CONTACTBABEL

The Inner Circle Guide to Chatbots, Voicebots & Conversational AI

Sponsored by

kerv

The Inner Circle Guide to Chatbots, Voicebots & Conversational AI (UK edition)

kerv | GENESYS Gold Partner

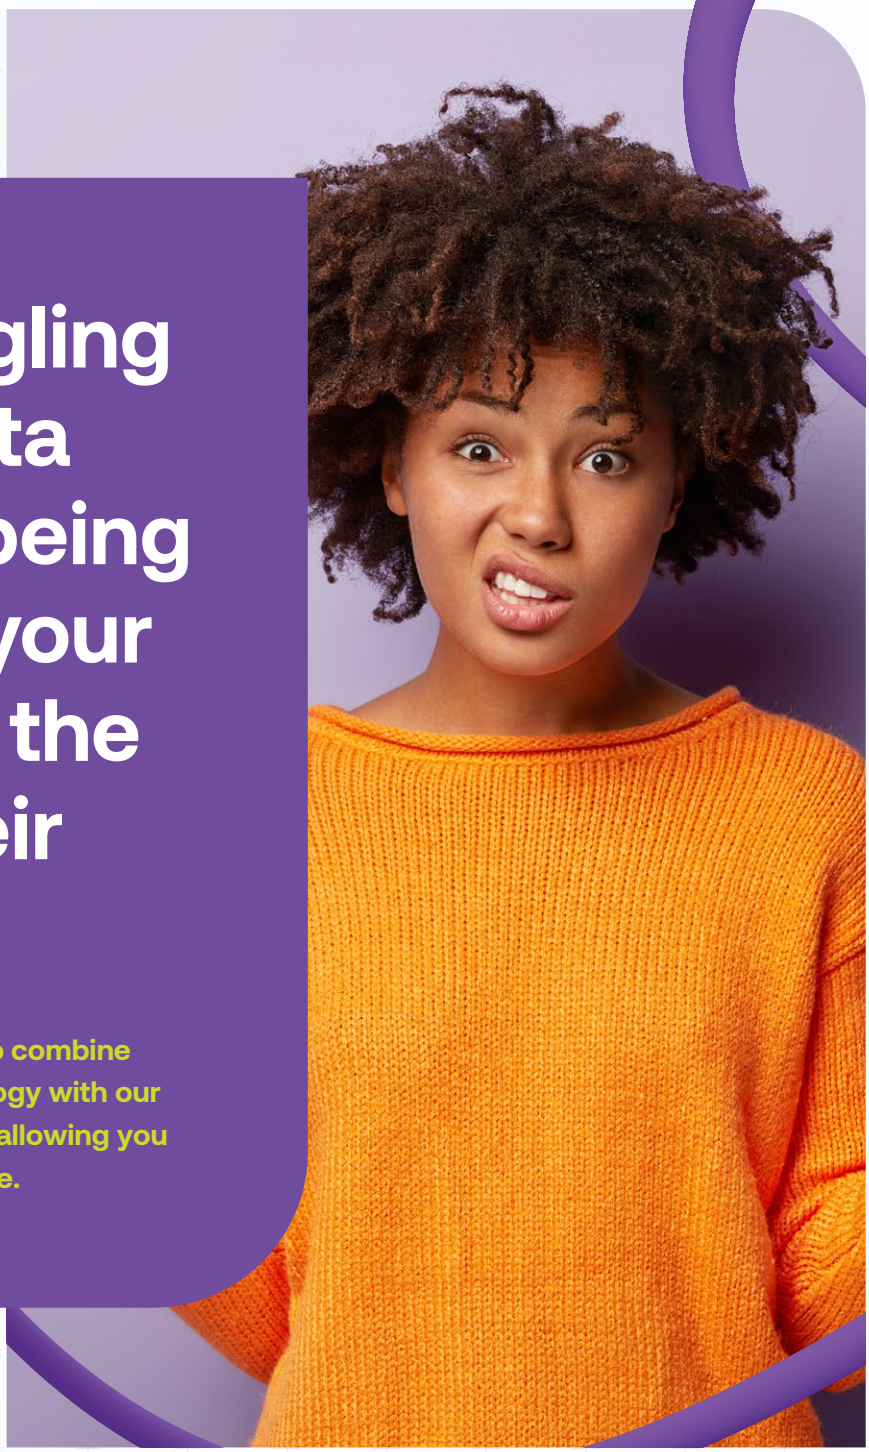

Are you struggling with siloed data systems and being able to serve your customers on the channel of their choice?

Kerv Experience partners with Genesys to combine industry-leading contact centre technology with our own applications, services and support, allowing you to deliver the CX your customers deserve.

SKIPTON BUILDING SOCIETY Since 1853

ZSL LONDON ZOO

Cyngor Sir CEREDIGION County Council

arvato BERTELSMANN

PUREGYM

GOOD SHAPE

mi hub

Want to optimise your CX strategy and drive better customer and agent experiences?

Contact us today for a demo or free trial

For more information please contact

sales.experience@kerv.com 0330 403 0000

kerv

CONTENTS

TABLE OF FIGURES

At Kerv, we leverage the power of technology to help our customers stay ahead, providing end-to-end cloud solutions and digital transformation.

Our specialist CX practice, Kerv Experience, are an established Genesys gold partner with 23 years' experience in the contact centre and CX space. Kerv Experience's deep expertise of Genesys Cloud helps you completely personalise your digital customer experience, improve customer insights, reduce effort with seamless integrations and automation and to deliver amazing service when it matters most.

As the 2021 Genesys EMEA Customer Success Partner of the year, we help deliver exceptional CX to our customers through the deployment of cloud, digital, data and AI technologies. We help our customers execute on their CX strategy by deploying the correct technology and ensuring they get the most out of the solution over time. Our clients include PureGym, HelloFresh, Heineken, Somerset County Council and P&O Ferries.

Our sister practice, Kerv Digital, work hand in hand with Microsoft to deliver bespoke services for Dynamics 365, Power BI and Power Platforms which integrates with Genesys Cloud to create a seamless customer and employee experience. Cloud Contact Centre as a Services allows you to access constant technology innovation and Kerv are the partner to help you maximise the benefits of that investment.

Contact:

w: www.kerv.com

t: 0330 403 0000

e: jas.bansal@kerv.com - Jas Bansal, Head of Marketing

e: sales.experience@kerv.com

LinkedIn

ABOUT THE INNER CIRCLE GUIDES

The Inner Circle Guides are a series of analyst reports investigating key business issues and the customer contact solutions that can help, along with various use cases, the reality of implementing and using these technologies and a view on what the future holds.

The Inner Circle Guides are free of charge to readers. Sponsors have **not** had influence over editorial content or analyst opinion, and readers can be assured of objectivity throughout. Any vendor views are clearly marked as such within the report.

There are Inner Circle Guides to:

- Agent Engagement & Empowerment
- AI-Enabled Agent Assistance
- Chatbots, Voicebots & Conversational AI
- Cloud-based Contact Centre Solutions
- Customer Engagement & Personalisation
- Customer Interaction Analytics
- First-Contact Resolution
- Fraud Reduction & PCI Compliance
- Next-Generation Technology
- Omnichannel
- Omnichannel Workforce Optimisation
- Outbound & Call Blending
- Remote & Hybrid Working Contact Centre Solutions
- Self-Service
- Voice of the Customer.

These can be downloaded free of charge from https://www.contactbabel.com/research.

As well as explaining these solutions to the readers, we have also asked the potential users of these solutions whether they have any questions or comments to put directly to the report's sponsor, and we have selected some of the most popular to ask. These branded Q&A elements are distributed throughout the report and give interesting insight into real-life issues.

The statistics within this report refer to the UK industry, unless stated otherwise. There is a version of this report available for download from https://www.contactbabel.com/research with equivalent US statistics and findings.

"Small" contact centres are defined in the report as having 50 or fewer agent positions; "Medium" 51-200 agent positions; and "Large" 200+ agent positions.

AI in the customer contact world is perhaps currently best known for chatbots. Conversational AI tools such as natural language processing (NLP), machine learning and analytics run automated tasks and simulate conversation with the customers.

The chatbot program using conversational AI may be given a human avatar and personality characteristics, and can include natural language processing, dialogue control, access to knowledge bases and a visual appearance that can change depending on who it is talking to and the subject of the conversation.

Chatbots – as the name suggests – are found in the web chat channel, but the functionality can be used in any other digital channel, such as social media, email and especially voice self-service in the form of voicebots.

However, chatbots do not always use conversational AI, and rule-based basic chat applications can be used in circumstances where there are a significant proportion of interactions about similar issues, where the chatbot can be tasked to answer these and pass more complex issues to a human agent.

Even without AI, rule-based chatbots can have a positive impact on customer care without requiring large investments or long implementation timescales. Chatbots using conversational AI use natural language understanding (NLU) and can ask questions to understand customer intent and improve the accuracy of the output, and may also use machine learning and analytics to predict and improve future outcomes.

A voicebot is an application made up from AI and natural language understanding (NLU). Voicebots convert speech to text, analyse it and respond appropriately using text-to-speech. It is integrated with CRM or a knowledge base in order to provide a greater accuracy and depth of response. It should be noted that a common use of speech recognition, such as keyword spotting in order to route a call, is not the same as a voicebot.

Like any machine learning application, voicebots require training, as well as volumes of clean data from which to learn. Voicebots are used to deliver full self-service experiences without requiring an agent, but are also increasingly used for customer identity verification and also call routing.

AI DEFINED

Understanding artificial intelligence is a complex matter. Definitions and terminology can differ slightly (or sometimes greatly) depending on who is describing functionality and the technology behind it.

Within the boundaries of this report, AI will be used as an umbrella term for solutions which appear to emulate human cognition through the 'understanding' of complex, natural language requirements, in order to reach its own conclusions, learn and thus improve itself.

Within the wider customer contact environment, AI involves technologies such as machine learning, speech-to-text, deep learning, analytics, chatbots/voicebots and natural language understanding, all closely integrated and working together, aiming to provide outcomes similar or even superior to those achievable by human agents.

Some of the typical characteristics of AI-enabled solutions include:

- An understanding of the customer's meaning and intent, rather than just accurately decoding the syntax of the request
- Use of multiple questions in a conversational format to improve understanding
- Using past outcomes to predict and deliver the likeliest most successful output
- The use of confidence levels rather than a binary right/wrong output
- The ability to learn and improve without constant human support
- The ability to improve future outcomes without constant human input or monitoring.

Here are descriptions and definitions that we will use throughout the report, although do bear in mind that other commentators may have different views.

CONVERSATIONAL AI

Conversational AI refers to software applications such as advanced chatbots (also known as virtual agents) and voicebots (which can be referred to as voice assistants). Customers communicate with these applications using natural language rather than simply keyword recognition, and the bots reply using context and in a manner that a human would.

Through being trained on large amounts of data and using machine learning and natural language processing (NLP), they are can recognise user speech and text along with context and intent. A key part of conversational AI is the ability to improve itself over time through a constant feedback loop which improve the AI's algorithms.

TYPES OF CHATBOTS / VOICEBOTS

Traditional, **rule-based chatbots** are trained to answer only a pre-programmed, specific set of questions, and while they can be extremely useful, are not seen as part of the conversational AI world. They cannot handle queries outside of their programmed scope, and do not use machine learning to improve their responses. Any changes to their output or flow needs to be done manually.

Conversational AI chatbots/voicebots use natural language processing – which includes speech to text transcription, allowing the voice channel to be automated – and have an understanding of the customer's context and intent. Machine learning allows them to improve responses over time. They can handle multiple questions in a single interaction and be deployed across numerous channels, including voice, email, social media and web chat.

MACHINE LEARNING / DEEP LEARNING / NEURAL NETWORKS

Through the use of pattern recognition, previous outcomes and other algorithms, machine learning enables systems to improve themselves without the need for continuous human user input (although supervision and guidance is often needed in reality). It relies upon extensive datasets and computational power in order to make predictions with continually improving levels of confidence.

Based on the workings of the human brain, neural networks consist of input and output layers as well as one or multiple hidden layers (Deep Learning uses multiple layers, each carrying out their own specific task), working to find patterns which will be too onerous or complex for humans to identify. Neural networks can be trained to spot patterns in data and provide accurate output, with programmers correcting any mistakes. Eventually the neural network can 'understand' whether it is producing accurate output with far less human correction.

Neural networks can be set up using supervised or unsupervised learning techniques. Supervised learning techniques involve giving the neural network a specific problem such as "is this customer likely to complain?". Programmers then provide the system with large datasets of customers who have or have not complained, and then the neural network will find patterns of characteristics that make some customers more prone to complaint. They are then able to predict which customers are likely to be dissatisfied, allowing the business to act accordingly. In the case of unsupervised learning, no specific output is given to the system, which will then find patterns in the data and classify groups accordingly. Supervised learning is by far the more common use of AI in businesses.

NATURAL LANGUAGE PROCESSING / UNDERSTANDING (NLP/NLU)

NLP refers to the branch of AI which enables computers to understand human language, whether spoken or written. It goes beyond speech to text processing – although of course accurate transcription is vital – and attempts to understand the actual intent and sentiment of the customer.

NLU is a subset of NLP which looks at the challenges of understanding human communication, such as mispronunciation, sub-optimal word order, slang and other elements which are a natural part of human speech but which can cause major problems for computers due to their unstructured and outlying nature.

One of the keys to successful automated service, whether via telephony or website, is for the user to be able to describe their issue in their own words, rather than feeling that they have to use specific terms or a stilted, incomplete account of the issue. Natural language processing-based systems encourage users to describe their issue more fully, asking follow-up questions if there is any degree of ambiguity in the initial request. One of the obstacles to overcome for NLP-based systems (whether through speech recognition or text recognition) is that many Internet users have been trained to use keywords, believing that simplifying the description of their issue will lead to greater levels of accurate response. In fact, NLP works best with longer and more detailed requests, and it is a challenge for businesses and solution providers to encourage and support users of the system in using the solution in an optimal way.

GENERATIVE AI

Generative AI refers to a category of AI algorithms / models that create **new** content based on the datasets that they have been provided, using deep learning techniques and neural networks to create similar types of content. Chatbots built on top of a large language model (LLM - a machine learning application) can provide answers to users' questions in a detailed and realistic manner. ChatGPT is an example of this.

It is important to consider which LLM(s) to train your AI tools on, focusing on topics most relevant to your organisation's products and services: there have been high-profile examples of generative AI going off-topic with embarrassing, incorrect and sometimes offensive results[1] [2]. It is equally important to decide which internal or external sources of information and knowledge the AI will have access to: solutions may require extensive mapping of subject categories to knowledge bases, whereas others may be able to learn to connect with the correct knowledge base after less training.

Generative AI is capable of understanding multiple languages, has a detailed knowledge of the information it has been trained upon, can carry out a certain amount of reasoning and uses language in a human-like way, including sentiment analysis.

Some experts predict that generative conversational AI – a combination of approaches – will emerge in the near future, blending the natural communication style of generative AI with the accuracy of response provided by conversational AI. The key to this is providing the right type of data and examples of successful interactions upon which to train the AI, providing it with only the relevant data specific to the work it is doing.

END-USER QUESTION #1: ARE CHATBOTS / VOICEBOTS BETTER SUITED TO BUSINESSES WITH VERY HIGH VOLUMES OF CUSTOMER CONTACT?

Businesses with frequent customer interactions stand to gain significant advantages from chatbots. These digital assistants can efficiently handle repetitive inquiries, freeing up human agents for more complex issues. Their 24/7 availability ensures constant support for high-volume customer situations. Additionally, chatbots can streamline operations by automating tasks and offering self-service options, potentially leading to cost savings.

However, the benefits of chatbots extend beyond high-contact businesses. They can act as "always-on" representatives, qualifying leads, answering general inquiries, and even recommending products. Furthermore, chatbots can personalize communication by tailoring interactions based on user information, even for companies with moderate customer interaction. Ultimately, chatbots can enhance customer experience by providing efficient and immediate support, regardless of the volume of contact.

[1] https://www.forbes.com/sites/marisagarcia/2024/02/19/what-air-canada-lost-in-remarkable-lying-ai-chatbot-case/?sh=65f98ea0696f
[2] https://www.bbc.co.uk/news/technology-68025677

VIEWS ON THE EFFECT OF AI ON THE CONTACT CENTRE

Survey respondents from over 200 UK companies generally do not believe that AI will replace agents: only 17% agree to some extent that this will be the case, with 56% disagreeing. It is worth noting that after a growing feeling five years ago that AI will replace agents, recent years' views are very much of the opinion that they will not.

Unanimity was found when the question was asked as to whether AI would support human agents, with all respondents agreeing or strongly agreeing that this would be the case, reducing risk, speeding up responses and providing customers with higher quality resolutions.

52% strongly disagree that AI will be irrelevant to their contact centre, with almost unanimous agreement that AI will affect contact centres of all sizes. This figure is growing year on year as AI becomes more widespread and the benefits better understood.

Figure 1: Views on the role of artificial intelligence in the contact centre

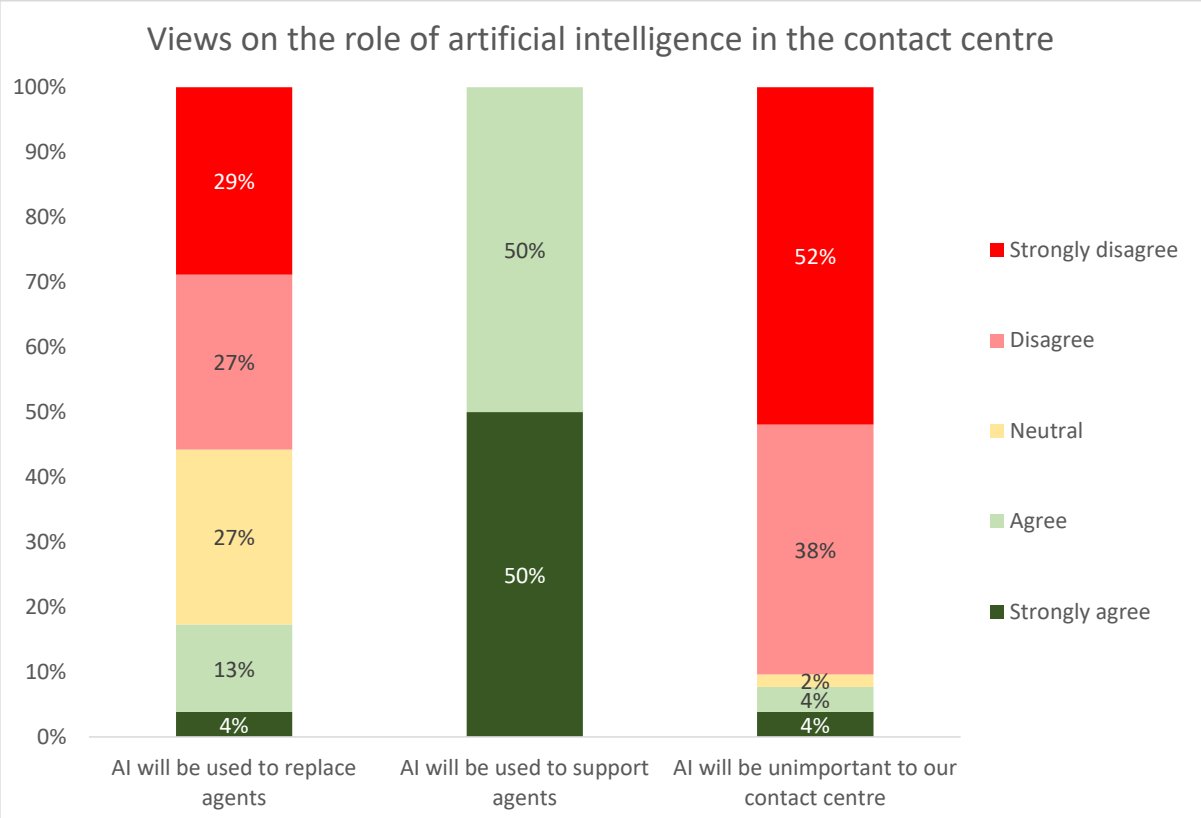

There was a unanimous belief that customers would not have a problem with AI if it helps them to resolve their issue as quickly and easily as possible. The uptake in web self-service and automated digital service suggests that customers will accept non-human assistance if it is most convenient for them, which is a positive sign for chatbots and voicebots.

There was also agreement that older generations will take a lot more persuasion to be happy with AI compared to a younger generation that is already used to dealing with AI in their everyday life (e.g. through smartphones or other virtual assistants in the home).

There was also a widespread feeling that AI should not need to be hidden from customers. To some extent, this finding goes somewhat against one of the main ideas about chatbots: that they should communicate naturally with customers, and can even have their own style and personality.

Respondents disagree to a lesser extent than usual about whether customers will always prefer human interactions: far more believe that customers will always prefer human interactions, with only 14% feeling differently about this. Our research has found that customers would prefer to speak with humans rather than self-serve, but this is only the case where the time required and outcome received is exactly the same. Where self-service is quicker, many customers choose this first.

Figure 2: Views on how customers will perceive artificial intelligence in the contact centre

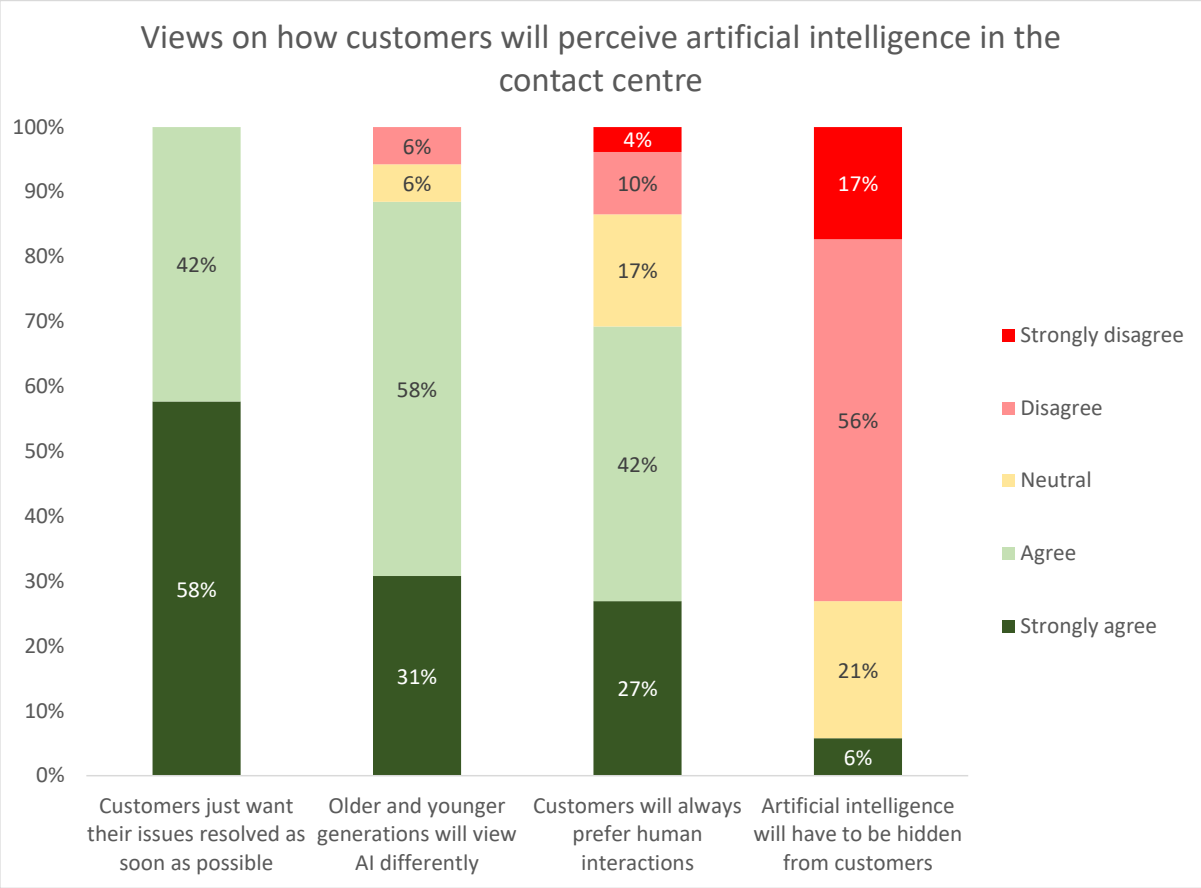

This report focuses on the use of AI to support self-service. There is another report – "The Inner Circle Guide to AI-Enabled Agent Assistance" – which looks at agent-related use cases for AI.

As such, at the most basic level, the business drivers for chatbots, voicebots and conversational AI are around trying to make self-service a better experience for customers while reducing the costs of unnecessary calls.

Specifically, the main business issues that these solutions can help with are:

- Excessive customer queue / wait times
- The uptake and effectiveness of web self-service is lower than desired
- Voice self-service is rigid and difficult to use
- Customer authentication costs too much.

CASE STUDY
PureGym

PureGym solves peak demand challenges while significantly improving customer satisfaction and the agent experience

Founded in 2008, PureGym Group holds market leading positions across the UK, Denmark, and Switzerland with recent entries into the US, Dubai and Saudi Arabia.

The rapid business growth came with a penalty. The company's contact centre in Leeds was predominately call-based and soon became stretched. In addition, agents were slowed down and couldn't perform at their best because of different systems.

Following an extensive tender, the company chose a proposal from Kerv Experience and Genesys and made the decision to make their contact centre 100% digital. By taking full advantage of self-service, deflection and automation tools, PureGym can now perfectly handle the seasonal peak demands. Agents feel empowered and engaged, as opposed to frustrated and demotivated by legacy systems.

> **"This January we had a 72% increase in demand on the previous month, which was 16% higher than the previous January. Despite this huge demand spike, we managed to grow our member base by 10% while reducing agent hours by 15%."**

Michelle Kaye
Head of Member Services, PureGym

Results

4% decrease
in contacts requiring agent assistance

~300% increase
in call deflection and auto handling

15% reduction
in the number of advisor hours

18% increase
in contacts handled per agent per hour

BUSINESS ISSUE #1: EXCESSIVE CUSTOMER WAIT TIMES

The chart below shows the importance of various customer experience factors, based on a survey of 1,000 UK customers. Aggregating the results allows an understanding of which factors were placed in the top three overall, while also providing insight on age-related opinion. Figures below are expressed as the percentage of each age group that expressed an opinion.

Having a short queue time is seen as the most important factor driving customer experience, with respondents of almost every age group rating this extremely highly.

Figure 3: What are the top 3 most important factors to you when contacting an organisation by phone or digital channel? (by age range)

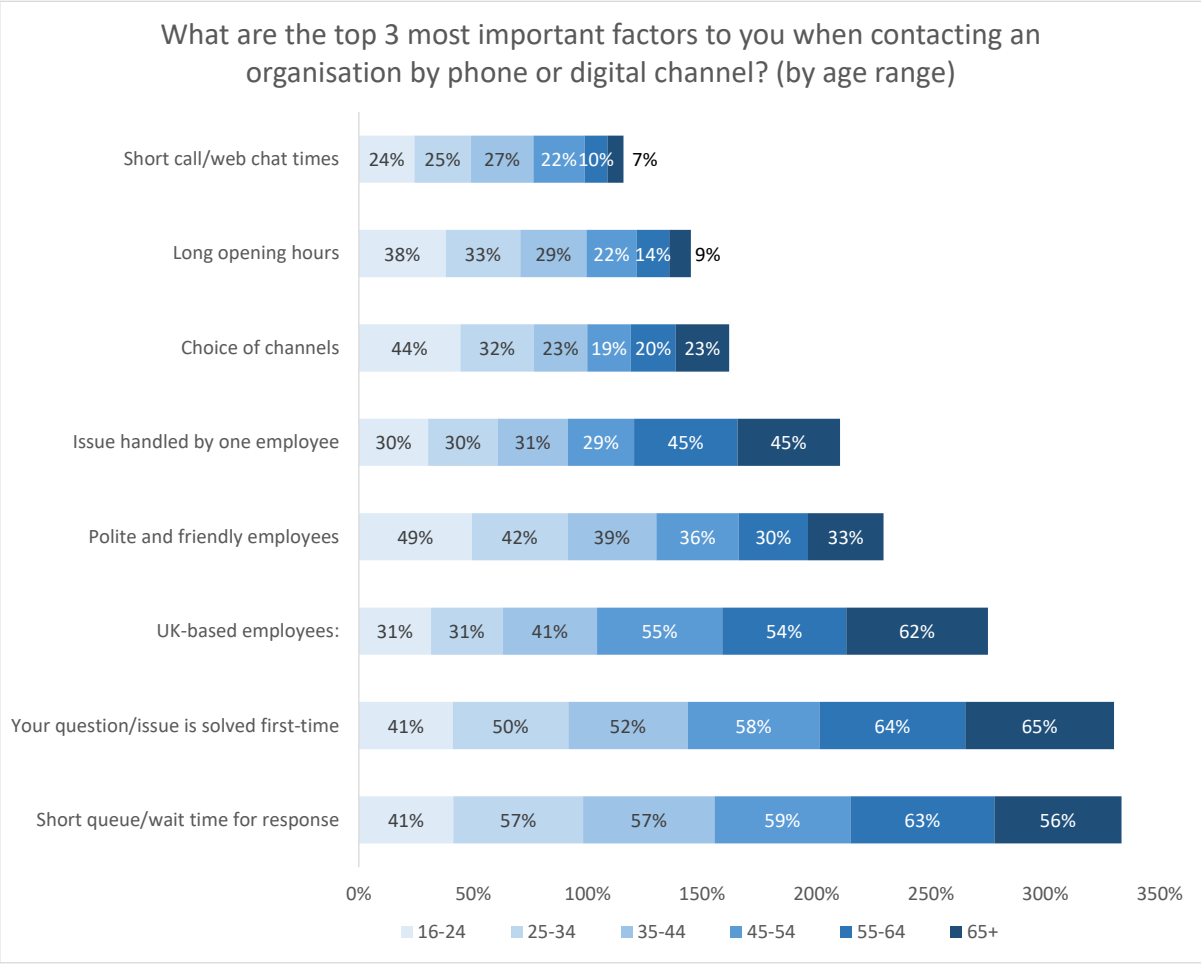

Yet the following charts show that queue time is at an unacceptably high level for most customers: in fact, the most widely reported customer issue is that the queue time / speed to answer is too long.

Looking at segmentations by age group, older customers are more likely to report long queue times happening very often, with almost half of over-55-year-olds stating this to be the case.

Figure 4: Customer Issue: "The queue time is too long" (by age range)

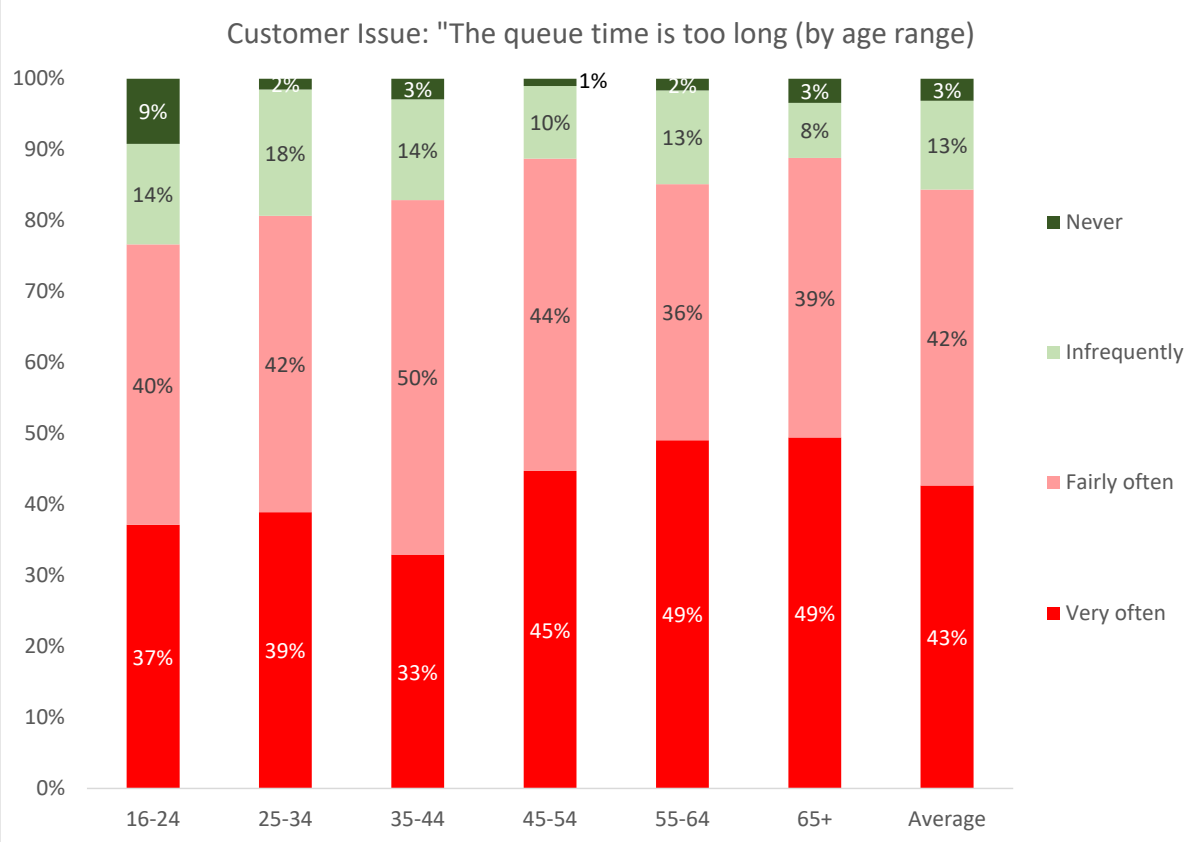

Customers are not imagining that the average speed to answer is excessive: the industry statistics bear this out.

The following chart shows historical figures for average speed to answer and call abandonment rate.

From 2009 onwards there has been a gradual increase in average speed to answer, with a spike to 36 seconds in 2014 which remained quite steady until 2018's jump to 41 seconds. There has been a huge uptick in 2020 and 2021 caused by pandemic-related working practices and an increase in demand experienced by some businesses, which is even stronger in 2022 despite a return to some sort of normality.

Call abandonment rate did not historically show any upward trend, with all but three of the 16 survey results showing data between 5% and 6% between 2004 and 2019. 2020's abandonment rate was up by 20% and is almost certainly linked to the changes in speed to answer, with 2022 the highest yet recorded and little sign of improvement being shown since.

This is clear proof that customers are not overexaggerating their frequent experiences with long queue times. Some sectors – utilities, public sector and housing in particular – are reporting extremely long queues, and this needs to be looked at urgently.

Figure 5: Historical average speed to answer & call abandonment rate, 2004 - 2023

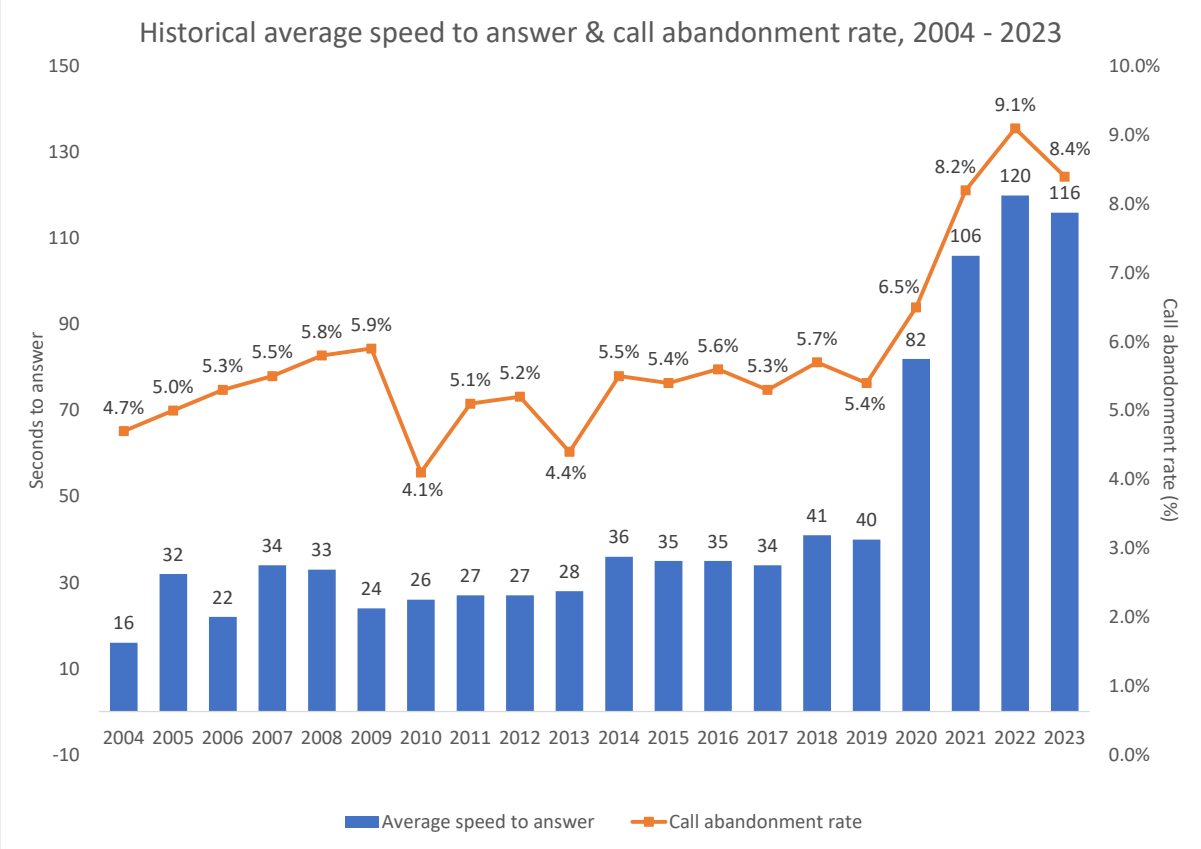

Excessive queue times are without doubt a major driver of negative customer experience, and the UK contact centre industry as a whole has lost its grip on this key benchmark.

Chatbots and voicebots can alleviate this through encouraging more customers to use self-service – whether through digital or voice channels – providing 24/7 service. Adding AI to the self-service mix means that self-service becomes more powerful, approachable and flexible, solving more of the issues that would otherwise have to be dealt with by a live agent.

Rules-based chatbots have done a fairly effective job of answering the most common FAQs, but there is a massive opportunity for more complex issues to be solved through conversational AI, which allows the understanding of context, slang and nuance, allowing customers to ask questions in their own way, rather than having to guess the right keywords.

The following two business issues look at web self-service and voice self-service to show how AI-enabled chatbots and voicebots can improve the take-up and effectiveness of automated service, which will impact positively on queue lengths, call abandonment rates and other CX-related metrics, as well as reducing live call costs.

BUSINESS ISSUE #2: UPTAKE AND EFFECTIVENESS OF WEB SELF-SERVICE IS LOWER THAN DESIRED

Many current self-service systems are inflexible and structured rigidly in their information flow, handling simple, unambiguous service requests by customers (such as account balances). Generally speaking, these are very successful at delivering this information, and customers will often choose a familiar and effective method of handling the simplest enquiries.

However, historical interaction volume information shows that the number of live calls received by contact centre remains steady: although the contact centre is the primary channel choice for only around one-third of customers, almost two-thirds of interactions with the business still come via live telephony. This suggests that the various methods of using self-service and the supporting knowledge base still have a very long way to go before customers rate them as highly for effectiveness and timeliness as they do the traditional contact centre.

New channels such as social media, email and web chat have grown rapidly in popularity, yet the large majority of interactions involving the former two channels still require a customer to make a request to a human agent. The use of chatbots has increased hugely in the past few years as shown later in this report.

While web chat, social media interactions and emails have somewhat lower costs than telephone calls, the differential between these is far smaller than between live and automated service, offering cost incentives to move interactions that do not have to be handled by live agents onto an automated channel.

Figure 6: Cost per inbound interaction

Channel	Mean	1st quartile	Median	3rd quartile
Phone	£5.58	£5.77	£4.18	£2.58
Email	£3.55	£5.27	£3.00	£1.52
Web chat	£3.05	£4.00	£2.00	£0.87
Social media	£2.25	£3.00	£2.25	£1.33
IVR	c. 30-50p	-	-	-
Web self-service	c. 5-10p	-	-	-

However, despite the widely acknowledged desire to move live interactions to automated channels, web self-service is not as successful as it could be.

Although around a third of business survey respondents state that fewer than 10% of their customers have tried to resolve issues online before calling the contact centre, a similar proportion state that more than 1 in 4 of their inbound calls come from people who have failed to complete their objective on the website first, and who may approach the call in a state of frustration.

Figure 7: Proportion of callers that have tried to answer own queries through web self-service before calling

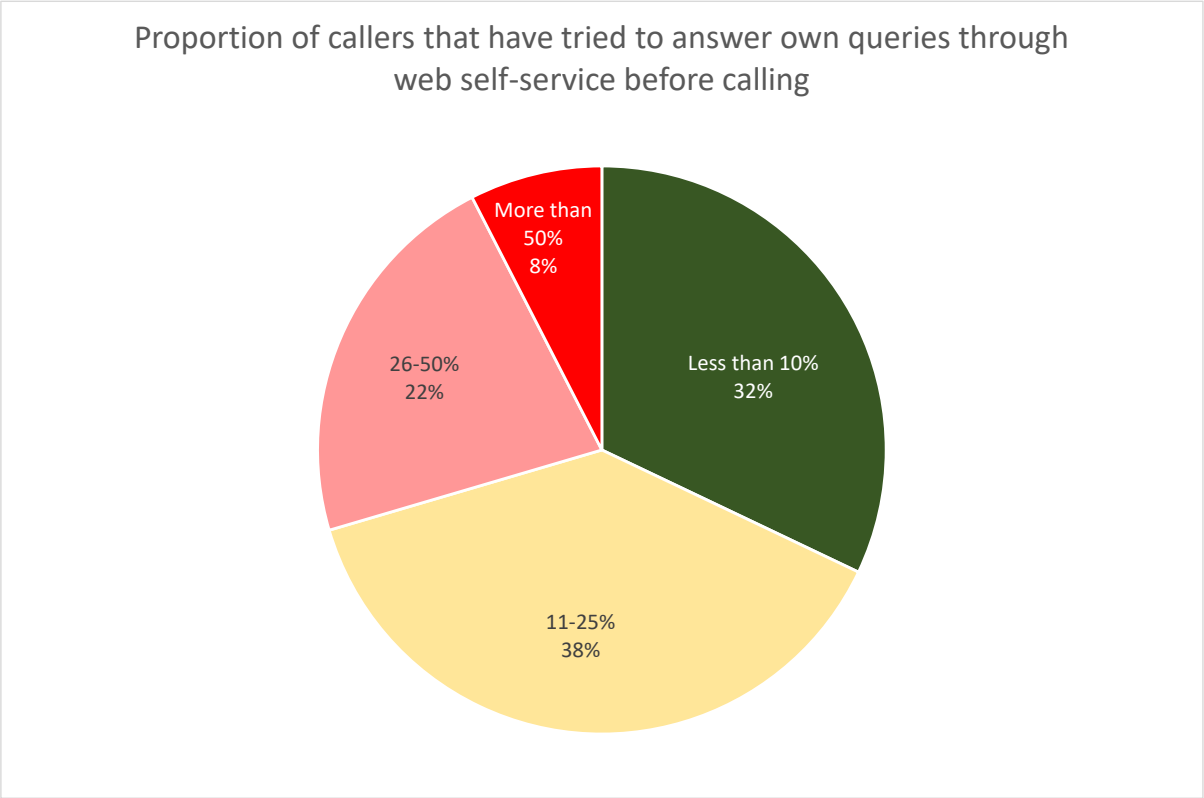

One of the two most important reasons for moving from web self-service to live telephony was that the escalation involved a complex issue requiring a live agent to complete successfully.

88% of respondents also felt that customers wanted the reassurance that a live agent brings to a conversation.

64% stated that the functionality that the customer required was not available online, but interestingly, 63% stated that they received calls about issues that could in theory be resolved online, but customers were unable or unwilling to do so. As such, businesses may consider that time spent educating customers in how to use self-service would pay benefits in the long term.

Relatively few respondents believed that website security authentication was an issue causing unnecessary inbound calls.

Figure 8: Why customers move from web self-service to live telephony

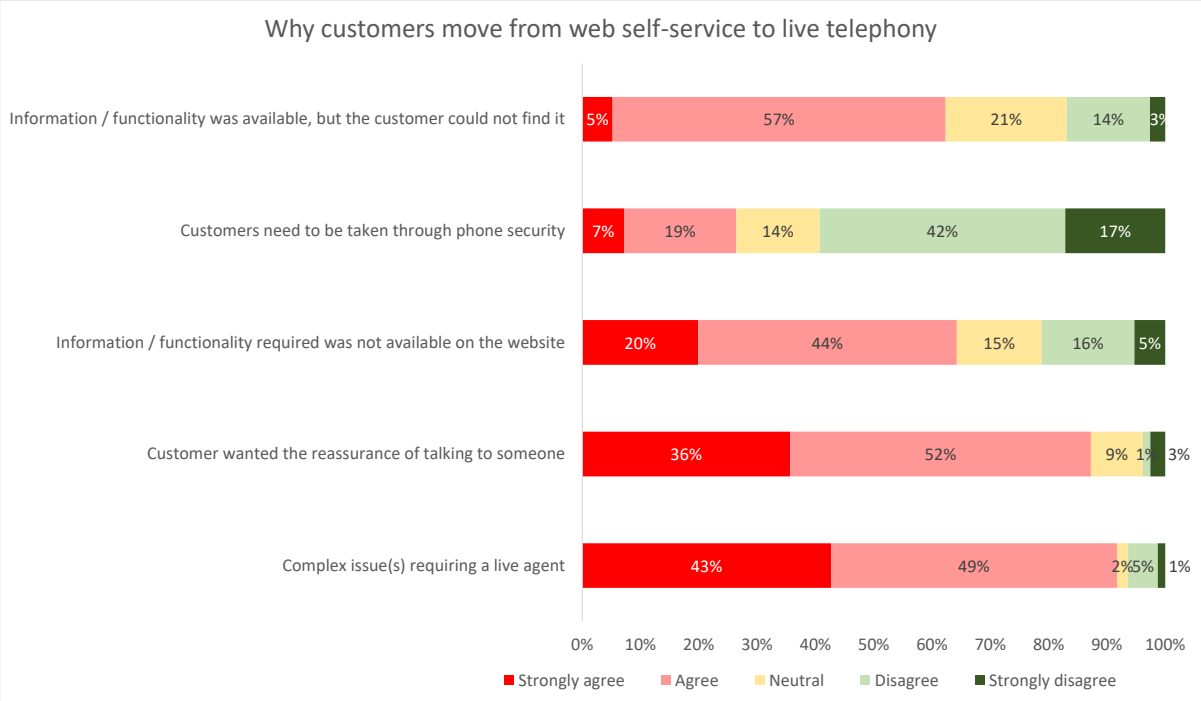

The use case for AI-enabled chatbots to help with web self-service is very clear: there are a very significant proportion of live calls that could be handled through web self-service if the customer were confident they could do so, and that the automation solution was able to handle the type of requests being given.

Clearly, the main issue to address here is that many customers – rightly or wrongly – feel that their issue is too complex to be resolved through a speaking with a chatbot or resolving it themselves elsewhere on the website. In some cases, this will be true and there is nothing to be gained by insisting that even the cleverest chatbots can solve every enquiry.

Yet even a relatively small increase in chatbot sophistication can bring more complex enquiries into play where self-service is concerned.

Conversational AI encourages the site visitor to engage with them using natural language, rather than keywords. Many applications now offer multiple languages, encouraging customers to at least try to resolve their issue online even if the website is not in their native language, increasing the possibility that bots can help.

Virtual agents parse, analyse and search for the answer deemed to be most suitable, returning this to the customer instantly. Many virtual agent applications will allow customers to give all sorts of information in any order, and either work with what it has been given, or ask the user for more detail about what they actually meant.

Sophisticated chatbot applications look for the actual intent behind the customer's question, trying to deliver a single correct answer (or at least a relatively small number of possible answers), rather than a list of dozens of potential answers that may contain a keyword. It may also try to exceed their brief by providing a list of related questions and answers to the original question, using machine learning and generative AI to predict what the next question may be and provide this answer as well.

The understanding of the context of what the customer is asking means the result can be more akin to that of an empathetic human who also has had access to what the customer has been trying to do. For example, if asked "When can I expect my delivery?", the context and the required answer will be different depending on whether the customer has placed an order and is enquiring about its status, or has only a hypothetical interest in turnaround times in case they decide to place an order.

Through 'listening' to what the customers actually say – perhaps through a mixture of large quantities of audio and text – the initial set-up configuration can achieve a good accuracy rate, which really benefits over time as a positive feedback loop is established. Solutions that gather and differentiate customer requests and results from multiple channels, noting the difference between them, have an even better success rate.

When the chatbot has low confidence that it has returned the correct result, it is able to escalate the customers query seamlessly to a live chat agent, who then has access to the self-service session history, enabling a greater chance of a successful resolution without repetition. (It is generally considered best practice that escalations to real agents are not hidden from customers). The eventual correct response can be fed back to the automated virtual agent (and the knowledge base underlying it), which will make it more likely that future similar requests can be handled successfully through chatbots.

Businesses should be aware that while the sophistication and capabilities of chatbot-type applications is varied, customers may not realise this and not even attempt to engage with the chatbot as they are used to basic rules-based applications.

Additionally, having been unconsciously trained over the years to provide their queries in a way which standard search functionality is more likely to be able to handle (for example, a small number of specific keywords), customers must be encouraged and educated to use natural language queries in order for virtual agents to be able to deliver to their full potential.

Implemented successfully, AI-enabled chatbots can provide a great deal more functionality and assistance than rules-based chatbots, going some way to addressing the main reason customers abandon the website to use telephony: that their issue is too complex.

Chatbots can also be used to address one of the other main reasons for channel migration: that the information and functionality is available on the website, but that the customer cannot find or use it.

Virtual agents can identify the reason for the customer query, and direct them to the correct part of the website, providing advice, links and ongoing support as needed.

Improving the sophistication and complexity of the issues that a chatbot can resolve will certainly help to deflect unnecessary live calls, but we should also consider the customer's emotion and experience in this. 88% of businesses agreed that one of the main reasons for moving to telephony was that the customer wants the reassurance of talking to someone.

While chatbots can be given 'personalities' and avatars, the general rule of current usage seems to be that bots are shown to be bots, and are not pretending to be human. It is debatable whether this will change in future as mimicry of human behaviour and language further improves and the interactions become more convincing, but one way in which the chatbot can give reassurance is by gathering the entire conversation and generating a summary in real-time where the customer can see exactly what has been done. If the customer is happy with the outcome and conclusions, this can then be emailed to the customer and assurances made that this has been kept on their file in case of any further issues or queries further down the line.

This will encourage customers to use chatbots again, and reassure them that the organisation understands and acts appropriately regardless of the channel used.

BUSINESS ISSUE #3: VOICE SELF-SERVICE IS RIGID AND DIFFICULT TO USE

A typical self-service IVR solution works on the assumption that the majority of customer enquiries can be solved by offering a limited number of solutions or options. This works well in cases where the caller has a simple request (e.g. speak to a sales agent or get a balance update) but often fails where there is a complex requirement.

Banks have commented that analysis of speech recognition-driven calls shows that customers can ask for the same information in many hundreds of different ways, which throws up significant problems for non-AI speech recognition that looks for keywords. Although touchtone IVR reduces the number of options and thus ambiguity, it is by its nature only useful in a limited number of calls and options.

Of those contact centres offering telephony self-service, a mean average of 16% are handled entirely by self-service without requiring an agent.

Figure 9: Overall proportion of calls handled entirely through self-service (only in respondents which offer telephony self-service)

	Proportion of calls handled entirely through self-service if offered
1st quartile	27%
Median	10%
3rd quartile	6%
Mean	**16%**

Even amongst those respondents for whom telephony self-service is a vital part of the customer contact strategy, it's no use trying to shift every customer service interaction onto telephony self-service, as if customers don't want to use IVR, they will "zero-out" (press 0 for a live agent, or try to find a similar shortcut).

If callers agree to try a company's self-service system rather than insisting upon talking to an agent, there is an implied contract that if the self-service session is unsuitable, the caller should be allowed to speak with an agent. Few things frustrate callers more than being hectored into using an unhelpful and irrelevant self-service system, and if businesses don't offer a live agent option to an irate and frustrated caller, they won't need to worry about providing customer service to them in the future as they'll have gone elsewhere.

Overall, a mean average of 14% of calls that go into the self-service option are "zeroed-out": instances where the customer decides that they in fact wish to speak with an operator, which is similar to the historical norm.

There is a broadly positive correlation between the size of the contact centre and the proportion of self-service sessions that are abandoned in favour of speaking to an agent: the larger the contact centre, the more often customers 'zero out'.

One possible reason for this might be that larger operations are trying to do too much with their self-service. There is some evidence to suggest that this is the case, as it is very noticeable that respondents from larger organisations tend to have far more options in the auto-attendant functionality of their IVR solution, and this tendency to offer a great deal of functionality and options may well also apply to IVR's self-service functionality as well. Overly complex or long-winded IVR functionality will tend to encourage session abandonment.

Figure 10: Reasons for abandoning telephony self-service sessions

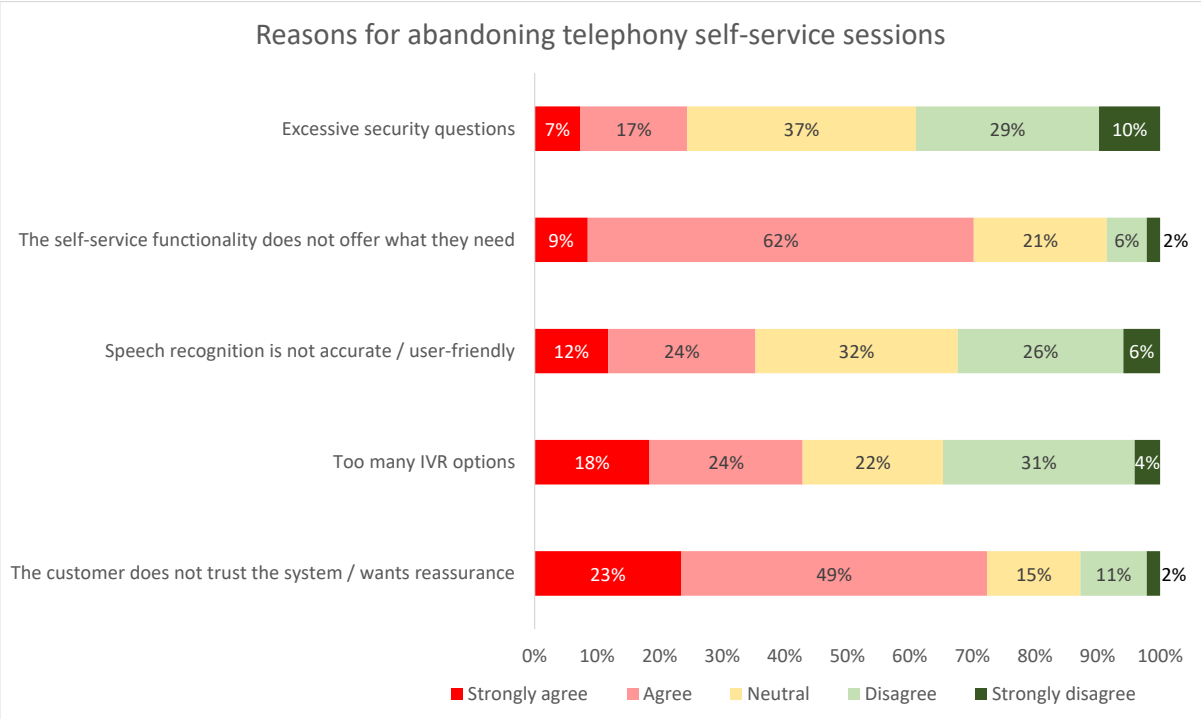

71% of business survey respondents agreed that customers abandoned self-service sessions because the self-service function simply does not offer what the customers want. While this at first glance may appear negative, it is the case that even in the most commoditised and transaction-driven environments a substantial proportion of customers will want to speak to a person, either because the system does not allow them to do what they want, there is a complicating factor involved, or simply that they wish reassurance or have multiple questions.

In such circumstances, it is the customer's choice to abandon the session, and this does not have to be a particularly negative experience as long as a clear exit path that leads to a live agent is marked early in the process. Situations where businesses hide their agents from customers, making them go around in IVR loops are the ones that give all telephony self-service a bad name.

42% of respondents agree that having too many options presented to customers is a major reason for them seeking human assistance. It is noticeable that 72% of respondents agree that the customer simply does not trust the system, preferring to have human reassurance that the request they have made has been carried out, or the information they are looking for is actually correct.

Of those using automated speech recognition, 36% of respondents agree that speech recognition is unpopular with customers due to lack of accuracy and user-friendliness. While this is high, it is a major improvement on past years, and it may be that customers are gaining confidence in how to use the system after many years of struggling. As customers continue to be encouraged to use natural language (both by successful interactions with corporate self-service applications, but also through digital virtual assistants such as Siri and Alexa), this issue should further decline.

Despite the wider and more powerful functionality that speech recognition – whether AI-enabled or not – gives to an IVR system, significant inhibitors are present. It is generally acknowledged that speech recognition can be considerably more expensive to implement than DTMF IVR, and is also likely to require significant, highly paid in-house resource to fine-tune and operate it going forward.

Some solution providers note that the majority of businesses' interest in moving from DTMF to speech recognition comes when the existing telephony self-service legacy system is approaching end-of-life.

The success or otherwise of speech-based IVRs is very affected by how callers are encouraged to use the service. It has been the case that some speech implementations have actually made life more difficult for the customer, who may not have the confidence that the system will understand their natural language request and provide very short, one-word answers in the same way that some customers use keywords when interacting with chatbots.

If nothing is given in the way of prompts or examples, callers may give too little or too much information as they are unsure of the sophistication or capabilities of the system, and this may be a reason for high self-service abandonment rates. Using prompts such as "describe in a few words why you are calling us, for example 'to start a new mortgage application'" can be extremely useful in setting ground rules for the successful use of a more sophisticated voicebot self-service application.

It is not just self-service that AI-enabled voicebots can help with: IVR is very often used for routing calls, but despite being in place for decades, this functionality is still often sub-optimal.

For many customers, IVR is seen as a way for the business to put up a barrier between them, involving a long and tortuous path before actually getting to speak with someone. Yet an IVR session should be used to capture information about the customer's identity and requirements that allows a business to provide an answer or route the call to someone who can actually help, rather than taking pot luck by dropping the call on the next agent available.

The IVR experience will often begin with a generic welcome announcement before offering various options for the customer to choose with a DTMF keypad (the vast majority of IVR is carried out with DMTF rather than speech recognition). More than one-third of initial IVR announcements take more than 30 seconds, which seems like a lot for customers to have to listen to, when all they want to do is talk to someone.

The audio-only nature of DTMF IVR places limitations upon how user-friendly the experience can be for a customer. There has always been a trade-off required between functionality and usability, which manifests itself in the number of menu options and levels that made available within the IVR system. The greater the functionality, the longer the announcements and the worse the customer frustration.

Looking at the number of levels used on a DTMF IVR (i.e. how many key-presses a caller must make to reach their destination), only 28% of respondents keep it simple with a single-level of options, e.g. "Press 1 for Sales; 2 for Service; 3 for Accounts".

24% of large operations present a possible four or more routing menu levels to their customers, a level of granularity that must appear daunting to their customer base.

Figure 11: DTMF IVR routing menu levels, by contact centre size

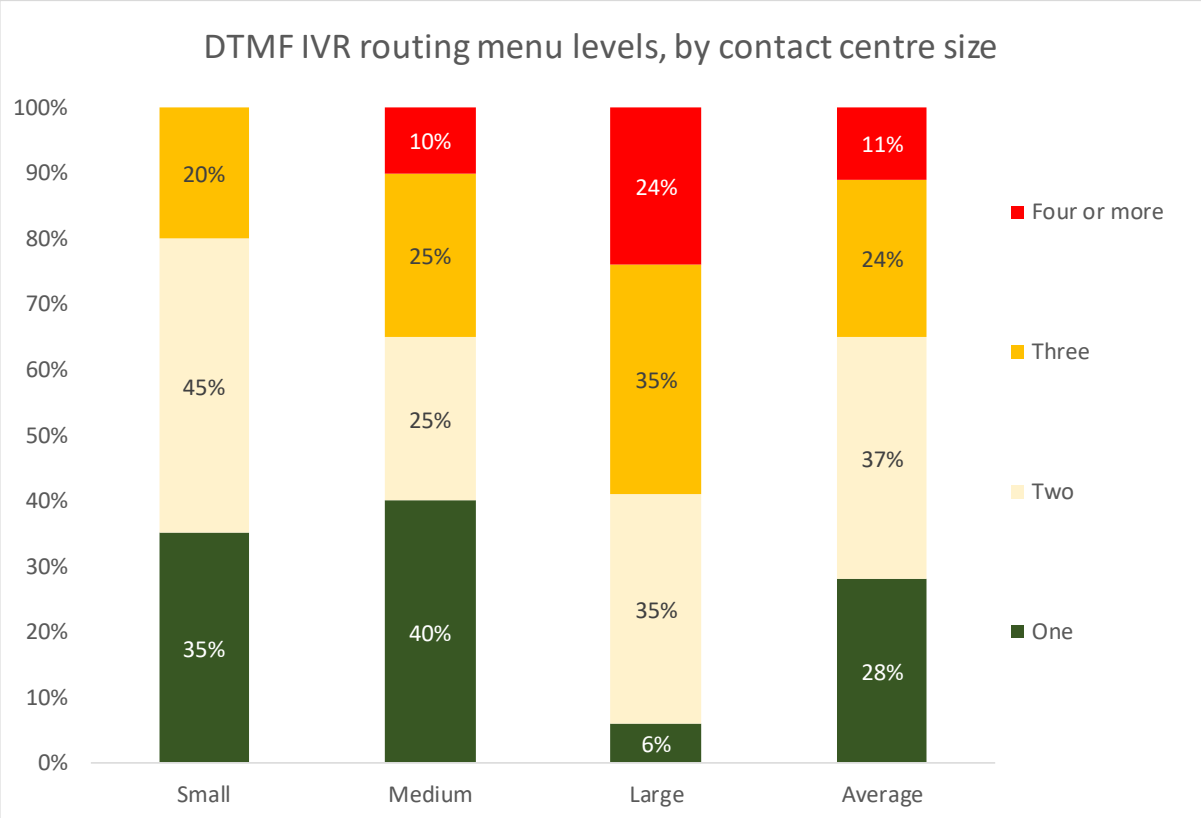

It is not just the number of levels in a menu that can frustrate customers, but also the overall number of options within each level. As the customer cannot see what the options are, but has to listen to each, it can be a very frustrating experience, and one which the movement to visual channels such as web self-service or visual IVR via a smartphone will go a long way towards alleviating.

Respondents report a median of between 6 and 7 options, which can still be a considerable number for a caller to listen to, especially if their preferred choice is the last one in line.

Logically, larger contact centres will tend to support larger businesses, which usually have more departments, offer a greater level of segmentation and have more products and services available to customers. Consequently, there are on average many more menu choices offered in the phone menu of large contact centres, with 70% of these respondents reporting offering seven or more routing options to their customers.

Figure 12: DTMF IVR total routing options, by contact centre size

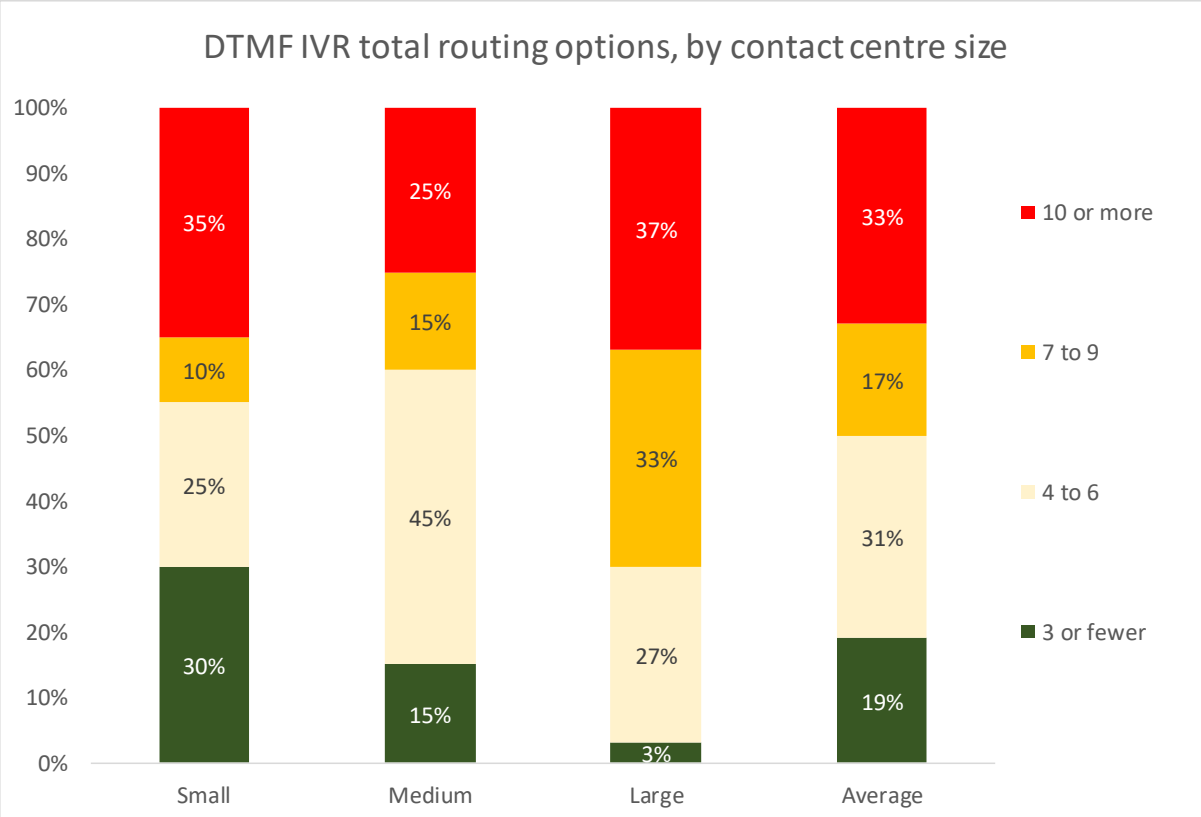

The report has shown that there are two main business issues around IVR: self-service is not as effective as it could be, and call routing is a painful (and sometimes ineffective) experience for the customer.

AI-enabled voicebots can provide a solution to both of these issues. A voicebot can be viewed as a chatbot that customers communicate with through voice rather than typed text. It needs the capability to handle natural language, to use artificial intelligence to determine intent and provide solutions, and to convert speech to text and back again.

Definitions differ, but whereas an IVR (whether speech-enabled or not) is restricted by programming to achieve outcomes based on specific patterns and inputs, an AI-enabled voicebot can have conversations with customers to determine even multiple intents and deliver a far wider range of personalised information and services. Crucially, it can learn from experience.

Advantages and use cases for voicebots include:

- Cost reduction - a typical self-service voicebot interaction costs around 30-50p per call, compared to over £5-6 for a live agent
- By asking questions and understanding the content and the context of what the customer requires, call routing can be far more quick and effective than touchtone IVR or keyword-driven automated speech recognition. In some cases, it is possible to turn a routing request into a self-service session, avoiding the need for a live agent at all
- Providing 24/7 service in multiple languages
- Create far wider opportunities for extending self-service capabilities through greater sophistication and understanding of the customer's intent
- Taking pressure off agents, particularly in times of crisis or volume surges
- PCI DSS compliance can be easier, voicebots can handle card payments rather than needing agents to do this, without customers having to type in numbers on a keypad
- Inbound and outbound activity can be linked to provide superior customer service: for example, an outbound SMS appointment reminder may initiate a change request from the customer which can be handled by a voicebot
- Encouraging customers to use natural language generates large amounts of data that can be used to further train AI models
- Voicebots can be used for outbound work as well, such as debt collection, reminders and surveys
- Personalising the voicebot allows greeting by name, and a change of voice / speaking style depending on customer preference. Customer identity verification is also possible (see next section of the report).

While CTI-like screen popping is useful for cutting time from the early part of a call, the insight that this functionality provides is often limited to the customer's name and a general idea about why they are calling.

AI-enabled voicebots offer instantaneous gathering and assessment of data from multiple sources to occur even before the call has been routed, which allows accurate prioritisation and delivery of the call if a live agent is needed, or initiate a self-service session if appropriate.

For example, a voicebot working in an airline contact centre may judge a call to be urgent if the caller:

- Has booked a flight for this day
- Rarely calls the contact centre, preferring to use self-service
- Is a frequent flier
- Is calling from a mobile phone rather than a landline
- Shares a similar profile with other customers who only tend to call for very urgent reasons.

In such a case, the AI may consider that there is a likelihood that the call is directly related to the flight that is happening today (e.g. there's a danger of missing the flight and the customer may need to rebook), and after listening to the customer's (possibly anxious and stressed) request, can either help through self-service or move the call to the front of the queue, routing it to an agent experienced in changing flights and whose communication style suits the situation and customer profile.

BUSINESS ISSUE #4: CUSTOMER AUTHENTICATION COSTS TOO MUCH

Customer security processes are about two factors: are you who you say you are, and are you allowed to do what you are trying to do?

Industry-wide, a mean average of 76% of UK inbound calls are stated to require caller identity verification, the highest on record.

92% of respondents who authenticate identity do so through human means, taking an average of 44 seconds to do so.

Around 10% of survey respondents use IVR or speech recognition, but many of these also use the agent to double-check in many instances once the call is passed through, wasting the caller's time and increasing the contact centre's costs.

Customer authentication not only costs businesses money, but also impacts negatively on agent engagement: an agent taking 80 calls per day may spend around 45 minutes of an eight-hour shift doing the mundane and repetitive task of taking customers through security.

Although in-call efficiency has improved, identity verification is slower than it has ever been: over the past decade, our surveys have found consistently that around 60%-70% of calls require identity checks, which take considerably longer due to more stringent testing (a rise in the length of authentication of around 90% since 2010).

Figure 13: Proportion of calls requiring caller identification & average time taken, 2010-23

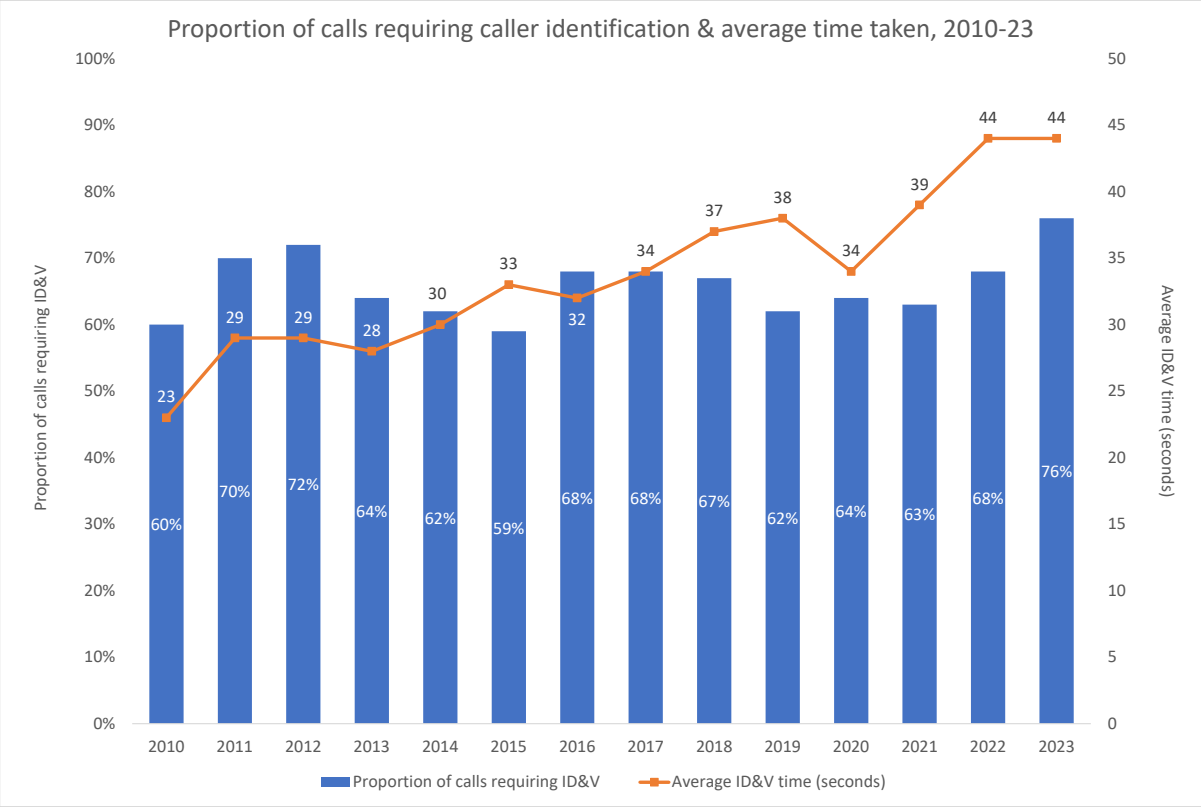

The unnecessary cost of caller authentication

Using figures from ContactBabel research, it is possible to estimate the industry-wide cost of customer identification authentication using an agent. Please note that as respondents change each year, this figure is an indicative estimate based on this year's survey and should be read as such.

76% of all calls require a security and identification process to be completed first. This year, 92% of calls were reported to be authenticated by agents. On average, it takes 44 seconds to go through security. Using these statistics, it is possible to estimate how much UK contact centres spend each year on screening customers by using agents.

Inbound calls per year (handled by agents): 5.38bn[3]

Proportion of inbound calls that require security and identification checks: 76%

Average length of agent-handled security and identification check: 44 seconds

Average call duration: 7m 1s (421 seconds), therefore 10.5% of the call is ID&V

Mean average call: £5.58

Cost of time spent on agent-handled security and identification check: 58.3p per call

Proportion of calls requiring ID&V: 76%, of which 92% require an agent

Therefore, overall cost of agent-handled security and identification checking: £2.19bn per year

While dedicated authentication solutions such as voice biometrics and call signalling analysis are increasingly being used, these are expensive and focused mainly on large businesses such as banks where fraud reduction is a key priority.

For many businesses, having a voicebot rather than an agent take customers through security will reduce costs while providing a similar level of security, and has real potential to cut costs and improve agent morale.

The security process remains the same as if it were a live agent taking these details, with the voicebot simply taking their place. If the voicebot notices undue levels of stress or anxiety, they can flag the call to the agent as potentially fraudulent for further security checks.

[3] ContactBabel, "UK Contact Centres 2023-2027: The State of the Industry"

It's important for businesses to understand that if they're not already using AI, then they haven't already missed the boat, and that even with unlimited budget and resource, there will still be many customer service activities that are more appropriate for a live agent to do.

For first-time AI deployments, the focus should be on the solution delivering a high-quality solution to a relatively small and clearly defined business process or issue, rather than taking on more complicated situations, even if there is a potentially higher benefit.

In many cases, it might be appropriate to start with a self-service-focused chatbot or voicebot project, and then look to roll out AI to other parts of the customer journey including call routing, back-office processing, analytics and agent assistance.

Apart from the dangers associated with an overly complex initial AI project, scale is also an issue to consider. To begin with, businesses may consider it wise to limit the number of concurrent customer users that AI supports (i.e. dozens rather than hundreds of concurrent users), in order to learn what works and what needs improvement in each use case, and to optimise processing performance by providing the right amount of processing capacity.

Over time, machine learning tends to require less processor power and running a relatively small scale AI implementation for a few months will provide a more informed view of what full-scale usage of AI will involve, meaning that the right amount (and cost) of processing power can be established.

If you're considering implementing AI, there are some questions that you should ask yourself first:

- Is there a specific pain point or issue within the operation that needs to be addressed? e.g. large numbers of simple calls about a small number of issues
- How does this affect CX, and how would the customer like this to be improved?
- Are there solutions available that have successfully addressed these issues already?
- How quickly can this be implemented, and what initial and ongoing resource will actually be required to make it run successfully?
- What upheaval would it create within the existing operation? What effect does it have on the customer experience?
- Are the improvements measurable? What are the cost impacts?
- Is there a sufficient volume of clean data in order to train an AI system effectively?
- Will infrastructure or existing platform need to be replaced?
- Is AI likely to be the most appropriate way of dealing with this issue?

It is likely that some senior decision-makers within the enterprise have an unrealistic expectation of how AI can help within the CX environment. It is important that the boundaries of the project are clearly understood, with relevant baseline metrics captured before the project, and clear and achievable outcomes signposted and agreed so that the eventual level of success of the project can be clearly understood.

Many contact centres may consider a limited, low-risk use case which can be implemented quickly and relatively cheaply in order to demonstrate a quick win and assert the viability of AI within a customer contact operation. For example, increasing the number of self-service interactions through improved AI-enabled website guidance in certain defined cases is an example of a project which has a clear and easily measured metric which translates directly into call and cost reduction.

Having said this, it is important for contact centres not to sell this to high-level management as being an opportunity to reduce headcount, as it is unlikely that this will be the outcome from any AI project, certainly in the short-to-medium term. It may be better for the project to be viewed as

improving CX through providing customers with an alternative to a frustrating web browsing experience, ending with an unnecessary and unwanted live call which then impacts on queue times and call abandonment rates.

While it is important for the initial AI implementation to focus on achieving success within its own terms, it is also important that this is not seen as a tactical point solution with a single end in sight. For example, while the initial implementation may be focused on increasing the effectiveness of self-service in a defined area, the longer term view may be to roll out AI into the agent's sphere, assisting them while on live calls.

As such, a roadmap of logically linked business cases can help to establish a long-term vision which can be shared with non-operational senior personnel to help them understand the strategic use of AI across the customer-facing parts of the business.

For example, a simple yet strategic rollout of AI may look similar to the following:

- Use a virtual assistant / chatbot to improve the take-up of knowledge held within the FAQ database, by improving the search mechanism and offering a two-way conversation interface in order to provide more accurate answers. Capture the phrases used by customers in existing human web chat sessions to understand the questions they will ask your chatbot

- Place this virtual assistant upon the agent desktop in order to provide them with more knowledgeable potential answers within the call

- Meet customer requests over voice and text through the use of natural language processing, in order to assess customer intent, and provide answers or optimal routing strategies

- Improve efficiency, consistency and effectiveness of back office processes connected with the contact centre through the use of robotic process automation

- Deploy analytical AI in order to discover patterns of data relevant to the business that would otherwise not be identified.

Once the process, objectives and outcome are clearly defined, the selection of a vendor and solution can then be approached. In a rapidly growing and heavily hyped market sector such as AI, it can be difficult to compare vendors with like-for-like solutions.

For example, in the case of chatbots, on the one hand these can be rule-based, have limited conversational capability and are unable to learn; on the other, they may use natural language processing, engage with customers in order to ask further questions to determine intent, and be capable of self-improvement. The development time, resource and cost associated with each of these types of chatbot are very different, and businesses must decide whether they are looking for a quick win, or whether they have a definite long-term AI strategy in mind.

Businesses should also consider the type of developer and implementation model that's most appropriate: some self-service chatbots can be based on off-the-shelf software which is then customised and implemented by an in-house development team, whereas some businesses may prefer to bring in third-party developers with greater experience in AI implementation. The rate of change within this technology sector is very high, so implementations that are measured in a handful of months rather than longer would seem to make more sense at this point.

At the request for proposal (RFP) stage, businesses may consider asking potential suppliers:

- What are the current capabilities of your AI solution and what does your product roadmap look like?
- How do you propose escalating interactions to live agents if the AI solution cannot handle it?
- What metrics do you propose using in order to judge the success of an AI implementation?
- What does the timeline of a successful implementation look like? Do you have a reference site?
- How do you propose to train the AI, and what will our training data need to look like?
- How do you propose to integrate AI with our existing systems, and how much customisation will be needed?

At the initial stage of the implementation process, datasets that the AI models will be learning from must be analysed, cleansed and curated to provide a solid basis for the AI solution to learn from. Vendors will have dedicated examples of neural networks that work for various business cases such as providing answers to queries or estimating the time taken for a process to be completed. These can be used as a starting point for training the AI model, and to enable it to start making predictions of its own.

While each vendor will have their own framework and architecture, they are likely to follow a similar path involving input, interpretation, action and improvement. Input is gathered by the system – often from a customer – and if necessary is then translated into a format which the system can understand (e.g. through speech-to-text or OCR). Once the data are converted, the AI looks for the customer intent behind the input using NLP as well as other metadata such as location or customer history. Once the intent has been decided with a certain level of confidence, various solutions are considered and presented to the customer. Finally, the loop is closed through gathering feedback about the success or otherwise of the answer, which is then taken into account in future interactions, with the AI learning what works best.

In theory, despite the often onerous effort involved in creating a clean pool of data, the implementation of a virtual assistant or chatbot should not have to change the existing IVR or web chat infrastructure materially, as the AI agent is treated as just another user of this technology.

The more data that the AI has to train on, the more likely it is to succeed. As machine learning works through pattern recognition, this can include metadata and context which may seem somewhat peripheral to the process, but there are likely to be patterns that have not been recognised by human users. This allows the AI model to understand customer intent and also to be able to discern which customers need to be treated in a way outside of the ordinary (e.g. in an emergency situation, if the customer is likely to defect, or if they have contacted the business multiple times in a short timeframe). Analysing the use of existing data shows the ways in which customers want to express themselves.

As with any IT project, testing is key to success but with AI implementations this is even more important. Most IT systems work on an input-processing-output basis, where the point of the implementation is that the same thing happens every time, reliably and predictably. As the processing element of AI involves elements of learning (and hopefully improvement), the output can change over time. This does not always end optimally: Microsoft's Tay[4] is only one example where the AI displayed inappropriate and offensive responses after interacting with users. Businesses should be aware that AI solutions, especially in the early stages, may require very close supervision and possibly intervention. Dedicated chatbot testing vendors offer services to make sure the chatbot is working properly before putting it into a live environment.

Measuring the performance and success of an implementation is always vital, but never more so when it is for a highly anticipated and poorly understood solution such as AI. There is likely to be far greater interest in and pressure from the higher echelons of the business than is the case for most contact centre technology implementations, and thoroughly understanding the outcome of the initial implementation is vital.

There is no baseline set of metrics that every AI implementation should be measured against, although in the widest sense, the impact upon customer experience, agent experience, cost and operational benchmarking should all be considered. Of course, it also depends on the area of the contact centre business processes that implementation is aimed at improving.

Some examples of AI-related metrics around self-service include:

- volume and success rate of successful chatbot / voicebot interactions, segmented by type of interaction
- customer satisfaction by self-service segment (particularly useful for comparing AI-enabled self-service with scripted self-service applications such as rules-based chatbots)
- proportion of self-service attempts that are handed off to agents, and following from this, the proportion of these which are handled by a single agent (showing the efficiency and accuracy of routing and the collection of relevant information in the initial self-service session)
- length of self-service session (this is related to customer effort, which has a strong influence on CX)
- change in inbound call volumes.

[4] https://en.wikipedia.org/wiki/Tay_(bot)

While each AI project is different, businesses may wish to consider following these initial steps:

1. Review operations, categorise them and if possible, quantify the cost, complexity, value and volume of each type of interaction to determine which are most suitable for automation. Use topic modelling to consider the subjects most under discussion as well as their relative complexity. Agents are likely to have a good idea of the types of questions that they are frequently asked, and which ones are most easily automated, so they should be included in this phase. AI, RPA and analytics can also be used in the discovery phase of the project to identify topics, processes and bottlenecks which may be causing excessive cost and customer effort

2. Consider whether processes and queries can be handled by a simple 'if x then y' type of rule which can be programmed and always adheres to the workflow, or whether customer intent and requirement is likely to be more complex, and therefore more suitable for AI

3. Develop AI user cases for specific processes and interactions, including giving each bot a 'job description'. Although starting small is often a good idea, have a roadmap that doesn't constrain you to follow any early decisions

4. Fully understand the necessary underlying workflows, systems, data and processes which support this function and which are supported by the action of the bot

5. Develop a clean pool of data for the AI to learn from, being closely supervised by human experts

6. Companies using voice data will of course need to implement speech recognition, as using text transcripts for analysis is far easier and more powerful than using audio data. This will allow a wide variety of applications, such as search, scoring of agent behaviour, monitoring and QA of agents, and predictive capabilities (e.g. predicted NPS or first-contact resolution) which are based on large quantities of calls rather than just a handful

7. Make sure that there is a reliable path to a live agent if escalation is needed, including providing the context and history of the interaction so far

8. Monitor, track and report on the success of the bot as if it were an agent, and continually look for improvements. Be aware that processes and data may change over time, so regular reviews are vital.

While a voicebot can deliver the same type of sophisticated AI-enabled functionality as chatbots, there are some extra elements to consider in implementation and usage:

Noise cancellation: the additional requirement of carrying out speech-to-text in order for the interaction to be analysed is potentially more difficult because of background noise, but there are a number of software-based, AI-driven noise cancellation solutions available that can provide a clearer and more accurate delivery of the customer's voice. This is particularly important when capturing names, addresses, payment card details and account numbers that could then be passed onto the agent if required.

Resilience: the dependency on multiple systems – sometimes external ones – means that integration may not be entirely straightforward, with components such as speech-to-text, natural language processing and noise cancellation all having to work together seamlessly and quickly, which risks system breakdown.

System latency: users of even basic automated speech recognition have been aware for years that it is not close to the experience of speaking with a live agent, with pauses of several seconds typical as the system processes the speech before replying. Solution providers claim that it is now possible to achieve latency of around half a second with an ideal voicebot set-up, but latency is something that businesses have to consider seriously when implementing AI-enabled voice assistants.

Other considerations around voicebot implementations include:

- Voice conversations don't follow linear paths: people speaking may repeat themselves, decide to change the phrasing halfway through a sentence to make it clearer, or interrupt. These are not issues that a text-based application such as a chatbot has to cope with.
- While accent recognition has improved greatly, it is still not infallible. There is also the issue with similar- or identical-sounding words (e.g. "1 / won"; "4 / for"; "2 / to") being misunderstood by the system.
- By its nature, a voicebot cannot show customers anything, unlike a chatbot which can send a clickable link or picture to the customer. As such, voicebot implementations will generally not be as wide-ranging or ambitious in their scope as chatbots.
- Whereas a pause in a conversation with a chatbot is expected – and allows the customer to consider their next statement carefully – this is not a part of spoken conversation. Customers can feel rushed to explain their issue, often sub-optimally, and can end up flustered, give the wrong information or want to start again.
- Intent detection can be more difficult with voice conversations, as people will tend to be more long-winded and include information that is not strictly relevant. There is also the necessity for systems to be able to break up multiple intents within a single utterance (e.g. "I want to book single train tickets from Newcastle to London for two people, travelling tomorrow and arriving before 11am.").

- Consider the voice brand you would like to use: many voice assistants sound robotic, overly produced or do not share the same national accent as the typical customer. It may be that you decide that certain types of customer (younger / older; male / female) should be communicated with through different voice personas.

- Voicebots should be able to cope with interruptions, and keep the context of the conversation so as to be able to answer additional questions while keeping the required dialogue flow.

KNOWLEDGE MANAGEMENT

One of the most central and critical elements to a company's service capability is the knowledge base, which is vital to the accuracy and consistency of both the self-service and assisted service experiences across channels. An AI project cannot succeed unless the data it learns from is clean and deep.

For many organisations, a knowledge base started off as a list of useful documents and files, which quickly grew into a wider, less coherent collection of information sources, requiring increased levels of expert management, amendments, editing, and deletion.

However, the resources required to keep these knowledge bases up-to-date are very scarce, as the people within the business that have the right capabilities and expertise to do so also have their own jobs to do. Very quickly, what started off as a useful and highly tailored information resource mushroomed into an expensive, out-of-date and increasingly less useful collection of information of wildly varying quality. AI can assist in the management of knowledge bases by feeding back successful outcomes, and noting when the answers provided did not meet the requirement.

On an ongoing basis, feedback from agents and customers will identify gaps in the knowledge base which will need to be filled by product experts. Some knowledge bases will require full-time, dedicated resource to manage them, whereas others will rely on automated systems making dynamic changes depending on callers' and agents' requirements.

It is often the case that large businesses with many products and services to maintain will have numerous editors across many departments who can make suggestions, although it may only be a small handful of people who will verify and publish this information. Businesses may want to consider allowing experienced contact centre agents to create new entries based on their communications with the customer. Understanding which documents are being used the most allows the maintenance efforts be focused on the most important areas.

While some knowledge base solution providers state that 80% of questions can be answered by 20% of content, it is each business's decision to decide how the remaining 20% of queries will be handled (but of course, even these 20% of documents will change over time as customers' requirements and the businesses' products will not stay static). Some will consider that this is a reasonable proportion to be handled by more traditional means, such as the contact centre, whereas others will leverage expert internal resource, as well as customer communities and forums to fill these knowledge gaps.

It is not just the publishing of information that is vital: it is feedback on its accuracy and success from the wider user community and any automated systems which will help the business to fine-tune the knowledge base. Processes to gather this feedback should be put in place, and continually revisited to check their effectiveness, and it is possible to add successful answers to the knowledge base very quickly if a response from an agent (for example, via email or web chat) has been marked to be successful.

By their nature, knowledge bases only contain information that a company thinks is relevant to its customers, who may disagree or find the sheer volume of data to be unmanageable when searching through a company's website. AI can be used to understand what customers are doing on a website (including any web chat questions) and provide answers in context, either directly to the customer or through providing agents with the likely correct response.

One of the keys to successful knowledge management is the continual monitoring, updating and publishing of the most accurate and in-demand information. Businesses should consider setting internal service levels for the knowledge base, for example only returning documents and suggested answers that have over a specific score for relevancy, and no more than a small number of answers per enquiry.

If customers are trained to expect a self-service or virtual agent experience that returns pages and pages of documents that bear little relevance to their original query, they will very soon abandon self-service entirely. It is also vital that the information contained in the knowledge base is available consistently across all channels, whether through a virtual agent or human agent.

AI cannot succeed or learn without clean, up-to-date and accurate information. It is tempting – and certainly the easiest route – to allow knowledge to stay only in the heads of the experienced agents and supervisors, who can share this with colleagues as and when needed. However, in a remote working environment this is sub-optimal, and even the most loyal agent will eventually leave the company, taking their knowledge with them. Although building procedures and allowing the time for these agents to update the knowledge base is an extra expense and complexity, it is vital for the success of AI-enabled self-service.

It is not just the publishing of information that is vital: crowd-sourcing of answers, and feedback on accuracy and success from the wider "super-user" community will help the business to fine-tune the knowledge base and train the AI.

Much of this may be unstructured data, which the AI is capable of handling and putting into a usable format. Processes to gather this feedback should be put in place, and continually revisited to check effectiveness, and it's possible to add successful answers to the knowledge base very quickly if a response from an agent (for example, via email or web chat) has been marked to be successful, and AI is an effective method of doing this regularly and consistently.

Those who contribute timely and useful information – whether a customer or an employee – can be rewarded and recognised accordingly. People **want** to share their knowledge with others, and enabling them to do so easily is beneficial for all parties concerned. Businesses could measure the success of the knowledge management system by measuring the return on investment from call avoidance, by the rating or score given by readers of recommended articles, or through targeted customer satisfaction ratings.

END-USER QUESTION #2: WHAT DOES SETTING UP AND TRAINING AI-ENABLED CHATBOTS / VOICEBOTS ACTUALLY INVOLVE?

The complexity of the process will vary depending on the desired capabilities of your chatbot but a general overview of what's involved includes the following:

- **Defining the purpose:** First, you need to determine what your chatbot will do. Is it for customer service, answering FAQs, or something else? This guides the training process.
- **Gathering data:** You'll need an amount of text data, like conversations, emails, or documents, relevant to your chatbot's purpose. This data helps it understand language patterns and user intent.
- **Training the model:** Different AI techniques are used, like machine learning and natural language processing (NLP). NLP helps the chatbot interpret user input and generate responses. The training data is fed into the model, allowing it to "learn" from the information.
- **Testing and refinement:** Once trained, the chatbot is tested with real users or simulated conversations. This helps identify weaknesses and areas for improvement. Based on the feedback, the training data and model parameters can be adjusted for better performance.
- **Deployment and monitoring:** Finally, the trained chatbot is integrated into your platform (website, app, etc.) and goes live. However, the learning process doesn't stop. You'll need to monitor user interactions and continuously refine the chatbot with new data to ensure it remains effective and up-to-date.

In any technology implementation there will be risks of failure: with AI covering a vast amount of territory and with the potential to be misunderstood by business owners, planning and expectations must be managed very carefully.

- Expectations of what the AI implementation can actually achieve must be closely managed. There may be the expectation from senior management that headcount will immediately begin to drop, but in the majority of instances this is not why AI is being implemented. Focusing on a tightly defined use case will reduce the risk of implementation delays and expecting too much, too soon from AI. However it is important not to see even a relatively modest implementation of AI as being a point solution, rather than a single strategic step.

- AI in the contact centre is relatively new, and with it being so popular, there is a shortage of skills, support and resource within the industry as a whole. In-house technology departments are less likely to have capability, expertise and experience, meaning that the risk of suboptimal deployment and the requirement for third-party assistance may be higher than with other more traditional IT implementations.

- Businesses data assets must be in place before implementation of AI, as this is a technology that relies upon having large, clean pools of data that it can be trained on and learn from. Without this in place, it will be virtually impossible for any AI implementation to get close to its potential.

  The preparation of data will involve having an organised, non-siloed data architecture, a consistent data vocabulary, the means of accessing this data securely and quickly, and the ability to access other pieces of relevant information (e.g. customer-related metadata) in order to include greater context. Without this, it will be more difficult for a machine learning process to train itself effectively, or for a chatbot to be able to use all of the relevant data in order to reach a correct conclusion.

- Always have a well-designed and clear path out of AI-enabled self-service to a human agent. Trapping a frustrated customer in a self-service session runs the risk not only of training them not to use self-service again, but also poisons the well for other companies using AI. This is what happened in the early days of email support: customers would try to communicate with one or two businesses via email, and when they didn't receive a response for days (or ever), they decided that the whole email support channel was unworthy of their time. It took many years to change this perception and to get them to trust the channel.

- In the AI world, knowledge management is not something that is a part-time job or that can be handled by amateurs. Consider developing more full-time, expert roles to support knowledge bases and to enable understanding of data models and flows across the entire enterprise. AI experts have to understand both data and also the real-life business / customer issues, and this resource can be difficult to find. Some businesses use 'superuser' teams of experienced agents who understand which requests are most suited to automation, and the process steps that are required for successful outcomes.

- Some vendors have predefined applications and bots which can be specific to vertical markets, but even these will need to be refined and developed over time.

- The business originally needs to identify the outcomes that are most important to them: reducing cost to serve by lowering handle times; improving CX; reducing call-backs about the same issue, etc., and the success or otherwise of these outcomes must be tracked closely to gauge ROI. Track service levels of AI processes in the same way as you would your agents.

- Discovering why customers contact your organisation, and identifying repeatable, low-risk tasks that can be automated is a good way to start, particularly if the cost of handling these issues manually can be calculated. Soon, the business will have metrics on how many customers were successfully helped by AI, the cost avoided and any effect on customer satisfaction ratings, which can be shared with senior management and will help with future use cases.

- Explain to employees why you're implementing AI, how it will affect them, what it means for them in their role and how they can help. There may be concern that their jobs are going to be replaced by bots, so educating them about what the contact centre of the future will look like, training them thoroughly and getting their opinion and feedback will help to get them onboard.

- AI and bots need to be entirely integrated into the tasks, flows and reporting. Choosing an approach that allows a bot to be built once and replicated across channels will save time and money, as well as standardising the customer experience.

- There have been a lot of media scare stories about AI and robots making people unemployed. It is important to emphasise to agents that any AI implementation is about making their jobs more interesting and effective by allowing AI to handle simple and repetitive requests, as well as providing them with more of the information that they need to serve the customer more effectively.

- While agents are experts on answering customer queries, it may be too much to ask them to spend significant amounts of their time on contact curation as well. As such, businesses should consider how to incentivise power user experts (both inside and outside the enterprise) to help with knowledge management and problem resolution.

Common pitfalls from conversational AI implementations include:

- A lack of clean data
- Low quality content in the knowledge base
- Little or no feedback from whether the bots' actions are successful, or even understanding what success looks like
- Not having a smooth handoff to live agents with full context
- High maintenance requirements from continuously training AI models.

END-USER QUESTION #3: IS THERE ANYTHING THAT SUCCESSFUL AI-ENABLED CHATBOT / VOICEBOT IMPLEMENTATIONS HAVE IN COMMON? ANY PITFALLS TO AVOID?

Successful AI chatbots share a winning formula. They have a clear purpose, like answering questions or generating leads. They excel at understanding natural language, thanks to powerful NLP capabilities. This, along with training on massive amounts of data, allows them to deliver accurate and engaging responses. They also continuously learn and improve through user interaction, becoming better over time. Additionally, they offer a user-friendly interface and transparently acknowledge their limitations, directing complex issues to human agents when needed. By excelling in these core areas, AI chatbots can become powerful assets for businesses and users alike.

The main reason for the slow response rate and excessive length of web chat and especially email is that in the past there has been very little automation used in the UK contact centre industry. This has also meant that the cost of an email or web chat has historically been similar to that of a phone call.

Digital channels may work quite well for customers, but many businesses are not seeing the cost savings that successful automation can bring. Even now, very few emails are handled entirely by AI, although this year has seen another increase in the proportion of web chats being dealt with by AIs working alongside agents, suggesting responses which agents can then accept or amend. This way of working is most likely to be the norm in the short-term at least, with the speed of automation and the emotional intelligence of humans providing superior service at a lower cost.

It is worth noting that the proportion of web chats handled **entirely** by automation has increased from 4% in 2019 to 20% in 2023, and as shown later, there has been a significant drop in the average cost of web chats in the past few years almost certainly caused by this.

Figure 14: Human / AI email and web chat handling

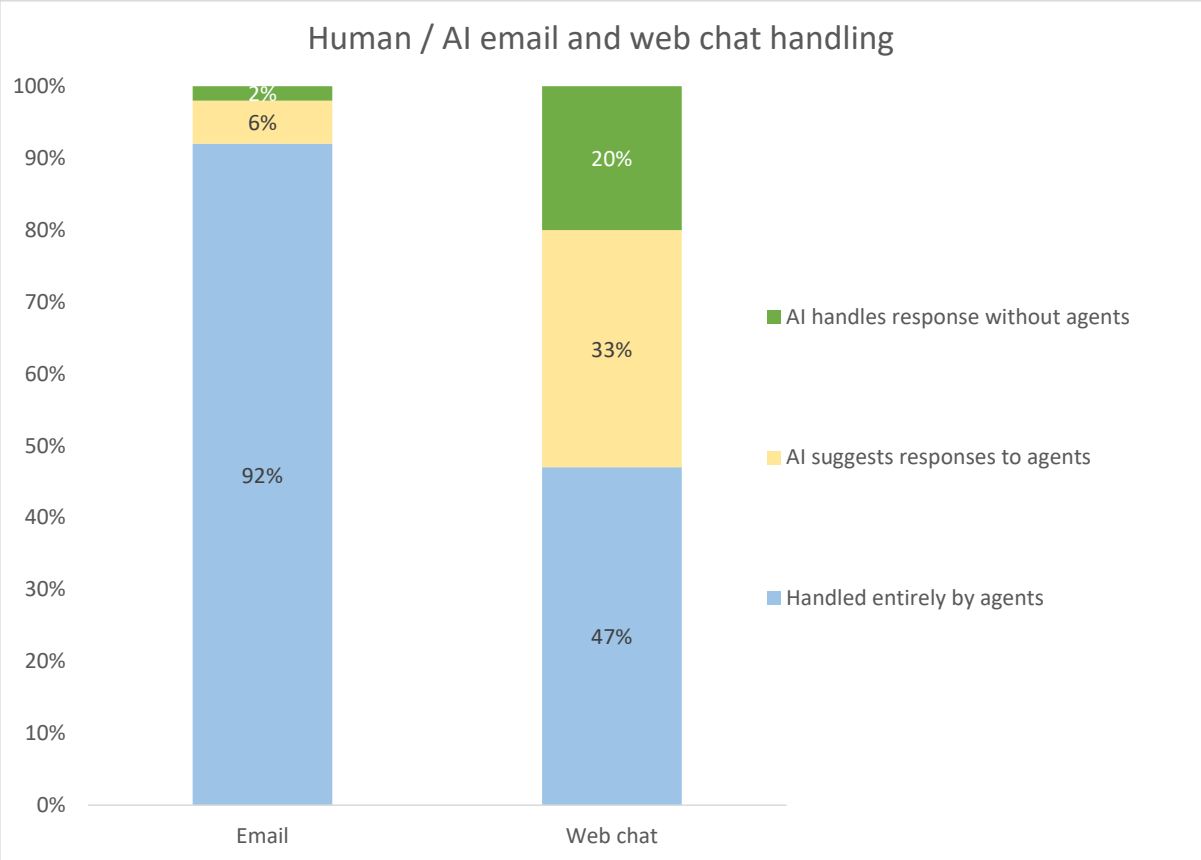

Whereas only 5% of web chats had any automation involved in 2015, this grew to 53% by the end of 2023, mainly as a result of initial handling by automated chatbots which may then hand off to live agents where appropriate.

It should be noted that fully-automated AI-enabled web chat has increased very significantly in recent years as well.

Figure 15: Level of automation used in web chat, 2015-23

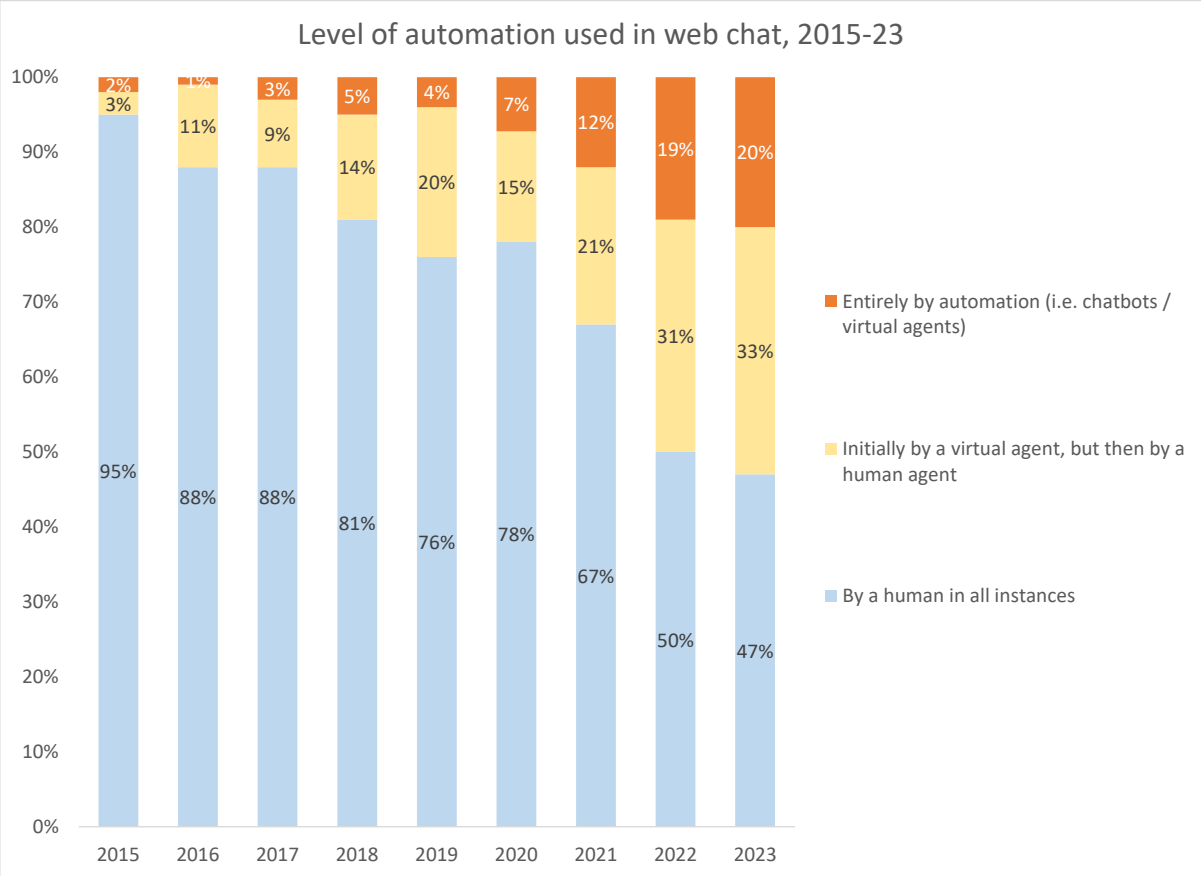

Despite a fairly low current use of AI across industries, there is widespread interest in implementing this solution, with 31% of respondents intending to implement AI within 12 months. While these figures are probably overly ambitious, it does show real interest from the contact centre industry.

The outsourcing, TMT (telecoms, media & technology) and retail respondents report the greatest current use of AI this year.

Figure 16: Use of AI / Machine Learning, by vertical market

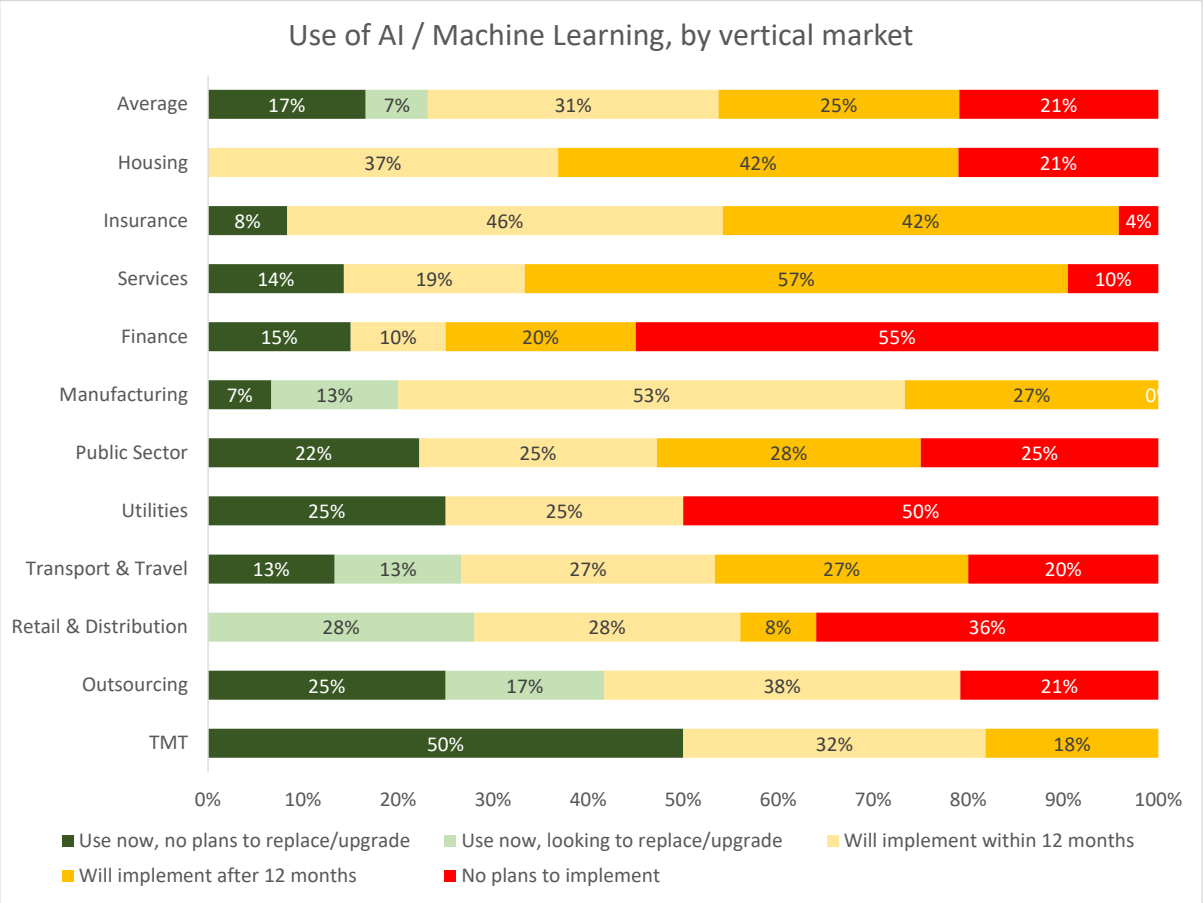

In recent years, larger operations were much more likely than sub-50 seat contact centres to have implemented some form of AI, although there had been very significant interest even amongst smaller operations.

This year, many more survey respondents from small contact centres state that they are using AI – it is not to say that these solutions will be as powerful or sophisticated as those used in large contact centres, but there has definitely been a noticeable increase in usage.

Figure 17: Use of AI / Machine Learning, by contact centre size

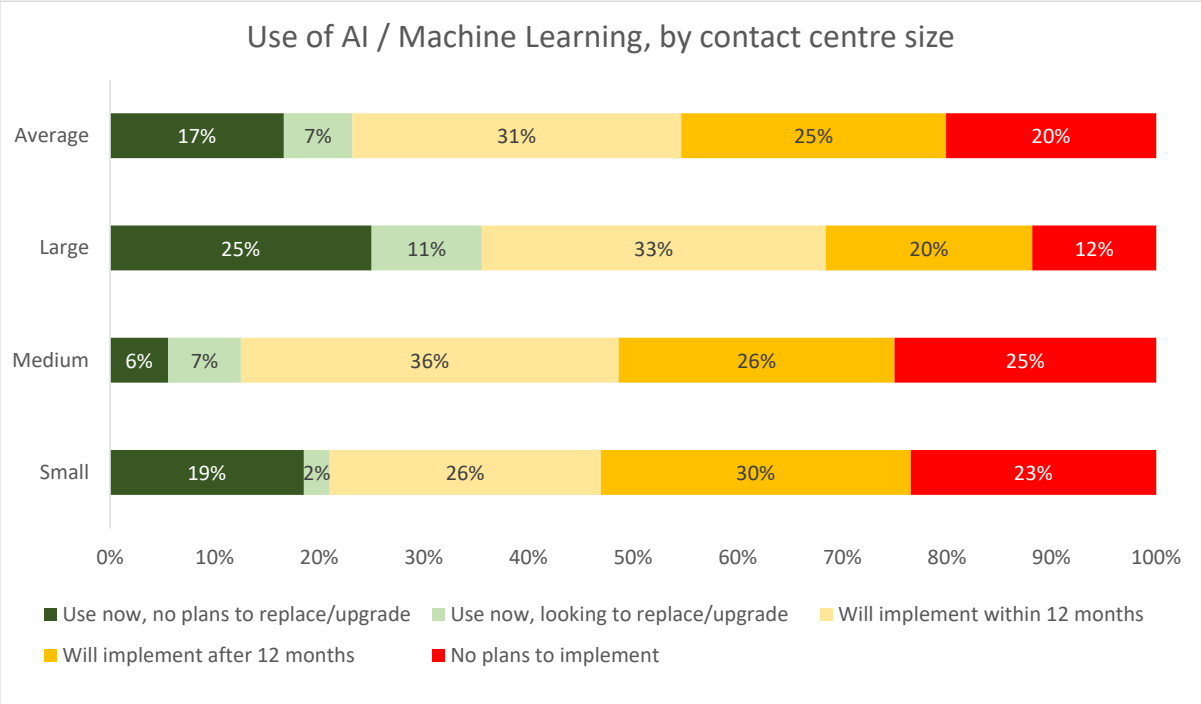

For a chatbot / voicebot implementation, what would success look like when considering the four business issues identified earlier in the report? What KPIs could be impacted, and what benefits seen?

Excessive customer wait times

Queue times are close to their historical peak, customers show no sign of leaving the voice channel and state that queue times are the no.1 driver of CX – whether good or bad.

AI-enabled chatbots / voicebots can impact queue times by encouraging the uptake of self-service, removing pressure on the queue. While rules-based chatbots have gone some way to soaking up some of the simpler queries, AI-enabled chatbots and voicebots can add further depth and encourage self-service by understanding customers' natural language and the context of the request, and responding conversationally, including turn-taking and further exploration of customers' issues.

KPIs impacted include: average time to answer; call abandonment rate; CSAT / NPS.

Uptake and effectiveness of web self-service is lower than desired

Despite most companies encouraging web self-service, many customers will abandon a self-service session if they have complex requests that they feel only a live agent can solve, want reassurance, or cannot find the answers or functionality they need on the website.

AI-enabled chatbots can encourage customers to stay on the website by asking sensible, context-related questions rather than just following a rigid rules-based script. Interacting conversationally and understanding customers' natural language will add reassurance and confidence, and being able to break up complex requests into multiple pieces while retaining the flow of the conversation will further encourage customers to interact more deeply with the chatbot.

This is not to say that all requests can or should be handled by AI-enabled chatbots, but that a greater proportion of them can be. It will take some time for customers to realise that not all chatbots are the same, and to invest time and effort in interacting with them rather than reaching for the phone, but the trajectory is in the right direction.

KPIs impacted include: % of chat interactions successfully handled; reduction in call volumes (and improvement in the associated phone queue KPIs); CSAT with chat functionality.

Voice self-service is rigid and difficult to use

While touchtone IVR has been around for decades, it is not particularly popular with customers nor is it generally used for much more than call routing or very simple self-service. Despite this, it is a real cost-saver for businesses, with around 5-7% of all interactions being handled through voice self-service across the entire contact centre industry, although these figures have been stagnant for many years.

AI-enabled voicebots can offer a deeper understanding of customer intent, helping route calls successfully without the need for customers to dig through long-winded IVR menus, and providing richer self-service functionality. Voicebots go beyond traditional speech recognition systems, which often act on keywords and tend to have a poor user experience.

The ability to understand and answer questions in multiple languages is also a strong positive point for many voicebots, which – like chatbots – respond to customers in a conversational style which can even be personalised depending on the customer and the brand they are communicating with.

KPIs impacted include: reduced voice self-service abandonment rates (zeroing-out), improved call routing (reduced call transfer rate); reduction in live call volumes (and improvement in the associated phone queue KPIs); CSAT with voice self-service functionality.

Customer authentication costs too much

At an industry-wide level, customer authentication by agents costs almost £2.2bn each year – over 58p per call. Quite apart from the cost, taking customers through the same security questions dozens of times each day is morale-sapping for agents, adds nothing to the customer experience and impacts negatively on queue times and call abandonment rates.

AI-enabled voicebots can carry out customer authentication just as quickly as an agent, with a much lower cost per interaction. AI-driven improvements in accent comprehension and noise cancellation mean that voicebots are able to understand more of what the customer is saying, reducing caller frustration and improving outcomes. It is possible for voicebots to gauge voice stress, sentiment and other indicators that suggest potentially fraudulent activity, and flag these to the agent for further security checks.

KPIs impacted include: reduced cost of customer authentication, reduction in call times (and improvement in the associated phone queue KPIs); improved agent morale.

The outcomes of using AI-enabled chatbots and voicebots include better customer experience as well as operational efficiencies and cost reduction.

An example of the latter is shown below where we have tracked mean average cost per phone call and web chat over a nine year period.

As can be seen, both phone and web chat costs rose until 2018. Phone costs kept creeping upwards, but web chat costs have diverged from that and are now quite flat at around £3, despite the average length and complexity of a chat increasing over that time.

While correlation does not prove causation, this chart seems to make a strong case for chatbots (whether rules-based or AI-driven) impacting positively on the cost of service, as the uptake in automated chat started at around the same time and has accelerated since.

Figure 18: Mean cost per inbound call & web chat 2015-2023

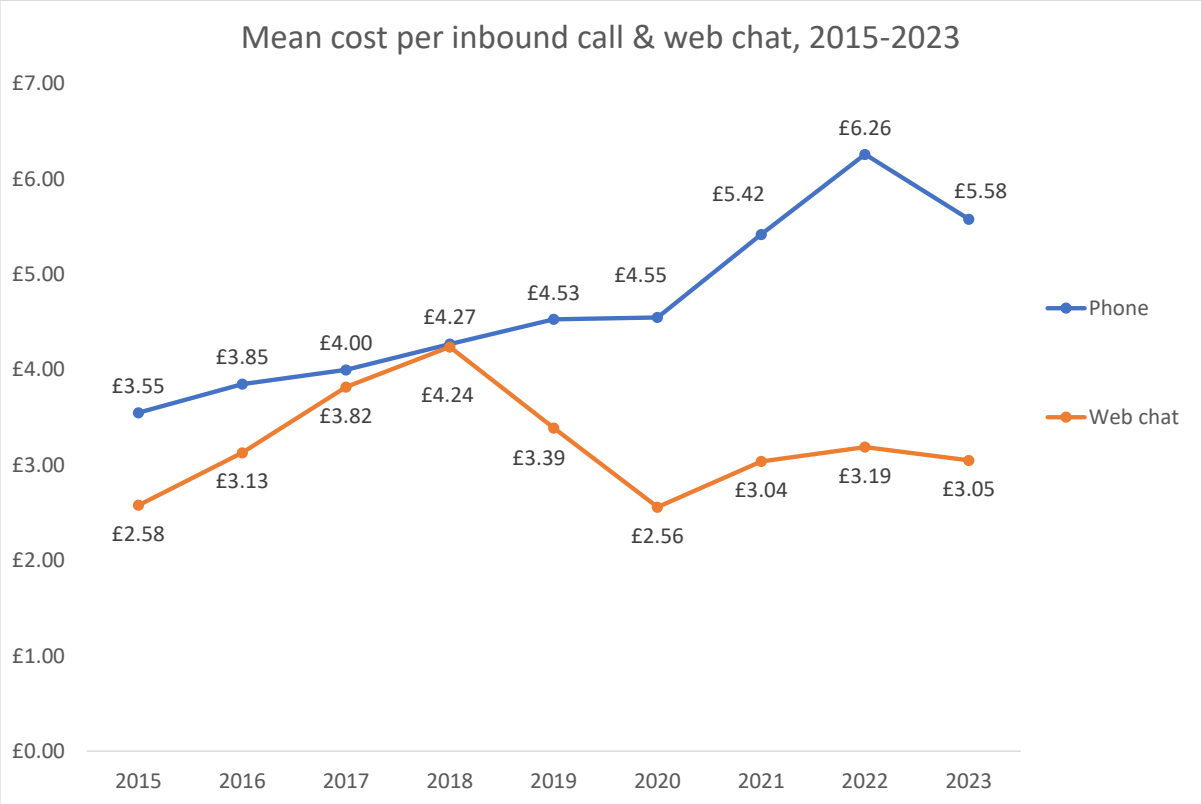

END-USER QUESTION #4: WHERE DOES THE ROI COME FROM IN AI-ENABLED CHATBOT / VOICEBOT IMPLEMENTATIONS?

AI-enabled chatbots offer a compelling return on investment (ROI) through a combination of cost savings and revenue generation.

On the cost-saving side, chatbots can significantly reduce the need for human intervention. They can handle routine inquiries and tasks, freeing up human agents to focus on more complex issues that require their expertise and judgment. This translates to reduced agent time and overall operational costs. Additionally, chatbots can offer 24/7 availability, answering questions and providing basic support around the clock, eliminating the need for extended call centre hours. Finally, by automating tasks and providing self-service options, chatbots can streamline processes, leading to increased efficiency and potentially lower processing costs.

AI chatbots also have the potential to generate revenue in several ways. They can qualify leads, answer product inquiries, and even recommend products, potentially boosting sales conversion rates. Furthermore, by providing a positive customer experience, chatbots can help build customer satisfaction and loyalty, leading to increased customer lifetime value. Finally, timely and efficient support provided by chatbots can help retain customers and reduce churn rates, ultimately resulting in higher overall revenue.

There seems little doubt that the eventual overall roadmap of the contact centre industry will lead to significant levels of AI involvement in all areas of customer contact. In the foreseeable future, AI is likely to be used to improve self-service and to assist agents, rather than having seismic effects on headcount.

An AI implementation whose success is to be measured mainly by the reduction in HR resource is unlikely to be a success from a customer's perspective. The combination of generative and conversational AI promises much from the perspective of providing a genuine-seeming conversation and lower set-up cost and effort, but deployment in a real customer environment may have risks, as recent high-profile generative AI incidents have proven.

Businesses' interactions with customers will continue to become a highly polarised mixture of the automated and the personal touch. Moving a large proportion of interactions onto self-service reduces business costs, and is increasingly popular with a customer base that is becoming more sophisticated and demanding in what it expects from self-service. AI takes this a step beyond, in an increasing number of cases offering personalised service without the need for a human agent.

Looking further ahead, we can expect to see personal technology applications seeking out the best deals on offer or actually interacting with a business without involving the customer at all. This leads to the conclusion that many customer-agent interactions will be exceptional, such as a complaint, an urgent or complex issue or a technical query that an FAQ or customer community couldn't solve. It is also likely that whole segments of the customer base who don't want automation at all will be handled directly by live agents in many cases.

Many self-service scenarios suggest a world in which customers speak directly to 'intelligent' systems, but an "e2e" world is becoming real, where systems talk directly to other systems without a human being involved at all. The customer will delegate many of their business interactions to an intelligent device, which will store information such as personal preferences, financial details and individuals' physical profiles. Customers will instruct the device to research the best deals for products and services, and to come back to the device's owner with the best selection. The personal AI would 'call' the relevant contact centre (which could in fact be either a AI or possibly a live agent in some cases) and even purchase the best deal without having to involve the owner in any way.

The same principle applies to customer service: using the 'Internet of things' means that, for example, utilities meters can send their own readings to suppliers on request, and a manufacturer can detect when a part on an appliance is about to fail, and organise a replacement part and engineer visit with the customer's permission.

At this early stage, most businesses might decide that implementing AI in a small scale on a clearly defined user case is the most appropriate action to take, building up their in-house knowledge and expertise while following a strategic implementation roadmap. Chatbots are usually the first move.

In the case of such a heavily hyped solution, expectations should be managed and care taken in identifying and forecasting the improvements that the initial AI implementation can bring, with the success of the project being clearly based around specific, easily understood metrics.

In the longer-term, there's no doubt that AI will be used as a key part of handling customer interactions in most businesses, but the question is: how? The use of AI should be focused on use cases where the AI does a better job than a human, whether that's being quicker, more accurate, available 24/7 or able to see patterns in data that no person could see.

It's our view that customers call live agents not necessarily because they want to hear a friendly voice, or that they're Luddites who won't countenance automation, but because they've found through personal experience that this is the most effective way of making sure their issue is resolved, even if there is significant effort involved.

So while AI-enabled automation will handle much of the simple work, customers will still seek out a live channel for complex or emotional interactions: probably voice, but perhaps digital or video too, as customer confidence in these channels builds up.

Yet even here, AI will be playing a part, identifying the customer's intent, gauging their sentiment, and understanding through past experience what the appropriate actions for the agent will be. Over a long period of time, AI will become thoroughly enmeshed in every element of customer interactions, but we are confident that the human element to customer interactions will remain vital in the times when it is most needed.

ABOUT CONTACTBABEL

ContactBabel is the contact centre industry expert. If you have a question about how the industry works, or where it's heading, the chances are we have the answer.

We help US and UK contact centres compare themselves to their closest competitors so they can understand what they are doing well, what needs to improve and how they can do this.

The coverage provided by our massive and ongoing primary research projects is matched by our experience analysing the contact centre industry. We understand how technology, people and process best fit together, and how they will work collectively in the future.

Email: info@contactbabel.com

Website: www.contactbabel.com

Telephone: +44 (0)1434 682244

Free research reports available from www.contactbabel.com (UK and US versions) include:

- The Inner Circle Guide to Agent Engagement & Empowerment
- The Inner Circle Guide to AI-Enabled Agent Assistance
- The Inner Circle Guide to Chatbots, Voicebots & Conversational AI
- The Inner Circle Guide to Cloud-based Contact Centre Solutions
- The Inner Circle Guide to Customer Engagement & Personalisation
- The Inner Circle Guide to Customer Interaction Analytics
- The Inner Circle Guide to First-Contact Resolution
- The Inner Circle Guide to Fraud Reduction & PCI Compliance
- The Inner Circle Guide to Next-Generation Customer Contact
- The Inner Circle Guide to Omnichannel
- The Inner Circle Guide to Omnichannel Workforce Optimisation
- The Inner Circle Guide to Outbound & Call Blending
- The Inner Circle Guide to Remote & Hybrid Working Contact Centre Solutions
- The Inner Circle Guide to Self-Service
- The Inner Circle Guide to the Voice of the Customer

- The Australia & New Zealand Contact Centre Decision-Makers' Guide
- The UK Contact Centre Decision-Makers' Guide
- The US Contact Center Decision-Makers' Guide
- The UK Customer Experience Decision-Makers' Guide
- The US Customer Experience Decision-Makers' Guide
- Exceeding UK Customer Expectations
- Exceeding US Customer Expectations

- UK Contact Centre Verticals: Communications; Finance; Insurance; Outsourcing; Retail & Distribution; Utilities
- US Contact Center Verticals: Communications; Finance; Healthcare; Insurance; Outsourcing; Retail & Distribution.