AI for Humans

A simple guide to help navigate the use of AI in the insight industry

discover.ai

At a recent insight industry conference, we were giving a talk on how AI can make us be more human. Afterwards, someone with years of insight industry asked us...

“Can you explain to me what AI is all about? I know it’s going to be important but it’s not easy to know where to start.” Indeed, someone else asked “who’s AL? Is that short for Allen?”

These two encounters, and others like it, make you realise that a reality where AI is a day to day part of your work and life is a bubble. It’s easy to forget just how marginal an area it is for so many other people. It’s not that people aren’t interested, but there’s a gap between what they need and want to know about AI and what’s out there and easily accessible. Whilst jargon-filled technical information and sensationalist media stories abound, it’s a lot harder to find the kind of simple and practical help that could take people through the basics of what you need to know to navigate the choice of AI innovations, and how they might help your business. And if you want something focussed on the Insights industry?

That’s even harder to find.

That’s the gap this paper is designed to fill.

Reading this should give you a basic definition of what AI and machine learning, how they are being applied, their benefits and potential drawbacks, and what to look out for to find the most relevant technologies for you and your business. You may think AI is something to be feared or mistrusted. You may see it as an exciting yet confusing new frontier. Or perhaps you’ve not though much about it at all as you get on with your day to day life. Whatever your current views on AI, we hope that reading this will leave you with a little more understanding, as well as a feeling of optimism for how AI can help us in the future.

So what is AI?

Artificial Intelligence is not quite what it sounds like it might be, or what popular culture often tells us it is. We might think that it is machines trying to mimic the intelligence capabilities of humans. The often-quoted Turing test (where a human interacts with an unseen entity and must decide whether it is human or machine) does much to perpetuate this impression. But it's not. Wikipedia defines AI as "any device that perceives its environment and takes actions that maximise its chance of successfully achieving its goal." AI is about machines in the real world, experiencing it, making decisions and acting in a way that influences that world. But machines can be doing this in a very different way to how we as humans do it. Machine intelligence can seem either very familiar and 'human', or completely alien to us, depending on how it functions.

AI generally isn't human intelligence in machines, it's machine intelligence in the human world.

Knowing this definition helps you realise that AI is a really broad catch-all term for any technology that helps machines to perceive the world, make decisions and act. It can mean almost anything. So, in practical terms, knowing that a person or company is working in AI tells us little or nothing about what they are actually doing. You need to dig deeper. Anybody can add something AI related to their software platform in minutes (more on this later) so it's hard to draw any conclusions from its use.

What most people mean when they say AI in the Insight industry right now is Machine Learning, making this a more important term to watch out for. So why do we talk about AI most of the time, when we mean machine learning? Probably because machine learning sounds dull and technical and doesn't shift the emotional needle in the way talking about AI does.

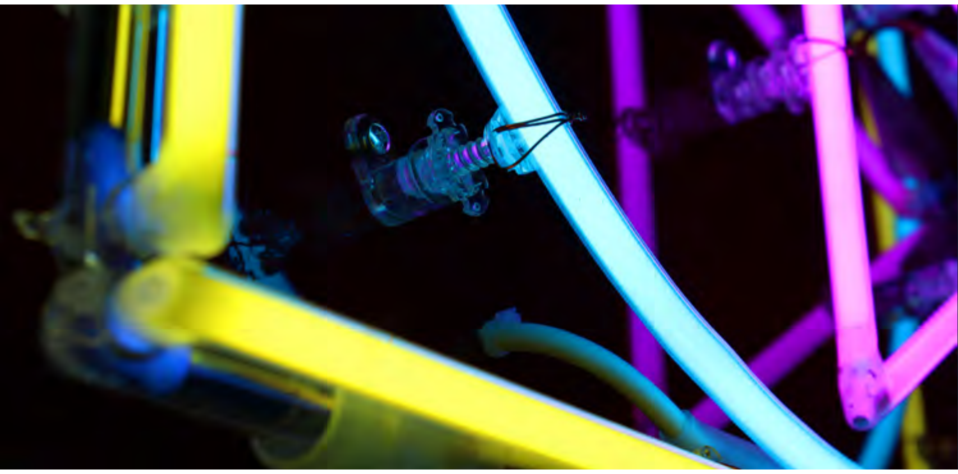

For this paper we will assume that you are really interested in machine learning, which accounts for most of the AI technologies in the insight space. If you want to explore other AI technologies, then you could start here

OK then – what’s Machine Learning?

In simplest terms, Machine Learning is about teaching machines how to do something rather than programming them how to do it.

Before Machine Learning, if you wanted a computer to do something you had to programme it to do it. That means trying to anticipate every scenario and writing a set of instructions (code) for what to do in that situation. Fine if you are talking about something simple, but the more complicated the scenarios become (and the more ‘intelligent’ you want to the machine to seem) the more complex are the instructions to manage. There will always come a point where the real world is more complicated than your instructions. For example, when trying to programme a machine to win a game of chess the rules are actually relatively simple to define and programme.

But trying to programme a machine to make decisions across every potential game playing scenario will become incredibly complex, very quickly.

Machine Learning in effect reverses the principle. Instead of writing instructions to get the outcome you want, like win a game of chess, you show the machine thousands of examples of a chess game from start to finish and let it create its own rules of how to play to win, based on these examples. It’s learning by experience. You ‘train’ the machine with live examples and the machine ‘learns’.

It took programmers decades to write the code that could beat a chess grand master, but modern machine learning algorithms (the clever technical bit that learns) can process thousands of games and learn to do the same thing in hours.

The unifying principle of prediction

At this point machine learning can start to get extremely complicated if you let it. There are all kinds of different algorithms our there which rely on a wide variety of analytical principles, from Bayesian statistics to genetic models of evolution and survival. I would suggest it's more practical to ignore all this and leave it to the technical experts.

Ultimately all machine learning systems are trying to predict something, it's that simple. Sometimes prediction is the end goal of the system, other times prediction is a means to an end, a way of achieving another goal. The better the machine learning system, the more accurate the predictions become.

+ Amazon is trying to predict products I will like to get me to buy more stuff
+ Driverless cars are trying to predict what car controls to apply in every situation to drive someone safely and comfortably to their destination
+ Machine marking systems are trying to predict what mark to grade a paper to make marking faster, more cost effective and more accurate
+ Medical diagnosis systems are trying to predict what ailment someone is suffering from based on all available data to make a treatment recommendation, improve health and save lives
+ Automated weapons are trying to predict where enemy targets are and kill / destroy them (yes, these weapons already exist)

All these examples, and countless more, are ways in which machines are trying to understand the environment, make predictions about outcomes based on the available data and take action to achieve a goal. The only question we need to ask ourselves when faced with a Machine Learning system is 'what is the machine trying to predict and how does that help me?' This principle can apply to just about any machine learning system or insight platform or service that is using AI.

From simple to intractable questions

It's important when understanding the potential role of AI and machine learning in your business to understand the true nature of the questions you are trying to answer. You need to know the complexity of the world in which you are trying to make predictions and so how well machines are able to take on this task and the role of AI vs. humans in finding the answers.

Simple questions

Some questions are based on simple logic and only need a programmatic approach to solve them. For example "my Queen is on square A1 of a chess board, where can it move to next?" There are a clear set of rules and limited variations for this question making it simple to solve using computational power getting to answers far quicker than humans can do.

For simple questions, automation can be achieved without AI

Complex questions

Other questions are complex and might seem like a challenge for even the most advanced machines. For example "how can I win a game of chess against a Grand Master?" But despite being complex these questions are 'machine solvable', because they are based on fixed rules, have a clear outcome and are repeatable. Give a machine learning algorithm enough examples of a complex question that has been solved and it can learn to solve that question.

For complex questions, automation can be achieved using AI

Intractable questions

But what if our questions isn't how to win, but instead "why do people like chess and how do we get more people to play?" Now our question has no fixed rules, is subject to the vagaries of human nature, is not repeatable and there is no clear outcome of historical precedent. Where do we find the thousands of examples to train a machine to solve a problem like this?

For intractable questions, AI alone will not find the answers

The principle of intractable questions is essential for the Insight world because these are the questions we typically face every day. Questions about human needs, culture, communication, emotional engagement, trends, growth strategies and more. People may try and tell you that machines finding patterns in big data sets will answer these questions for you, but the reality is that they won't.

In the insight industry, AI and machine learning will help us automate many of the questions that rely on answering simple and complex questions. But where we face intractable questions, we need to be exploring how these technologies can accelerate the human expertise needed to solve them. Machine working hand in hand with human.

Returning to our chess example, it's interesting to note that while machine can beat the grand master, the grand master playing in combination with machine can still more often than not beat the machine. The future of AI in the Insight world is about how we apply this principle of partnership. Human and machine working together.

There is a belief that given a big enough data set and a smart enough algorithm, that a machine finding patterns in this data can give us the answers to all our questions. Sadly this isn't true. The real world is far more complex than any game.

Knowing the limitations of AI

Machine learning can achieve some amazing things, but it's not without its limitations and drawbacks. This is particularly true in the insight industry, where people, culture and emotion play such as important role. It's important to understand what these limitations are.

Beware the Black box

The first drawback is the 'black box' nature of how machine learning works. Because the rules that the machine is following to make decisions have been learned, not programmed, it is often impossible to know what these rules are and why it has chosen the solution it has. There are some circumstances where this is not an issue, but for the Insight industry, where understanding is key and evidence matters, this can be a problem. Brands and businesses don't just need to know the 'what' of any question, they need to know the 'why'. Using an AI solution can often feel like a step back to the dark days market research reports with data points dumped on an audience with no context or human understanding.

Pattern isn't truth

Recent scientific research has begun to cast doubt on the results of machine learning algorithms and the patterns they find in big data. A recent BBC article stated, 'according to Dr Allen, the answers that they (machine learning algorithms) come up with are likely to be inaccurate or wrong because the software is identifying patterns that exist only in that data set and not the real world'. The same article also refers to a 'reproducibility crisis' and 'an alarming number of research results that are not repeated when another group of scientists tries the same experiment'. So not only are machine learning answers 'black box' in nature, patterns in big data are held up to be a truth about the real world when there is a significant chance that they are not.

Be mindful of bias

This is an important one as machine learning analysis of big data sources is often seen as the search for an objective truth. But these algorithms are trained, and therefore prone to the biases of the people who train then and of the historical data sets used for the training. This bias can be of a social nature and therefore a cultural concern. For example, a 2015 study found that women were less likely to be shown high-income jobs by Google's ad sense. This is historical bias compounded into future decision making because of biased data sets. But the concerns aren't just social, they are practical; the same principles will apply to algorithms trained to look for insights. Always remember the results of any analysis you see are only as good as the data set used to train and will contain all the biases of the people that created them. Always ask 'who trained this algorithm? With what data? And how does this influence the results I will see?'

Past doesn't predict the future

Using historical data to predict the future can be seen as a bias in and of itself. There is an inbuilt and potentially misleading assumption in any research that what you can learn from the past will help you in the future. When you are trying to beat a grand master this isn't such a problem, as the rules of the game don't change (although strategies employed by your human adversaries might evolve). But in the insight world this can be more of an issue as the world in which new brands, growth strategies and communication are launched might be very different from the ones in which your historical data has been collected. Context can be everything and the solutions of tomorrow may need to look very different from the success stories of the past. Or, essential learning for success in one market or industry might prove irrelevant for another.

These limitations have always been true of traditional research methodologies, but the danger is that the innovative nature of AI and big data sets can blind people to the fact that they still apply and lead them to put more faith into the predictive powers of AI solutions than is sensible. When people pitch the benefits of their AI platforms what they promise can often sound magical, but the right questions can often help uncover the limitations of their technology

Being clear what AI people are selling

It’s incredibly easy to create a piece of software that has an AI element to it, and it won’t be long before all software platforms will include AI in one way or the other. So it’s helpful to understand what level of AI is being used by any given platform and its role in providing the platform’s solution.

Pre-trained algorithms

The most accessible form of AI is pre-trained algorithms. This is where someone has used an algorithm or used a data set to train it to do a specific task and then made this available for others to put into their software. Some very common examples of this are sentiment analysis, automatic translation and natural language processing (NLP), but there are hundreds of others. It’s very simple to pipe these services into your software via an Application Programme Interface (API). Anyone can access these services from providers like Amazon, Google or Microsoft and claim their software AI enabled in minutes. But if this is all they are doing then they are providing quite simple commoditised AI services built on top of a programmatic solution. This isn’t a problem, but it can help make you aware of the scope and limitations of the technology.

Self-trained algorithms

The next level of AI solution is when a unique data set has been used to train an algorithm to perform a new and focussed task. This takes a step up in investment and expertise and will give that platform an AI service that is in some way unique and tailored to their particular use case. But in these cases, it’s important to understand how the model has been trained – by whom, with what data set and with what relevance to the task trying to be solved. There is a difference between models created by high quality data sets and expert trainers versus people renting off shore ‘training farms’ with low-cost, low-skilled training. So, it’s important to understand not only what the AI is trained to do but also how it is trained, the quality of that training and the potential biases influencing the solutions it creates.

Unique algorithms

At the far end of the spectrum is where people are not just training algorithms to solve new challenges, they are created entirely new AI technologies. This will be at the cutting edge of AI technology and typically companies doing this will be investing millions in building their solutions. Most companies in the insight space won’t be doing this, and indeed this level of technological investment is unnecessary for most challenges we are trying to solve in the insights industry right now. Pre-trained or self-trained algorithms are well placed to solve our challenges at a much lower level of investment in time and cost. But the breakthrough algorithms of today are the ones we will likely be using to train our solutions of tomorrow. These technologies typically filter down rapidly.

Sector vs. Technical expertise

Another area to look out for when evaluating AI-based solutions is the relationship between sector and technical expertise and the balance in the team. It's essential to get this balance right.

AI experts often talk about the quest for the universal algorithm. This would be the one AI algorithm intelligent enough to predict the right outcome in any situation the world or life could throw at it. This sounds great in principle, but it is accepted that any best we are probably decades off creating such an algorithm, and it may not be possible at all.

Instead current algorithms are trained to solve relatively narrow sector specific tasks. So the algorithm that will drive your car is not the same one that will diagnose your illness or choose you a film to watch.

Each of these sector specific challenges requires not only an algorithm dedicated to that task, but also a team that combines the best of technical and sector expertise and then an organisation capable of bringing it to market and engaging with the relevant industry customers in a meaningful way. This is not an easy task. Many AI solutions are lone technical innovators struggling to make their breakthroughs relevant in their chosen fields.

When looking for AI solutions for your business it's essential to look at the breadth of skills and experience that any company brings, from technical to sector specific, and how well these skills are integrated to bring out the greatest benefits of AI.

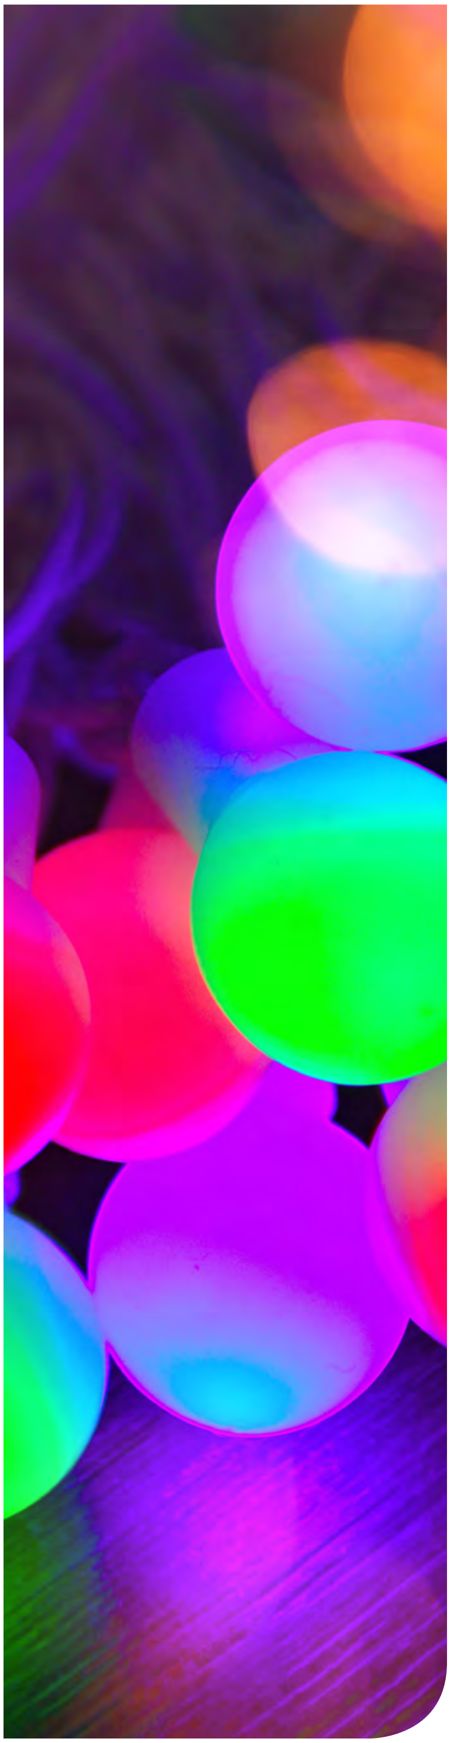

AI and Machine Learning are in their infancy and as an industry we are finding our way, with new innovations appearing at a rapid rate.
It can be an overwhelming space to try and understand and explore. Ultimately though, the proof of any AI innovation is in the experience rather than the pitch and the ability of any AI platform to deliver on its promises and add value to your brands and business.

If you want to learn more about using AI to discover opportunity, then get in touch.

jonathan@discover.ai
discover.ai

---

discover.ai