


HuDeE

**HUMANITY
IN THE
MACHINE**



Contents

- **Why AI?**
- **The promise of chatbots**
- **Humanity in the machine**
- **How to build a chatbot**
- **Conclusion**



Why AI?

Artificial Intelligence – self-learning programmes solving problems – is the next major technology that will empower the world.

While today it is attracting huge investment from both VCs and the tech giants, it has a long history going back to the 1950s when Alan Turing developed the Turing Test to ascertain whether a machine could think.

After many false dawns, the underlying driver of Moore’s Law and advances in machine learning techniques mean the power of AI is now becoming manifest in the world around us – search, facial and voice recognition, drones, self-driving cars and robotics are all powered in one way or another through AI.

Google’s CEO Sundar Pichai now explicitly talks of an AI led world:

“Over time, the computer itself – whatever its form factor – will be an intelligent assistant helping you through your day. We will move from mobile first to an AI first world”.

Alongside the proponents like Pichai and Mark Zuckerberg, we have seen some notable voices of concern from Elon Musk, Bill Gates, Stephen Hawking and others about the dangers of runaway AI super

intelligence. Hawking has even said *‘the development of full artificial intelligence could spell the end of the human race’*.

AI is set to be the defining technology of this century and will impact our lives, whether for good or ill, in ways that we struggle to imagine from today’s vantage point. We will all need to get used to interacting directly with AI, whether in physical or virtual forms, and one of the earliest incarnations of this type of human to AI interaction is the chatbot – the subject of this report.

AI chatbots

AI chatbots are starting to emerge as a transformative way of interacting with businesses and brands. Chatbots can be simply defined as artificial intelligence programmes that conduct conversations with humans through chat interfaces.


They can exist within brands’ own digital platforms, on their websites or embedded within apps. But increasingly, bots are appearing where people are spending more and more of their time – in messaging apps. Already 60bn messages are sent every day within WhatsApp, Facebook Messenger and the like – 3 times the number of SMS.

As Wired magazine has put it *“there’s a new UI paradigm emerging, and it’s based around messaging.”*

We believe we’re at the start of the next great platform shift. We’ve moved from browsers to apps as our primary means of accessing information, and now we’re moving towards bots. The launch of *‘Bots for Messenger’* at f8 (Facebook Developer Conference) in April is testament to Facebook’s commitment to this view.

Brands need to start thinking about a world in which millions of independent AI agents are interacting with their customers and making decisions on behalf of their organisation.

The promise of chatbots

A man with a beard and long hair, wearing a white button-down shirt, is looking down at a smartphone in his hands. The entire image is overlaid with a semi-transparent green filter. The background is a plain, light-colored wall.

What do chatbots promise for brands?

The rise of bots presents a number of opportunities for customer – centric organisations in terms of both an enhanced customer experience and cost efficiencies.

Primarily, bots will let brands get closer to their customers by being present in their preferred environments (messaging apps) and interacting in what is increasingly their preferred mode (chat). Ultimately, this means the ability to give customers a better, faster experience – less time spent waiting for an operator to become available or navigating through call centre menus; problems being resolved more quickly.

Critically, it also leads to operational efficiencies. IBM Watson believe in most companies a relatively small number of queries occupy the majority of call centre operators' time. If bots can handle basic queries about billing, payments or deliveries, call centre operators can be freed up to handle the more complex queries adding greater value to the business overall.

What did we set out to find?

Technologists all too often focus on the technical capabilities at the expense of the end user's human (and therefore invariably irrational) needs – as Google Glass can bear witness...

The vision of a world in which people switch seamlessly from chats with their friends to chats with an AI bot about their grocery delivery is built on the premise that people will feel comfortable talking to machines.

But how do consumers really feel about interacting with brands through chatbots?

And how can brands build chatbots that enhance users' humanity and deliver a positive human experience?

Humanity in the Machine sets out to answer these questions. Through consumer led research this report provides a set of guidelines for brands developing their own customer facing bots, in preparation for an AI first world.



Humanity in the machine

To understand how AI chatbots can be built in a way that delivers a positive human experience, we designed a very hands-on research approach.

We needed to move from the theory and give consumers as much of a first hand experience of the potential of AI chat as is currently possible.


Working with our academic partners at Goldsmiths, University of London, we took the following approach:

Month long co-creation phase with 12 participants aged 18-35 including:

- Auto-ethnography: Participants logged their everyday lives particularly noting customer service experiences and moments of interaction with AI
- Experience sampling: Participants were given a range of hardware and software to help them imagine a future AI driven marketplace including Amazon Echo, your.md, Quartz and BotWorld
- Co-creation workshops: These were used to explore the participants' experiences and collectively develop rules of engagement for brand chatbot interactions

Customer Experience Trials

- Partnering with IBM Watson, we ran trials testing the impact of tone of voice on consumer satisfaction when carrying out a task with a banking AI (powered by Watson)
- Working with Extreme Biometrics, participants ran through a controlled experiment exploring changes in their stress levels when completing a task using AI (SwiftKey) versus a human



Alongside this field research, we conducted interviews with eight subject matter experts (SMEs) who provided us with valuable insight into the current AI chat landscape, and the future trends to follow.



Dr. Kathleen Richardson
Senior Research Fellow
in the Ethics of Robotics
at De Monfort University



Prof. Mark Bishop
Professor of Cognitive
Computing, Goldsmiths
University of London



Prof. Blake Morrison
Professor of Creative and
Life Writing, Goldsmiths
University of London



Duncan Anderson
CTO,
Watson Europe at IBM



Jonathan Crane
CCO
at IPSoft




Pete Trainor
Founding Partner & Director
of Human Centred Design,
at NexusCX



Matteo Berlucci
CEO,
Your.MD



David Roth
The Store WPP



63%

would consider messaging an online chatbot to get in touch with a business or brand

61%

agree 'It would be more frustrating if a chatbot couldn't solve my problem than a human' .

75%

agree that 'I'd prefer to know whether I'm chatting online with a chatbot or a human'

79%

agree that 'I'd need to know a human could step in if I asked to speak to someone'

48%


agree 'It feels creepy if the chatbot is pretending to be human'

60%


agree that 'it would feel patronising if a chatbot starting asking me how my day is going'

52%

agree that 'I would be happy to give quick feedback after a chatbot conversation' .

A top-down view of a person's hands typing on a laptop keyboard. The image is overlaid with a semi-transparent green filter and white geometric lines forming a grid pattern. The text "How to build a chatbot" is centered over the keyboard area in a large, white, sans-serif font. The person's left hand is on the keyboard, and their right hand is on the trackpad. A silver watch is visible on the left wrist. To the left of the laptop is a smartphone, and to the right is a black leather wallet. In the top right corner, a camera lens is partially visible.

How to build a chatbot



1. Define the role

Brands seeking to develop AI chatbots firstly need to go back to core marketing principles, and ask what need in the consumer journey the bot is trying to meet.

AI chatbots can offer a specific function (eg dealing with delivery queries, offering product suggestions) or can focus on being a general gateway offering 'navigational triage' - an initial place that people can go to from which they can access anything they want.


In the more sophisticated 'gateway' role, the bot directs people to where they can get what they need and manages the efficient allocation of resources behind the scenes. The bot can deal with the query or serve up a link, connect to an API, or to a chat response, and, along the line, to a human.

Ask yourself: are you building the entry point bot or the bot someone is triaged to from within the entry point?

The AI chat ecosystem is getting increasingly complex and includes things like general purpose AIs (Siri, Cortana, Viv), chat platforms and channels (Slack, Telegram, WhatsApp, WeChat, Kik, Line),

individual branded AIs (Luvo) AI licensing (Watson, Amelia), vertical AIs (Your.MD, Cleo, Alexa), RoboAdvisors (Penny, Digit), smart hardware (Echo, Jibo), Human AI augmenters (Facebook M staff), social chatbots (Xiaolce, Tay), and bot stores (BotList, KikHotBots, Facebook).

Working out a distribution strategy for a brand bot is critical. Sitting within the brand's owned platforms (website and in app) is likely to be a good starting point. But it is increasingly likely that the bot will sit 'on top of the stack' and act as an omni-channel brand concierge interacting with customers across multiple platforms, including the physical in store environment.



‘In a retail environment, it could be next to the mirror in the changing room. Or the watch on your wrist is your conversational interface, because it knows where you are, it knows you are in a shop, and therefore the conversation is about what you were doing last night online, and then it helps you find the product. I’m trying clothes on in a changing room, and you are basically saying to the AI that’s there, is this what I was looking at online last night? Where can I find that item that we were looking at last night? So I think where the screens appear in a retail environment, either on your wrist, or around it, create those links of continuity that create a very nice omni-channel experience. In a very un-minority report way.’

Pete Trainor

Founding Partner & Director of Human Centred Design,
at NexusCX



More and more though, brands should be extending their bots to where their customers are spending more of their time – messaging apps. Distributing within the messaging environment has distinct advantages:

Building conversational bots costs less and happens faster than building and maintaining cross-platform apps

There's less friction in the experience for users as they no longer need to "download and install" an app – they can just invite a bot to a conversation and interact with it as they would with a person

Critically, messaging apps know the user's identity and other contextual information (eg location and social data) which along with payment capabilities make 'conversational commerce' a strong opportunity



2. Focus on establishing trust

One of the key challenges any brand will face in building a bot is the issue of trust. It lies at the core of good customer service (our research found that 76% of people say trust is key to good customer service) and it has to be established before users can be expected to take part in deeper engagements.

We have found that people are often more trusting of bots around sensitive information than they are with human customer service operators. Only 37% say they are happier to give sensitive info over the phone to a human than to a chatbot. For 'embarrassing medical complaints', twice as many people prefer talking to a chatbot than a human than for 'standard medical complaints'.

We found that users consistently preferred to receive behavioural suggestions from an AI than from humans, and many of our participants at the Watson trial reflected this:

“I would prefer the interaction with the AI over the human for things like budgeting and managing my money. I think it would be quite intrusive if an actual person was trying to advise you on your spending, but if you could set an AI goals to manage your money better, and you could get advice and tips, that would be a really nice feature”.



‘It’s better if a computer knows about you than if a person knows about you’

This is largely because there is a feeling that the bot will never judge you in the way a human might. As Jonathan Crane, the CTO of Amelia puts it, “You are talking to a technology, providing information, and that technology isn’t judgemental. It can’t go further. It can’t misuse your information. And so we think there is going to be an increase in trust on the artificial intelligence”.

However, crucially, people are less forgiving of mistakes made by a bot than a human operator. In our field trials with Extreme Biometrics, participants ran through a controlled experiment exploring changes in their stress levels when completing a task using AI (SwiftKey) versus a human. We found consistently higher stress levels when people encountered mistakes made by the bot versus the human. This was echoed in our quantitative survey where we found 61% agree ‘It would be more frustrating if a chatbot couldn’t solve my problem than a human’.



So in their early iterations, brands need to make sure that their bot doesn’t get much wrong. As Duncan Anderson, the CTO of Watson, argues, “The essential part of establishing trust between the human and machine is having a machine which probably first of all does a good job understanding what the human is looking for, secondly gives a clear and consistent response.”

They should use specific tactics in how the bot delivers responses to help build confidence. This could be adding pauses before delivering a response which gives the illusion of the bot ‘thinking’. It could also be asking additional questions – Your.MD requires three questions to achieve its optimal probability score but have found that users are more comfortable with the responses if the bot asks five or six questions. As Matteo Berlucci their CEO explains, “You don’t do it because you’re trying to pass as human, you do it because you want to make it easier for the user to interact with the service.”

It is also important to be clear when the user is chatting to a ‘machine’ and when they are chatting to a human, as the uncertainty undermines trust. 75% agree that ‘I’d prefer to know whether I’m chatting online with a chatbot or a human’.

One of the key considerations in building trust is identifying the right conditions for handing off to a human operator. 79% agree that ‘I’d need to know a human could step in if I asked to speak to someone’.

When implementing Amelia in organisations, IPSoft typically toggles the ‘confidence’ the AI needs to have to directly address a customer question. If the AI is not sufficiently confident it passes the question on to another resource in the organisation like a human call centre worker. The AI can then follow the outcome of the interaction and learn how to handle it more directly the next time. It is typically better to err on the side of caution and offer frequent hand offs in the early stages of implementation.



3. Align tone of voice with category & brand

In building a bot, one of the critical considerations is what tone of voice to give it, and in particular how 'human' it should be.

Working with IBM Watson we set out to test this issue, by building two alternative banking bots with very different personalities: one was chatty, informal and conversational; the other was more straightforward, with a serious and functional tone of voice. We asked 20 respondents to carry out a simple banking task of ordering foreign currency with both bots.

While some people enjoyed the chatty version, there was a small majority in favour of the more straightforward version. Many found the chattier version unnecessarily off-putting, patronising or even creepy. As one respondent put it 'The chatty one is like my dad when he uses emoticons, it's creepy'.

In our survey we found that 47% agree with the statement that 'chatbots can never seem like humans. So they shouldn't even try' and 48% agree 'It feels creepy if the chatbot is pretending to be human'.

As AI becomes more sophisticated, building a bot with ever more 'human' characteristics will become increasingly appealing to technologists. But, in our view, brands shouldn't be trying to make the bot sound as lifelike as possible.

They should instead focus on delivering the customer need in a simple, straightforward manner and using a tone of voice which is flexible, contextual and personalised to different users and different situations. This will also mean using copywriters alongside programmers to create consistent style and tone.


47%
agree with the statement that 'chatbots can never seem like humans. So they shouldn't even try'

48%
agree 'It feels creepy if the chatbot is pretending to be human'



“There certainly seems to be a group of people who are very happy with a machine that is just a machine, that doesn’t feel like a human, is not trying to pretend to be a human. I think some of our implementations, where the machine isn’t trying to be a human, sometimes they can come off slightly robotic, but it seems that people don’t get offended because they know it’s a machine, and their expectations are that it may be slightly robotic. It doesn’t surprise them at all, they seem to feel comfortable with that idea.”

Duncan Anderson
CTO Watson



4. Make people feel human

Much of the research into AI focuses on what it means to be human, as this invariably informs the design and development of AIs. Throughout our research we too asked this question, but we found a different answer.

What defines a 'human' AI does not depend on how human the AI appears to be, or how life-like its interactions are. What defines a human experience is the experience itself - it is measured by how a person feels when dealing with the AI, and not some intrinsic humanity in the technology.

This may or may not mean engaging with an AI that has an avatar, name, gender, or politics but what makes an experience human, is its ability to make people feel human.

A 'human' experience is defined by how the user feels, not how 'life-like' the bot is.

Bots should aim to use context and emotional understanding (empathy) to deliver a 'human' experience by meeting the user need. In doing this, the style of the bot should ideally go unnoticed. If it feels too 'robotic', then interacting with it

leaves the user feeling dehumanised. If it's too 'life-like' then the user can be left feeling patronised or even disturbed.

60% agree that 'it would feel patronising if a chatbot started asking me how my day is going'.

As Prof Blake Morrison puts it "I know I'm not speaking to a real human being, and I don't want this. It's sort of a patronising thing, you know, 'I'm gonna speak like you, I'm gonna try and be chummy with you'. You don't want too much of that, but you don't want to feel like you're speaking to a robot. Even though you are, you want something a bit more."

The challenge is to get the balance right and leave the user feeling as though they have had a human experience. And, crucially, to avoid falling into the 'uncanny valley'* through creating a bot that feels creepy by attempting to emulate humanity.

[FOOTNOTE* the term coined by Japanese robotics Professor Masahiro Mori to describe the feeling of revulsion that people get when a robot becomes so realistic that it comes close to resembling a human but feels like a zombie or a corpse.]



5. Build for continuous improvement

As the AI learns from people, it develops its understanding of people's needs and wants. Building this bank of knowledge will enable brands to provide more efficient and engaging interactions with consumers. This necessitates a 'feedback loop' where people help inform the AI for the future.

The challenge here is to create easy feedback mechanisms to aid machine learning, but to do so in a way that isn't irritating for users. 52% agree that 'I would be happy to give quick feedback after a chatbot conversation'. Importantly, it doesn't need to be a huge proportion of users.

Matteo Berlucchi, the CEO of Your.MD explains,

“You don't need everyone to give you feedback, you just need a reasonable percentage. The interesting thing with the bots is that there is a learning curve, and it plateaus quite quickly, so once it does that it doesn't need as much customer feedback.”


Berlucchi finds it analogous to the 1-9-90 principle of content interaction on the Internet where 1% of users create content, 9% comment on or help evolve the content, and 90% of users consume the content as their only form of interaction.

It's also important to nurture community activists who will help develop bot capabilities proactively, in the same way that volunteers edit and moderate Wikipedia and the like. Your.MD has up to a hundred UK GPs who have volunteered to help adjust and improve the service.

But feedback needs to be managed carefully as unfortunately there are plenty of people online who will be more than happy to sabotage your bot as Microsoft found out recently with their social bot Tay which was taught racist obscenities within the space of 24 hours.

52%

agree that 'I would be happy to give quick feedback after a chatbot conversation'



Conclusion

AI is set to be the defining technology trend of this century. With its advocates arguing it could help solve humanity's greatest issues such as cancer or poverty, and its critics arguing it could end humanity, it's clear that AI's impact will be utterly profound. While the visions of both its proponents and detractors are many years and decades away, we will all need to start getting used to interacting directly with AI, whether in virtual or physical form.

All businesses need to interact with their customers directly, and the chatbot is the first manifestation of AI powering that interaction. Understanding how that can be done in a way that not just suits the business through efficiency savings but also offers the customer a faster, better, more satisfying experience is critical. And the key to that is being respectful of the end user and delivering a positive interaction that respects their humanity.

Interacting with a self-learning machine has the potential to be a dehumanising experience, and no doubt many mistakes will be made by businesses along the way. But the brands that start experimenting first will be the ones that can learn fastest and reap the rewards from the upcoming AI revolution.

At Mindshare, we're keen to be part of that learning journey and we will be offering consultancy services to help brands start designing interactions with chatbots in a way that puts the Humanity in the Machine.

A special thank you to our contributors

Mindshare UK:

Jeremy Pounder - Futures Director

Rob McFaul - Business Director

Sam Barton - Account Manager

Goldsmiths:

Dr Chris Brauer - Research Director

Dr Yael Gerson - Research Lead

Daniela Rumlova - Research Assistant

Marton Gaspar - Research Assistant

Cleary Ahern - Research Assistant

Ricardo Leizaola - Multimedia

Huddle is Mindshare's festival of thinking and ideas. Every November since 2011, Mindshare has gathered media, cultural and tech players, entrepreneurs, academics and interested amateurs for one day of intimate and no-holds-barred discussions about the future of media.

In 2016, we've brought all of Mindshare UK's work on inventing the future of media research under the Huddle umbrella. Huddle will still be the chaotic, vibrant event our clients and partners love, but it is also now our R&D facility, our lab and our incubator of new ideas.

For more information, or to enquire about how we can help your organisation explore the potential of Artificial Intelligence, please contact:

Jeremy Pounder

Futures Director, Business Planning
E: jeremy.pounder@mindshareworld.com

Louise Richardson

Director of Marketing
E: louise.richardson@mindshareworld.com

Central Saint Giles, 1 Saint Giles High Street
London WC2H 8AR

www.mindshare.co.uk

Follow us on:

[@mindshare_uk](https://twitter.com/mindshare_uk) [#mshuddle](https://twitter.com/mshuddle)