

“You just can't differentiate between a robot and the very best of humans”

- Isaac Asimov, “I, Robot”

In recent editions we've featured a series of articles from the Natural Language Understanding experts at ContactEngine. The more we found out about the company, the more interesting we found their slightly leftfield take on the role of Machine Learning and AI in the customer experience, so we sat down with their charismatic CEO Prof. Mark K. Smith to find out more about ContactEngine, proactive conversational AI, and his view of the future of customer experience. Along the way we'll pick up some crucial insight as to where AI does and doesn't fit in your customer journeys.

The problem with Chatbots

When you talk about AI and the customer experience, many people think immediately of Chatbots. I'm sure we've all encountered them, and I'm equally sure we've learned to be

suspicious of vendors who claim that theirs are indistinguishable from humans. Those vendors, it seems to me, must know a lot of very stupid, boring, people.

Used in the right way, Chatbots serve a purpose. When you're honest about what they are, which is essentially a user-friendly skin built on top of a FAQ, that's fine. The mistake is to see them as an alternative to a proper conversation. As Mark comments,

"Why do you have chatbots? Why do they exist? It's containment for overspill of people going to websites, stopping them reaching call centres. They even call it 'containment', as if customers have a virus. It's just wrong."

You may be wondering why someone who runs a company specialising in conversational AI is so dubious about the benefits of Chatbots. The answer is that Mark believes he has identified a unique niche in which AI can be used to enhance the customer experience, rather than to save cost at the risk of making the customer experience worse. It's a niche where customer experience, business efficiency, and the strengths of Machine Learning line up to allow automation to help everything flow more smoothly.

AI and the Customer Experience

Based on our conversation, I think there are 5 crucial elements that make this kind of automated communication work, which we can take as generalisable rules for where automation makes sense in the customer experience. I believe AI makes sense when it is proactive, focused, conversational, learning, and context-aware. Let's look at each of those in turn.

Proactive

One of the problems with Chatbots is that they are reactive; they respond to a request or enquiry from a customer, and that request or enquiry could

be almost anything, worded in a massive variety of ways. ContactEngine's approach is different.

"Because we asked the question, we know the context of the reply. We might ask the question about a loan application, or an insurance product, or a washing machine, but we know what was first said."

"Use humans to do what humans are best at, and then machines."

Focused

ContactEngine sells itself on using communication to improve the small moments of inefficiency that bedevil so many businesses: the missed appointments, the unhappy customers who need an opportunity to be heard, the information updates that prevent inbound calls.

"We start a conversation with somebody that says something like, 'we're coming to your place in three days' time, is that still on?' And when somebody says 'yes', we'll say, 'we're coming to number one, the high street, is that the correct address?' and then we carry on the conversation."

The point is that, because you know so much about the context for the customer's response, and because you started the conversation, you have naturally constrained the possibilities for what they are going to want.

"The intents, when you ask a question, are reduced by quite a lot, particularly if it's a single question. You can maybe say 15 intents cover 98% of the objectives, something like that. Machine learning algorithms fly when they are fed training data like that."

What about the 2%, then?

"There will always be the need for human beings to deal with exceptions, but machines are better at a lot of that sort of work."

This is a crucial point. When we use AI well

we take away unengaging tasks that humans would rather not be doing, and we give those tasks to machines that not only don't mind that they're boring, but they actually perform them better and more reliably (not to mention more cheaply). We can easily think of situations which we wouldn't want to automate, for instance when someone calls to make a claim on a life insurance policy. As Mark points out, that logic may apply only to the initial call, and automation may well have a role later on in the journey:

"A machine won't be the best to do that, because those calls are long, and dealing with grief. The first call is counselling, this person is in bits, so that's where it has to be humans. After that, the machine is fine, but initially you need a human being because machines can't do empathy. Use humans to do what humans are best at, and then machines."

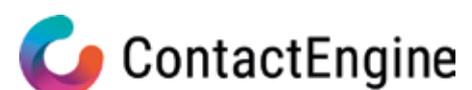
Conversation

Of course a lot of this kind of communication is already automated, but what is relatively rare is for an organisation to automate conversation in this context, so that the customer can get an SMS or email and interact intelligently with a computer at the other end of it.

"We're dialogue, not monologue. The technical challenges of conversation are vast, of course, but if you are connected into what the company wants and the service the customer wants, then you could make massive cost savings."

Learning

AI is a frustratingly vague term. Even if we restrict our definition to Machine Learning (ML), the plethora of algorithms, approaches,



and implementations often makes it very difficult to know what vendors are talking about. I suspect, though I can't prove, that the much-hyped AI solutions of some vendors are often very simple algorithms applied with varying degrees of cleverness to very simple problems.

One of the trademarks, it seems to me, of true ML, and one that is rare because it's relatively difficult to do, is ongoing learning. As Mark says,

"It's got to be learning, it's got to get better with time, and that's really rare. By labelling the data you arrive at a point where you can outperform a human agent very rapidly. The learning bit comes from when you take the exceptions, deal with them, and then that's added to the algorithm. So it gets better, and better, and better."

Context

Making outbound contact to a specific customer about a particular event means that the context for the conversation is well understood. That has benefits in terms of language understanding, as we've already seen, by narrowing the scope of likely responses. It also opens up the ability to personalise the conversation.

That opportunity can be a risk—there's a very fine line between intelligent personalisation and creepiness—but there are cases in which it clearly makes sense.

"It's a really fine line. You have to travel very carefully through that, and you have to make sure your GDPR compliance and all those things are right there, but there are things that you can do. Take telco as an example: for some processes someone has to have something before the next thing can happen, like receiving something in the post before the connection can be made. If you choose not to connect those two events, there will be 10-15% people where it will not have happened, in which case the second communication makes no sense. So what you need to do is confirm that they've got it before the second communication happens. That's a very logical sequence and it's not creepy, it's just sensible."

Judging that line between personalisation and creepiness can seem difficult, but a good starting point is to ask who benefits from the use of the data that we've got. If, like the telco example, it's 100% in the customer's interest, then it falls on the right side of the line. We can even make a good argument, as Mark does, that judging the timing of a sales message is ultimately showing respect for the customer's feelings:

"In the world of financial services, where someone has a successful mortgage application, and then is surveyed on NPS - if they give a 10 out of 10, then it's perfectly reasonable to offer them an additional product, maybe home insurance. If the answer was zero, then don't do that right now. That's rapport as well, because you're looking at patterns in the data to make an offer at an appropriate time, which isn't irritating."

That's obviously in the organisation's interest as well, but I think it's fair to argue

that it's common sense not to try to upsell a customer while they're unhappy, and that they'd rather you didn't!

The ethics of AI

There are times when AI can slide from creepiness to impacts that are downright unethical. Some of the key issues, all of which are related, are interpretability, bias, and the impact on society.

Interpretability

Perhaps the majority of AI at the moment rests, directly or indirectly, on people tapping into big "AI as a service" providers, especially a few key players such as Amazon, Google, Apple, Microsoft and IBM. Mark is glad that ContactEngine decided early on to develop their own algorithms in-house:

"What they do is not open, and it's a GDPR nightmare. We recognised that some years ago and decided to build our own, which really went against the flow. We were lucky we made that decision, because there's now a big kickback against the black box AI solutions that people use."

Interpretability is a big challenge in ML, because it's often the case that we can train the machine to get the right answer, but we don't know how. If we can't explain how, then there is always the possibility that the machine will make unexpected mistakes*, or bake in bias. Developing explainable AI is important to Mark:

"There is an argument you're not ever arriving at singularity or sentience, but you are absolutely performing like a human and getting better with time. Therefore, by doing this, you can not only out-perform the agent, but you can explain it as well. You can visualize it. You can actually say, 'we made this decision because of that'. So we're not trying to make a life or death decision, we are living in a simpler world than that, and that is proper AI; applied, and white box, and explainable."

*There's a great apocryphal story about an early neural network that the US Army trained to spot camouflaged tanks, but which was really detecting photos taken on a cloudy day. Sadly it's not really true: <https://www.gwern.net/Tanks>

“The way I see it is that computers take away jobs that humans simply don't want to do, and they make them happen better.”

Bias

Most ML applications work by working with a set of training data, and learning to replicate the label a human would apply by looking at patterns of association between features of the data and the label applied. If there are systematic biases in the way that humans apply those labels, then the algorithm will learn those too, which has the potential to introduce biases. Importantly, the machine doesn't do this on purpose,

“I dislike intensely the notion that the AI itself possesses human traits of bias. Algorithms are not racist, or sexist, or homophobic, or antisemitic. The data reflects society. It is not the computer's fault.”

In fact, there's an interesting parallel between the ideas of algorithmic bias in machines and unconscious bias in humans—both reflect structural problems in society that probably need to be addressed at a societal level. It's not really fair to expect AI developers to address these issues, but I think it is fair for them to be expected to engage with the issue, and at least not make the situation worse. Explainable AI means that the biases and the model are there to be checked and talked about and discussed. If it's a black box, you can't.

Robots in society

AI opens up the potential to use the data that we hold about customers in ways that

are simply not possible with traditional approaches (although, frankly, we were never making the most of our data anyway!). Before we dive into it, we need to stop and think about what we should and shouldn't be doing with the data with which customers have trusted us. Mark gives an example:

“You could imagine a situation where you were trying to do inferred importance of value to a client based on the quality of the language that's coming back to you. We don't do that, but there is quite a lot of work that suggests you can work out people's educational background based on the way they write. So you could make that inference. Humans do it all the time.”

That last point is really interesting, isn't it? Here we are wringing our hands about algorithmic judgements, but what about the judgements that our human staff are making every day? It's true that algorithmic biases can scale in a way that an individual human's wouldn't, but again there seems to be a wider point here about the ways in which we make decisions about how to deal with individual customers. People are nervous about self-driving cars, but what about the human drivers who are killing 2,000 people a year on British roads? As Mark comments,

“The autonomous vehicle is held to a higher standard than the human.”

What about the impact of AI on jobs? When should we expect to be replaced? With a few very specific exceptions, we should probably take the more extreme predictions with a pinch of salt:

“I think there's a tremendous arrogance from the tech community to imagine that computers will cross into sentience. It's just ridiculous. I also think that every 10 years there will be cataclysmic predictions about the end of humanity because of AI.”

As far as Mark is concerned, the most effective use of AI is in very specific, limited, domains. Jobs that a machine can

do better than a human, and that humans find unengaging. ***“The way I see it is that computers take away jobs that humans simply don't want to do, and they make them happen better. I know one large telco who have 12,000 people in a call centre dedicated to taking a call when broadband goes down. The call centre churn is a hundred percent, every eight months. No one wants this job. You need automated proactive communications in that situation. We know your broadband has a problem, or an imminent problem, so I will give you all the information about what's happening when it's happening, and keep you informed across all available channels until such time as the situation is resolved. And that reduces the anxiety of the customer and lets them know what's going on.”***



Mark K. Smith
CEO
ContactEngine

Mark is a serial entrepreneur who IPO'd his first business on the London Stock Exchange in his early 30s. He is credited with inventing online conferencing in the 1990s, built the first Content Management System for blind people in the 2000s, built 'Parasport' to help talent spot disabled athletes in the run-up to the London 2012 games, and invented a live streaming audio product that allowed commentary from anywhere in the world via phone. Mark is now CEO of ContactEngine, a conversational AI technology used by large corporates to automate customer communications. The company employs linguists, behavioural scientists, mathematicians and software engineers to design machine-learning algorithms that automate human-like conversations. The company began as an idea in Mark's head 10 years ago and is now a multi-£million company. Throughout his career, Mark has relentlessly applied science over instinct and believes technologies like AI can be a force for good.

No one, I think, can really object to machines replacing humans in a job that sees that kind of churn, and this is exactly the kind of interaction that a machine can handle better than a human. Not only that but, by handling it effectively, the machine is able to create a better emotional experience for the customer. This is a really crucial point—don't imagine that customer emotions can only be influenced by human-to-human contact. Proactive automated communication, like this example or even Amazon's simple delivery status notifications, can do a lot of work to reduce customer anxiety.

And if the automated interaction can't handle a particular customer's needs, or if they just want to speak to a human being, then there is always the option to escalate those cases to the call centre. Those cases which, almost by definition, will be more unusual, and more interesting for a human to handle.

“The call centre person would normally be a long-serving staff member, because they have a better job dealing with actual problems that humans deal with. That will normally be better paid as a consequence of that loyalty, and they will stay longer because we've got rid of all the crap that otherwise would have made them leave after six months.”

Coders & linguists

Effective Natural Language Understanding (the work to teach computers to understand human language as it is really used) happens at the intersection between linguistics and machine learning. ContactEngine employs a variety of specialists from different disciplines to work together at this point of intersection and, with one of their offices at Bletchley Park, Mark sees a parallel with the code-breaking teams assembled during WW2:

“There were four types of people that were employed at the time: there were men and women that were the equivalent of dev ops, they were programmers, there were people putting the tapes in the machines, so the equivalent of software engineers, there were mathematicians looking at the statistical patterns of data, and there were linguists. They are exactly the disciplines we employ now. What we do is a little less important than stopping a world war.... but it's intriguing that 75 years later, it's the same group of people, addressing very similar challenges.”

Getting machines to understand humans speaking or writing naturally is extremely difficult, and it's not something that you can expect mathematicians or programmers to solve on their own. These are problems that need to be solved with real world knowledge, and by testing the impact of approaches with real customers.

What we do is a little less important than stopping a world war... but it's intriguing that 75 years later, it's the same group of people, addressing very similar challenges.

“The language that you use in communication can massively affect response rates, and you can personalise that as well, based on additional information. The next generation of what we're doing we call human-computer rapport, which is a phrase we had to invent. You can market to individuals as individuals based on the patterns of what they do and using a concept of rapport means that you learn ways of communicating better over time, by building up an understanding of their communication needs.”

The future of customer-facing AI

The niche that ContactEngine has found is extremely revealing of an

opportunity that exists in a huge number of customer journeys across most consumer sectors and not a few business to business ones. Despite all the hype around AI and the potential for machine learning to improve the efficiency of many business processes, nowhere near enough attention has been paid to the potential that it offers to not just save costs, but also to improve customer journeys.

By focusing on proactive, outbound, communications (backed by smart conversational AI), rather than reactive enquiry handling, ContactEngine has built a very successful business which is demonstrably saving its clients money. More importantly, I think this is a great example of the way in which AI should be approached, not as an alternative to humans which is cheaper and “nearly as good”, but as an enhancement. In ContactEngine’s case, they’re adding conversation at a point in the journey which currently has either one-way communication or nothing at all.

Should you build AI into your journeys? This, for me, is the acid test: will it make the customer experience better?