FRONTIER TECH INVESTOR

ISSUE NO. 2 JUNE 2016

Superintelligent, superfast machines

Artificial intelligence, quantum computing and the #1 tech firm to take advantage

Eoin Treacy

"If I have seen further, it is by standing on the shoulders of giants"

Isaac Newton

I'm "doubling up" in this month's issue of *Frontier Tech Investor*. I want to introduce not one but two of the most exciting – and important – themes developing in the technology world right now: the quest to create artificial intelligence machines (AI), and the race to develop quantum computing.

It's fascinating stuff. But we're not here just to gaze in wonder at the incredible things technology could help humanity achieve in the coming years – we're here to take advantage of them financially too. As our mantra goes, don't fear rapid change – master it, by first understanding it and then making the right investments at the right time to profit from it.

With that in mind, I want to introduce you to a company at the heart of this story. It's a core investment that I believe should be present in any tech investor's portfolio. It's also a company that you'll doubtless be familiar with – but don't let that put you off. There's a compelling case that this company could play a major role in developing and commercialising both AI and quantum computing.

The firm I'm talking about, of course, is IBM.

From floppy disks to intelligent machines

IBM celebrated its centenary in 2011 so to say it has been around for a long time is an understatement. A survival record like that in the technology sector is unheard of considering the threat of obsolescence and the increasingly low cost of competition.

The secret to IBM's success is simple. It has an astonishing track record of coming up with cutting edge technologies that push so far

Intelligence bulletin

Supercharged machine learning
Last month Google announced a new breakthrough in its effort to create machines capable of learning – seen by many as a forerunner of AI. The secret? To spend more time working on newer, faster hardware – instead of just more powerful software. As the company itself put it: Great software shines brightest with great hardware underneath. That's why we started a stealthy project at Google several years ago to see what we could accomplish with our own custom accelerators for machine learning applications.

The end of oil is coming
You've heard of peak supply – that one day, we'll run out of oil and that supplies are therefore in terminal decline. But what if we're approaching peak demand? According to consultancy firm McKinsey, we could hit that point by 2030. Lower GDP, a shift away from energy intense industry and the rise of renewables could all put the nail in the oil coffin.

beyond the realm of the possible that they become indispensable as soon as they hit the market.

I'm thinking about the impact inventions like the ATM, floppy disk, hard drive, financial swaps, dynamic random access memory (DRAM) and a host of others have had. The company's employees have racked up Nobel Prizes, Turing Prizes, National Medals of Technology and National Medals of Science among a host of other awards. But despite all of this success it is in the past, and markets are all about the future. It's what IBM is turning its considerable expertise to next that we're interested in.

What anyone can see from the above list of products IBM created is they are all hardware related. However almost all hardware is now manufactured in Asia, and the fungibility of designs is such that competition and obsolescence are almost taken for granted in the modern era.

Management recognised this years ago and has been selling off hardware divisions for years. The most recent glaring example of a prescient disposal was the sale to Lenovo of IBM's desktop and laptop divisions. Lenovo was initially able to bring China's vast manufacturing economies of scale to bear which allowed the company to flourish, but that didn't last as tablets have killed off growth in the desktop market. Lenovo has since been through a number of failed product launches that has seen its share price collapse.

As one of the largest employers in the US, IBM's efforts to sell off units at risk of becoming outmoded or declining in the face of intense competition means it has had to lay off workers; while at the same time hiring legions of people with the requisite skills to do the important work it is now pursuing. That gives rise to a lot of online commentary about how disloyal the company is to past employees. But the reality remains that IBM has embarked on a massive overhaul of its operations which reflect a renewed focus on services, software and cloud computing.

Let's put some flesh on what this transformation has meant for the company. First off, revenue is down from $104 billion in 2012 to $81 billion for 2015. That's not a great outcome and highlights just how much of a transition has already taken place. Its Global Technology Services are in serious decline, and the company is betting that its software and network/cloud solutions will eventually compensate for that loss. It's a big transition, and it has to work if IBM is going to prosper in the next century.

The beating heart of the tech world

In the digital era we are creating data at an ever increasing rate. If you have kids you know how quickly the storage on your phone fills up with games, screen grabs, photos and videos. Everyone, everywhere, is doing this all the time and the desire for storage is booming. As the world becomes ever more advanced technologically speaking, it creates more and more data.

More importantly, the desire to parse this data to distil nuggets of information about our lifestyles, choices, needs, wants and longings is driving a wave of companies to invest in big data programs. In the decades before now there was a problem with accumulating sufficient data to run programs, but today the issue is the complete opposite. We now have oodles of data being created on a daily basis, and the devices we rely on to create the data, namely mobile phones and desktops, don't have even close to the computing power to parse the data.

IBM sits very squarely in the nexus of that surge in demand. Other major technology firms like Google, Facebook and Apple among others invest heavily in this kind of technology for their own purposes. On the other hand, Amazon and IBM are at the cutting edge of providing technology as a service for big data applications. Between those two, IBM is way ahead because it's what it focuses its attention on whereas it's an evolving side business for Amazon.

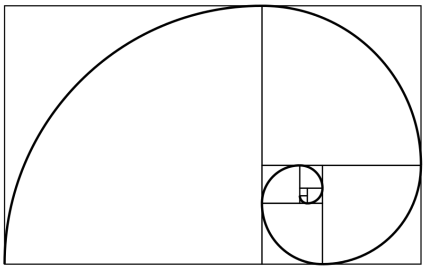

I first came across the Fibonacci spiral while at university when it was used to describe how the numbers of rabbits in a population can expand exponentially. With greater knowledge of exponential growth we can now see the pattern in all sorts of natural phenomena. For example, the nautilus shell is almost an exact match for the illustration to the left. Meanwhile the ratio is found all over nature in everything from the distribution of leaves on a stem to the family tree of honeybees.

At this stage you might be wondering where on earth I'm going with this, but the answer is simple. We can see these patterns in nature now because we came up with a sequence to characterise it. What if there are other sequences that exist in nature and consumer behaviour that would lead to a much greater understanding of our ecosystem, but we do not yet have the tools to make those discoveries? What if there is a pattern hidden in all the data being created by technology – the reams of consumer, medical, financial, climate, transport and economic data – that we are yet to understand but could yield incredible results? This is the goal of big data: to find the meanings hidden in the mountains of information.

Mainframe and supercomputers are now testing that hypothesis with datasets so large it's difficult to get one's head around. That's both a business opportunity and a challenge. Think about it, we can use superfast computers to explore how bacteria evolve to generate resistance to antibiotics and experiment with ways to slow them down. The global climate is about the most complex puzzle facing humanity right now, and yet our ability to predict the weather is woefully inadequate. Each of the people on earth has about 100,000,000,000 individual nerve endings in their brain. When they start firing the number of possible permutations it isn't quite infinite, but it's so close it doesn't really matter.

The opportunity is that if a company has a computer so powerful that it can derive an answer for just about any question you might wish to ask, then it has an enormous amount of commercial potential. It can simply lease, replicate or sell that resource out. That's just what IBM is planning by pioneering cloud as a service business model. Previously the company sold mainframe computer networks which were expensive to buy, install, maintain and staff. It was a great business but is likely to be dead in less than a decade. It is being replaced by the much cheaper mobile cloud where companies can have their data securely warehoused for employees to access wherever they are. When those employees need to beef up with the data usage, the company just leases a little more for the time it is needed. That means they don't have to go out and buy whole new servers.

Turing mountains of data into artificial intelligence

But IBM isn't stopping at simply storing and analysing data.

It's also developing ways of turning that data – and the patterns, hidden meanings and other as yet unknown discoveries within – into artificial intelligence.

Teaching a computer to think like a human or to respond to natural language is an enormously time consuming process. In many

respects it's like teaching a child to recognise pictures. Anyone with kids will remember holding up a picture of a dog and saying "What does the dog say?" From about 12 months old the child will say "woof". Teaching a computer to do the same thing requires millions of pictures to be inputted into a library and then some complex code to tell the computer to identify what is common to all dogs and what isn't. Every one of us comes up with those answers instinctively and with a little help from our parents, but a computer has to be painstakingly taught.

A running joke in our home is to say "Dunno phone" whenever we don't know something. My daughter derived the term when she was a year and half, and her favourite book was a book with lots of pictures of random objects. She would sit between my wife and I before bed and recite what each was but would always get stuck on saxophone and would say "Dunno Phone." It's a happy memory for us, but it's a good example of substitutional learning that computers are not good at. Humans have little trouble identifying what we know about something and what we don't know. Computers need a very complex set of rules and algorithms to come up with approximate answers. In a world of 1's and 0's, yes and no answers are easy – "maybe" requires a lot more subtlety.

In order to overcome that challenge the number of libraries an artificial intelligence program has to consume is vast. The tricky part is teaching the computer to learn from its mistakes. The good news is you only have to teach a computer once.

Computers only need to learn a set of rules once. It might take time to teach them but they never forget a lesson once learned. Developing artificial intelligence libraries has been enormously time consuming, but it has been going on for a long time already. It began accelerating at an exponential rate with the invention of the internet, because the availability of digitised information to train on is increasing exponentially.

IBM has form in the AI world, too. In 1997, its supercomputer Deep Blue defeated Garry Kasparov at chess. That was an amazing feat of processing power. We still talk about it today. And yet the computing power of mobile phones is now greater than Deep Blue. You are literally walking about with a 1990s supercomputer in your pocket.

After Deep Blue beat Kasparov, IBM began looking for new challenges. That led to the next generation of intelligent machines. Specifically, Watson.

Meet Watson

Watson runs on today's supercomputer s and is a natural language machine designed to answer questions posed by regular people, and it succeeded by defeating the reigning champion of the US gameshow Jeopardy in 2011. It can process about 500 gigabytes of information per second or the equivalent of a million books per second, and it will only get better at it as it is taught how to do so even more economically and with additional power.

Demonstrating party pieces by beating humans in a highly structured game environment is all well and good, but it doesn't exactly make money (unless of course IBM was to let Watson loose on a poker site). IBM chose the healthcare sector as the first field in which Watson could have real-world applications.

Why? Well about 12 million people seeking outpatient care are misdiagnosed in the US every year. That amounts to a great deal of needless suffering. Unleashing an analytical computer program on that problem has the capacity to really improve patient outcomes, reduce the risk of medical practice lawsuits and lower costs all round. That's exactly what IBM set out to do when it rolled out Watson as a diagnostic tool from 2011, and the number of hospitals and universities developing questions for the system has increased rapidly since. It's not too difficult to imagine that with success rates improving, Watson will become a ubiquitous medical assistant in every hospital in the world. (For a fascinating industry perspective on this, see the feature on page X for our exclusive interview with Jeremy Howard, CEO of Enlitic – a company dedicated to using machine learning to diagnose illness.)

Of course the system's applications are not limited to healthcare. Processing client data and requests is an enormous job most companies are incapable of completing internally. IBM is working with partners to develop applications in the retail sector to provide advice on shopping. An app developed with Welltok is due to launch on the Apple App Store next year and will be the company's first digital health concierge for the Apple Watch.

Not content with that, Watson has also turned its hand to fashion design. IBM teamed up with designers to produce a dress for the Met Gala this year, which was influenced by the conclusions made by Watson to judge how it would be received by audiences.

And if that doesn't do it for you, how about letting Watson loose on the legal industry. Law libraries are already a thing of the past, but having a cognitive assistant by your side that happens to have an encyclopaedic knowledge of case law is invaluable when on a tight deadline. By entering the legal services field, IBM is taking on established companies like LexisNexis, Elite, Eclipse, Thomson Reuters and Bloomberg but it's fertile ground for the type of service it brings to the table.

Would you prefer to have a Watson as your banker than a person? I'm sure your answer is likely to depend heavily on your experience at various banks and how easily you have been able to secure a loan when you wanted one. What if a cognitive loan officer had the capability to look at more factors than your credit score and come up with a truly informed decision about how reliable a debtor you might be? It's already starting to become a reality for Westpac New Zealand customers though their partnership with IBM.

When your car is truly hands free, having an ability to communicate both with a central information hub as well as the vehicles around you will be invaluable and will require the kind of computing capability and rapid sifting of inflation Watson has been bred to deliver.

But all that is only the beginning.

IBM could make billions from AI

Quantum computing is coming – and IBM is leading the charge

What does the digital age really

Quantum computing is coming

mean? It's a term bandied around a lot, but at its absolute simplest all it means is that each piece of information is being encoded onto a computer system so that it can be catalogued, analysed and retrieved. Look a little deeper and everything is recorded as a series of 1's and 0's on a vast network of computers that spans the world. Despite the fact that the number of libraries and coding languages continues to proliferate, the basis of everything they do is still 1's and 0's. What if that was all about to change?

Richard Feynman, the granddaddy of quantum computing, first theorised in 1981 that quantum theory had matured, at least to a point, and that computer science was almost exclusively devoted to physical computation. What he desired to see was a computer capable of simulating the relationships between subatomic particles which do not follow the rationale of the physical world. He was talking about a quantum computer, and today the very first working prototypes are up and running. It's only been 25 years.

We're in a hurry because Moore's Law, since first observed in 1965, has seen the number of resistors on a chip double every 18 months since. Today the smallest process has 14nm (nanometres) wide spaces between its resistors and a silicon atom is 0.2nm wide. With such tight conditions and an increasingly voracious appetite for smaller, quicker and lighter products the race is on to come up with the next big thing. If Moore's Law continues we need to either come up with a different cheap material (possibly graphene), much better manufacturing processes and/or we need to look inside the atom for how to make faster computers.

That's exactly what quantum computers do. They achieve speeds that make even the fastest computers look prehistoric. They achieve this lofty goal through a quantum mechanics effect known as quantum parallelism. By existing across billions of planes, on a subatomic level, concurrently, the computer can run the problem across all iterations of itself at the same time with the result that answers pop out a lot faster.

That's not to say they are the answer to every problem or that we are going to have personal quantum computers any time soon or ever. The one thing they are really good at is processing vast amounts of data instantly to come up with answers to problems we didn't even know existed. IBM just happens to have a five qubit machine, which it developed with the aim of leasing its computing power via its cloud services – so companies and governments can make use of it to conduct the experiments and research that were literally impossible until now. This is the first of its kind and is absolutely scalable.

How IBM could make billions from AI and quantum computing

This is all exciting stuff, but where the rubber meets the road, so to speak, is in hard cash. As we discussed earlier IBM is in a major transition as it moves from providing hardware and consultancy, to software and computing power as a service. In an age where information and data and the ability to refine it for profit are driving major corporate profitability, there are obvious advantages to the business model IBM is now pursuing. However this is not coming without a price. IBM's revenue has been falling for the last four years, dropping from a peak of $104 billion in 2012 to a low of $81 billion last year. The losses have not been confined to any one business segment but have been broad based, suggesting the company has challenges that is has yet to resolve.

Some other important facts worth knowing about IBM are that it is an S&P 500 High Yield Dividend Aristocrat, which means it has been raising dividends for at least 25 consecutive years. Dividend growth has continued at a 14% to 15% rate since the early 1990s and the yield is currently 3.67%. With that kind of record all you have to do is hold the share to outperform gilt or Treasury bond yields.

IBM has an inexpensive P/E ratio of 10.18 which is estimated to rise to 11.28 by the end of the year. However the valuation does not tell the whole story. IBM shares are becoming collectors' items because it buys back so many of them. The board voted in 2015 to extend the programme which will purchase $4 billion this year. By reducing the number of shares in issuance it boosts earnings per share, which helps to artificially reduce the P/E ratio. While $4 billion might sound like a lot, IBM is not included in the PowerShares Buyback Achievers ETF constituents because it does buy back enough (5% of free float) to merit membership. That suggests that while we might argue management could find better uses for that money, the reality is it can afford the buyback programme.

With such a long history, IBM has a very different capital structure from other household name technology companies in that it makes abundant use of debt. With an A+ rating it has therefore benefitted enormously from the low interest rate environment we now find ourselves in, because debt that previously required a coupon of 7% or 8% can now be issued at less than 3%.Its most liquid bonds now yield less than the shares themselves do. That suggests equity investors might come lower down on the pecking order, but they are being amply rewarded.

All of which adds up to mean the conclusion I consider IBM a "must own" core investment in any tech focussed portfolio. I'm placing it in the Technology category of the *Frontier Tech Investor* portfolio. It's a perfect counterbalance to some of our higher risk plays (such as our "moonshot" company, SolarWindow).

The share was among the better performers following the credit crisis, but it hit a medium-term peak in 2013 and almost halved before finding support in February. It has since rallied to close the overextension relative to the trend mean, represented by the 200-day moving average and has been consolidating above it since March. I rate the share BUY between $140 and $160 and believe a stop would be appropriate around $135.

The most potent risk is that it fails to capitalise on the impressive suite of new technologies it has developed to prosper in the modern age. If that is the case then it will extend what has already been a three-year downtrend. The potential for the Federal Reserve to raise interest rates also represents a threat, because with so much debt IBM will no longer have the luxury of refinancing at progressively lower levels. Higher debt repayments would take a toll on the company's ability to continue to fund share buybacks, which would also act as a headwind.

Name:	**International Business Machines Corporation**
Ticker:	**IBM LN**
Exchange:	**New York**
Price close of 07.06.16:	**US $152.98**
Market Cap:	**US $146,855m**
Buy Between:	**$140 and $160**

Five Year Performance:

2011 +25.29% | 2012 +4.17% | 2013 -2.08% | 2014 -14.46% | 2015 -14.22% |

How AI could revolutionise the medical world – a conversation with an industry expert

Nick O'Connor
Publisher

One of our goals here at *Frontier Tech Investor* is to go a step further than simply helping you invest successfully – it's to help you understand exactly how world changing technology is developing. To do this my strategy is simple: go out and speak to the most respected and informed experts in the field, and report their insights to you.

With that in mind, I'd like to share a conversation we had with Jeremy Howard, CEO of Enlitic. Enlitic is a firm that's essentially

a competitor to IBM in the field of using AI and machine learning to diagnose illness (in that sense, take his comments on IBM with a pinch of salt – he's a competitor after all). But his insights into how AI could work both medically and commercially are invaluable.

Q: Where did you get the idea to create your company, Enlitic, in 2014? It seems odd for someone with no medical background to set up a company with exclusively medical uses.

A: I started out my career in management consulting. I was at McKinsey & Co. I was kind of unique there, in that my approach to solving problems was based on rigorous data analysis. Other people at McKinsey were more reliant on interviewing people, and on subject matter expertise.

It was partly because of my limitations that I had to invent this different approach. I was only 18, so I didn't have any subject matter expertise. I had to find other ways to get good answers. I started from the very earliest time being focused on data-driven problem solving.

Q: How did you apply data to your job at McKinsey?

A: Everybody has their own way of answering whatever questions they have to answer on a particular study. I just created a database and put all the information I could find in the database, analysed it, and used linear programming to find optimal approaches – a much more data-driven approach. It was just how I happened to go about answering the questions. I discovered that when I talked to the project director, they would say: "This is a really unusual way of tackling the problem." It was the only way I knew.

Q: Did you study computer science?

A: No, I didn't. I studied philosophy. Because I started work full time when I was 18, I didn't have a degree. I was self-taught. I got a degree without going to the lectures -- just doing exams. When you study philosophy you can do that!

Q: It's fascinating to see someone like you studying philosophy. One would imagine you doing computer science.

A: I guess so. That's why I've always been a bit of an outsider, and I think that's a good thing. I'm a total outsider in medicine now. I've always been an outsider in computer science.

I started a company which was entirely based on analysis. It was a new approach to pricing insurance. It was called Optimal Decisions Group, and based on a new algorithm. I also started an email company. We tried to build the biggest email business we could with the smallest number of people we could.

When Kaggle came along – this international community of machine-learning practitioners who compete against each other to come up with lots of different algorithms – I entered the competition with the expectation that I would come last, because I didn't have the background that everybody else had. I was very surprised that I actually came first. That was when I realised that my unusual background means I have skills that other people don't have.

I ended up joining Kaggle full time. I became the president of the company. I got to know most of the world's best machine-learning practitioners through that. I discovered that many of them had unusual backgrounds like I did.

While I was there, I discovered that one particular algorithm was increasingly winning the hardest competitions. The hardest competitions by definition are the most challenging data problems of the day: things like finding the Higgs boson, or mapping the dark matter in the universe, or big difficult questions that people hadn't been able to solve before. This particular algorithm was deep learning, and heavily related to something called neural networks, which I did a bit of 20 years ago.

Neural networks allow computers to learn from examples

Neural networks are a way of doing machine learning, and machine learning is anything which allows a computer to learn to solve the problem from examples, rather than you having to program it step by step.

In the past, computers have always been very limited by the fact that they only work on structured data -- data that neatly fits into rows and columns, as you might have in a database or a spreadsheet. It works well for things like Amazon's book recommendations, or Google's search engine. It hasn't worked for things like an MRI scan or an ECG. Deep learning allows us to apply machine learning to unstructured data. The neural network is the underlying mathematical function that deep learning uses. It's a function that is a simplification of the network of the human brain. Therefore it's flexible enough to be able to handle anything we channel it.

I got to the point at Kaggle where I had a sense of déjà vu.

Back at McKinsey, I was there at the birth of the commercial internet. Everywhere I looked, I thought, "The internet is going to change that, that and that." I talked to a bunch of people at McKinsey about it, and they all said, "No, don't be silly. This is business, and the internet is students and hackers." So I ignored my gut feelings, and I stayed there for another eight years. In the end, that feeling of "the internet is going to be everywhere" was right.

I got that same feeling about deep learning. One of the reasons I left Kaggle was because I didn't want to have this happen again: to be sitting on the sidelines while something changed our world.

The other piece which influenced this was that after I sold my second company, I had a sense of disappointment. I'd achieved everything I had set out to achieve, and yet I didn't feel like I had made a significant dent on the world. I didn't feel proud. That seems a bit wasteful after ten years of hard work. I was determined that the next time I embarked on something seriously, I would pick something that I would spend 25 years on, and if it was successful, that I would feel proud, that it would be important.

So I spent a year doing nothing but research into the impact of deep learning on the world. The way I did that research was, I prepared a presentation about it, which to some extent became the Ted Talk.

I asked a lot of Nobel Prize winners and CEOs and government leaders: "This is what we can do thanks to deep learning. What would it mean if you had access to that technology?" And every one of them said, "That would totally change everything in my business," or "That would totally change everything in my research."

Their response totally confirmed my gut feeling. I thought, what's the biggest of these opportunities? To some extent, I left medicine aside, because if you're not involved in medicine, it seems like something other people do. But I started to see people saying, "Unstructured data -- images and sounds and words – are the thing that doctors rely on nearly entirely. They're what deep learning lets us analyse with the computer."

I started to realise that there was a whole different way of doing medicine: what we now call data-driven medicine, where we figure out how to diagnose people and make them better by using rigorous computer analysis of their medical data. That's totally different to the artisanal way which has happened for the last 1,000 years.

In the end, I started Enlitic because, after that year of research, I realised that (a) deep learning is definitely going to change the world, and (b) medicine is the place where it's going to have the most impact. It's the world's largest industry, and the one where the social impact is highest, because there are four billion people that don't have access to medicine, and that's largely because of the shortage of expertise. Finally, because the type of data that doctors use is exactly the kind of data that deep learning is good at analysing, it seemed like a perfect match -- although it's pretty scary to get into that world knowing nothing about it.

Q: At the moment you've got this technology that allows radiologists to read scans more effectively. You're not replacing the radiologist, but rather the computer is assisting the radiologist in arriving at a diagnosis. How foolproof and

reliable is this technology? Is there a risk that an erroneous diagnosis by the computer would mislead the radiologist?

A: Actually, the algorithm doesn't do any diagnosing. The main purpose is to guide the radiologist to the interesting part of the scans, faster and more accurately. The worst thing that could happen would be that we should fail to guide them to an interesting part of the scan. So the technology is only as good as the radiologist who's using it. They would never want to rely 100% on assuming that the computer catches everything.

Humans miss things, and computers miss things. The search results we have so far show that computers miss fewer things than humans, but they still miss things. So if humans and computers are both looking for things, they're going to see even more things.

Radiologists still need to be thoughtful and careful, but by leveraging this thing which is saying "Hey! Did you look over here?" they can focus on the important bits.

Q: You're not actually aiming to replace the radiologist?

A: There's not really much purpose in doing that. What I'm aiming to do, ideally, is to fill the huge expertise gap we have in the world. One statistic from the World Economic Forum – that it's going to take 300 years to train enough doctors to meet the needs of the developing world -- is the scary thing that drives all of this. If we can make radiologists initially five times more effective, then that means that in the developing world, there will be a demand for twice as many radiologists as are available, rather than ten times as many. They can start analysing data that has been gathered from remote Yunnan province, or from the slums of Mumbai, or anywhere else.

The idea would then be to replicate that success across all of the other types of medical data, so that everybody in the world can have access to effective medical diagnostics, and know what needs to be done to treat their current problems.

The nice thing about deep learning is that it actually works equally well for everything. So we can handle anything we can see: an MRI or an X-ray or a CT scan. Cancer is just one of the areas of operation. Other things would include bone fractures, or aneurysms, or kidney stones, and so on and so forth.

Q: You've said that you want your company to tackle more data types like genomic and pathology data.

A: Absolutely. That's the next phase.

There are constraints which are stopping us from being able to do that right now. Some of those constraints are kind of crazy. The pathology constraint is basically that there is no FDA-approved digital slide scanner, which basically means that you're not legally able to look at a digital photo of a slide and make a diagnosis. You have to have the physical specimen under the microscope.

You and I both know that a digital photo of something is the same as looking at the thing. But it's just a technicality. There's no point in us creating pathology solutions when people aren't allowed to use digital data for making those decisions. That's maybe three to five years away. It's the FDA process.

As for genomic data, the problem there is technological. The genome is just too big, and we don't currently have the ability to analyse that much data in this deep learning. It might be a few more years before we can fully use a full genome sequence. In the meantime, we can use some very simplified pieces of the genome, which is what people do nowadays. They basically say, "Do you or don't you have this gene?" That's all we can do at the moment. It's very, very simple.

The other thing about genomics is that the prices do not come down far enough that using genomic

The nice thing about deep learning is it actually works well for everything

analysis is universal for diagnosis and treatment planning. There are many, many times more medical images than genome sequences, so there's a shortage of data.

Q: Let's talk money now. I read that you recently raised $10 million for your company, and previously raised $5 million. Is that the total money you've received?

A: Yes, that's the total. The first $2 million was venture capital money. The next $3 million was angel money from people who wanted to get involved in healthcare. The last $10 was a full-service radiologist company that wants to be the first company in the world to fully integrate this technology in their practice.

Q: Are these capital injections enough to pay for the research you're doing?

A: Building software and algorithms is not a very expensive exercise. Our initial seed investment was $2 million. Our team size was ten for the first year. That was the right number of people to work on the R&D side of things.

The additional money is to support the commercialisation, not to support the research. It's convincing a hospital to use some new kind of technology, to go through the validation process and the integration process that's a slow and expensive thing.

Q: You're aiming for a majority share of the $10 billion US market for radiology image interpretation. IBM has just bought Merge Healthcare, and they spent $1 billion for the medical imagery platform. How much of a threat are they to your market leadership objective?

A: It's too early to say. Even with the Merge acquisition, it's not at all clear the particular kinds of medical images that Merge had. It's not at all clear that it's actually what IBM would need to build a successful deep learning algorithm for medical imaging.

I would say IBM has recently become interested in medical imagery. They're pretty late to the party. They're not the people that would probably keep me up at night, but I keep an eye on them.

Q: Your sphere of interest will remain human health, right? You're not going to go off and use big data for other subjects?

A: Exactly. This is the place that we can make an impact, what I really want to focus on for the next 25 years of my life.

Q: As you know, big names such as Elon Musk and Bill Gates have warned against the perils of artificial intelligence. Do you think AI does more good than harm? How do you feel about those warnings?

A: My biggest concern is not so much the robots taking over the world and killing us. It's the robots taking our jobs and causing an extremely unequal wealth distribution, where wealth is even more concentrated in the hands of the people who have the data and the algorithms.

If deep learning is going to do what we think it's going to do, there's going to be a lot of jobs which just aren't necessary anymore. We need to have a plan for that. People tend to assume that certain new jobs will come along, and people can work in the new jobs. We don't actually have any reason to believe that. There's nobody saying, "What are those new jobs going to be?" That's my biggest worry.

Energy: do as I do not as I say

Eoin Treacy
Investment Director

I was invited to Jeff Gundlach's "Strategy and Timing" talk in Los Angeles a week last Thursday. DoubleLine Capital, the company Gundlach founded in 2009, had just closed the month with $99.9 billion under management so the event was something of a victory lap. I'm not going to bore the readers of a technology newsletter with the gory details of bond valuations, but I thought a couple of comments he made about the energy sector are well worth repeating.

The first is that the number of bankruptcies in the energy sector is set to rise. There are a large number of wildcatter operations that got overleveraged, caught out by the drop in prices and

have been running on fumes over the last few months. Even the almost doubling of oil prices since the January low is unlikely to save them. That's going to have a disruptive effect, but the oil is still in the ground and better capitalised companies will take it over. What happens after the debt cycle runs its course is way more important.

One of the more notable anecdotes Gundlach recounted was focused on what sovereign wealth funds are doing with their money. He said that he had spent a lengthy afternoon with the chief asset allocator of the largest sovereign wealth fund in the Middle East the week before. The allocator said they were making all investment decisions based on the assumption oil prices could be $10.

The allocator gave three reasons for this view:

Hydraulic fracturing

Hydraulic fracturing is considered unconventional, because it is so much more expensive than regular drilling. The technicals of getting the process right are complicated, but the methodology is reasonably simple.

A well is drilled using horizontal drilling techniques so the profile of the holes looks like an inverted umbrella. A set of explosive charges are set at intervals along the holes, and when they are detonated they create fractures in the rock. Next, a sludge of water, sand and chemicals is sent down to the blast site to keep the fractures open so oil and gas can flow. It's a process that requires a skilled team, and drillers just pick up the rig and move to the next site once the well head has been capped off. As a leaseholder you want to get as many holes drilled as possible before your drilling rights run out, so a lot more holes were dug than brought on stream. When prices get back to where it is economic to bring new supply to market, all the leaseholder has to do is install a pump jack and the majority of the hard work has already been completed. When that happens, prices will start to come back down.

You might have heard of hydraulic fracking, but have you heard of refracking? The nature of shale or tight oil deposits is they are locked in impermeable rock. The oil and gas only flow because the rocks near the bore hole have been fractured. That means the initial flow rate is incredibly prolific, but supply from unconventional wells quickly peaks and declines sharply. Rather than going off to drill expensive new wells, there is a growing body of knowledge that suggests existing wells can be refracked which would create new fractures and return the well to its former supply trajectory. Companies have been cautious so far, not least because prices are low but as refracking could dramatically reduce the cost of unconventional supply – which would be an additional negative influence on prices.

Preparing for $10 oil

Electric vehicles

China set an aim of putting five million electric cars on the road by 2020. That represents a 30% increase per annum between now and then if the objective is to be achieved. China is the world's largest car market, but if the demand for the Tesla Model 3 is any guide then the US is likely to have plenty of electric cars on the road too. European car manufacturers are already well on the way to having a wide variety of EVs on the road in the hope demand will mirror growth in Asia and North America, not least following the diesel emissions cheating scandal.

I'm sure I'm not alone in wanting to own an electric car as my primary mode of transport, but the reality is I don't want to have to worry about recharging half way to Las Vegas because I've been running the AC too heavily in the hot desert weather.

A Tesla Model S has a maximum range of about 250 miles; subject to road conditions, weather and air conditioner use. These are still very new cars, so with innovations in battery packs the potential for that range to double in the next few years is non-trivial. Taking the range

of an electric vehicle up to that level would eliminate range anxiety, and demand for cars that don't create emissions and don't have negligible fuel costs would increase at a startling rate. Jeff Gundlach quoted the Middle Eastern asset allocator as saying that it could represent as much as $4 million barrels a day in lost demand from the US.

Renewables are the future

The introduction of the supply unlocked by hydraulic fracturing represented about 1.5 million barrels a day, and just look at the impact that has already had on prices. If electric car volumes come in as expected the oil industry is going to be in serious trouble.

Renewables intermittency

Coal is still the primary source of fuel for electricity generation globally. Natural gas is the fastest growing global fossil fuel market. However, the improving efficiencies of solar and the potential for batteries to close the loop of intermittency holds out the real possibility that fossil fuels will lose the top spot in electricity generation over the next couple of decades – and not just in the Sun Belt. For more than a decade already, German policy makers have been floating the idea of running high tension cables to Algeria and importing solar energy rather than coal, oil or natural gas.

The dream of replacing dirty coal with clean non fossil fuel alternatives is not based on a single potential solution. Energy security depends on an energy mix, so more than one solution has a place in the overall market.

For example, Toyota, having been an early advocate of hybrid electric vehicles, is no longer producing the Prius for the North American market and is only now selling the model in China at the government's insistence. Instead, it is betting on the hydrogen fuel cycle with its new Mirai model. Right now the most expedient way of producing hydrogen is from natural gas, but Toyota has a view on the future where fusion will enable vast quantities of hydrogen to be produced at a very affordable cost.

That's just two – but with transparent solar cells, wind turbines, geothermal, biomass and batteries as potent potential supply alternatives, and greater efficiencies driven by better batteries, autonomous vehicles and the Internet of Things on the demand side, it means there are a large number of potential replacements for oil. Is it any wonder Middle Eastern oil despots are planning for a future where the product they depend on for their entire economies is no longer in such high demand?

Energy

Company	Ticker	Rec Date	Price Then	Price Now	Gain/ loss %
Orocobre Ltd	ORE. AX	07/04/16	A$ 2.92	A$ 4.63	58.56

Technology

Company	Ticker	Rec Date	Price Then	Price Now	Gain/ loss %
Cyberdyne	7779	30/05/16	JPY 2,353	JPY 2,549	8.32
IBM	IBM LN	07/06/16	152.98	152.98	0

Medical

Company	Ticker	Rec Date	Price Then	Price Now	Gain/ loss %
Acorda Therapeutics	ACOR	07/04/16	$ 27.98	$ 28.92	3.35
Intuitive Surgical	ISRG	07/04/16	$ 626.12	$ 633.35	1.15

Moonshot

Company	Ticker	Rec Date	Price Then	Price Now	Gain/ loss %
SolarWindow	WNDW	07/04/16	$3.96	$4.05	2.27

For the full portfolio including live prices, please visit the *Frontier Tech Investor* subscriber area. **You can view that by following this link.**