I'm going to let somebody else answer that.

Janet Bastiman: I fundamentally disagree with taxation of innovation because it's just going to throttle us as a country. I think there are different solutions that need to be investigated.

Nick James: Let's wind it up there. Thank you to the panel, that was really interesting discussion and next, we got the healthcare panel.

Part 4

Sally Taylor: Good morning. I'm Sally Taylor from the Numis Healthcare Research Team. I'm joined today by Ken, Mark, Ben, and Jonny, and I'll hand it over to introductions in just a moment. Firstly, I thought it was worth highlighting some of the early trends that we're hearing in terms of applications within the AI space. We're hearing the potential for greater ability to access primary healthcare and we just heard from Babylon in terms of doctor-based applications. Will surgery outcomes improve through the applications of robotics? Will remote surgery be possible one day? [00:00:30] We're also seeing significant improvements in imaging and that's paving the way for deep learning, and the potential to be able to diagnose certain conditions such as heart disease or skin cancer at levels at, or above, the leading global specialists in the field. Other areas of focus include within behavioral health, and analyzing large quantities of data, to try and predict outcomes, and ideally, be able to intervene to prevent patient relapse.

All of these areas could ultimately improve outcomes for patients, and also save healthcare costs in the longer term. [00:01:00] There are still some challenges surrounding data, particularly around bias who owns that data, and we'll touch on those, hopefully, as part of our discussions today. Our companies today are focused on both the acceleration and improving success within drug development where only a few percent of drugs actually enter clinical trials, are ultimately commercially approved and launched. Also, within social care, we're attempting to help with the NHS bed-blocking crisis that costs the NHS around £900 million a year. The panel is open to questions [00:01:30] from the audience, so after introductions, if you do have a burning question, please do raise your hand and let's make this an interactive session. Just to hand over for introductions, if I may, and perhaps you could provide a kind of elevator pitch if you like, and for your company. Also, the big problem that you're trying to solve, perhaps we'll start with you, Jonny?

Jonny Wray: I'm Jonny Wray. I'm the head of informatics at E-therapeutics. We're a small, let's say, [00:02:00] biotech based in Oxford, and we're focused particularly on improving the discovery of [inaudible 00:02:08] drugs. I would say if we addressed the big problem, the big problem facing the drug industry as a whole is productivity. As you mentioned, a really small percentage of drugs that go to clinical trials get through. This has been analyzed quite a lot in recent years. I wanted to focus on one particular potential reasons. [00:02:30] There are lots of reasons with that productivity because we're addressing, I think, one reason. There are lots, reasonable, amount of evidence that a number of drugs that get to late-stage trials, so the very expensive ones, fail because effectively, the wrong biology, the wrong targets has been optimized against. That goes back to the early stages, so basically, in the early stages, the wrong thing has been addressed.

You can address that two ways. You can come up with better techniques to come up with [00:03:00] the right target in the first places. Big pharma is

Part 4

definitely doing that. We're addressing it in a slightly different way, which is effectively, to address early-stage stroke discovery in what we call a target, agnostic way. We're using computational approaches to model disease biology and look for optimal interventions into that disease biology. To put it into the context of today, we might be slightly different, I don't know. We've hadn't augmented intelligence [00:03:30] before. I would put us in that area, so our modeling approach isn't pure AI, it's much more mechanistic. We meld all of the mechanisms of disease biology, but we use techniques from AI and multiple places within that modeling to augment the modeling, augment the data, address biases, et cetera.

Sally Taylor: Thank you. Then to you, Ben?

Ben Maruthappu: Hi, my name is Ben Maruthappu. I'm a doctor by background. I've spent the past three years advising the CEO of the NHS on innovation technology, [00:04:00] establishing a number of our key flagship programs in that space, and now I run, and am building, Cera, a textile startup that is looking to transform social care using digital and AI. We deliver home care services across a population around six million people, partnering with the public sector to deliver services that are both paid by the public sector, and independently, to deliver services around now because we've accumulated quite interesting data on all the people who have multiple co-mobilities, which is a relatively [00:04:30] rare and unique dataset. We are trying to determine how we can identify deteriorations in their health and act earlier on to prevent unnecessary A&E attendances, and hospital admissions across the system.

Sally Taylor: Thank you. Ken?

Ken Mulvany: I'm Ken Mulvany. I'm from Benevolent, like you. What Benevolent is doing is really a follow-on from my last company, which was a biotechnology company, focused mostly on neurodegenerative diseases. One of the things that I've found [00:05:00] working with that group, is there's a tremendous mismatch between humanity's ability to generate information, and our ability to use that information. When starting this group, I had two fundamental goals. One is to establish what all the facts are, and they'll come from databases, but the vast majority of our information comes from [00:05:30] literature. We ingest about 10,000 new publications every night, we extract every fact from every sentence, of every paragraph, of every paper that's published, and we relate that back to everything else that's in there to form a consistency of truth in the facts that are there. The reason we're doing that is one, establishing the facts. The second is being able to reason on top of that fact. Imagine a scientist who's able to [00:06:00] read everything that's ever been written, and everything that's known about human health, has perfect recollection of that, and is able to reason on top to find the should-be-knowns in that.

It's surprisingly rare, actually, there was a publication in January, a major publication, that mapped the actual number of novel disease-target relationships that have been made in the history of [00:06:30] drug

Part 4

development. There's about 550, so 550 times, in the history of humanity, we have mapped accurately, a target disease. Our system has done that eleven times in the last 18 months. Probably the most widely published one was the AOS society saying that there's a potential here for the disease. That came from a machine brain, rather than human [00:07:00] brain, and then we make a chemical that modulates that target.

Sally Taylor: Thank you.

Mark Swindells: Hi, I'm Mark Swindells. I'm an exscient here based in Dundee and Oxford. We are concentrating a slightly different area that's been introduced so far. We're concentrating on the area of drug discovery where we actually have to design the compound that will go into the clinical trial. Of course, you have [00:07:30] to select the correct target, or take in another target agnostic approach for that, but ultimately I think it comes down to getting the right molecule. These molecules are extremely complicated. They don't just have to have potency, but they must be absorbed, it must stomach, if they're a psychiatric disease, they must get across the blood-brain barrier. You are trying to optimize many different characteristics into the same compound. This [00:08:00] has been incredibly unproductive in the whole pharmaceutical industry. The current averages for just getting that one compound is the typical company explores, in other words, they make and they test about 1,600 compounds to get to that stage. When you consider that each one costs about 5,000 pounds, dollars, euros, to synthesize in tests, you're getting to enormous amounts of money, and it takes [00:08:30] a long time. It takes about four to five years.

What we're trying to do is use AI design, and more specifically, what we call a centaur approach where we have chemistry experts driving the algorithms to reduce both the number of compounds made, currently down to about a quarter, 400 compounds. Therefore, by virtue also, reducing the time from about 40 years, down to about 12 months, [00:09:00] and that's our objective.

Sally Taylor: Great, thank you. If I could just pick up, perhaps with Ken and Mark on this topic, we've had a loss around the likes of Google, IBM, Watson, Apple, investing heavily within the healthcare space. Do your businesses compete directly and with the tech giants, and how easy is it for others to replicate both the technology and the know-how that you're building within your businesses?

Mark Swindells: You want your chance, Ken?

Ken Mulvany: We have a pretty good idea what happens within both of the groups. [00:09:30] A lot of what's happening, in public domain where we collaborate with each of those groups in a business case. As it turns out, the head of our tech division is the former head of AI at Watson. We have kind of an understanding of some of the bottlenecks that each of the groups have and how we may complement the work that they would do. We don't actually sell our technology though, so we don't partner it. We [00:10:00] consume our own technology, we make

Part 4

medicines. We have about 20 programs that are in development, and our business model is to sell the medicine, not the technology.

Sally Taylor: I guess given the datasets that are publicly available, how much of a lead time do you have when someone catching you?

Ken Mulvany: I think some of it is a matter of specificity. If you look at benevolent as a whole, it's effectively two different companies, [00:10:30] three different companies at the moment, where you have what looks like a biotech. It's got a chief exec, it's got a chief medical officer, it's got regulatory affairs, chemists, pharmacology, all of that. Then you have a technology company with a CEO, the requisite people, and they work very closely together to define the things that are interesting to them. It's kind of a purpose-built [00:11:00] system for drug discovery and development.

Sally Taylor: Perhaps you could talk to that, Mark, and perhaps linking in some of the deals that you've got?

Mark Swindells: Just to give an analogy, just trying to think of one whilst you were talking. I guess if I was given the choice of an expert surgeon using AI, or an expert AI company that thought they could get into surgery, [00:11:30] I think I know which one I might select, and I think that's the core of the problem. I think some of the big companies are trying to bite off too much. They do have a lot of budget, far more budget than we have, but I get the feeling that somebody above them is saying, "You must be able to sort out every single problem." Even within the drug discovery area, we're very lucky. Our chief chemist is Andy Bell. He was the co-inventor [00:12:00] of Viagra, and he also was a key part of the team that got voriconazole, which is a less well-known, but a very important anti-fungal to market. Not many people have got a drug to market, nevertheless, we've got two. It's trying to instill the knowledge that you get into the algorithms and apply those in an intelligent manner.

That's why we use this reverse scent or analogy. The machines won't work perfectly [00:12:30] all the time, but what they can do is they can empower a specialist to be roughly at four times more productive. We can conduct that for projects in parallel. Whereas, in the old days, you would have had all your mind on a single issue. That's why I don't really think that we have competitors, really. The likes of the big companies, because I just think they're spread too broad.

Sally Taylor: That's helpful. Thank you. Jonny, I think when new therapy suits is being focused and [inaudible 00:12:59] [00:13:00] and you've got a deep background in this space, perhaps you can give your insight over 20 years in terms of what the changes has been. Maybe the recent breakthroughs and how e-therapy suits is changing its approach to encompass here?

Jonny Wray: I guess some of it follows on, I think in previous comments, as you mentioned, I've been in competitional biology for more than 20 years now. I think there's

Part 4

been a numerous peaks [00:13:30] and troughs in both AI and the application of competitional approaches to biology where the competitional approaches are over-promised and under-delivered. Systems biology maybe 10, 15 years ago was probably the most recent one that's in my head. I think where we are now, though, is maybe slightly different. I think it's a combination, actually, I think it's mainly driven, at least in my [00:14:00] narrow area, basically by the data. If you take all of the technologies of post-eunomic biology that come on next generation sequencing, et cetera, that have the capabilities to generate a lot of data, a lot of valuable potential knowledge about the biological system, these data sets are too big and complex for individuals to interpret on their own. More advance competitional techniques [00:14:30] are needed to understand that. I think that, in our industry, is becoming very apparent and accepted.

Sally Taylor: Thank you. Ben, perhaps you can discuss how Cera is competing with encumbrance within the social care market and you are generating significant data, granularity of that data, and at what point you'll use that to potentially predict, and how easy is it to sell within the entrenched NHS and social care system, and how willing is it to change and embrace new [00:15:00] technologies?

Ben Maruthappu: Sure. I think if you look at the social care space, it's one where you have thousands of very fragmented care providers who are largely offline. It's a cottage industry, mom-and-pop shops, so there aren't really any pioneers who are trying to push the boundaries, innovate, explore digital, let alone AI. I mean, a number of these companies don't even use e-mail yet. When it comes to our competition, they're probably quite analog or flying-in-nature. That presents really opportunity combined with the fact that social care is [00:15:30] in crisis at the moment and has tremendous knock-on effects on the health service. I think you rightly asked how it is to work with the NHS. I mean, when I was on the opposite side of the fence, people found the procurement cycles extremely tedious, life science companies who wanted to strike up new collaborations or drug deals, also found it quite tricky given some of the funding constraints the health service was facing. We've been all right, actually, simply because we're kind of pushing at a semi-open door, which is bed-blocking, [00:16:00] as you rightly pointed out.

Because social care, the moment it's poured from a liquidity and agility standpoint, there aren't enough care workers and care cannot be organized quickly enough for people as they move from hospital to home. This is causing tremendous operational issues in trust up and down the country. It's why we saw the A&E crisis last year. It's very likely why we will see it this year as well. People are really now trying to focus on this at the local level, but also at a national level with the chancellor [00:16:30] announcing an extra £2 billion to be focused specifically on this problem. That has meant that tackling the problem, and partnering with the health service to address it has actually been relatively straightforward.

Part 4

Sally Taylor: Just stepping back a bit, there's a number of investors here today and certainly from Lamis' perspective, this area is so fast-moving. What should we be asking all of you in getting to know your businesses? What's the killer question if you like, and maybe if you can answer that, I don't know who wants to go first?

Ken Mulvany: I [00:17:00] think it's demonstration of the utility of your system with investors. There's so much hype around AI. There's so much hype around it, and being able to deliver on what you're saying that you're going to deliver, I think is ... I invest as well, I invest in a lot of companies, and what I want to see all the time, is okay show me this booking.

Sally Taylor: [00:17:30] I guess that's the challenge within healthcare because timelines are a bit longer than perhaps in other sectors. What are the KPIs that we should be looking for in the nearer term within this space?

Jonny Wray: I guess I would like to just extend on that. Does it work? Is certainly the first question. I think there's some other issues, especially, it comes back to my previous answer to why some technology solutions over the years appear to have worked, and then haven't [00:18:00] been taken up. That's because their practical feasibility in the area that they've been trying to apply to hasn't been addressed. This typically will come out of people who don't know about the industry. Can you get a hold of the data you need, et cetera, et cetera. I think maybe a more fine-grained answer to: "Does it work?" There are different aspects to the "does it work". In terms of how we address progress, it's a good question in drug discovery. 10, [00:18:30] 15 years will be the ultimate promise of what a lot of us are working on. Obviously, that's too long. Our industry has standard ways of measuring progress, and it's progress into clinical trials and getting through. Not much more than that, I don't think.

Sally Taylor: That's great. Ben, you could answer that for your business, and KPIs that we should be thinking about.

Ben Maruthappu: I will answer that. I just wanted to touch on the previous question, which is the investor one, because I've got both [00:19:00] current and prospective investors who are actually in this room. They're probably wanting to know about that. I think a question that commonly comes up here is, "How do you monetize the data?" Nick rightly pointed out in the last panel that lots of AI algorithms are actually open sourced, and will probably continue to be. The value chain is moving towards being more data-centric in some regard. I think broadly, it's either using the data to provide a high-quality service that gives you a conducive edge, or selling it, which is [00:19:30] somewhat more controversial, but you need to think about how you can anonymize the data, who has the capital to deploy it? Actually purchase it, and analyze it in a respective way, it could be research units, insurance companies, pharma, but of course, there are lots of points there regarding privacy, confidentiality that need to be taken into account.

Part 4

	Roughly, on the monetization of data, for some of the people I know out there, it's using it and selling it. In terms of KPIs, as we heard, it can be difficult. Essentially, [00:20:00] it's identification of surrogate markers, just to be concise in that, where in time. That's how you classically solve that problem.
Sally Taylor:	That's very helpful. I guess stepping back a bit, we traditionally think of the West Coast as the leader in tech, where do you think the UK can compete? This is possibly a leading question given it's a healthcare panel, but does the UK's domain expertise in science, or the unique structure of the NHS, does it lend itself for the UK to be more innovative in certain areas? Mark?
Mark Swindells:	[00:20:30] Well, in my sector, I think the UK has always been very good at extracting knowledge from data. I think one of the reasons for that, I think back 20 years, way back then, we couldn't afford these very expensive computers, and therefore, we had to be smarter with the amount of money that we were given. Therefore, we have become, at least in my area, world leaders on how to extract data [00:21:00] and to, I want to say repackage it and give it higher value. I think that there's enough people who have worked in the UK, and probably Europe as well, over that period, that that's where our strength is. I think if we can, Ken's company, we're all trying to do very much that same sort of approach. Of course, fortunately, the markets come to us. Now, of course, computing is far cheaper than it was 20 years ago. I think we've got a very, [00:21:30] very strong platform, and that we should be really strong in touting on differentiators from the West Coast, because we do things differently and we should be proud of it.
Sally Taylor:	Absolutely. Has anyone got anything else to add on?
Jonny Wray:	Your second piece about the NHS, I think there is an advantage there of having an integrated healthcare system, especially when it comes to questions like personalized medicine, and how we address this drug discovery [00:22:00] or doing clinical trials, et cetera, in that context. It's a data collection issue. I lived in the US for 20 years, their healthcare system is horribly un-integrated, and I'd be very surprised if they managed to get the amount of clinical data associated with the potential molecular, et cetera, data, anywhere near to the level that we could potentially do it if it was done correctly [00:22:30] using the centralized health system that we have.
Sally Taylor:	Just to open up to the audience, has anyone got any questions? We've got a question here, gentlemen.
Male:	I should state up front, I'm not a healthcare specialist, but I do study innovation. Listening to you, something strikes me as an interesting thought, which is that difference between the computational up-stream methods that we now have, and the history of medicine and accidental discovery. If [00:23:00] I look at the people I know in their 70s and 80s, they take five, six, seven different kinds of medication. There must be a lot of accidental discovery data available on how these drugs actually interact in real use if we could access it and work

Part 4

	backwards to discover things. How does that contrast with what you're trying to do up-stream in discovering things?
Jonny Wray:	The approach you call it, there's a competitional academiology approach out there that people use. [00:23:30] I think one of the biggest barriers is access to the data, again, through the NHS. You can't just get that data unless you go through the issues of privacy, et cetera. From our point of view, that would be a statistical approach, people who take drug X are less likely to get cancer, et cetera. Personally, we're standing at the other end and trying to come up with mechanistic hypothesis of potentially how cancer starts in the first place to look at intervention. [00:24:00] From our point of view, it's a mechanistic verses a statistical, and at the very different ends of the pipeline.
Sally Taylor:	Ken, did you have some thoughts?
Ken Mulvany:	Yes, it's an interesting question. The way our industry, the way that scientists look at the scientific innovation piece of that can be tracked through literature. Everybody uses the scientific method, every paper is written the same sort of way. It's interesting because you can use that to [00:24:30] track how good your hypothesis are in predicting. I'll give you an example. I could stop my literature ingestment in 2005 and see whether or not I can predict what happens in 2015. What is the epic of information that exists in the 2005 that's never been made before? That's important. No one has ever made this particular hypothesis before. Can I predict whether or not that hypothesis has born out to be true? The [00:25:00] information is there. You can train your algorithms to do that sort of thing. I think the way that we've done it now, I think we're about 75% accurate in predicting. We're not 100%, we're 75% accurate in predicting those hypothesis. Actually, industry is 6%, so we're significantly better than industry in making these predictions. The obvious thing to do is ingest everything up until today and tell me what I could [00:25:30] predict what would happen in 2023. That's actually what we do now.
Male:	In terms of why you have a drug discovery and a tech because big pharma have been around for ages, they know how to bring these drugs to market and get them through pre-clinical up to phase three and into market. You're kind of replicating something that already exists. Do you add value [00:26:00] to that part of the chain as well?
Ken Mulvany:	We could do. About two thirds of the programs we develop, we will license after a phase two. About a third of those in the areas of diseases of the brain, central nervous system, and inflammation, we keep and will take those to market. Our most advanced programs are phase two B, and then we have other clinical assets that we'll be partnering. [00:26:30] For that exact reason, we have an immunology, stroke, allergy franchise, we're not going to build a sales force to attract that, but in areas where it's specialty, CNS, hospital potentially, those are areas that as a company, we can attack. That was the intention of my last company, was a biotech before the company was sold to a US pharma company. To hold [00:27:00] onto assets to bring them later in development, and

hopefully to market. Still flying that flag for the UK. I was really outspoken. I will never sell a UK company, but then I did. There's probably a whole new thought to it.

Mark Swindells: They're really good at doing the clinical trials, which are very expensive and they've got a big sales force, but there's no data to suggest that they're remarkably good at doing the rest. If they were fantastic, it wouldn't need 20 efforts of which 19 fail, [00:27:30] to get to a job. There's clearly enormous room for maneuver within that space. How long have you got? There's a lot of challenges. Wrong concept, wrong target, wrong linkage, and wrong compound. Put those all together, [00:28:00] you can go wrong in many, many ways.

Ken Mulvany: You still have 50% of drugs failing late stage development because of the wrong biology.

Mark Swindells: One, very, very recently.

Ken Mulvany: It's extraordinary. You don't get the biology right to begin with.

Mark Swindells: You're doomed.

Sally Taylor: That's the last question, thank you.

Male: We had a joke, I think it was at MERC, where three weeks of work in a lab will save you one hour in a library. One of the things that we struggled [00:28:30] with, I was a portfolio director at MERC, and then during the show, plan merger, and then more laterally, GSK. One of the things we struggled with was the economics with the pre-2A space, particularly because effectively, our net present value basis, it's negative. Effectively, a cost of doing business. I guess the question I would be interested in, I think Ken, you're starting to touch on, I was wondering if you could expand on it, which is the idea of how are you reducing the capital intensity of the early space, you're actually turning into a value creation center as opposed to a cost. [00:29:00] Have you got any early reads on return on capital in that space? I know it's early experience.

Ken Mulvany: I do. I've got some statistical, but I also have some anecdotal stuff as well. My last company, I said it before, it was [Proximigen 00:29:16], the largest biotech exit here in New York about 10 years when it went. I had a pipeline roughly what mine looks like now, Benevolence is a bit bigger, it's more advanced, [00:29:30] but the point is, it took me 10 years to build that pipeline, 10 years. I have this pipeline in 18 months. Some of the things that we're talking about here in terms of speeding up the chemistry design process, predicting whether it's the right biology, I called them alpha-go moments because like Alpha-go from Deep Mind, you're not a go champion, [00:30:00] you play a go-champion, you beat a go-champion, it's probably because of the tech. In our space, you have a disease, which humanity has never been able to treat a cure, a machine comes up with that disease, we obviously then have to have it tested, and we

Part 4

have it tested the way you normally would, and lo and behold it works. It was because of the tech. We don't have any particular experience in gleo-blastoma or in ALS.

[00:30:30] I think it's a really exciting time for us right now deploying this type of technology. From a pharma perspective, I'm sure everyone will be on board. There's a lot of cultural resistance within pharma, but some of the proof is just evident. It's there to be seen.

Jonny Wray: I guess to emphasis that from our anecdotal, process now, basically, the tech side, the prediction [00:31:00] of what the biology is, isn't a bottleneck at all. The bottleneck now at the early stage of discovery has become going into the lab and testing it. We've basically managed to eliminate that, coming up with novel biology and novel hypothesis from it being a bottleneck at all.

Sally Taylor: That's great. Thank you very much, just conscious of time. Thank you very much to all the panelists today. It's been very insightful and it is very exciting times for the future of AI applications in healthcare. I think well summed-up by the late Steve Jobs who towards the end of his life in 2011 [00:31:30] said that he thought that the biggest innovations of the 21st century will be at the intersection of biology and technology, a new era is beginning. Wise words indeed. With that, I'll hand over to Gary for closing comments.

Male: Thank you for your comments, just to say thank you very much to all our speakers and all the panelists, who [inaudible 00:31:51] conversation and for good reasons, we've got all of your data [inaudible 00:31:55] emailing you with [inaudible 00:31:59] running over the next [00:32:00] six months and hopefully, we'll get them out when the recording of this at some point this afternoon, if not, Monday morning. Thanks to you all for coming.