


10 things they don't tell you about **Artificial Intelligence**



A Humanotics White Paper | Dr David Naylor



10 things they don't tell you about Artificial Intelligence

(and you probably didn't think to ask)

Do you want to know the inside track on the state of AI in the customer management industry today? We will share a few home truths that will help you navigate the hype and hopefully make some decisions that will put you on the right track.



1. You've been using elements of it for years.

That's right. You may remember Fuzzy Logic in washing machines in the late 1990s. That was a form of AI. Chase Bank was using it even before that for reading cheques. Anyone who has used speech analytics tools has been using AI. There is a

very broad definition of AI being applied now which vendors are taking advantage of. We've been using 'traditional' algorithms to make predictions for years and many application features are simply being rebadged as 'AI-driven'. For example, logic engines and decision trees have supported knowledgebase systems for many years and now they are called AI. The latest AI developments really centre around the use of artificial neural networks. Again, these are not new but found little commercial interest until Google started developing the self-driving car. These models can be trained with the large real-world data sets that have now become available rather than having to be told how to work based on a hard-coded formula that defines the relationship between an input and an output. Hence the term Machine Learning. Today you may flag a customer as a high churn risk simply if they are out of contract and offer them a promotion. With ML, you can look for more complex patterns in many data points simultaneously to identify hidden risks and relationships. For example, people who have complained twice in the last 6 months and have called about their bill 3 times this year are 80% more likely to churn on average. Therefore, the real question to ask of any vendor is "how do you use machine learning to enhance the performance of your solution?"



Think about what you want to achieve and data you need before selecting an approach.

2. It's still not that intelligent, really.

Self-driving cars are clearly at the cutting edge of the AI we see today. While we're fully expecting these to be on our roads in the next few years, they are still lacking some key capabilities of our brains and the learning process has been a long one. Apparently, Google cars have driven for over 300,000 hours on public roads so far – but still not passed the driving test. My daughter has driven for around 30 hours and is taking her test in about a month. Ok, I'm stretching the comparison but it's also important to realise that AI is good at very specific tasks and there isn't one AI 'brain' in any of our systems. Specific learning algorithms are applied to specialised tasks such as interpreting the image from one camera or monitoring environmental data. A central management system, which may have further learning algorithms as well as

traditional logic, then reviews the data it is getting before making decisions. If you want to know more without going too deep watch the talks by Yann LeCun, Director of AI Research at Facebook (LeCun, 2017).

In our contact centre applications, learning algorithms play similar specialist roles – doing one thing, hopefully really well – before delivering an output that can be actioned by a human or the next step in the system. However, while the average contact centre application doesn't need 300,000 hours of data, train the algorithm with incomplete, skewed or insufficient data and you will end up with the common problem of garbage in and garbage out.



3. Many chatbots contain very little AI

Not every chatbot application is created equally but there is a common theme to most solutions available today. Understanding what the customer wants to do is called 'establishing the Intent'. Typically, this will be done using a cognitive application that feeds the text into a natural language processing system which responds with a code that represents the Intent. At that point a specific script is executed which processes the request. It may have steps to gather more information from the customer, access a database or knowledgebase and then provide a response. That script is built by the business team who know the process and needs to be very robust to handle as many as possible of the scenarios customers create. It will never be complete before launching and needs

ongoing work to tune the process. Humans complicate the tuning as we are good at going off script and using historical context to infer meaning. The new Lidl wine advisor on Facebook Messenger is an example of a chatbot that can only 'remember' what you asked for in your last request. Try to ask it to 'show me red wines from Australia' and then 'show me some from New Zealand' and it fails.

Some chatbots that handle more conversational requests such as FAQs are being built without scripts and learn from analysing lots of previous live chat interactions – if this data is available. Some vendors target agent assistance in a chatbot as they don't believe the technology is capable of delivering effective customer interactions, yet. It's therefore important to think about what you want to achieve and what data you need before selecting an approach.

4. The basic AI tools you need may be available online at little cost

In the last 3 years IBM Watson, Microsoft, Amazon and Google have all released extensive toolsets for cognition (i.e. speech and text analytics), chatbots and machine learning within their cloud platforms. There are others which are less well known but equally capable such as Floydhub. The best part is that they are free or amazingly cheap to get started with and typically have extensive online guides and support communities. Building your first machine learning algorithm or simple bot is something you can do almost overnight. Of course, you will need the right team (see more below) to build something in-house that is robust enough to deploy and it is less likely to be something that can be managed easily by the business rather than IT. The management and reporting framework is a critical component that needs to be created to ensure tuning is effective and the applications can evolve with customer needs. An in-house build may not fit with your strategy but if it does then it's worth exploring the benefits.



Could machine learning predict Lifetime Value or Churn better than current approaches?

5. There isn't one answer and the answer may not always be AI

AI is powerful and there are many questions it could be used to address – most of which are not clear to us yet. But there are also many AI solutions. As I've mentioned, lots of data science techniques are now being classified as AI and there are many variations of machine learning.

Machine learning is a computationally expensive approach which involves an element of discovery to determine if a solution exists that improves on traditional models. For example, could machine learning predict Lifetime Value or Churn better than current approaches? If it can deliver a 10% improvement, then it could be worth it, but what will it cost in time and effort to achieve?

So, what are the elements that affect the discovery process? The size and quality of the dataset are key factors. Then, the initial training parameters and size of training steps for the algorithm are important. But critically, different algorithms work better on different types of data. For example, Convolutional Neural Network models work best with image processing and Recurrent Neural Networks are good for understanding context in language. The amount of data needed for training and evaluation will vary by type of model used and how you train it. I won't go on, but you get the point – machine learning algorithms can do *some* tasks very well and much better than we could ever expect from a traditional computing approach. Knowing if you are applying machine learning to the *right* task is a business problem as much as a data science problem.



Data creation, extraction and consolidation is time consuming.

6. AI won't magically make sense of all your data

The data manipulation stage of any AI project is typically the most labour intensive. Why?

As anyone who has been through a system migration when implementing a CRM platform will testify, the quality of the data you extract can be very poor. Spelling mistakes, missing fields, data in different formats and so on. Basic cleansing always takes much longer than ever expected.

Your data may also be skewed. If you are looking at time series data as you would in a forecasting model, poor service levels will drive more repeat contacts rather than genuine new contacts from different customers. You may also have less understanding about the needs of customers who interact with you less frequently perhaps. The data could also contain human biases – there are many articles on that (Tobias Baer, 2017) (Giannandrea, 2017) (Devlin, 2017). Machines that learn can handle missing

data reasonably well as long as the data that exists is representative. You need to spend time validating that this is the case.

Also, it is rare that you will be accessing only one system. The data needed may come from different systems, may not be stored digitally or may be in people's heads – especially when dealing with customer process exceptions. It may not even exist and so you may need to create it, if you don't track a particular metric, for example. Data creation, extraction and consolidation is time consuming.

Chatbot vendors may just ask for your live chat transcriptions but that is often just the start of weeks of data manipulation behind the scenes to make it fit for use in training and testing. When starting with machine learning, it is therefore wise to consider data quality very early. You may not be able to solve the biggest issues until you have gathered more reliable data.



7. You need to think about RPA if you are going to optimise the AI you deploy

You will often hear the terms Robotic Process Automation and Artificial Intelligence talked about in the same context – is that just the marketeers and conference organisers getting more bang for their buck? Well no, it makes sense to think about RPA with AI. RPA provides the glue that allows the business to automate repetitive aspects of customer processes. Often these are the tasks queued to back-office teams, but RPA can also be used to improve performance at the desktop by reducing rekeying of data between systems.

For example, a current client is looking to implement a chatbot to take a request from a customer so they can issue a replacement loyalty card. While it is possible to use an existing web API to link

into the CRM system to capture or validate customer details, the request to send the new card is issued through a standalone system. In a live agent chat scenario this is inconvenient but not an issue as the agent can initiate the request manually. What you don't want to be doing with a chatbot is providing a self-service process that does not complete the transaction but instead queues the work for a back-office team – bad process ! This is where RPA can help.

Clearly, the benefits of RPA need to be assessed relative to the option of building a dedicated interface. However, as AI evolves it is going to provide more of the 'decision making' glue between automated processes tasks, with only exceptions being handled by humans. Typically, vendors focus on RPA or AI solutions so you will need to raise your own thinking up to the business process or, better still, customer journey level, in order to evaluate where RPA and AI can jointly streamline your business.

8. You will need to add new skills and processes to your business to manage AI

If you choose to go down the DIY route and use the tools available online to build your own AI applications, then you should clearly already understand the technical capabilities needed. You don't just need great Data Scientists. Real world AI problems need a team. This team may often consist of business process experts, product owners, hardware engineers, programmers, cognitive engineers and a commercial lead focused on the outcome. Unfortunately, beyond designing chatbot scripts, most business analysts who have come up through the contact centre with increasing proficiency in Excel just don't have the mathematical skills needed. I've therefore seen many businesses building up their Business Intelligence teams with very expensive Data Scientists who are detached from the operational world. That's why a team is essential. This combines the business knowledge of the Analysts and Commercial lead with the specialist skills of the Data Architects and Scientists.

If you don't already have a digital customer transformation strategy in place then deploying a chatbot is probably not your next step

When it comes to developing solutions, adopting Scrum principles (often called

'Agile') is almost essential, whether it's for building a chatbot or developing a predictive up-selling model. This is because the development of machine learning models requires a highly cyclical approach and regular feedback from experts in the business to ensure the solution is evolving in the right direction. Scrum requires a different mindset from the business and a highly empowered team.

Finally, an operation that uses machine learning models needs to be continuously evolving the model as the environment changes. If you introduce new products or contact channels, how does that affect what you can up-sell to customers or affect the accuracy of your call forecasting, for example? The models need retraining and constant monitoring to evaluate effectiveness. Someone must take on the responsibility for the decisions that the algorithm makes, and own the metrics that track performance – a critical new role.

This sounds like a lot of business change but this is part of every technology implementation. Unfortunately, it's probably the area where most projects apply the least attention.

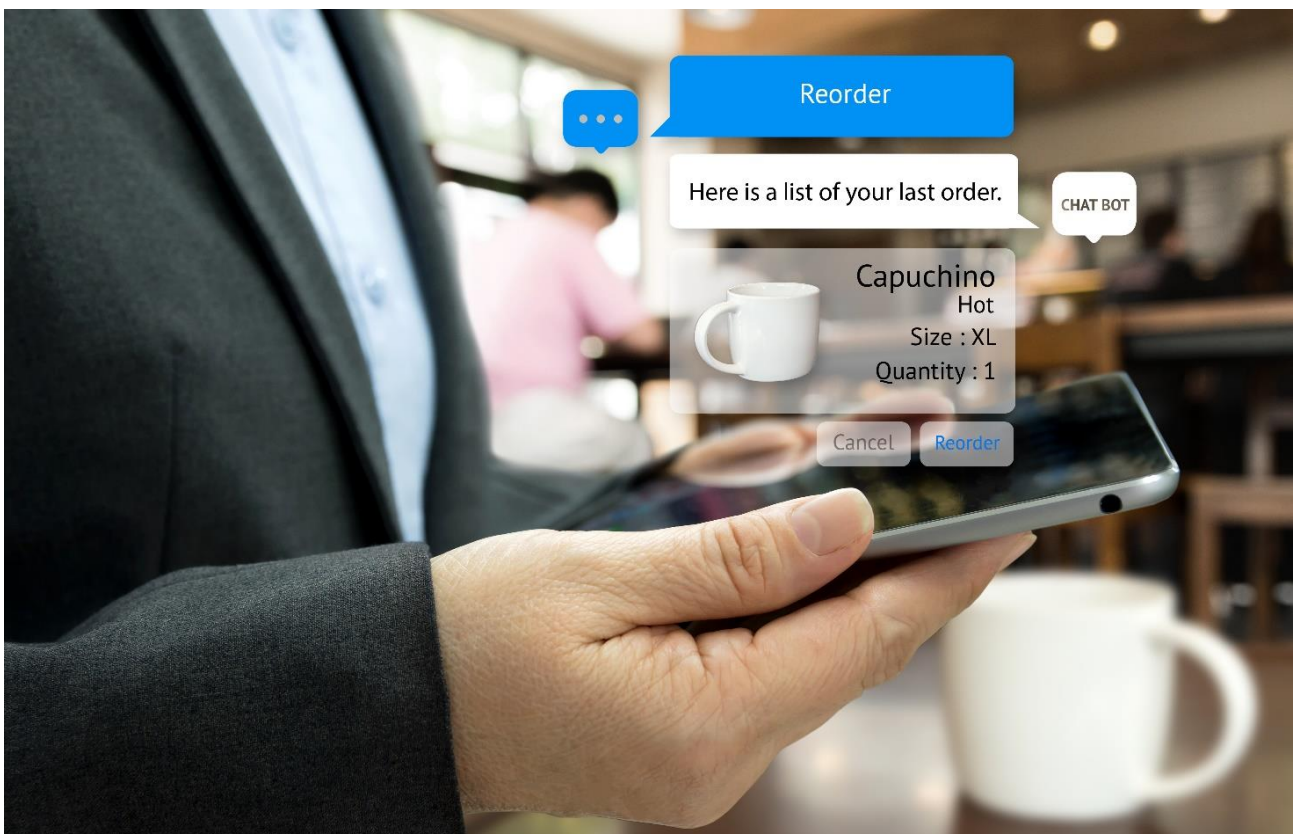
9. It might not revolutionise your business today

You've read this far, and you are probably thinking, "why would I even bother right now?" McKinsey estimate that investing in the right AI technology will give benefits that outweigh the costs by up to 10 times (McKinsey Global Institute, 2017). But if you don't already have a digital customer transformation strategy in place then deploying a chatbot is probably not your next step. Work out your strategy first and think about where AI could play a part – you may look for help with that.

You are not missing the boat. No vendor has a magic piece of technology that is going to reduce 50% of your costs in one go. At present you can buy packaged chatbot solutions or you can invest in Business Intelligence tools (or managed services) powered by machine learning algorithms. These tools allow you to

explore your data in new ways and identify patterns that were previously hard to see. But the question has to be asked before you start on this path – "what are we trying to achieve?"

New applications are being developed that will embed the technology at key points of the interaction with customers. For example, making smarter decisions on routing of contact, next best actions for sales or to improve agent quality and performance monitoring. Salesforce.com has embedded tools called Einstein into its platform and is claiming its clients see reduced attrition and increased sales conversion rates as a result. Amazon.com also has a lead on AI in service given the expertise it has gained through its own operations. I predict that the Amazon Connect CRM platform is 'one to watch' in the next 12 months as much of the extensive AI capability now offered through its web services platform could be surfaced in Connect (Amazon Web Services, 2017).



10. AI is changing fast. Instability is opportunity not just a threat.

According to government statistics, a new AI company has been launched in the UK almost every week over the last 3 years and we are small compared to the advances being made in China and the US. The amount of private and public funding available continues to rise and so the rate at which research moves into commercial applications is continuing to speed up. A new technique called Generative Adversarial Networks (GANs) was proposed by academics only in 2014. It has been highlighted as “the most interesting idea in the last 10 years in machine learning” by Yann LeCun, Facebook’s AI Research Director. GANs involve two neural networks competing against each other to help the system self-learn with the intention of being able to mimic rather than classify patterns that most networks do today. General applications are seen in areas such as art and music where the network could compose a piece of music in the style of Mozart or paint a Van Gogh. These applications are moving us towards the AI of the future where the system can create something new. They will be inventing a game like chess rather than just playing it. How far they will be able to go in mimicking the characteristics of your best sales advisor or

top performing customer service representative, we will have to see.

While that is all some way off, the current challenge is ensuring you are taking the right steps on the AI journey when you select services, solutions and vendors. Has a product simply been rebadged? Is there any real AI in there? Is the AI ‘vapourware’ and are you being used as the guinea pig? Will they be around in 12 months’ time?

The answers to these questions are important of course, so enter a relationship with a supplier with your eyes wide open. There can be some great risk-reward opportunities available not to mention competitive advantage in a brave new world.

Dr David Naylor, November 2018.



About Humanotics

Our mission is to demystify complex Artificial Intelligence technology. We help organisations to understand and deploy its capabilities in an AI-powered Operating Model, that creates real financial benefits and better customer experiences.

If this technology is to deliver on what it promises for intelligent customer engagement, solutions must be business-led and be embedded in a holistic customer management strategy. Humanotics will help you achieve this.

We adopt an end-to-end approach on your AI journey. Our experienced team has a unique balance of skills to support each step from customer experience design and operational change to data science, speech analytics and integration.

We will deliver for you more of the benefits that AI promises faster than you can do alone.

For more insight or assistance please call David Naylor on +447824397097 or email him at david.naylor@humanotics.co.uk.

References

- Amazon Web Services. (2017, September 21). *Maximising the Customer Experience with Amazon Connect and AI Services*. Retrieved from Slideshare: <https://www.slideshare.net/AmazonWebServices/maximising-the-customer-experience-with-amazon-connect-and-ai-services>
- Devlin, H. (2017, April 13). *AI programs exhibit racial and gender biases, research reveals*. Retrieved from The Guardian: <https://www.theguardian.com/technology/2017/apr/13/ai-programs-exhibit-racist-and-sexist-biases-research-reveals>
- Giannandrea, J. (2017, October 3). *Forget Killer Robots—Bias Is the Real AI Danger*. Retrieved from Technology Review: <https://www.technologyreview.com/s/608986/forget-killer-robotsbias-is-the-real-ai-danger/>
- LeCun, Y. (2017, November 19). *YouTube*. Retrieved from The Power and Limits of Deep Learning: <https://www.youtube.com/watch?v=0tEhw5t6rhc&t=1354s>

McKinsey Global Institute. (2017, January). *Harnessing automation for a future that works*. Retrieved from McKinsey&Company: <https://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works>

Tobias Baer, V. K. (2017, November). *Controlling machine-learning algorithms and their biases*. Retrieved from McKinsey&Company: <https://www.mckinsey.com/business-functions/risk/our-insights/controlling-machine-learning-algorithms-and-their-biases>