

Botman vs. Superagent: Man vs. machine in the future of customer experience

Dr Nicola J. Millard

Head of Customer Insight & Futures, Global Services Innovation Team, BT



The march of the machines

Machines are getting smarter, so should we fear for our jobs? Find out how far machine-learning has come and how we can use it to make both our jobs, and customer experience, better.

The fear of “robots taking all of our jobs” is nothing new. In 1811 the people who gave us the term for resisting technological progress, the Luddites, destroyed the new automated looms that they felt would result in the loss of their factory jobs. In 1928, the New York Times printed the headline ‘March of the Machine Makes Idle Hands’, and in 1930 economist John Maynard Keynes coined the term “technological unemployment”. Variations of the same headline have appeared as technology has evolved. But are things different this time? Or are our mechanical overlords about to make us all redundant?

Deep machine learning capabilities, combined with more powerful processors and enormous amounts of data, make things different to the last wave of machine learning development in the 1980s, and are already drastically changing the world of work. Economists¹ have suggested that the march of technology has already started to “hollow out” middle level jobs and cause wage stagnation. As digital business races ahead, technological progress may leave many people behind, and may mean that a lot of us will need to adapt, reskill and retrain.

Some jobs have already succumbed to the march of the machines. Forty years ago I would have either hand written or dictated this report for it to be typed up by the ‘typing pool’, now a victim of word processors, sadly I now have to do it myself. Are taxi drivers, lawyers, doctors, pilots, and contact centre advisors going to follow the typing pool into the history books?

Before we get too depressed, let’s get a few things straight about artificial intelligence.

1. **Artificial intelligence doesn’t exist** – much of what we call “AI” is more accurately termed “machine learning”.
2. **Machine learning is very unlikely to “eat your job”** yet. Unless you’re unlucky enough to have a job with a lot of good, machine-readable data, that can be purely defined by rules. Machine learning is only as good as the data that fuels it.

So, let’s take a look at how automation is going to affect jobs and skills, and how machine learning can be harnessed to improve the customer experiences we deliver.

Is it really a titanic struggle between **Botman** and **SuperAgent**, or is it more of a powerful alliance?



Machines vs. humans: what does the data tell us?

One of the most quoted studies about the future of jobs is the 2013 Oxford Martin 'Future of Employment' paper². It calculated that 47 per cent of jobs were at risk of being automated in the next ten years. For telemarketing that rose to 99 per cent, and for customer service 55 per cent.

A later study by the McKinsey Global Institute³ took a more activity-led view, on the grounds that within every job category there are activities that may have a different technical potential for automation than others. The study estimated that five per cent of jobs will be fully automated in the next ten years, but that 60 per cent of jobs could have 30 per cent of their activities automated. Another activity-centred study from the OECD⁴ has also estimated that nine per cent of jobs will be automated in the next ten years.

We can see a future when a minority of jobs – somewhere between five and nine per cent – have been fully automated, with the rest having elements with the potential for automation. This means that jobs are more likely to change than be automated out of existence entirely. On the positive side, this means that some of the tedious tasks that we often don't enjoy, such as managing our inboxes, logging timesheets and entering data into multiple systems, have high potential for automation, freeing up time for you to do more challenging and interesting things instead. Rather than destroying jobs, this could increase productivity.



"It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility"

Hans Moravec, Robotist

What activities are more likely to be automated?

The first question we need to ask is 'what tasks are humans good at, that machines find difficult?

Robotist Hans Moravec⁵ suggested that there was a paradox at work here. He noted that "it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility". This paradox can be illustrated by a recent feat of advanced machine learning by Google DeepMind's AlphaGo program, which defeated professional player Lee Sedol at a game of Go in 2016. Despite its relatively simple rules, the ancient game of Go is extremely

complex, with more possible moves than the total number of atoms in the universe. This however, played to AlphaGo's strengths, as machines are generally good at learning in areas where there are complex rules and large amounts of data. AlphaGo was self-taught, learning winning Go strategies by learning from observing examples of successes and failures, and by playing the game millions of times against itself.

The problem, is that AlphaGo's algorithm can't then instantly become a chess master, a poker player, a great conversationalist, or drive a car, because that's not what it is trained to do, and Lee Sedol, without a significant reprogramming of his brain, will.

The problem with Artificial General Intelligence

So-called “artificial general intelligence” (AGI) is extremely difficult to achieve. Getting machine learning to transfer skills from one task to another is a challenge: Google’s DeepMind team are currently looking at the possibility of creating a network of neural networks (called PathNet) to do exactly this.

Playing games, interpreting the stock market, writing articles about football, spotting patterns in large, messy data sets, and performing activities in highly structured and predictable environments are all relatively easy for machines to learn to do. Having a sensible conversation, negotiating with a difficult customer, having a gut feeling about a situation, and being creative about solutions are all things that a human is typically better at than a machine... at least for now.

Consider the value that these skills contribute to organisations. Jesuthasan and Zarkadakis, suggest that we need to look at the distinction between two types of role⁶:

1. **Proficiency roles:** requiring high levels of skill, but don’t deliver more value once people perform beyond a certain standard.
2. **Pivotal roles:** also requiring high levels of skill, but higher performance yields more value for the organisation, e.g better customer service.

The example that Jesuthasan and Zarkadakis give to illustrate this point is that of pilots, who, once they achieve a certain high standard in flying a plane, don’t deliver more value by getting even better at flying the plane – so they have a proficiency role.

The cabin crew, however, have pivotal roles as they are responsible for delivering customer service.

Higher levels of achievement from the cabin crew can have a huge value impact on customers. Investment in value-producing pivotal roles should be higher than in proficiency roles, but this is frequently the other way around. Taking some of the routine responsibilities off from the cabin crew, and making sure that technologies help them to deliver better customer care, is likely to be a wiser investment than automating them out.

In reality, much of the pilot’s role can be automated, since it is generally rules based and data rich, however it’s when things go wrong that their skills are essential. Few passengers would want to step onto a plane without a pilot (yet), but many planes are already flown extensively by auto-pilot, albeit with humans there to supervise.

This brings another element to the debate – how much do we trust machines to fly planes, drive cars, diagnose illnesses, or manage our financial affairs unsupervised?

This becomes exceptionally problematic as the machine learns, rather than being

programmed. Even the designer of the algorithm might struggle to explain the output of its child’s reasoned actions.

The moral dilemma that is often discussed here is that of the self-driving car. Should the car swerve to avoid a pedestrian if it thinks that there is a high likelihood that its actions will kill the driver? This is also a legal grey area – who is responsible?

Is it the car manufacturer, the driver, the designer of the algorithm, or the makers of the multitude of sensors that are involved in the operation of self-driving vehicles? Statistically it’s probably safer not to have a human at the wheel. However, there are a number of cultural, legal and logistical issues that need to be resolved before we see large numbers of self-driving cars on roads.

The debate extends to robo-advice in the financial services industry, and medical diagnosis. If a person is given a personalised recommendation based on the output of a machine learning algorithm, how is that advice regulated if the learning algorithm can’t show us how it came to that particular conclusion?



Past, present and future: where are we now with machine learning?

1980s first generation machine learning



My first practical exposure to machine learning was at BT Laboratories (now Adastral Park) in 1990. Machine learning and expert systems were all the rage, so we created a neural net based expert system called ESCFE (Expert Systems for Customer Facing Environments).

ESCFE helped customer service advisors to navigate complex diagnostics of international telecoms network faults. It was designed to make call handling more accurate and consistent by supporting both the business customer and customer service advisor through the process, regardless of the levels of expertise on either side.

The system used an inference engine, based on a neural network and natural language processing ("NLP"); a multimedia / multimodal

dialogue system at the customer advisor interface; and screen scraping technology, to get information about products, services, policies and selling opportunities to and from a number of several back end systems.

ESCFE was trialled in BT's International Customer Service Centre in Burnham-on-Sea in the West of England. It spawned a number of spin-off variations, using neural nets and natural language processing to listen to dialogue and proactively push knowledge to advisors. These were trialled in other regional BT call centres in Thurso and Newcastle.

The trials were successful proofs of concepts but there were some significant issues with this approach – some of which have been solved by 21st century technology, but some that haven't.



Figure: The Expert Systems for Customer Facing Environments (ESCFE) architecture.

The issues with first generation machine learning

The first issue was **data** – notably the kind of data that a ‘neural net’ needs in order to learn. Deep learning in the days of ESCFE was done mostly by interviewing various experts and putting their knowledge into the expert system. It was little more than a glorified decision tree that needed to be updated manually. It was also difficult to codify how experts made decisions. “Gut feeling”, built up through years of experience, wasn’t something you could easily express in a neural net.

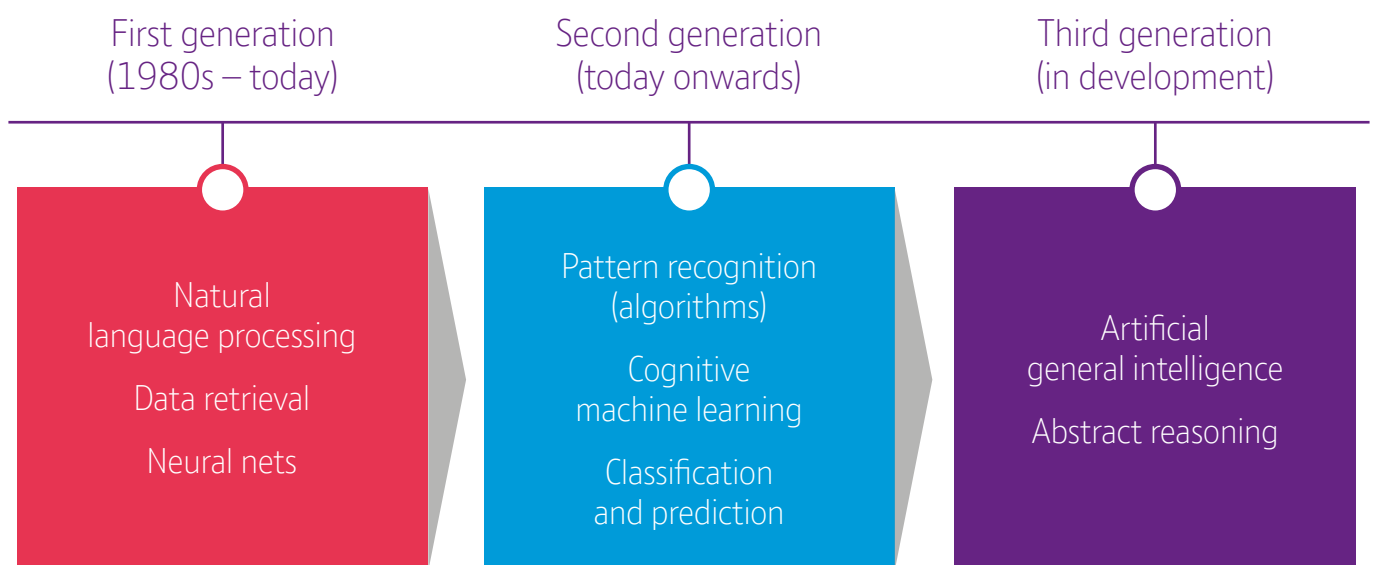
Now, machine learning has a lot more machine-readable data to fuel its algorithms, and processing power is a million times more powerful than it was in the 1990s. The cost of maintaining the system goes down if the system can learn, rather than needing continual human intervention. Gut feeling is still difficult to codify, though.

The second issue was **screen scraping** – pointing the system at specific data fields on multiple back-end databases so that it could extract it. The issue was that the method fell down if any one field on any database was shifted even a millimetre to the left.

This was a stumbling block that added greatly to the maintenance costs of the system. This has now been partially solved by robotic process automation or “RPA” – that is, using a software robot to follow processes that extract, fill in, and duplicate data into multiple back end systems, cutting out the need for humans to re-enter it.

The third issue – which has still to be definitively solved – is the **reliability of natural language processing, or “NLP”**. NLP is a technology that has come on significantly since the 1990s, but still falls down on one thing: the ability to process the context of the language, both in terms of the knowledge needed to process customer conversations, and the emotional situation. Because NLP was notoriously unreliable in the 1990s, advisors tended not to trust and ignore the output from the system. Interestingly, when advisors were more involved with the development of the expert knowledge base, they trusted and used it more.

Timeline of machine learning:




Machine learning today

Twenty-six years later, we are largely still in the **first generation**, formative years of machine learning. Most of the applications of machine learning we are familiar with today are based on pattern recognition (algorithms), used by Google, Netflix and Amazon for example. These algorithms can reason over narrowly defined, structured problems, but have only limited learning capabilities. There's still a fair amount of human intervention needed to maintain them, and they can't handle ambiguity or uncertainty.

A **second generation** of machine learning has however, now started to emerge. It learns from large sets of machine readable data, usually in specific areas, such as playing Go, or diagnosing cancer, and is sometimes referred to as "cognitive machine learning". Classification and prediction is possible, but it has limited ability to reason, or process context. As with the first generation of machine learning, it is only as good as the data available to it.

The **third generation** future of machine learning is likely to be closer to the artificial general intelligence ("AGI") that we talked about earlier. AGI will be able to both learn and adapt to context, but it is still unlikely to be able to do things like abstract reasoning.



Machine learning and its role in customer experience

The hype around "chatbots" (programs that can simulate conversation) means that many people think that this is the only technology using machine learning as part of the customer experience. The truth is that machine learning has been shaping digital customer experiences for a very long time.

Machine learning and analytics can make customer experiences personal, proactive and predictive.

1. Personal – using customer data to create personal experiences, such as Amazon and Netflix. This relies on the customer sharing their data. Of course, they may not necessarily know that they are doing this, as the biggest lie we all tell on a daily basis is that we "have read the terms and conditions," however, if it makes it simpler to navigate, or inspires us to buy things we never knew we needed until "other people like us" wanted them, we probably don't mind.

Much of this consumer behaviour is about the "me" economy. In other words, how much of my data am I willing to share in order to get a better experience, a more personalised service, make things easier for me, or get a better deal? One industry that is capitalising on this "me" economy is insurance.

Crippled by the spiralling costs of car insurance for younger drivers, telematics boxes in cars are being used to significantly reduce premiums when the driver drives sensibly. The driver sacrifices some of their privacy in return for a better deal. Wearable health technologies are carrying out a similar role in the health insurance sector.

Retailers are also doing their bit to personalise experiences by bringing digital into the physical store. They are using the consumer's smart phone to identify them (with permission, of course), and are deploying technologies such as facial recognition, geofencing (virtual perimeters) and beacons (devices that broadcast to nearby portable electronic devices) to guide consumers to special offers based on their past preferences.

Loyalty schemes have been doing this in various ways for years. Supermarkets across the world have been quietly gathering data from in-store and online purchases, and crafting offers “just for you”. It has been well documented that some of these offers cross “the creepy line” for example they have been predicting pregnancy, divorce etc. This is the point that people sometimes decide that this is too “big brother”, and stop sharing their data.

For those of us opting in to the “me”conomy, we need to be able to see tangible advantages to doing so – otherwise, we may start to conscientiously object. As we explode into an era of the Internet of Things and omnipresent connectivity, we may all start to have a sneaking suspicion that our technologies are talking about us behind our backs.

Machine learning could be used to predict credit ratings for people with little or no credit history.

2. Proactive – the next step beyond personalisation is proactivity. This is about telling customers that there are problems or issues with their service (or solving them) before they have to do anything themselves. Noticing that there is an issue in the first place may not be obvious to the naked eye, but it may well show up in the data.


Security and fraud detection is a major application for proactive machine learning. This can include simple things, such as your bank contacting you because there is unusual activity around your account, to far more large-scale ‘denial of service’ attacks on systems. BT’s ‘Saturn’ tool is an example of a proactive machine learning algorithm that can warn of potential attacks, failures, anomalies and threats, by drawing on a fusion of data feeds.

Similar tools, such as BT’s ‘Aperture’, can look at patterns in large sets of customer and other related data. It can spot customers that are likely to churn or complain, and also look at incidents and scenarios that are likely to cause customers issues and identify the root causes of them. In the case of the churning customer, if they are a valuable one that you’d like to keep, you might want to contact them before they contact you to cancel their service.

3. Predictive – this is more uncertain territory for machine learning. Most intelligent systems don’t have a ‘crystal ball module’, but they can run scenarios based on big data sets to predict the most likely outcome of a situation.

This is largely how organisations like the UK’s Met Office currently forecast the weather. The machine gathers data from multiple data points, and then creates a number of potential forecast scenarios before, typically, a human meteorologist takes an educated stab at the most likely forecast. They are not failsafe, but the amount of data and processing power is such now that they are far more accurate than they used to be.

But can this be used to shape customer experiences? Atom Bank in the UK says it wants to become the “world’s first psychic bank”, using machine learning through their app to predict what customers might need next. Barclays have trialled a chatbot to help customers to manage their future “financial health”, based on their past transactional data. Similarly, machine learning could be used to predict credit ratings for people with little or no credit history, by using a multitude of online behaviours including how they search and browse⁷.





How will machine learning change the contact centre?

The contact centre has been at the forefront of the march of automation for many years. Digital customer care, self-service technologies and the rise of the app have already pushed a lot of the simple, repetitive stuff out of human contact queues.

The contact centre has been at the forefront of the march of automation for many years. Digital customer care, self-service technologies and the rise of the app have already pushed a lot of the simple, repetitive stuff out of human contact queues. Overall contact across all channels has gone down, but contact handling times into many contact centres, before machine learning technologies have been widely adopted, have gone up by as much as 40 per cent.

Our SuperAgent 2020 research⁹ found that the contact centre's role has shifted from "processing lots of simple, repetitive calls, really, really fast", to taking on much more complex and emotive issues than ever before. 54 per cent of those asked in the survey believed that 'complex problem solving' would be the core function of the contact centre by the year 2020.

To examine the effect of machine learning on the future development of the contact centre we need to consider five critical factors³:

1. Technical feasibility
2. Cost of development and deployment
3. Availability of skills in the labour market
4. The economic benefits
5. Regulatory and social acceptance

1. Technical feasibility

The machine learning algorithm's effectiveness will depend on the data it has available to it (no data, no point), the suitability of the activity, and the maturity of the machine learning technology available.

We must start with data, because “garbage in equals garbage out”. The centralisation of customer, product, and service data, combined with stable and well-documented processes, and increasing amounts of real-time data, can all provide fuel for the machine learning algorithm. If you have that, machine learning is a viable option – if you don't, you need to look at how you get to that point before machine learning can get you anywhere.

In the contact centre, it's highly likely that contact data is freely available. If you can go beyond queues and durations, and figure out what people are contacting you about in real-time, you can use this to be more proactive. This means that contact centres can better manage demand, rather than be managed by it. If there is a sudden flurry of contact about the website, it may indicate that there are issues with the website that need to be quickly addressed, before the flurry becomes a blizzard.

We also need to consider the technical feasibility of a machine doing a human advisor's job. A typical customer-to-advisor interaction will consist of:

- **Triage and routing** – A fair proportion of the up-front triage work can be done by the precision routing engine, based on customer input by voice into interactive voice response (IVR) on a device, by text into a chatbot, analysis of social media content, customer activities on the website or app, and any other data that may be relevant, such as complaints. This will “speed date” the contact to the advisor with the skill set most likely to match the contact (see the section about advisor skills later on in this paper).
- **Identification and verification** – Next is identification and verification of the customer's identity. Rather than remembering strings of numbers or passwords (or, in many cases, not remembering strings of numbers or passwords), much of this is likely to become automated. It could be replaced with biometrics such as voice, selfie, thumb/finger/vein/iris print; or with pre-authentication through apps and visual IVR, a mechanism where customers are presented with a list of personalised options, based on their profile before clicking to call or to chat.

- **Conversation** – If everything has gone smoothly, the advisor with the right skillset will then pick up the contact and have a conversation with the customer. A critical, and often underestimated, part of the advisor role here is to:

1. Listen and understand the customer issue.
2. Translate that into the language that the organisation uses.
3. Use a combination of their knowledge and experience, plus data from back end systems, to solve the issue.
4. Negotiate with the customer to reach a mutually acceptable outcome.
5. Make sure that the customer understands what the outcomes and next steps are, and that they are happy with them.

Although some of these activities are automatable, there is one critical element that falls foul of Moravec's paradox. Conversation is a surprisingly complex activity to automate. McKinsey's study³ estimated that it has a low potential for automation, at just 20 per cent. This is because conversation typically requires not just linguistic and cognitive capabilities, but also social and emotional ones.

Natural Language Processing has come on leaps and bounds in terms of recognising what's being said, however, it's still a victim to the quirks and ambiguities of human language – which evolved somewhat chaotically as a way for humans to communicate and build relationships with other humans, not with machines¹⁰.

This brings us to the subject of chatbots.



Rise of the chatbots

The growth of chatbots is largely fuelled by the popularity of chat as a communication tool for customers. Some of the most successful chatbots to date have been used on messaging platforms such as Facebook Messenger, Viber, Kik, Skype, Line, Slack, Google's Allo messaging, and WhatsApp. Brands such as Uber, Edwardian Hotels, CaixaBank, Bank of America, Mastercard, KLM, Dominos Pizza and Burger King are all experimenting with chatbots.

One of the most successful pioneers of chatbots is WeChat in China. It supports 10 million businesses, including Nike, Burberry and Pizza Hut. WeChat allows people to hail a taxi, order food, pay a bill, and book a doctor's appointment. Chatbots on WeChat can take advantage of application programming interfaces (APIs) with over 100 "bubble message" types, including text, voice, video and speech. They also have APIs that allow the integration of multiple data sources, including access to direct messages, voice messages, user IDs, and customer location data.

From a developer's perspective, chatbots have several advantages over apps, including the facts that they are less data intensive, and use less memory. Apps need to be downloaded from an app store, while chatbots don't need to be pre-installed. This means that we can quickly pay for a car parking ticket without having to go to the trouble of downloading an app we'll probably only use once.

From a brand perspective, the decision is around where to deploy a chatbot. Is it over one of the many chat platforms, over webchat from a webpage, or over an in-app chat? Integration with existing CRM, customer service systems, and databases also needs to be considered because, as we've already covered, the chatbot is only as good as the data behind it.

First generation chatbots work well for simple queries and transactions. Rather than allowing full conversational natural language processing, many chatbots restrict dialogue by offering multiple choice options, or simple yes/no answers.



Chatbots are able to steer the customer through the "known knowns", without the customer having to plough through lists of FAQs, however, they tend to struggle with complaints and complexity. If we get angry, we tend to tell long and rambling stories which are hard for machines to break down into something meaningful.

Sarcasm may be the lowest form of wit, but it can also confuse algorithms. Sentiment analysis can fail horribly on complaints coming through on digital channels, such as social media and email. Here are a few of my favourites, that would probably be interpreted by a machine as a positive comment: "Thanks, @TrainCo for my free sauna this morning"; "It was so good to see that your maintenance department hadn't spoiled things by making unnecessary repairs"; "I would have cheerfully strangled him".

Some of the best chatbots don't pretend that they are human. No machine has yet to convincingly pass the Turing Test, the test that is passed when a human interacts with a computer, and they don't know they

are interacting with a computer, or the Winograd Schema Challenge, which is similar to the Turing Test but uses a set of multiple choice questions that require common sense reasoning to successfully answer. When faced with ambiguity, complexity, sarcasm or a dead end, the best bots act as "IVR for digital" and quickly triage the customer to an advisor with skills that are most likely to get the customer to their goal. The advisor has the advantage that they can see the conversation that the bot has had with the customer, so the customer shouldn't need to repeat themselves.

Undoubtedly, the ability for machines to process natural language will improve – especially with the evolution of virtual assistants such as Siri, Cortana, Alexa and Google Now. Second generation chatbots, such as the impressive, Loebner prize winning Mitsuku (see <http://www.mitsuku.com/>) are getting better at free-flowing dialogue, but aren't yet glittering conversationalists. Humans find it far easier to have a conversation – as long as they aren't constrained by too many rules, or by a script.

2. The costs of development and deployment

The business case for investing in machine learning is unlikely to stack up if it costs more to develop and train the algorithms than it costs to pay the employees currently doing the job.

This was the main reason that it didn't work in the 1990s. It's also why people have argued that it simply isn't worth automating some lower wage jobs, while it might be worth automating higher wage ones, however this argument fails to take into account the value of the work to the customer – the pivotal vs. proficiency role argument⁶.

Customer service work is a pivotal role because of the value it can create in the relationship with the customer. By taking the routine, transactional activities, which add little value to the relationship, away from contact centres, the advisors are able to focus on the more complex and emotive tasks that do add value.

Dixon, Ponomareff, Turner and DeLisi¹¹ have pointed out that companies that have focussed on new self-service technologies have tended to then underinvest in the human frontline, and while the self-service experience has evolved, the contact centre has changed very little. This has created a yawning gap in the service that customers expect, versus what they get from the contact centre.

The cost of putting untrained, unprepared and disempowered advisors in front of a customer with complex and emotive problems is a false economy that may, ultimately, drive costs up. In one of our studies, we compared the top and bottom quartiles of contact centre advisors based on pay and experience. On average, we found that it took the bottom quartile eleven minutes longer to deal with complex customer issues than the top quartile, with a resulting negative four-point decline in customer survey measures.

It makes sense to put the most experienced, talented and innovative advisors in front of challenging customer issues, but to do that, technology needs to take away the simpler issues that often eat up their time. Automation can also free up time for the less experienced advisors to train and upskill.



3. Availability of skills in the labour market

Some hard to get, hard to find skills that are highly rules-based, proficiency activities, such as some medical or law roles, may well be desirable to automate because there aren't enough people available with them. Similarly, jobs that are classed as “dull, dirty or dangerous”, such as bomb disposal or factory assembly, for which it is difficult to recruit and retain people, might also be ripe for automation.

Recruiting “SuperAgents” into contact centres who can handle all channels and all problems is likely to be challenging⁹. They may not need X-ray vision, but we may need to look at some very different skills for the future, as automation changes the demands of the traffic going into the contact centre (see section on SuperAgent skills later).

4. The economic benefits

Economic benefits are found more in adding value and productivity, rather than making human employees redundant. By automating routine tasks in the contact centre, advisors can add more value on more complex contacts, especially seeing as they have a pivotal role, where higher performance adds more value to the customer.

Similarly, technologies such as robotic process automation can improve consistency and quality, whilst reducing error rates in data intensive back office functions. This means that employees have more time to upskill both themselves, their less experienced colleagues, and the machine learning algorithms.

5. Regularity and social acceptance

The key to implementing machine learning technologies well is to make life easier for your customers. Do it badly and you'll create a backlash.

For example, if chatbots are IVR for digital, then we need to ensure that we follow what we've learned from traditional IVR design. Too many chatbots have been deployed with little attention to customer interaction. The conversation often results in dead ends, or the dreaded "sorry I didn't understand your request". For chatbots to work properly, options must be simple and in the customer's language. There also needs to be an option, initiated either by the customer or the chatbot, to connect to a human advisor. If that happens, the conversation with the chatbot needs to be sent across to the advisor, so that the customer doesn't

need to repeat themselves, and the human advisor can pick up where the machine left off.


From a regulatory perspective, transparency is a key issue with machine learning. If a customer acts on a recommendation made by a chatbot and it goes wrong, where does the liability lie? This is especially difficult when dealing with second generation chatbots, which learn rather than act on programmed behaviour.

Take the case of Tay, the Microsoft chatbot that was corrupted by Twitter in a period of only 24 hours. It quickly became what the UK's Daily Telegraph called "a Hitler-loving sex robot"¹², because it was taught some dubious habits by some unsavoury characters on that particular social network.

Similarly horrifying was Google's machine learning photo classification app that incorrectly tagged a photo of an African American woman as a "gorilla"¹³.

These behaviours could be damaging if the technology is seen as representative of the brand. That's why machine learning algorithms need to have some responsible parenting from humans as they learn.

Several of the chatbots that have been developed for customer service, including Digital Genius, learn appropriate responses from the human advisors who sit alongside them. The bonus here is that the advisors become part of the development process, and they tend to become more supportive and positive about the technology as a result.




What skills do “SuperAgents” need?

With a combination of self-service technologies, and chatbots and machine learning at both the customer and advisor interface, the fabric of the contact centre is changing. This means that the skills that we need are changing from rules based, constrained and scripted environments, to more empowered, empathetic and creative ones.

The things that are harder for an AI to do in contact centre work are those that make us uniquely human: conversation, empathy, creativity, intuition and negotiation. These are the pivotal skills that the human “SuperAgent” needs to enhance the customer experience, and make it unique to the brand. Complex contact will be in the realm of a skilled, empowered human advisor, working alongside a number of different types of machine learning. Our research tells us that this kind of advisor is likely to inhabit the “contact centre” as we head into the future.





Meet Sally, the “Swiss army knife” advisor of the future

Sally is an experienced advisor. She is adaptable and has proved to be an exceptional communicator across multiple channels. Her role is to take over once the technology has handed the contact over for further triage. The bulk of her work is dealing with chat. She’s able to read through the conversation with the chatbot before picking it up seamlessly with the customer.

Many of the customers Sally deals with have fallen between processes, got frustrated, and need a bit of advice or redirection. A stable and user-friendly systems support helps her to deal with this. It includes a machine learning engine that helps her to navigate through internal processes, and make sure that she is compliant with regulation. Robotic process automation also enters data on multiple back-end systems on her behalf, so she doesn’t have to. This means that Sally doesn’t spend all of her time training on products, processes and systems: it’s the customer that’s important.

Sally knows that her emotional intelligence, empathy and creativity set her apart from the automated aspects of the contact centre. She has regular assessments from the

sentiment analysis engine that looks at the positive or negative impact of interactions on both Sally and the customer. Her work can be extremely challenging, as customers are often angry or frustrated, so she knows she needs to take a break if she gets too stressed, as it could negatively impact on the next customer.

Sally can see everyone’s performance ratings and goals across the whole contact centre. Transparency is key: she knows how they all contribute to the centre’s overall purpose.

On more complex issues, Sally has access to a skills registry that helps her to find the right expert to pass the customer to. She also has access to a lot of collaboration tools – including internal chat and enterprise social media – to help her engage with others around the company who might be able to help her out with customer problems. Rather than simply responding to demand, much of Sally’s job is also about being proactive with customers.

Data coming in from sources such as the Internet of Things allows the analytics engine to identify issues affecting the customer

before they notice. Much of the proactive interaction is handled by the chatbot – but Sally gets involved on the more emotive issues.

In her spare time, Sally acts as a coach for both the chatbot, making sure that it responds to customers appropriately, and other rookie advisors who are learning the job. She got value from a mentor when she was getting up to speed with the job, and wants to give the same back to new advisors. They no longer get to cut their teeth on the easy stuff anymore, as the automation takes care of the simple things for them.

Sally actively feeds back suggestions for improvements to Sheila, her contact centre manager. Sheila is part of the panel of managers who act as guardians of the customer experience. She feeds back issues daily to the Chief Customer Officer, along with the results of real-time insight on which contacts are coming in, the root causes of their referral, and other customer measures. Because both Sally and Sheila can see how they influence the customer experience and can actively feedback improvements, they are both highly engaged in their work.

Meet the “networked experts”

These advisors are experienced specialists who can take on complex issues. They may live beyond the boundaries of the contact centre, and may not be considered to be “contact centre advisors” in the traditional sense. In this model, the traditional contact centre is a bottleneck – so complex issues are triaged by a combination of machine learning and, where necessary, the Swiss Army Knife advisor, and then precision routed to the most appropriate and available expert, regardless of channel.

Networked experts have typically built up knowledge over years of practice, and much of it resides in their heads. Knowledge tends to be a child of the past – you can solve known problems with it. If it is documented somewhere, it can become part of the dataset for machine learning. This is why networked experts are so valuable. Mastery is needed for new, or emerging, problems.

The issue here is that mastery and expertise are rarely visible in the organisation chart, pay structure, or job descriptions. Every organisation has “secret knowledge” about people with mastery. We tend to know who these people are, but is only through exploiting human connections that gets us to them.

For this model to work, companies need a much better view across their business (or, indeed, their supply chain) as to who and where these experts are, whether they are available to talk to customers, on what channel, and what training they have, such as in customer service or brand language.

This is why, on a technical level, embedding tools like Unified Communications (UC), and collaboration into the contact centre makes a lot of sense. Advisors need to be able to search for skills, see presence information and collaborate over channels like instant messenger (IM), or enterprise social media with people who can help solve complex customer issues. Leadership in this model is less about silos and command and control, and more about connection and collaboration.



Networked experts can cover a number of ‘hard’ and ‘soft’ areas of expertise, all of which are difficult for machine learning to achieve. Here are a few examples:


Paula the problem solver

Paula loves problems. Her role is to take on customers who have fallen through the gaps in process, and project manage them on their behalf. Paula acts as a single point of contact for these customers. Once she has taken them on, they can get in contact with her over various channels. Her job is to navigate the spaghetti of internal processes, question decisions made by the machine learning and work with other networked experts, to attempt to come to a solution that works for both the customer and organisation. Success measures for her are around the quality and effectiveness of her problem solving, and customer outcomes.

Natalie the negotiator

Natalie works with other networked experts on especially emotional contacts. She trained as a lawyer and is extremely articulate and persuasive, however she is also patient and a good listener. She knows that rules based negotiation with no leeway, empowerment, or empathy usually results in a “computer says no” moment, whether it’s a human or a machine.

Natalie picks up the complex interactions, where straightforward, rules-based negotiation won’t work. She knows that the key to negotiating mutually acceptable outcomes is for her to be open, empathetic and flexible, but to know where the company’s limits are. She tries to find areas of common ground with customers and work from that: understanding where they are coming from, being open to suggestions, but being assertive when they are demanding too much.



Tony the technical master

Tony is an experienced engineer, with a whole career behind him. Retirement is not something Tony likes the sound of so, at 71, he puts his skills and expertise out in the customer space to solve complex technical issues. All he needs is a network connection and the ability to fly himself into the cloud, so he does much of his work flexibly from home.

Tony has worked with customers all his life, so he knows how to talk to them. He is often asked to mentor some of his other technical colleagues about how to talk to customers. All new technical networked experts have to earn customer service skills accreditation, so that they don't blind customers with science and confuse them.

Tony is often networked into conversations across multiple channels – including phone, video, chat and social media – by people such as Sally, the Swiss army knife advisor, and Paula, the problem solver. A machine learning algorithm alerts him when one of his products are being talked about on social media and forums, and he can choose whether to get involved in the conversation or not. Other algorithms may also be used to spot patterns in unstructured data and bring this to the attention of Tony, so that he can identify specific issues, such as hack attacks.

Tony often works on new products and services: in these areas, there are always teething problems and very little in terms of a stable knowledge base. Tony is responsible for working with other experts in product and service lines to make sure that the data that any machine learning algorithm may be working from is up-to-date and accurate, and reflects any new or emerging issues that he has come across. He also involves “crowdsourced advisors” (usually enthusiastic customers), on the development of more customer friendly processes. The network of expertise can go beyond the organisation, and customers may also become part of it.



Chris the crowdsourced advisor

Chris loves helping people and solving problems. He is a technical consultant making a good living in the gig economy, but he also gives a lot of his time away for free on social networks and forums, trouble shooting for other people. Chris is occasionally paid for his trouble, but he gets a lot of other things from companies he “works” for. He gets kudos points ratings for his solutions from other customers and, because he is reliably good, he often gets a lot of free and exclusive stuff from companies he works with as a thank you.

Chris is also on a number of customer panels and goes to a number of hothouses, as involving customers in innovation is becoming increasingly important to many organisations. This lets him get exclusive access to beta trials of new products and services. He loves to feedback any issues he’s found with them, or feed in suggestions for improvements.

On some issues, Chris works with technical networked experts, such as Tony, on “communities of problems”. Because he is recruited specifically for his obvious skills as a community trouble-shooter, he is usually paid for these contributions.



All these experts:

- Are self-motivated and more experienced
- Are collaborative in creating solutions with other advisors and experts
- Enjoy the kudos of being an expert, or a master.

In a world of automation and self-service, it’s clear that talented advisors can make all the difference. We don’t just want lovely people who are great conversationalists and empathisers, we want people who can deliver results for customers in a more and more complex world, with the help of intelligent technologies.

Advisors cannot be drones who follow rules and processes, because machine learning can easily do that. They need to deliver the “human factor”, think on their feet, tackle the tough stuff, and get involved in improving the customer experience. This also means that the contact centre, probably a very distributed, rather than centralised “centre”, is far more strategic in terms of understanding customer issues, and mobilising the whole organisation to solve them.

This relies on customer experience being given priority in organisations, including visibility of customer experience at board level. Investing in machine learning technologies shouldn’t be about cutting costs. It should be about getting machines to take away the simple things, to free up human advisors to add value to the customer experience.



Botman vs. SuperAgent: The verdict

We like pitting man and machine against each other. It's the 'bot verses the agent in a struggle for human survival'. Dystopian science fiction stories like 'The Terminator' loom large in our minds, with machines crushing humans under their feet. This obviously isn't a great scenario for us – but we do have a choice.

The 'Iron Man' scenario looks far more exciting. We want to be able to prolong our lives, fly, and have cool head-up displays floating in front of our eyes. We want machines to enhance us as humans, not replace us.

Machine learning has come a long way since the 1990s, but there are still some formidable technical challenges ahead. While machines can be trained to perform a range of cognitive tasks, they remain limited. They're not good at putting knowledge into context or improvising. They have little of the common sense that is the essence of human experience and emotion. They struggle to operate without a pre-defined methodology. They are far more literal than people, and therefore poor at picking up social or emotional cues, such as sarcasm or irony.

One big clue to the future of man and machine can be found in Garry Kasparov's experiences with IBM's Deep Blue. The chess master was beaten by the machine in 1997. At that point, he could have simply have thrown in the towel and given up, but from the ashes of that defeat came 'advanced chess' – where human and machine in combination are pitted against other humans and machines. As a result, an average chess player can become a supreme champion. Machines play chess in different ways to humans, but the combination of both can create better outcomes.

In the customer experience space especially, the most likely scenario isn't about Man vs. machine – it is man plus machine. We need "cobots", not robots, to allow us to become more productive, valuable, and more human. If we create customer service people who act like robots, we are missing a huge opportunity to add the value of a human brain and its ability to be empathetic, creative, and intuitive.

Simply put, if the combination is right, smart people partnered with smart machines have the potential to superpower us.

It's not Botman verses SuperAgent; it's Botman plus SuperAgent.

References

1. McAfee, A. & Brynjolfsson, E. (2016), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton & Co.
2. Frey, C.B. & Osborne, M.A. (2013), *The Future of Employment: How Susceptible are Jobs to Computerisation?* Oxford Martin, http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.
3. Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P. and Dewhurst, M. (2017), *A Future that Works: Automation, Employment & Productivity*, McKinsey Global White Paper. <http://www.mckinsey.com/~media/McKinsey/Global%20Themes/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx>
4. Arntz, M., Gregory, T. and Zierahn, U. (2016), *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*, OECD Social, Employment and Migration Working Papers, No. 189, OECD Publishing, Paris. <http://dx.doi.org/10.1787/5jlz9h56dvq7-en>
5. Moravec, H. (1990), *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press.
6. Jesuthasan, M. & Zarkadakis, G. (2016), *Automation Will Make Us Rethink What a “Job” Really Is*, Harvard Business Review, October 12.
7. Tett, G. (2017), *Donald Trump’s campaign shifted odds by making big data personal*, Financial Times, 26th January, <https://www.ft.com/content/bee3298c-e304-11e6-9645-c9357a75844a>
8. Byrnes, N. (2017), *An AI-Fuelled Credit Formula Might Help You Get a Loan*, MIT Technology Review, February 14th, <https://www.technologyreview.com/s/603604/an-ai-fueled-credit-formula-might-help-you-get-a-loan/>
9. Millard, N. & Alcock, T. (2014), *SuperAgent 2020: The Evolution of the Contact Centre*, BT White Paper, <http://www.slideshare.net/btletstalk/bt-whitepaper-superagent-2020-the-evolution-of-the-contact-centre>
10. Zarkadakis, G. (2016), *In Our Own Image: Saviour Or Destroyer? The History and Future of Artificial Intelligence*, Ebury Digital.
11. Dixon, M. Ponomareff, L., Turner, S. & DeLisi, R. (2017), *Kick-Ass Customer Service*, Harvard Business Review, January-February.
12. Horton, H. (2016), *Microsoft deletes ‘teen girl’ AI after it became a Hitler-loving sex robot within 24 hours*, Daily Telegraph, 24th March, <http://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit/>
13. *Google apologises for Photos app’s racist blunder*, BBC News, 1st July, 2015, <http://www.bbc.co.uk/news/technology-33347866>

With thanks to Daniel Ballin and Sandeep Raithatha from BT Technology, Service and Operations for their comprehensive review of chatbots.

Offices worldwide

The services described in this publication are subject to availability and may be modified from time to time. Services and equipment are provided subject to British Telecommunications plc's respective standard conditions of contract. Nothing in this publication forms any part of any contract. © British Telecommunications plc 2017. Registered office: 81 Newgate Street, London EC1A 7AJ. Registered in England No: 1800000