



The Inner Circle Guide to AI, Chatbots & Machine Learning

Sponsored by



The Inner Circle Guide to AI, Chatbots and Machine Learning (2nd edition) – UK edition

© ContactBabel 2021

Please note that all information is believed correct at the time of publication, but ContactBabel does not accept responsibility for any action arising from errors or omissions within the report, links to external websites or other third-party content.

CONTENTS

Contents 3

Table of Figures..... 6

About the Inner Circle Guides 8

AI: Definitions and Terminology 9

 Definitions..... 11

 AI (Artificial Intelligence)..... 11

 Chatbots / Virtual Agents / Virtual Assistants / Conversational AI 11

 Voicebots 12

 Machine Learning / Deep Learning / Neural Networks 12

 Natural Language Processing / Understanding (NLP/NLU) 13

 Current Use and Future Plans for AI 14

Use Cases for AI in the Contact Centre 16

End-User Question #1: Does AI only really make sense for large contact centres? 19

 Self-Service 20

 Web Self-Service and Chatbots..... 22

 IVR, Speech Recognition and Voicebots..... 28

 Email Assistance and Automation 32

 Assisted Service 34

 Digital Interaction Assistance..... 34

 Voice Interaction Assistance 35

 Sentiment Analysis..... 36

 Predictive / Intelligent Routing 38

 AI & Remote Working 39

End-User Question #2: How can AI help with remote working? 40

 Robotic Process Automation (RPA) & AI 41


 Interaction Analytics 46

Discovery	46
Customer Journey Analytics	47
Identity Verification and Fraud Reduction	47
Workforce Optimisation	49
Real-Time Analytics	50
Improving Routing and Outcomes	52
AI and CX Metrics	52
Implementing and Using AI	53
End-User Question #3: How can AI best be used for a quick win? Where can ROI be found?	58
Initial Actions in an AI Project	59
End-User Question #4: What are the steps we need to take to use AI in our contact centre? Where would be the best place to start?	60
Knowledge Management	61
End-User Question #5: How much initial and ongoing effort/resource will AI require? Do we need a dedicated AI professional to keep everything running?	64
Agent Replacement or Augmentation?	65
Customer Opinions of AI	66
Making AI a success and avoiding pitfalls	68
End-User Question #6: Is there anything that successful AI implementations / projects have in common? Any pitfalls to avoid?	70
The View from CX Decision-Makers	71
The Future Role of AI in the Contact Centre	72
About ContactBabel	74
Further ContactBabel reports	74



CallMiner leverages AI and machine learning across voice, chat, email, social & more.

Your daily contact centre interactions are a goldmine of operational and customer experience insights just waiting to be unearthed. CallMiner's Eureka platform brings these to the surface for you with the help of AI and machine learning. Gain an unparalleled view into both voice and text-based conversations, identify and measure CX through not only the words spoken, but the manner in which they are spoken. Discover the answers to important questions you would have never thought to ask!



Improve your visibility into, and understanding of Voice of the Customer.

- Analyse 100% of your interactions across channels
- Identify, classify, measure & analyse VoC
- Gain insights about CX, product, sales, marketing & more
- Turn those insights into actions to improve enterprise performance

Take a **FREE** CallMiner analytics test drive starting with your own phone calls!

Visit us at: <http://callminer.com/free-speech-analytics/>



Download with our
Compliments

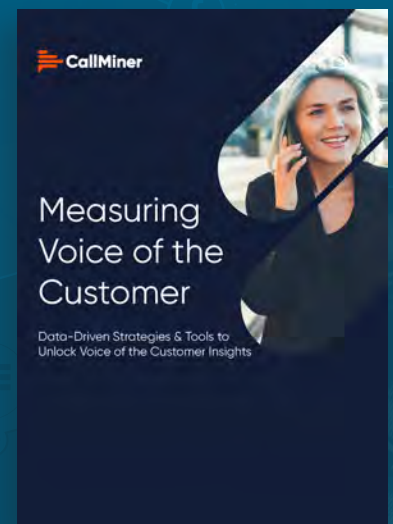


TABLE OF FIGURES

Figure 1: Use of AI / Machine Learning, by vertical market 14

Figure 2: Use of AI / Machine Learning, by contact centre size 15

Figure 3: Preferred method for contacting a company (high urgency interaction), by age range 20

Figure 4: Cost per inbound interaction 21

Figure 5: Web self-service methods, by contact centre size 22

Figure 6: Proportion of callers that have tried to answer own queries through web self-service before calling..... 23

Figure 7: Why customers move from web self-service to live telephony 24

Figure 8: Level of automation used in web chat, 2015-20 25

Figure 9: Web chat & inbound call lengths – a comparison 26

Figure 10: Overall proportion of calls handled entirely through self-service (only in respondents which offer telephony self-service)..... 28

Figure 11: Proportion of self-service sessions ‘zeroed-out’ to an agent 29

Figure 12: Reasons for abandoning self-service sessions..... 30

Figure 13: What proportion of emails are answered successfully and completely within these timescales? (2009-20)..... 32

Figure 14: Level of automation used in email management (2015-2020)..... 33

Figure 15: Human and AI email and web chat handling..... 34

Figure 16: Proportion of agents homeworking, 2019 - 2021 39

Figure 17: Selected performance metrics, by number of in-call applications used..... 44

Figure 18: Usefulness of real-time analytics 51

Figure 19: In-call access to computer-based knowledge sources for agents 63

Figure 20: Views on the role of artificial intelligence in the contact centre 65

Figure 21: Views on how customers will perceive artificial intelligence in the contact centre 66

Figure 22: Would you prefer to speak with an agent or use automation, if the outcome and time were identical? (by age range)..... 67

Figure 23: Importance of CX developments in the next 2 years 71

 **CallMiner**

CallMiner is a speech analytics platform that drives business improvement by connecting insight to tangible action. We use the power of A.I. to scale human understanding, analyse interactions at the deepest levels, identify patterns and root causes, and reveal opportunities.

We believe that business improvement starts with a deep interest and curiosity in people. How do we detect a customer's true emotion, and how do we act on it to shift a business's culture and steer it on the most successful path? We aim to close these gaps through innovation, but the heart of our work lies in humanity: understanding, followed by action. We apply this same principle within our culture, promoting an attitude of kindness, compassion, genuine interest, and respect for one another.

Highlighted by multiple customer achievement awards, including six Speech Technology Magazine's Reader's Choice Awards, CallMiner has consistently ranked number one in customer satisfaction.

Learn more about our customer engagement and speech analytics solutions to help your business:

Contact:

w: callminer.com

[LinkedIn](#) / [Twitter](#) / [Facebook](#) / [YouTube](#) / [Blog](#)

User Community: EngagementOptimization.com

e: marketing@callminer.com

ABOUT THE INNER CIRCLE GUIDES

“The Inner Circle Guide to AI, Chatbots & Machine Learning (2nd edition)” is one of the Inner Circle series of ContactBabel reports.

Other subjects include:

- Cloud-based Contact Centres
- Customer Interaction Analytics
- First-Contact Resolution
- Fraud Reduction and PCI DSS Compliance
- Omnichannel
- Outbound & Call Blending
- Remote Working
- Self-Service
- Voice of the Customer
- Workforce Optimisation.

They can be downloaded free of charge from [here](#).

The Inner Circle Guides are a series of analyst reports investigating key customer contact solutions. The Guides aim to give a detailed and definitive view of the reality of the implementing and using these technologies, and a view on what the future holds.

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, and we have selected six of the most popular to ask to the report’s sponsor. The answers to these are distributed throughout the report and give interesting insight into real-life issues.

Statistics within this report refer to the UK industry, unless stated otherwise. There is a version of this report available from the sponsor with equivalent US statistics.

“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: DEFINITIONS AND TERMINOLOGY

Within the customer contact space, there is a great deal of interest in how artificial intelligence (AI) can work to deliver a superior customer experience at every hour of the day, across channels, leveraging the vast amounts of data that are available to many large organisations. Supported by the speed and availability of affordable processing power, and the enormous amount of structured and unstructured data available, the opportunity exists for AI to take customer contact far beyond what is feasible now.

Although we are still in the early days of the AI revolution, there are already numerous well-known examples widely used by the public, including Amazon's Alexa and Apple's Siri. These virtual assistants 'understand' unstructured natural language requests and deliver the solutions in a manner similar to a live personal assistant.

As AI can be given access to all of the relevant data a company holds on its customers, as well as unstructured data held elsewhere (for example, forums or social media channels), it has a far wider source of knowledge from which to draw, compared to human agents. In theory, an AI with sufficient sophistication could make human agents all but unnecessary, but for the foreseeable future, AI will usually work alongside its human colleagues.

The usage of the term 'AI' in the contact centre covers an enormous area, and has often been used by solution providers, media and businesses to refer to functionality that may only very tenuously be said to be linked to true AI, which is itself a wide-ranging term for technology solutions which appear to emulate human cognitive capabilities through the 'understanding' of complex, natural language requirements, in order to reach its own conclusions and improve itself.

Rather than arguing about semantics, the umbrella term of AI will be used descriptively rather than prescriptively within this report. Its use within the contact centre will be linked to four broad types of linked functionality – the "4 A's of AI" – analysis, anticipation, augmentation and automation.

Analysis:

Whereas for humans, enormous, fast-changing datasets make understanding and action more difficult, AI requires extremely large sets of data in order to find patterns and work optimally. Tools such as speech-to-text and optical character recognition (OCR) enable the AI to normalise data and compare like with like, and machine learning allows systems to improve accuracy and the effectiveness of outcomes without constant input and tweaking from human users. This self-learning element is a key part of what makes some solutions truly artificial intelligent.

Anticipation:

Based upon the customer's history, the context of the interaction, and the factors influencing successful outcome of similar interactions in the past, AI will be able to predict the next best action to take. This may be in the form of an answer taken from the knowledge base, the correct prioritisation and routing of a call, or the prompting of an agent to ask a specific question or make a relevant sales offer.

Augmentation:

The AI is able to gather relevant information from numerous sources in real-time in order to provide enhanced information to human agents or the self-service system, increasing the likelihood of a successful outcome. The AI is also tasked with updating relevant systems and initiating the correct business processes.

Automation:

In circumstances where there is a high level of confidence that the solution presented by the AI is correct, human intervention may be circumvented altogether. The AI system may monitor the interaction in real-time, using sentiment analysis to determine whether there is a need for a live agent to collaborate.

DEFINITIONS

AI (ARTIFICIAL INTELLIGENCE)

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 contact centre, 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.

CHATBOTS / VIRTUAL AGENTS / VIRTUAL ASSISTANTS / CONVERSATIONAL AI

As with so much in the world of AI, there is disagreement about definitions. In the case of chatbots, virtual agents, virtual assistants and conversational AI, it is better to focus on the functionality and ‘intelligence’ powering it, rather than the phrase in itself.

AI for customer contact is perhaps currently best known for chatbots, applications that run automated tasks and simulate conversation with the customers. It may be given a human avatar and personality characteristics, and includes 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 are often found in the web chat channel, but the functionality can be used in any other digital channel, such as social media, email or even voice self-service (in the form of voicebots).

Chatbots are not always fully AI-enabled, and these are often 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, which can have a significant impact on customer care without requiring large investments or long implementation timescales. Other chatbots use NLP and can ask questions to understand customer intent and improve the accuracy of the output, and may also use machine learning to improve future outcomes.

In this report, “chatbots” and “virtual agents” are used interchangeably and refer to the same functionality. Virtual assistants (VAs) are not dedicated to a single task (such as customer service), and can assist in numerous ways such as taking notes, carrying out web research, setting alarms, communicating with smart devices, etc.

Both chatbots and VAs are conversational interfaces, but the level of AI involved can differ greatly.

VOICEBOTS

Like a chatbot, 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 authentication.

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 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, with a 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.

CURRENT USE AND FUTURE PLANS FOR AI

Despite a low current use of AI across industries, there is widespread interest in implementing this solution, with 35% 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 utilities, insurance and outsourcing respondents report the greatest current use of AI.


Figure 1: Use of AI / Machine Learning, by vertical market



As AI works most effectively with very large pools of data to learn from, it is of no surprise to see large contact centres being at the forefront of implementation. These operations are also the most likely to have large numbers of self-service sessions which can be optimised through AI, although there is very significant interest even amongst smaller contact centres.

It should be noted that respondents from smaller operations may consider their use of chatbots to qualify as being AI, although this is not necessarily the case.

Figure 2: Use of AI / Machine Learning, by contact centre size



USE CASES FOR AI IN THE CONTACT CENTRE

There are numerous use cases for AI in contact centres, and these are growing all of the time. In order to structure this report, we have chosen four wide areas in which AI is actually helping customer contact now:

1. Self-Service
2. Assisted Service
3. Interaction Analytics
4. Robotic Process Automation and AI.

There can often be overlap between these areas, for example when a self-service interaction is passed to an agent, who then hands off to robotic process automation.

It is important that the various elements comprising AI should not be viewed in isolation. When they work together, the benefits can be much stronger and unlike siloed applications, AI's capabilities can be made available to any channel, whether live or automated.

For example, AI-enabled discovery may find that a specific demographic group will purchase multiple products from a supplier. This could change the company's strategy to target this demographic with cross-selling and upselling offers. When a customer from group calls or is successfully offered web chat, the AI can route this interaction to an agent that it has calculated has the right experience and capabilities to cross-sell successfully. The agent is guided through the conversation by the AI, which follows the dialogue while looking at what has worked best with customers of this type previously, personalising offers without being intrusive to the agent. The outcome is analysed and used by machine learning to develop future customer identification, routing and agent support.

It should also be noted that some of these use cases can also apply to non-contact centre staff. For example, staff in physical stores can be given access to AI reasoning capabilities in order to answer technical questions, and can even be used for internal staff to answer their questions on products or services.

Example Self-Service Use Cases

- Replace or augment IVR which is rigidly structured
- Website navigation assistance: help customers around the website and provide answers to FAQs. Feedback from customers can help to measure the success or otherwise of the AI. A path to live assistance should always be offered
- Develop and support true omnichannel capabilities: ContactBabel research has consistently shown that siloed systems and processes are the single largest factor holding companies back from being able to offer a true omnichannel experience to their customers. AI encourages the decoupling of data from the channel of engagement as it draws data from many sources in order to recognise patterns and act accordingly, with AI's capabilities being used on any channel, whether in a self-service environment or through assisting agents
- Offer 24/7 self-service assistance

Example Assisted Service Use Cases

- Use real-time analytics on calls to improve the quality of interaction, agent behaviours and customer outcomes
- Identify cross-selling and upselling opportunities (this can be used in either an assisted or automated environment)
- Augment agent capabilities in real time to improve outcomes
- Use sentiment analysis in real-time to support outcomes.

Example Interaction Analytics Use Cases

- Recognise patterns of customer behaviour: e.g. use sentiment analysis to analyse large blocks of data and show that customers trying to carry out a specific action are more likely to be stressed and annoyed, allowing the business to investigate and reengineer the process
- Customer behaviour predictions: look for which customers are defecting, and identify any patterns in their behaviour or experiences. Learn which actions are most successful at keeping them loyal
- Predictive routing, inc. CTI-like information popping to allow better matches and outcomes
- Monitor, quantify and improve agent performance through post-call analytics
- Predict the actions and requirements of the customer, improving cross-selling / upselling opportunities
- Discover areas for improvement and automation in association with RPA
- Use customer journey analytics to identify bottlenecks and processes which frustrate customers and create unnecessary callbacks
- Improve identity verification and customer authentication and decrease fraud
- Improve workforce optimisation and training
- Quality assure 100% of calls and predict QA and CX metrics (e.g. NPS)
- Share intelligent real-time alerts with agents
- Predicted CX and QA metrics

Robotic Process Automation Use Cases

- Use RPA/AI to populate fields and take some tasks away from agents within the call
- Use RPA/AI for post-call wrap-up and to kick off and monitor back-office processes.



CallMiner

Intelligence
from Customer
Interactions

How AI, ML and Automation Can Help Humans Deliver Success in the Contact Centre

As businesses tackle mounting organisational challenges, in terms of efficiency, scalability and more, many have adopted AI, ML and automation technologies to level the playing field. But where do humans fit in?

The emergence of artificial intelligence (AI) and machine learning (ML) has played an important role in improving organisational efficiency in a variety of ways. AI offers up a wealth of productivity-enhancing features fit for use in organisations of all sizes, and when it's applied to an individual business' use case it can reveal unique functions based on specific industry, vertical and more.

Contact centres are one area of a business, in particular, that can **benefit from AI and ML technologies**, including streamlining internal processes. Offering AI-enhanced processes and recommendations to contact centre agents does more than make their work easier and help them be better at their jobs, it also enhances the customer experience through improved speed in resolution and outcomes. In fact, **Accenture predicts AI will increase business productivity** by over 35 percent before 2040 in the US alone.

So, it's no surprise that investments in AI across the contact centre industry are on the rise. **A recent report** estimates that the call centre AI market will grow to \$2.8 billion by 2024, an increase from \$800 million in 2019. What's more, **Gartner predicts** that 50% of enterprises will spend more annually on **chatbot** development compared to traditional mobile app development by 2021. "In the 'post-app era,' chatbots will become the face of AI and bots will transform the way apps are built," Gartner explains. "Traditional apps, which are downloaded from a store to a mobile device, will become just one of many options for customers."

The potential for AI at organisations and in contact centres extends far beyond AI-driven chatbots. Identifying where AI fits into your organisation comes down to assessing your own internal operations and determining which processes would benefit most from AI tools and automation.

The most effective place to start with AI in the contact centre is usually in areas where you have manual processes that are repeatable – it's easier to automate what you already know. In a contact centre, this could be around agent performance and training, quality assurance, data capture and more.


However, for AI to deliver the fullest benefits to organisations, employees and the customers they serve, it's critical to consider **how humans and AI work together**.

AI works best when it plays a supportive role for contact centre agents. That's because today's AI technologies aren't capable of navigating the complexities of human conversations entirely unsupervised. Instead, AI is the most successful when it improves how agents engage with customers, such as helping agents in identifying next-best responses or when to route a customer to a supervisor. While there are places where AI can stand alone, including in self-service instances where customers completing repeatable tasks, like checking a balance, the real power comes from when AI enhances human interactions, not replaces them completely.

Creativity, **empathy** and spontaneous judgment are still human traits, and with the support of AI, organisations can ensure their entire agent base can focus their attention where it's needed most in real time during customer interactions.

AI, ML and automation technologies have the potential to revolutionise the way businesses develop and adapt to shifting consumer trends. When implemented correctly – by managing repeatable tasks and supporting humans in more complex interactions – organisations have the opportunity to increase contact centre efficiency, make agents more scalable, improve customer experience and achieve significant ROI.

END-USER QUESTION #1: DOES AI ONLY REALLY MAKE SENSE FOR LARGE CONTACT CENTRES?

 **CallMiner** No. Small contact centres and large contact centres will likely have different problems to solve with AI, but there are many AI approaches that are just as successful regardless of the size of the contact centre.


SELF-SERVICE

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 10-15% of customers, 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.

As the following chart shows, self-service is rated particularly highly for urgent interactions (as opposed to complex or emotional interactions, where the personal touch still takes precedence).

Figure 3: Preferred method for contacting a company (high urgency interaction), by age range



New channels such as social media, email and web chat have grown rapidly in popularity, yet the large majority of interactions involving all of these channels still require a customer to make a request to a human agent.

Although web chat, social media interactions and emails tend to have slightly lower costs than telephone calls, the differential between these is far smaller than between live and automated service.

Figure 4: Cost per inbound interaction

Channel	Mean	1st quartile	Median	3rd quartile
Phone	£4.55	£4.50	£3.49	£2.89
Email	£4.10	£4.30	£2.89	£1.70
Web chat	£2.56	£3.12	£2.00	£1.10
Social media	£2.37	£3.80	£2.40	£1.20
IVR	c. 30-50p	-	-	-
Web self-service	c. 5-10p	-	-	-

Of course, not only are businesses missing out on huge potential cost savings, but one of the main customer experience problems still exists: that of having to wait until an agent is available to answer the query.

Expanding the boundaries of self-service outside the simplest and least ambiguous requests is one of the main challenges for most contact centres. Success in this will mean not only greatly reduced costs for businesses, but also improved customer experience through higher real first-contact resolution rates through the customer's channel of choice.

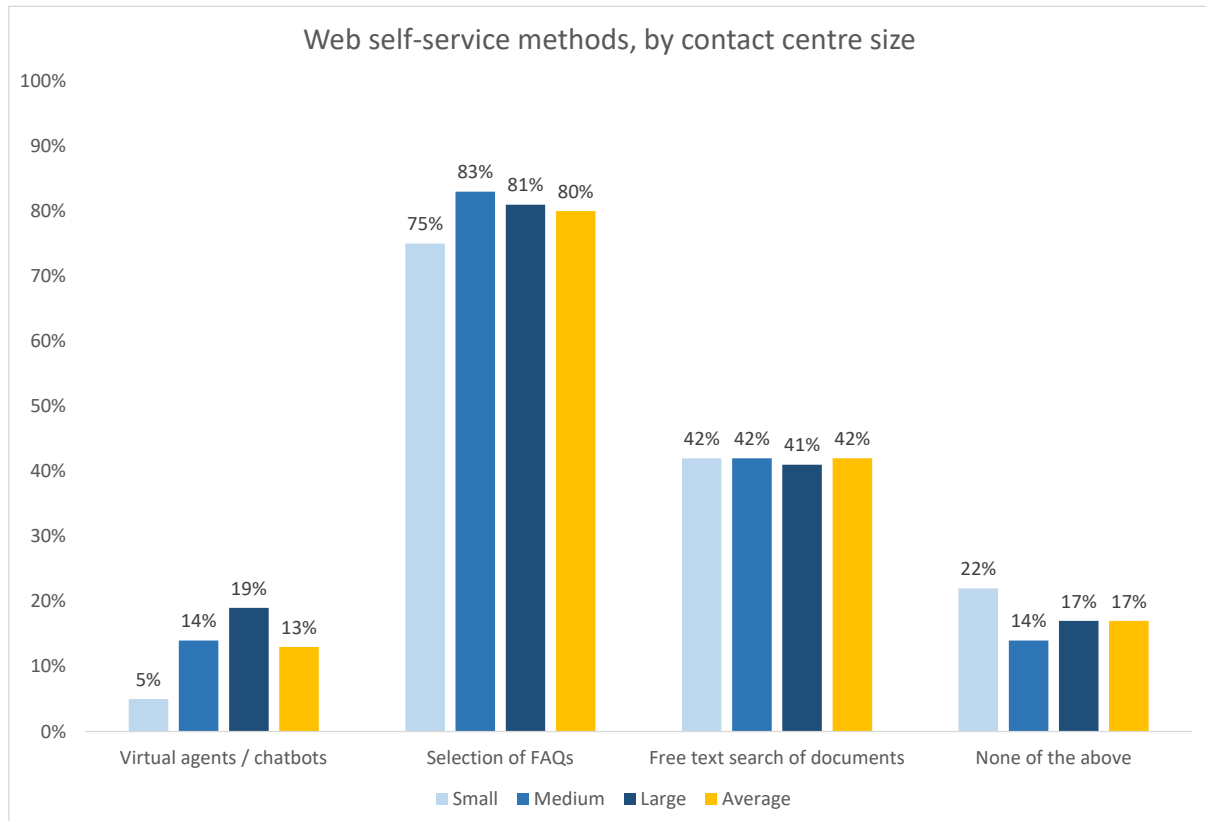
Self-service can provide many advantages:

- Significant cost savings
- Redirected demand frees agents for more complex work
- Greater agent satisfaction and less repetitive work
- 24/7 coverage
- Disaster management capabilities
- Handling of large volumes of similar requests for information e.g. power outage
- Acting as triage, meaning that bots and agents can work individually and together on the interactions that are most suitable for each.

WEB SELF-SERVICE AND CHATBOTS

Most businesses will use at least some form of web self-service. By far the most prevalent form is that of the FAQ (frequently-asked question), which is used by 80% of 2020's survey respondents. Free text search of the document library is somewhat less well supported, at 42%. Virtual agents (chatbots) are employed by 13% of respondents (up from 9% in 2019), more often those within large enterprises. 17% of respondents still offer no web self-service at all.

Figure 5: Web self-service methods, by contact centre size

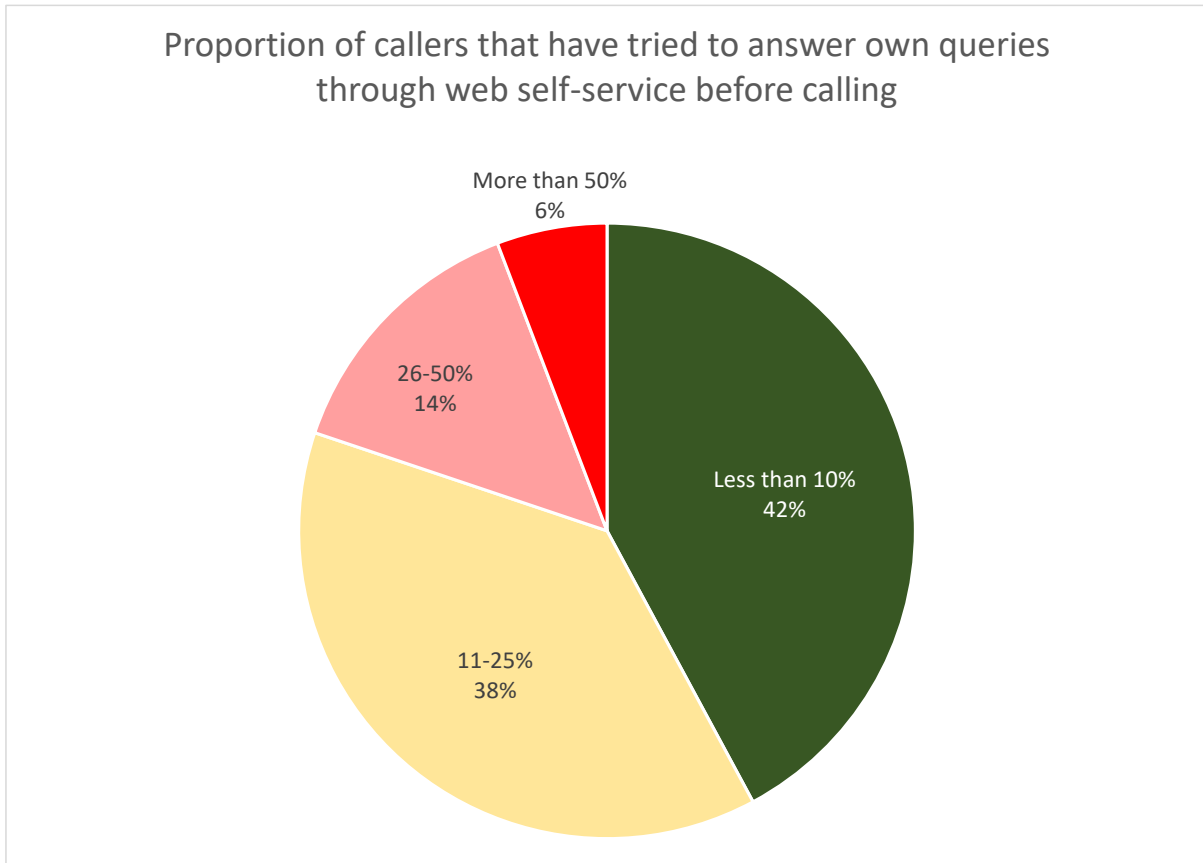


For many businesses, the customer is given free rein to search through documents, pre-written answers and archives, hoping to stumble across the right answer for themselves. This often proves time-consuming and ultimately frustrating for the customer, who will then go elsewhere or call the contact centre, already with a negative sentiment. An AI guide would be a valuable aid in improving CX and deflecting unnecessary calls.

Although 42% of survey respondents state that fewer than 10% of customers have tried to resolve issues online before calling the contact centre, 20% 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.

Although not shown here, 30% of respondents using web self-service do not have any idea of its success from the customers' perspective.

Figure 6: Proportion of callers that have tried to answer own queries through web self-service before calling



Machine learning can analyse the patterns and behaviours of website visitors, which is particularly useful in determining any differences between what eventual buyers do compared to those who abandon the site, as well as noting what customers were trying to do before they gave up and called the contact centre. This allows targeted and personalised offers and information through the right channel to occur at the right time.


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.

91% of respondents also felt that customers wanted the reassurance that a live agent brings to a conversation.

59% stated that the functionality that the customer calling in required was not available online, but interestingly, 55% 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, or the implementation of a chatbot or other guide to help them find the right answer would pay dividends.

Relatively few respondents believed that the need to take customers through phone security was an issue in receiving inbound calls, with the wide availability of password-protected online accounts and easy automated password resetting being a factor in this.

Figure 7: Why customers move from web self-service to live telephony




AI and chatbots can assist with some of these issues, especially assisting the customer to find the right information through conversation and questions. Putting agent-handled web chat on the website allows businesses to collect the data with which to train the chatbot, and for businesses to understand fully the processes and steps required for the large variety of questions and actions that customers may require.

Perhaps the most obvious potential use of AI/chatbots in the customer contact environment is in handling digital enquiries, where web chats generally take far longer than phone calls (due to agent multitasking, and typing time) and some email response rates can still be measured in days.

As the cost of web chat is not dissimilar to other channels such as email, voice and social media, there is still considerable room for increasing efficiencies and lowering costs through automation.

Whereas only 5% of web chats had any automation involved in 2015, this had grown to 24% in 2019, mainly as a result of initial handling by automated chatbots which may then hand off to live agents where appropriate. This figure dropped slightly to 22% in 2020, but the proportion of chats being handled entirely by AI / chatbots has risen from 4% to 7% over the past 12 months, meaning substantially lower costs.

Figure 8: Level of automation used in web chat, 2015-20




It would perhaps be a mistake to compare the content of web chats directly to that of phone calls: web chats tend to be about single, simple matters, whereas phone calls are often reserved for complex or multiple issues, and far more can currently be achieved in a five-minute phone call than a five-minute web chat. Non-AI chatbots can be of real benefit in answering simple but time-consuming customer enquiries and often do not require a major amount of time or investment.

Some solutions offer chat agents the opportunity to see what the customer is typing in real time, and enabling the agent to get a head start, while at the same time linking to the contact centre knowledge base in order to provide a list of most likely answers, which will increase the accuracy of response and decrease the overall time to serve.

Further comparing the experience of web chats with telephone calls, the survey finds that 41% of web chats take longer than 5 minutes to complete fully, compared to 65% of phone calls. While agent multi-tasking and the time taken to type differs from the experience of handling a phone call, web chat is still often a shorter experience for many customers. This may well be because the subjects of web chats will tend to be simpler than phone calls.

It is worth noting that the proportion of web chats handled entirely by automation has increased from 4% in 2019 to 7% in 2020, and that there has been a significant drop in the average cost of web chats in the past 12 months. It will be interesting to see if this becomes a pattern in future years.

Figure 9: Web chat & inbound call lengths – a comparison



The most sophisticated chatbots or virtual agents encourage the visitor to engage with them using natural language, rather than keywords. The virtual agent will parse, analyse and search for the answer which is 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. 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 couple of quite 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.

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 its brief by providing a list of related questions and answers to the original question, using machine learning to predict what the next question may be and provide this answer as well.

Solution providers and users train the system to pattern-match the right words or association of words with the correct result: the application, unlike older forms of web search techniques, does not simply guess what the customer wants, or how they will express themselves. 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.

Virtual agent functionality 'understands' the context of what the customer is asking, with the result being 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.

When the virtual agent application (whether AI-enabled or not) has low confidence that it has returned the correct result, it is able to escalate the customer's 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 automated agents.

Virtual agent functionality is of interest to most sectors, however the commercial reasoning and business drivers differ greatly. Banks have an appreciation that they need to understand their customers to keep them loyal in a highly commoditised and competitive environment, and as such there is considerable interest in using virtual agent functionality within Voice of the Customer initiatives. For example, using real-time analytics, such organisations can learn that customers are talking about a specific issue, which can feed into wider commercial decisions in business areas unconnected to customer service. Other sectors which are heavily focused on cost reduction may focus business cases on contact avoidance. Online retailers, which want to cross-sell and reduce their shopping cart abandonment rates, will have yet another strategy. Chatbots should be given their own "job description", in the same way as every agent will have one. It should be clearly defined and be linked to a desired business goal (e.g. a 10% reduction in agent-handled chats) so its success or otherwise can be quantified.

There are various levels of extra functionality that can be used to support webchat more effectively. If the customer has logged in, it is possible to identify them, and take into account past channel preferences, purchase history and other relevant information in order to personalise the experience, (for example including details of relevant offers to that customer).

It's important to reiterate an earlier point: not all chatbots or virtual agents are powered through AI and machine learning as many use programmer-defined rules and scripting in order to retrieve answers from a knowledge base. These types of chatbot can have great value to businesses, especially where there are a relatively small number of questions being asked by customers, and they can be very cost-effective too.

IVR, SPEECH RECOGNITION AND VOICEBOTS

A typical IVR solution works on the basis 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 complexity to the requirement. Even in the latter example, banks have commented that analysis of these types of call show that customers ask for this information in thousands of different ways, which throws up significant problems for non-AI speech recognition. Although touchtone IVR reduces the number of options and thus ambiguity, it is by its nature only useful in a limited number of calls.

Unlike touchtone IVR where an specific input is acted upon exactly the same way, natural language processing and machine learning techniques allow AI-enabled speech recognition to understand customer intent and give different outputs depending upon what has been proven to be successful in the past.

Of those contact centres offering telephony self-service, a mean average of 29% were handled entirely by self-service without requiring an agent, which shows that telephony self-service is used frequently by customers, although the majority of calls to these contact centres will currently still require an agent.

Figure 10: 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 <u>if offered</u>	
1 st quartile	40%
Median	20%
3 rd quartile	10%
Mean	29%

Many calls are not currently suitable for self-service, as they may require multiple requests within the same call, be of a complex nature or be from a caller who feels for whatever reason that they need to speak with a person. Additionally, some small businesses may have such a low volume of calls that it is not cost-effective to implement self-service.


Even amongst those survey 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 such 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 go elsewhere.

It is worth reiterating that 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 can frustrate callers more than being hectorred into using an unhelpful and irrelevant self-service system.

Overall, a mean average of 13% 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.

NB, 1st quartile performance for 'zeroing-out' is 3%, the median is 5% and the 3rd quartile is 18%.

Figure 11: Proportion of self-service sessions 'zeroed-out' to an agent

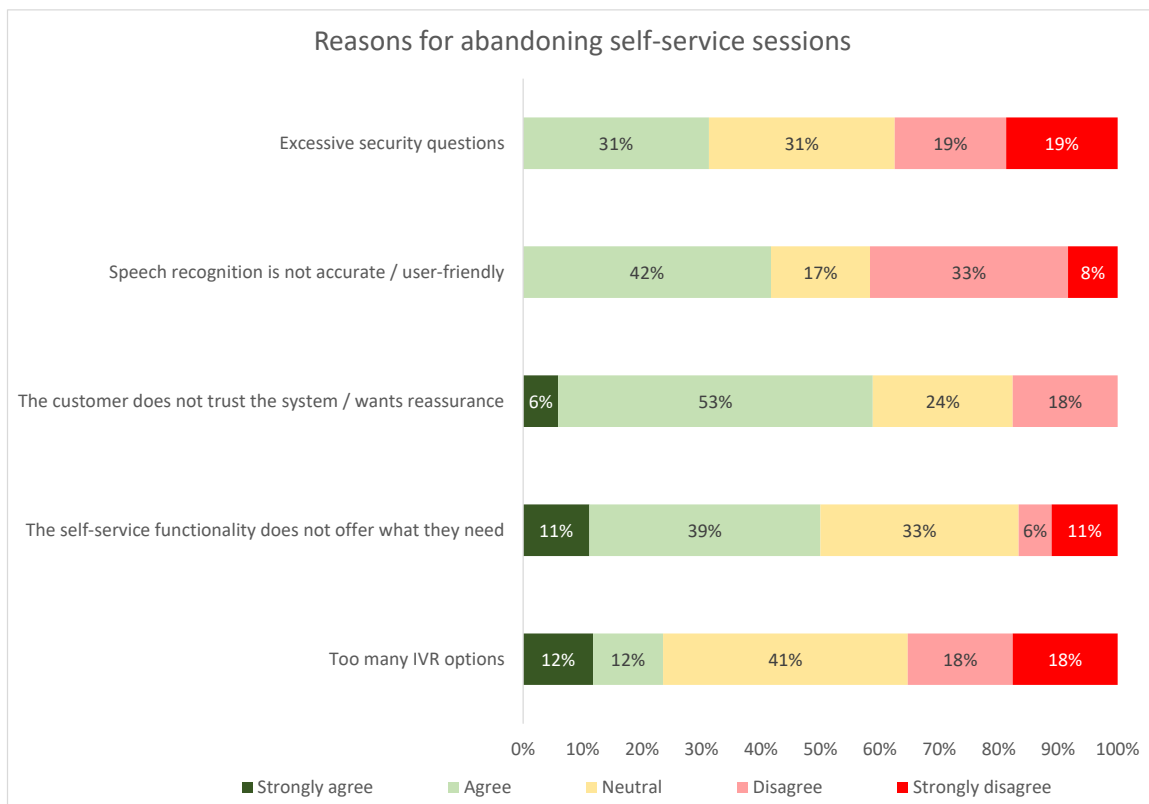


By a considerable margin, respondents agreed that the main reason for abandoning self-service sessions was that the self-service function simply does not offer what the customers want, with 80% of respondents stating that this is a factor. While this at first glance may appear negative, it is the case 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.

While 12% of respondents strongly agree that having too many options presented to customers is a major reason for them seeking human assistance, this is not widely seen as being an issue. It is noticeable that 59% of respondents agree to some extent 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.

Using AI-enabled natural language recognition can alleviate the first of these issues, as there is no fixed menu to navigate and no limit to the number of options a customer has to explain their issue. The onus is placed upon the system to understand the customer’s intent, rather than forcing the customer to shoehorn their request into a format allowed by the predefined rules and format of the business. Using NLU-enabled voicebots will go a long way to improving the effectiveness and user-friendliness of telephony self-service.

Figure 12: Reasons for abandoning self-service sessions



A voicebot should 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 voicebot interaction costs around 5-10p per call, compared to over £4 for a live agent
- Providing 24/7 service
- Create far wider opportunities for extending self-service capabilities
- Taking pressure off agents, particularly in times of crisis or volume surges
- PCI DSS compliance can be easier, as live agents are not required to handle card payments
- Encouraging customers to use natural language rather than to press numbers on a keypad generates large amounts of data that can be used to further train AI models
- 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
- Voicebots can be used for outbound work as well, such as debt collection, reminders and surveys. It should be noted that while customer response to obvious robocalls is generally negative, some voicebot solutions have an extremely convincing humanlike voice
- Personalised IVR allows greeting by name, and a change of voice / speaking style depending on customer preference. Authentication also possible.
- Using IVR as triage to decide who gets automation and who agents, based on customer profile, requirements, circumstances and past outcomes.


EMAIL ASSISTANCE AND AUTOMATION

For businesses that handle substantial volumes of email, while it is not suggested that they should aim to answer an email in the same amount of time that it takes to complete a phone call, it is desirable to manage all interactions closely to consistent business rules, and to act quickly if service levels slip.

In 2015, reported email response handling times reversed the improvements of previous years, especially in the all-important ‘less than 1 hour’ segment. In 2020, the proportion of emails answered the same working day is only 56%, while those taking more than 1 day has grown to 38%.

Taking longer than one day to answer an email runs the risk of the customer losing patience, and going elsewhere or phoning the contact centre, placing a greater cost burden on the business than if they had just called in the first place. Clearly the pressure placed on businesses as a whole and the phone channel in particular in 2020 does not help, but this chart shows that email response rates have shown little improvement for many years.

Figure 13: What proportion of emails are answered successfully and completely within these timescales? (2009-20)




AI can be used for email to create responses that look as though they have been written by a person rather than a machine, using natural language processing to write content as well as understand it. Emails can be tailored based on the customer’s history and behaviour, optimising marketing messages as well as service, sending emails at a time when they have been calculated that they are most likely to be opened.

Personalised emails can be sent, based on similar subscribers’ past email outcomes so as to maximise the success of the interaction. This is a way in which AI can outperform human agents, who do not have the time or expertise to find patterns or draw conclusions from huge amounts of data.

As the chart below shows, the most popular methods of answering inbound email are to use agents, rather than rely on automation. Around half of emails are answered by agents who start with templatised, editable responses and change them accordingly, thus not having to compose every email from scratch, but also being able to draw from a common pool of knowledge. Starting with a blank email and letting agents complete it themselves is not only likely to take longer, but also leads to an increased risk of poor grammar, spelling and punctuation, as well as a less consistent response.

Only 4% of emails have automated responses, (these statistics do not include simple automated acknowledgements), and of those, half are checked by agents before sending.

Figure 14: Level of automation used in email management (2015-2020)



Survey respondents state that around half of their inbound emails are queries about products or services that have already been bought, with only 1 in 6 being from prospective new customers, who have queries about products or services which they are considering buying.

Complaints represent around 15% of inbound email traffic for our respondents, compared to the telephony figure which is usually below 10%.

ASSISTED SERVICE


AI-assisted service provides agents with real-time information that is tailored and relevant to the customer interaction they are having. Findings later in this report suggest that businesses currently see AI as being more useful for augmenting agents' capabilities rather than being used to replace them. AI assistance is particularly useful for helping inexperienced agents, for remote working scenarios where asking colleagues is more difficult and for more complex interactions which may require multiple systems and database to be accessed.

DIGITAL INTERACTION ASSISTANCE

Digital channels may work quite well for customers, but businesses are not generally seeing the cost savings that automation can bring. Very few emails or web chats are handled entirely by AI, although a growing proportion of web chats are 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 foreseeable future, with the speed of automation and the emotional intelligence of humans providing superior service at a lower cost.

15% of web chats in 2020 involved an agent being guided or advised by AI on what to write, with a much lower proportion doing the same with email.

Figure 15: Human and AI email and web chat handling



VOICE INTERACTION ASSISTANCE

The use of AI to assist agents in real time within a call offers the chance of a real paradigm change: by the nature of the job, an agent-customer interaction has always necessarily been between two people, and the level of support that an agent can actually receive within a call is very limited. Advice on learning points have been restricted to post-call reviews, rather than realistically being able to improve the outcome of the interaction in real-time.

AI offers an opportunity to provide timely and effective support to every agent as necessary, actually within the call. AI can provide the agent with suggestions about next best action, pull up relevant information from the knowledge base, make suggestions based on customer history and sentiment about optimal cross-selling and upselling opportunities, and even the style of conversation that this customer may prefer. This has a positive impact on first-contact resolution as well as customer experience, and is of particular use to less experienced agents and for unfamiliar subject areas.

Bots monitor the real-time desktop and voice data, triggering processes such as information provision and back-office processes. They can also provide coaching or alerts if there's a lengthy pause in the conversation or anything has been done wrong. Agents can also use specific phrases, such as "I'll just look that up for you", triggering the bot to take action and putting the information on a single agent desktop application.

AI can be trained to understand intent and recognise patterns through immersion in vast quantities of historical data, so that when a call is taking place, it can draw upon this knowledge and provide advice or action that has proven successful previously, moving towards the actual provision of real-time analytics.

AI assists in real-time speech analytics through applying the results of machine learning that have been carried out on large quantities of previously recorded conversations, providing:

- agents with the understanding of where their conversational behaviour is falling outside of acceptable and previously successful norms (such as speaking too quickly or slowly, or in a monotonous fashion)
- an assessment of the meaning of non-verbal cues such as intonation, stress patterns, pauses, fluctuations in volume, pitch, timing and tone in order to support sentiment analysis
- understanding the actions and information that have been seen to provide successful outcomes in previous similar interactions, and relaying this to the agent within the call.

A combination of customer feedback and interests can be used to develop a customer profile, adding metadata around purchase history, demographics and lifetime value. Past customer outcomes with similar customers can be used to predict the best offers, communication method, channel and actions with that customer.

AI can work alongside agents to provide relevant knowledge that may be otherwise take a long time to find, and update the knowledge bases available to humans and AI self-service systems using an automated feedback loop that is constantly improving based on actual outcomes.

It's possible to fix customer service problems before they occur: for example, sudden numerous requests about the same thing is likely to indicate a breakdown in a specific business process or the occurrence of an outside event. AI can quickly recognise that this is an issue, and deliver information solutions to an agent's screen, to the chatbots and note that changes should be made to the IVR announcement.

SENTIMENT ANALYSIS

Sentiment analysis is a way of quantifying customer and agent emotions within interactions, whether on the phone or through an alternate channel, for the purpose of uncovering processes, behaviours and situations which cause strong levels of positive or negative sentiment that could affect business outcomes and customer experience. Using analytics and large data sources, datasets can be searched to identify and inspect the types of interaction that have major impacts on customer sentiment.

Agents, especially those with higher levels of empathy and experience, should be able to identify the emotions of the callers, so using technology for sentiment detection could seem at first glance to be an unnecessary elaboration. However, the use of analytics means that the sentiment and emotion of millions of calls can be assessed against their ultimate outcome in order to identify in real-time situations that have a higher likelihood of a negative outcome and to act before it is too late.

While language models can identify ostensibly positive and negative words and phrases, they cannot in themselves identify sarcasm or other less straightforward forms of communication, and they are less likely to identify the actual meaning in a series of conflicting positive and negative comments (e.g. "I'm happy that the product has **finally** arrived – I mean, this is good, but not exactly great, you know?"). Sentiment models are further trained to notice changes in tone, volume and speaking rate, instances of agent/customer talkover and the detection of laughter, silences or sub-audible noises expressing emotion, such as a snort of disgust.

Each interaction can then be scored on a sentiment scale from highly positive to highly negative, with nuances such as conversations which start positively and then turn negative able to be selected for root cause analysis. It is interesting to note that sentiment expressed towards the end of the call is a much better prediction of customer satisfaction than any emotions expressed at the beginning of the call: this makes sense, as a customer could easily be stressed if they have had to wait in a long phone queue for an urgent matter that they are stressed about, but if the agent resolves the query to the customer's satisfaction, there is likely to be a positive sense of relief and gratitude expressed, which is likely to indicate a good customer experience.

While sentiment analysis captures and analyses every interaction, it is generally thought to be of most use at an aggregated level rather than in judging particular individuals. Sentiment analysis can identify those processes, interactions and subject areas that are causing customers the greatest stress and negativity, and can view trends over time which allows the business to gauge whether any business or process improvements that they have made as a result are actually working. Some businesses do decide to look at sentiment at a team and individual level, but great care must be taken not to attribute negativity to a specific agent rather than the topic or product under discussion.

Sentiment analysis is potentially a very powerful tool and in common with the rest of the interaction analytics functionality it has many potential applications:

- **Discovery and categorisation:** by analysing thousands or millions of interactions, sentiment analysis is able to show the products, processes and topics which most often provoke the strongest negative or positive reactions, categorising them automatically for root cause analysis
- **Quality assurance:** interaction analytics is often used to analyse 100% of calls, rather than having a supervisor listen to a random, small selection which may not be representative of agent performance, and which may miss major opportunities to improve. Sentiment analysis plays a part in quality management, but an expectation of a correlation between poor agent performance and negative sentiment should not automatically be assumed. Analysing metadata such as the topic under discussion should indicate whether this negativity arises from a specific agent performance or is more likely to be linked to the subject matter
- Having said this, sentiment analysis can be a useful tool to use in order to rank agents by capability, in order to understand the behaviours and characteristics of top performing agents so that underperforming employees are able to be coached on these effectively
- As mentioned above, negative sentiment may be linked to a particular topic product or process. A dataset analysed by a sentiment model can be searched by product, giving a rapid answer to whether it is seen by customers as being broadly positive or negative. Delving further into the data – for example, looking only at the negative sentiment associated with a particular product – may identify areas for improvement (e.g. while the product performance itself scores highly for positive sentiment, the instruction manual scores negative, identifying an area for improvement)
- Some businesses use sentiment analysis to consider factors such as agent morale and motivation. This can be particularly useful in a sales environment, where the enthusiasm or otherwise of the agent can make a significant difference to the outcome
- Real-time sentiment analysis may be useful for offshore agents who have a different cultural and first-language background to that of the caller
- Sentiment analysis can identify stress in real-time, which may be an indicator that fraud is taking place, prompting the agent to take the caller through more detailed levels of security in order to prove their identity. This can be used in association with voice biometrics and/or phoneprinting, in order to identify the callers requiring stronger authentication.
- Sentiment analysis has been shown to be useful in predicting NPS, and is also useful in targeting customer satisfaction surveys. For example, for interactions with negative sentiment around a specific topic, a survey can be sent that asks customers specifically what went wrong with that issue, rather than relying upon a broad-brush general NPS approach with an open-ended question.

Some solution providers have recently noted that it is not only what we might consider the keywords within the conversation that indicate sentiment (e.g. “upset”, “disappointed”, “recommend”), but also the filler words (for example, if the inclusive “we” changes to “you”, which may indicate estrangement from the brand).

PREDICTIVE / INTELLIGENT ROUTING

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. AI enables an 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.

For example, an AI 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 is able to move the call to the front of the queue and route it to an agent experienced in changing flights, and whose communication style suits the situation and customer profile.

Taking this a step further, the AI is able to augment the conversation with suggestions based upon what the agent is doing on the screen and also, through listening to the details of the conversation, is able to provide relevant information without the need for the agent to search for it, such as the next flight to the customer’s proposed destination or the refund / transfer options. At the end of the call, the AI can then email or text the agreed solution to the customer without the agent having to do this manually.

AI & REMOTE WORKING

Remote working has been a necessity for most contact centres in 2020 and 2021. Even as the pandemic recedes, many companies are planning to continue a remote working model (usually in part rather than whole), and are considering better ways of managing this.

Figure 16: Proportion of agents homeworking, 2019 - 2021



While many agents like working at home, with generally positive outcomes for agent absence and attrition in particular, management of teams and individuals is a challenge. AI assistance for agents where supervisors or experienced colleagues aren't available (such as next best action) can help, along with real-time sentiment analysis which can flag a supervisor remotely who can break into the conversation or whisper coach.

AI also can show inexperienced agents the approach that typically provides positive outcomes, and also offer tailored cross-sell and upsell opportunities.

END-USER QUESTION #2: HOW CAN AI HELP WITH REMOTE WORKING?



AI can be used to make sure that agents have easy access to the right resources while remote, as well as helping managers monitor call quality or identify places for improved training from afar.

Furthermore, organisations can use AI to automate or streamline mundane tasks, the same as in the office. For example, a prediction model could give an agent more information about a call – suggesting the most likely reason a customer is calling – so agents can formulate a response while verifying customer information.

ROBOTIC PROCESS AUTOMATION (RPA) & AI

Robotic process automation (RPA) consists of digital software agents that handle repetitive, rules-based tasks at high speed, with great consistency and accuracy. RPA does not replace existing systems, simply sitting on top of existing logic and applications, using them in the same way that human contact centre agents, chatbots or back-office workers would do. In this way, it does not require complex integration, meaning roll-out of the robots can be relatively quick and flexible.

Processes and the necessary steps to perform a task are defined, put into a queue and the controller assigns various tasks to the robots. These robots can be monitored for speed and accuracy in the same way that a human workforce would be managed, with exceptions being flagged to human supervisors who can investigate why a particular task could not be completed as designed. RPA agents can input data, trigger processes, pass work onto other robots or humans as rules dictate and replicate data across multiple applications without making any copying mistakes.

87% of UK contact centres require agents to use multiple applications, often complex and only loosely linked, which require skilled and experienced agents to navigate, let alone to manage interaction with customers successfully at the same time. Even after the call is completed successfully, each system may need specific inputs from the agent in order to start the required back-office processes, and there is always the danger that even if the call has been completed successfully, opportunities to maximise revenues have been missed. There are also significant dangers around forgetting to key in information, not asking for the required information, starting the correct processes or failing to type in consistent data. The use of multiple applications will have a negative effect on training times and accuracy rates for new agents as well.

Using live agents to handle multiple, complex applications can have significant negative effects:

- Increased training costs
- Higher staff attrition caused by inability to complete tasks successfully
- Inconsistent data caused by keying errors or missed procedures caused by manual wrap-ups
- Increased call handling times
- Lower customer satisfaction caused by long queues and unnecessarily long calls
- Missed opportunities to cross-sell and up-sell
- Multiple open applications on the agent desktop can lead to system instability and lower performance.

The result is that even though a contact centre may be staffed with experienced, hard-working and skilled staff, its overall performance is suboptimal, leading to low customer satisfaction, unnecessary costs and decreased profits. RPA offers a way in which agents can be supported to assist customers through optimising the agent desktop without needing to rewrite systems or integrate deeply with multiple applications and databases.

In most cases where complex, multiple applications are used, they are necessary for the agents to do their job, so the question is not “How can we reduce the number of applications?”, but rather “How can we improve how the agent uses the applications?”.

RPA-assisted integrated desktop solutions can remove the need for agents to log into multiple applications, assist them with the navigation between applications within the call, and make sure that customer data is gathered from the correct places and written consistently back to any relevant databases without the need to navigate through multiple systems.

RPA can assist contact centres and back offices in numerous ways, including:

- Handling routine activities, such as the actions associated with a particular task such as change of address, including automated login to specific systems, field completion, screen navigation, copy and paste after a single entry is placed by a human agent in one application
- Triggering of processes based on call or digital interaction outcomes
- Record processes in ticketing systems
- Review and pass them onto the next stage in the workflow
- Validating customer account information
- Proactively sending updates to customers depending on the stage of the process.

The benefits of RPA in the contact centre include:

- **Accuracy and consistency:** if the data fed to the RPA robot is correct, and it has been programmed successfully, it will carry out its tasks consistently and accurately, with no possibility of mistyping or forgetting to complete a field
- **Reduced cost and improved customer experience:** agents may no longer have to do complete onerous post-call activities, reducing call handling time and increasing their availability to the next customer
- **Speed:** RPA robots carry out tasks many times quicker than even the most experienced contact centre agent. They are also available to work 24/7
- **Improved compliance with regulations:** if programmed correctly, RPA robots will never forget to ensure complete compliance by completing required fields and sending out correct documentation
- **Improved agent experience:** by handling some of the agents' repetitive work, there is less risk of boredom and low employee morale.

Unlike AI, RPA robots carry out their tasks in a consistent manner each time, neither altering nor learning from their behaviour or past outcomes, and are unable to improve their performance over time.

To automate more complex tasks and processes, AI needs to be involved which can create a more autonomous process which may require less human input. AI is important when looking at unstructured data such as free text (e.g. in emails or letters) or natural language, whereas RPA can handle structured data such as those found in drop-down fields. Where AI lacks the necessary confidence level to move the process on, it can refer to a human agent who can either accept or amend data, handing it back to automation afterwards. Crucially, the AI system over time will learn why human agent made that decision, and will be more likely to be able to handle this autonomously.

One of the major applications suitable for RPA is assisting front office agents (attended RPA). The agent desktop lies at the heart of the integrated contact centre, with data and processes flowing to and from it. An attended RPA solution looks after some of the work that an agent has to do within a call, and also afterwards in post-call wrap-up. The bot can create the call summary from the relevant data gathered and email this to the customer, as well as disseminating any relevant data to various databases and kicking off required processes (e.g. sales fulfilment). Within the call, the attended bots can be triggered by the agent's desktop activities, rather than requiring them to activate the bot manually. This can include populating fields, compliance reminders and cross-selling opportunities.

RPA has historically been about automating high-volume events, taking the weight from agents so as to eliminate errors and increase efficiency. The new generation of attended AI-enabled RPA takes this further, assisting agents in real-time through augmenting their desktop with relevant data and suggested actions. Through speech recognition and NLP, AI can identify what's happening within the conversation in real-time, and put next best actions in front of the agent, personalising these depending upon what the customer has done before and what they want from the interaction. The result can be a superior and timely cross-sell / upsell offer, or simply having the relevant answer to their query presented quickly, even if the agent was unaware of what that might be. This is clearly of great use to inexperienced agents and those working remotely, who may not easily be able to ask a team leader or a more experienced colleague. The result can be higher levels of customer and agent satisfaction, and improved cross-selling outcomes.

AI can work in association with other process automation solutions (which may in themselves not fall under the category of AI). For example, in the case of unstructured data such as customer emails or letters, optical character recognition can assist the entry of the customer requirements into the business system. Using natural language understanding, AI is able to discern the intent of the enquiry, using a knowledge base and assessing the previous best responses to similar enquiries in order to provide an agent with a recommended solution. It is very likely that the agent will be given the option to add or amend this response before sending to a customer. Any feedback from the customer can be assimilated in order to gauge success and fine tune future responses. Robotic process automation can also be used in order to update customer records accordingly.

Within the call, AI can seamlessly link with multiple back-office applications and databases, providing only what is relevant onto the agent's screen. Depending on the experience or profile of the agent, what the customer is trying to do and any regulatory inhibitors, on-screen buttons can be enabled or disabled, or access to fields limited according to business rules. Furthermore, adherence to business processes can be assured by making the agent complete all of the required steps in the transaction (for example, adding call notes, reading disclaimers, etc.).

The following table shows some key contact centre performance metrics that were analysed in the context of the proportion of time that agents spend navigating through multiple in-call applications. It is important to note that although there appears to be a correlation between superior performance metrics and less screen navigation, this does not necessarily demonstrate causality: this pattern of statistics do not mean that it is possible to say **definitely** that the use of fewer applications within a call will in itself improve contact centre performance.

However, it can be supposed that not having to navigate through multiple screens or spend significant periods at the end of the call typing out notes or making changes to multiple databases, and being given access to dynamic scripting that provides the correct information without having to search for it will encourage shorter calls, improved agent availability, and lower call abandonment rates.

Figure 17: Selected performance metrics, by number of in-call applications used

Metric	Respondents using 1 or 2 agent applications within a call	Industry average
Average speed to answer	55 seconds	81 seconds
Service call duration	241 seconds	368 seconds

It is logical to hypothesise that using complex, multiple applications without any specific agent support will often lead to longer calls. However, this is not the end of the problem, as this type of work also tends to initiate requests for processes to be carried out within the back-office (e.g. initiating an engineer or sales visit, sending out literature, moving a customer request onto the right department with the right information, flagging a customer as a hot prospect for a specific marketing campaign, etc.).

This, as well as the need to enter information in multiple applications, will tend to increase post-call wrap-up to a point where the agent spends a great deal of their time unavailable to take more calls. Historically, 10-15% of an agent’s time is spent on post-call wrap-up. It is in this stage that a lot of time and effort is wasted by sub-optimal manual processing of data. For example, a simple change of address request could take many minutes in a non-unified environment, with several separate databases having to be altered, which is itself a process prone to error, with a negative impact on the customer and business, as well as at least one extra unnecessary future phone call from the customer. Reducing wrap-up time through optimising the agent desktop is not simply a matter of writing consistently to the correct databases, although this is a key element. The contact centre also kicks off a number of processes elsewhere in the enterprise: it is the prime mover for sending out documents, instructing the warehouse to release goods, arranging deliveries, taking payment and many other key elements to a successful customer-business transaction. RPA is set up to handle these processes in a consistent, accurate and rapid manner.

Additionally, manual inputs involved in transferring data during wrap-up commonly lead to data entry and processing errors, causing an adverse effect on operational efficiency, contact centre cost, performance and customer satisfaction. Cost per call rises, productivity per agent declines and first-call resolution rates slip as more calls are escalated due to the complexity of the systems hindering agents, rather than helping them. So we can see that poor application integration and presentation at the desktop level has a direct and negative effect on those long-term contact centre strategies deemed most important and desirable, such as customer satisfaction, lower first-time resolution and reduced escalation levels.

Ai-enabled RPA can use machine learning and NLP to recognise products and processes that have been recently added, ‘understanding’ that while it may be unfamiliar with a new product, it should treat it in the same way as any other product, recognising the type of datum or process for what it is and acting accordingly. NLP and OCR may be used to identify and understand exceptional written notes on an order – a special request, for example – and be able to process the work without having to mark it for manual intervention.

RPA can also be used in the ‘discovery’ phase of an AI implementation to decide which processes should be automated by gathering information from agents’ desktops (e.g. field completion, mouse clicks, processes initiated, etc.) for various actions. This will give businesses a good idea which processes occur the most, which require the same actions or information, and which require the most time and effort from agents.

INTERACTION ANALYTICS

A key part of the AI picture, customer interaction analytics solutions offer huge opportunities to gain business insight, improve operational efficiency and develop agent performance.

DISCOVERY

'Discovery' is a term often used within the customer contact analytics industry, and refers to a deep, automated analysis of trends, patterns and results which are identified by the analytics solution rather than the knowledge or insight of the human operators. Discovery will help users to find interactions that are similar to each other, perhaps through similar groupings of words or phrases, and explore these links to discover the issues driving them.

The ability to see trends – to know that the instances of the words 'website' and 'password' have increased by 2,000% this week compared to the norms of the past 6 months – quickly identify likely pain points for the customer and potential broken processes. The continual tracking and analysis of similar information or categories over time also allows a business to see whether the remedial action that they put into place has actually worked.

Many analytics solutions offer AI-enabled automated discovery and this is an area that will always be improving and becoming more subtle and effective, having huge potential benefits for businesses.

Of course, any analysis where the direct beneficiary is not the contact centre must be properly aligned to the organisation's objectives and strategy, encouraging changes to be made to areas that have already been earmarked as needing improvement. Otherwise, if the focus is not aligned with strategic goals, information merely becomes 'nice to know', rather than actionable.

Customer interaction analytics has the ability to tear down the virtual wall between the contact centre and other areas of the business, meaning that the business intelligence extracted can be shared and valued by parts of the organisation that otherwise have little to do with the contact centre. With the historical and ongoing difficulty in getting the business to value the customer contact operation fully, this can only be a good thing politically.

Text analytics (using NLP and OCR) can be used within discovery mode to assess not only data held within the company, but also in unstructured, third-party environments, such as social media, letters and comments on websites and public forums, in order to learn and deliver proactive service before it is even requested.

CUSTOMER JOURNEY ANALYTICS

AI can be applied across the entire customer journey, including sales, marketing and service, helping organisations understand customer behaviour, intent and anticipating their next action. For example, an AI solution may find a pattern amongst previous customers that they are likely to search for specific information at a particular point in their presales journey, and proactively provide this information (or an incentive) to the customer before they have even asked for it. AI can also help with customer onboarding through predicting which customers are likely to require specific assistance, and when.

Machine learning allows analytics solutions to go beyond simply what they have been programmed to do, seeking out new opportunities and delivering service beyond what has simply been asked of them. Through understanding multiple historical customer journeys, AIs will be able to predict the next most-likely action of a customer in a particular situation, and proactively engage with them so as to avoid an unnecessary inbound interaction, providing a higher level of customer experience and reducing cost to serve.

Also, use AI to gauge customer effort at each stage of a buying or service process through analysis of what that customer has said (e.g. "I've tried to do this twice online but it won't let me"), which can identify bottlenecks and sources of frustration for the customer, to avoid churn, abandonment and low customer experience.

IDENTITY VERIFICATION AND FRAUD REDUCTION

Voice verification systems use spoken words to generate a voiceprint, and each call can be compared with a previously enrolled voiceprint to verify a caller's identity. Systems generate a voiceprint by using spoken words to calculate vocal measurements of a caller's vocal tract, thereby creating a unique digital representation of an individual's voice, as well as other physical and behavioural factors, including pronunciation, emphasis, accent and speech rate. These systems are not affected by factors such as the caller having a cold, using different types of phones, or aging.

A significant advantage of voice biometric verification is that both enrolment and verification can be done unobtrusively – in the background during the natural course of customers' conversations with an agent – using text-independent and language-independent technology. Real-time authentication significantly reduces average handle time and improves the customer experience by utilising voice biometrics to authenticate customers within the course of the conversation.

With this advanced technology, contact centres can:

- Voiceprint the vast majority of customers for seamless passive enrolment: in the course of a conversation, a voiceprint is created for that customer which lies on record for them to be authenticated against on the next call
- Securely authenticate customers with zero customer effort: the first few seconds of a call will be enough to match the customer's voiceprint against those on record

- Cut seconds off average handle time: no need for customers to answer numerous security questions as the conversation they are having provides enough information to identify them
- Significantly reduce fraud risk for all customers, and deter fraudsters when combined with other layers of security, for example, phoneprinting, which analyses the background audio of the call.

Some biometric solutions use neural network speaker recognition systems, which run in the background throughout the length of the call (including the IVR phase), stripping away background noise which can be used to mask identity.

Voice biometrics, while an excellent authentication tool, is not enough to deter fraud attacks. In fact, in 2015, researchers at the University of Alabama¹ found that a fraudster armed with just a few minutes of recordings of a person's voice, could build a model of the victim's speech patterns and successfully pass voice biometric security. As voice is a characteristic unique to each person, such attacks essentially give the attacker the keys to that person's privacy.

One alternative method of customer identity verification that can be used alongside biometrics is 'phoneprinting', which is perhaps focused more on identifying and preventing fraud than on simply authenticating genuine customers. Phoneprinting collects information about the call being made, such as location, the type of phone being used (VoIP is far more likely to be used in fraudulent calls), CLI (the calling number), the phone number's history and the chances it has been 'spoofed', levels of voice distortion, etc. These factors can be scored, and after assessing the likelihood of the call being fraudulent will then impact upon the security processes and questions that the agent is required to ask the caller, speeding up the process for genuine callers, and focusing the tightest levels of security on potentially fraudulent calls. Some solutions allow fraudulent phone numbers to be gathered and shared with other businesses, red-flagging likely fraudsters. Data from various sources can be added, such as consumer complaint sites, spam calls databases, detecting attack patterns and improving suspicious call identification.

Many businesses are moving from password-only authentication to multi-factor authentication (MFA). This often requires the customer to demonstrate something that they have (e.g. a one-time passcode to their phone), as well as something that they know (such as a password). NB: voice biometrics involve something that they are, and can form part of the MFA process. However, phones can be stolen, and one-time passwords intercepted. MFA can be improved by adding context, such as the IP address or location, which can flag anomalies to the usual methods the customer has used in the past to contact the business. This 'risk-based' authentication requires AI to consider a large number of factors in real-time, compare these to the customer's profile and deliver a risk score which may indicate a fraudulent interaction. The business can then ask for further proof of identity or deny access altogether. Using sentiment analysis to observe stress levels in the voice can also help to determine fraudulence.

¹ <http://www.biometricupdate.com/201509/uab-researchers-find-that-automated-voice-imitation-can-spoof-voice-authentication-systems>

WORKFORCE OPTIMISATION

Analytics solutions offer the capability to monitor and score 100% of calls automatically, and machine learning can use this large pool of data in order to analyse patterns of agent behaviour and characteristics connected with best outcomes, in order to develop performance and training programs always down to the individual agent level. Specific gaps in knowledge or capabilities can be identified and addressed based on thousands of calls, rather than relying upon manual evaluations which can only process a handful of calls from each agent.

Being able to score every call through an automated AI process means that the quality assurance team is able to review specific calls that have been flagged up as being potentially important, rather than hoping that they stumble across them in a random assessment. This may include calls where specific language is used, has long pauses, or where the agent or customer raises their voice or talk across one another. The AI system can be trained to understand which calls are “normal”, and which are outliers more likely to require input from the quality management department.

AI also enables marketing departments to predict the likely success of their campaigns, and translate this into how the contact centre or wider customer contact arena will be impacted, providing predictions on how resourcing and training will need to be changed to meet customer requirements. Being able to use peripheral data such as the specific level of agent experience available, the day upon which the marketing campaign is launched and even the time of the year will improve resource forecast accuracy.

REAL-TIME ANALYTICS

Many solution providers have worked hard to bring to market new or improved solutions to assist with real-time analytics, and recognition of key words, phrases, instances of talk-over, emotion and sentiment detection, pitch, tone, speed and audibility of language and many other important variables can be presented on the agent desktop within the call, triggering business-driven alerts and processes if required.

The speed of real-time analytics is crucial to its success: long delays can mean missed, inappropriate or sub-optimal sales opportunities being presented; cancellation alerts can show up too late; compliance violations over parts of the script missed-out may occur as the call has already ended. However, it is important not to get carried away with real-time analytics, as there is a danger that businesses can get too enthusiastic and set alert thresholds far too low. This can result in agents being constantly bombarded with cross-selling and upselling offers and/or warnings about customer sentiment or their own communication style, so that it becomes a distraction rather than a help.

Real-time analytics can be used in many ways:


- monitoring calls for key words and phrases, which can either be acted upon within the conversation, or passed to another department (e.g. Marketing, if the customer indicates something relevant to other products or services sold by the company)
- alerting the agent or supervisor if pre-specified words or phrases occur
- quality checking the agent for speech clarity and speed and notifying them of any changes to be made
- offering guidance to the agent on the next best action for them to take, bringing in CRM data and knowledge bases to suggest answers to the question being asked, or advice on whether to change the tone or speed of the conversation
- escalating calls to a supervisor as appropriate
- text analytics can also be used on inbound interactions such as emails, running an AI triage system to assess the priority and urgency of each request in order to handle these more effectively and in an appropriately timely manner
- detecting negative sentiment through instances of talk-over, high stress levels, negative language, obscenities, increased speaking volume etc., that can be escalated to a supervisor
- triggering back-office processes and opening agent desktop screens depending on call events. For example, the statement of a product name or serial number within the conversation can open an agent assistant screen that is relevant to that product
- making sure that all required words and phrases have been used, e.g. in the case of compliance or forming a phone-based contract
- suggesting cross-selling or upselling opportunities.

The effectiveness of real-time analytics may be boosted by post-call analytics taking place as well. For example, by assessing the outcomes of calls where specific cross-selling and upselling approaches were identified and presented to agents in real time, analysis can show the most successful approaches including the use of specific language, customer type, the order of presented offers and many other variables (including metadata from agent desktop applications) in order to fine-tune the approach in the future. Additionally, getting calls right first-time obviously impacts positively upon first-call resolution rates, and through picking up phrases such as "speak to your supervisor", can escalate calls automatically or flag them for further QA.

Combining AI-enabled real-time and post-call analytics, along with sentiment analysis, can quickly identify large groups of dissatisfied customers that may be complaining about the same thing, providing actionable insight to the business within minutes or hours, rather than days.

Real-time analytics offers a big step up from the traditional, manual call monitoring process, and is particularly useful for compliance, debt collection, and for forming legally binding contracts on the phone, where specific terms and phrases must be used and any deviation or absence can be flagged to the agent's screen within the call. Finance, telecoms and utilities companies – and indeed, any business where telephone-based contracts are important – are particularly interested in this.

Figure 18: Usefulness of real-time analytics



IMPROVING ROUTING AND OUTCOMES

Predictive behavioural routing uses insights gathered from historical calls and the analysis of customer communication types in order to choose the agent whose skills and characteristics are most likely to achieve a positive response from the next caller in the queue.

Predictive behavioural routing uses millions of algorithms to decode the language used by agents and customers, in order to understand sentiment, personality type, preferred method of communication, emotional intelligence and transactional attributes (such as ability to overcome objections and willingness to sell).

Each customer can be allocated a specific personality style, and when calling again, are routed through to an agent whose performance when interacting with this specific personality type has generally positive results. The matching of agent-customer communication styles has enjoyed significantly higher sales closure rates and better customer satisfaction scores than the average. Agents who are skilled at handling many types of callers' personality styles can be saved for callers whose character type is unknown, perhaps as this is the first time that they have called.

AI can also be applied to IVR interactions, asking a series of questions to customers using natural language processing to understand their intent. Depending on the customer requirements, it may be possible to answer the query without using a live agent, or in those cases where agents are needed, the prioritisation and routing of the call can be optimised, decreasing call transfer rates and increasing first-contact resolution. Over time, routing strategies will move away from being rules-based and towards cognition, which will also feed forecasting and scheduling processes.

AI AND CX METRICS

CX metrics such as NPS (Net Promoter Score) are viewed as vital by many businesses, especially their top executives. Yet most customers who are presented with an NPS survey will not complete it: there is usually a polarisation effect that means those who are very unhappy and those who are very happy will be much more likely to take part, meaning that the accuracy of results is compromised.

It is also the case that complex interactions, involving a lot of different steps, may not yield the type of insights that are useful. For example, a customer completing a mortgage application may have great difficulty in completing a certain form, but find other elements very easy. They may give a lukewarm NPS result (e.g. 6 or 7), whereas they may have wanted to give a 1 or 2 for the part they had difficulty with, but 9 or 10 for other parts. Collecting freeform verbatim comments is crucial for learning why customers are happy or otherwise, but this can involve a great deal of manual effort for large-scale surveys, which may fail to notice patterns and trends in the data that an AI would be more likely to spot.

Fewer than 5% of customers will complete a NPS survey, so AI can help to predict NPS and other CX metrics for this hidden majority through a mixture of sentiment analysis, NLP and speech-to-text which can then score the call / interaction and identify potential root causes.

IMPLEMENTING AND USING AI

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 are many contact centre activities that are more appropriate for a person to do.

For first-time deployments, the focus should be on the AI 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. It might be appropriate to start with a self-service-focused chatbot 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 or agent 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 in order 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. lack of available resource to handle existing enquiries, suboptimal business processes, inability to analyse large datasets, etc.
- How does this affect the customer experience, and how would the customer like this to be improved?
- Are there solutions in the marketplace that have successfully addressed these issues already in live environments?
- 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?
- Is there a sufficient volume of data in order to train an AI system effectively?
- Will our infrastructure or existing platform need to be replaced?
- Is AI definitely the most appropriate way of dealing with this issue?

It is likely that senior decision-makers within the enterprise have a sketchy or unrealistic expectation of how AI can help within the contact centre environment. As such, 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 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 very unlikely that this will be an appropriate response to the success of an AI project, certainly in the short-to-medium term at least. It may be better for the project to be viewed as improving the customer experience through providing customers with an alternative to a frustrating web browsing experience, ending with an unnecessary and unwanted live call.

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 roll-out of AI may look similar to the following:

- Use a virtual assistant 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 is then translated into a form 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. 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 Tay² is an example where the AI displayed inappropriate and offensive responses after interacting with Twitter 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 as 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 and operational change should all be considered. Of course, it also depends on the area of the contact centre business processes that implementation is aimed at improving.

² [https://en.wikipedia.org/wiki/Tay_\(bot\)](https://en.wikipedia.org/wiki/Tay_(bot))

Some examples of AI -related metrics include:

Self-service:

- volume of self-service attempts, segmented by type of interaction
- customer satisfaction by self-service segment (particularly useful for comparing AI-enabled self-service with scripted/programmed self-service)
- 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)
- change in inbound call volumes.

Assisted service (i.e. the AI's impact on live agents):

- first-contact resolution rates, with and without AI
- proportion/volume of simple/complex interactions and by agent
- call duration (is the AI assistant cutting down on unnecessary screen navigation by providing the answer on the agent desktop?)
- if using RPA, effect upon downstream business processes and reduced call durations / post-call wrap-up times
- impact upon customer satisfaction scores.

END-USER QUESTION #3: HOW CAN AI BEST BE USED FOR A QUICK WIN? WHERE CAN ROI BE FOUND?



The fastest ROI in call centres is commonly performance modelling. It starts by using a success metric at the agent level, and all the data you think may correlate, such as tenure, QA scores, workforce optimisation, and any other data that can contribute to success. From there, a model can be built to assign attributes to each decile of performance. Once a model determines these attributes, they can be used in a variety of ways. The fastest ROI would involve using the attributes to group similar interactions together or sort interactions into predetermined buckets based on how these attributes interact.

For example, group interactions with a high percentage of silence or low sentiment scores and then focus on those subsets. You could also use these attributes to create an ordering or ranking of the interactions. Once interactions are ordered in a meaningful way, you can improve the lower ranked interactions by training on what the higher ranked interactions do differently.

Another place where ROI can deliver quick value is finding the most frequent call drivers in a call centre, which can provide an overall picture why people are calling. It doesn't have to find every single call driver, but the top 10 would be useful to an organisation. With AI, this can be achieved using a simpler prediction model for whether a sentence is a call driver or not.

INITIAL ACTIONS IN AN AI PROJECT

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 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.

END-USER QUESTION #4: WHAT ARE THE STEPS WE NEED TO TAKE TO USE AI IN OUR CONTACT CENTRE? WHERE WOULD BE THE BEST PLACE TO START?



The best place to start with AI in the contact centre is what you already are doing manually. This is sage advice for any sort of automation – pick an area where you have a manual process that is repeated and use AI and ML to solve some or all of it. It is far easier to automate what you already know than it is to try to automate what is unknown.

Once you've decided on those processes, here are some key steps:

- Ask a specific, unbiased question. Consult with a team to help refine your question, especially if there is a risk of bias or unintended results. More brains are better, and everyone can draw on their lives to help create a responsible task for AI.
- Collect the data. This should include call transcripts, meta data, and more.
- Determine a metric for success. For example, if you are creating a prediction model, is it important that you always predict correctly (sale/no sale) or that no instances slip through the cracks (fraud)? Or is the cost and resource efficiency of the model as important than the accuracy?
- Choose the best approach based on the data you've collected and the question you are trying to answer and how you measure success.
- Train the model and check that the model is responsible. A responsible model is one that has an explainable process (even if it's black box, use a metric to get a sense), one that is trained on a diverse and reflective data set, and one that is answering a question about behaviour and not demographic information. It's easy to forget to look at the training, creation, and output steps of the process, but doing so will ensure a model that makes reliable decisions/predictions on every data instance. You should also continue this process of reflection and evaluation throughout the AI lifecycle to make sure that a responsible model stays that way.
- Monitor the model, don't set and forget.

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 the self-service and assisted service experience for both agent and customer across channels. An AI project cannot succeed unless the data the AI 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 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 has 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 certain 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 the AI.

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.

Within a call, an agent is likely to have to use multiple knowledge sources, which will also take longer and run the risk (especially for new agents) of missing vital information that is available but perhaps hidden away. AI-enabled RPA can gather knowledge sources and provide them to the agent in a unified manner, and any updates to this information can be shared automatically across applications and systems, providing an immediate, up-to-date and consistent source of information. RPA can assist with agent tasks in the background, provided guided assistance at specific stages of the call, including dynamic scripting and compliance hints.

The following table shows the knowledge resources that agents have within a call. Finding, reading, assimilating and using information actually within a call as very difficult and is rarely done seamlessly. An application such as case-based reasoning, which prompts the agent to ask specific questions, drilling down to find the right answer, is very useful but only 25% of agents have access to this sort of dynamic application and only 6% have access to an AI-enabled virtual assistant. Most agents have to search around on a company website or FAQ page, or rely on a wide, unsupported search of knowledge bases or the wider Internet, hoping to get lucky.

While case-based reasoning has been around for a long time, rigid rules-based processing can easily break down as enquiries become more complex. AI reasoning offers greater flexibility and power for agents who are inexperienced, or for whom the query is too complex to handle without assistance. AI encourages agents to ask the right questions to understand the customer’s actual intent and to guide them to the correct response. Any feedback from this – whether successful or not – increases the chances of the next interaction being successful, and this is disseminated across the whole agent pool, rather than being held in the minds of a few expert agents. This can also be used in complex compliance scenarios.

Figure 19: In-call access to computer-based knowledge sources for agents



END-USER QUESTION #5: HOW MUCH INITIAL AND ONGOING EFFORT/RESOURCE WILL AI REQUIRE? DO WE NEED A DEDICATED AI PROFESSIONAL TO KEEP EVERYTHING RUNNING?



Depending on the algorithm, model training takes a significant amount of data, because a model is only as good as its training data. It's also important for organisations to think critically about questions and outputs to make sure that models are behaving responsibly.

The use of AutoML providers can really help here, such as DataRobot or Microsoft Azure.

Also, partnering with the right software providers, particularly those that are AI-forward, at the beginning of adopting AI is also a great option. This can help organisations leap ahead, because the work has been done and the output is open for consumption – meaning the data is structured in a ML-friendly way.


AGENT REPLACEMENT OR AUGMENTATION?

AI is often seen by the general public and media as being used to oust humans from their employment. Our survey respondents generally do not believe that AI will replace agents: although 32% agreed to some extent that this would be the case, 68% disagreed. Respondents from large 200+ seat contact centres were more likely to feel that AI would replace human agents, with those in small and medium operations doubting that this would be the case. However, it is worth noting that the belief that AI will replace agents has strengthened over the past few years.

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 disagreed that AI would be irrelevant to their contact centre, with general 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 20: Views on the role of artificial intelligence in the contact centre



CUSTOMER OPINIONS OF AI

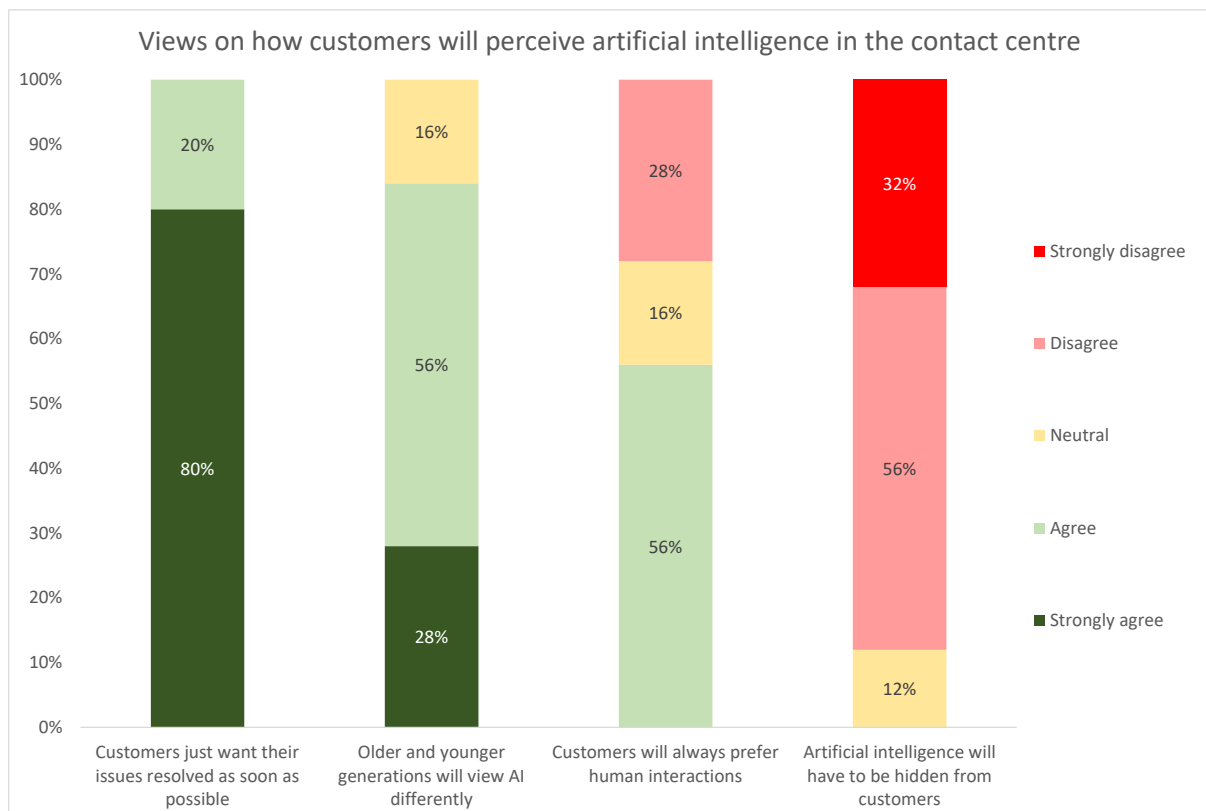
There is a very strong widespread belief amongst businesses that customers will not object to being served by AI if it helps them to resolve their issue as quickly and easily as possible. The recent uptake in web self-service suggests that customers will accept non-human assistance if it is most convenient for them.

There was also agreement that older generations would 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.

Respondents disagree about whether customers will always prefer human interactions: although more believe that customers will always prefer human interactions, there are few strong feelings about this. The next question looks at this issue from the perspective of the customer.

Figure 21: Views on how customers will perceive artificial intelligence in the contact centre

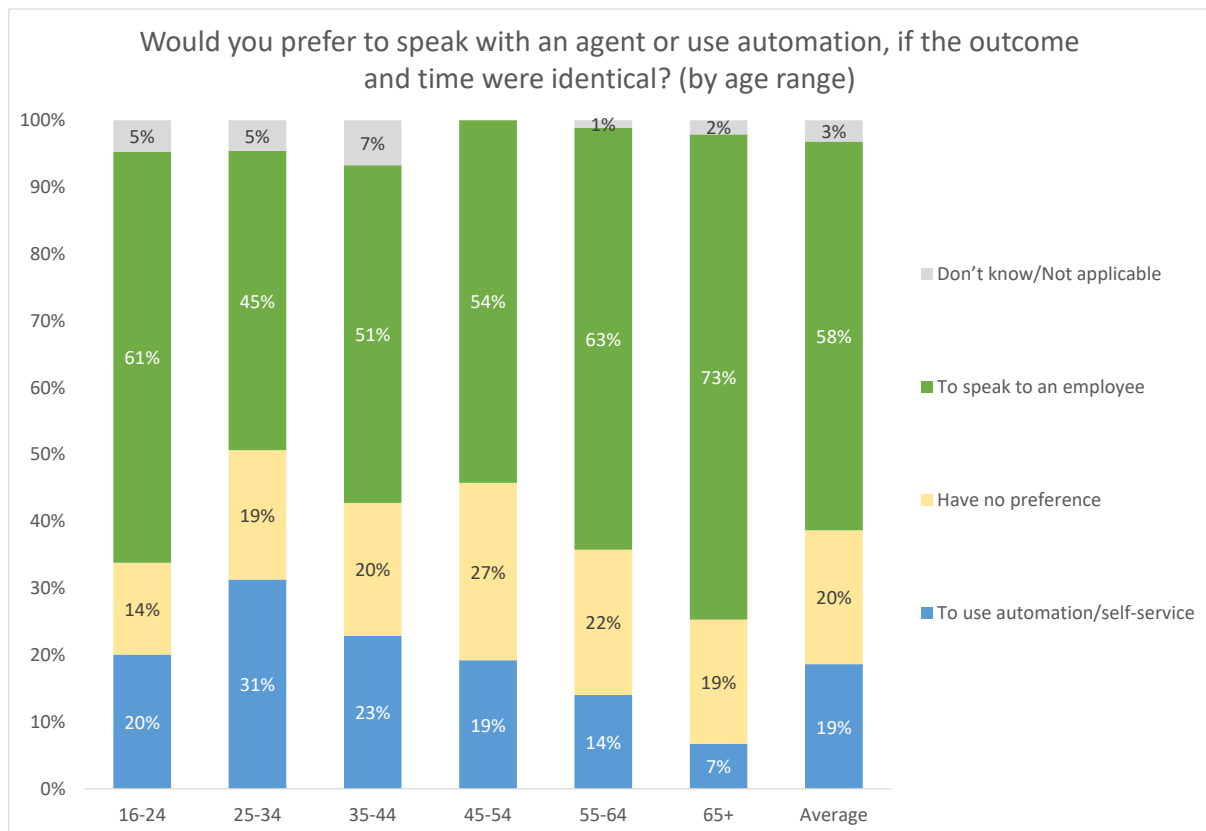


In order to gauge the level of acceptance and expectation around fully automated customer contact, 1,000 UK consumers were asked whether automation or human assistance would be preferable to the customer base if the customer effort, time and outcome were **exactly** the same. Bearing in mind the rapid advance and uptake in digital channels, the findings were quite surprising.

Looking at the age group of the customer base, older demographics feel more strongly about human contact, with 25-54 year-old customers being the age groups most likely to have no preference or to choose to use automation. This fits in with the previous findings that this section of the customer base places more value on their time, whereas the older demographic prefers to have their issue resolved first-time by a single employee. Interestingly, the youngest demographic still feels the need to speak with an agent, perhaps as they are relatively less experienced in dealing with businesses.

Bearing in mind that this question emphasised that the outcome and customer effort/time **would be identical** in each case, the results show that the customer base at present is not yet at a stage where automation is generally seen as being on equal terms with human contact, let alone the preferred method of contact with a business.

Figure 22: Would you prefer to speak with an agent or use automation, if the outcome and time were identical? (by age range)



Further analysis of this data showed that 58% of men preferred to speak with an employee, compared to 59% of women. At a socio-economic level, 21% of the AB class would prefer automation, compared to 15% of the DE group.

MAKING AI A SUCCESS AND AVOIDING PITFALLS

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 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 again
- 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 call 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
- 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 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
- Not having a smooth handoff to live agents with full context
- High maintenance requirements from continuously training AI models.

END-USER QUESTION #6: IS THERE ANYTHING THAT SUCCESSFUL AI IMPLEMENTATIONS / PROJECTS HAVE IN COMMON? ANY PITFALLS TO AVOID?



The best AI implementations are focused, beginning with the end in mind – think and start with a simple solution. Not all good solution involves AI. If a simple verification question could save a lot of time, ask that first. If linear regression solves 80% of the cases, don't use neural networks.

It's also important to make sure the solution applies to majority of the cases. Don't use AI to solve a problem that's only applicable to a small percentage of customers.

Here are some tips that should be considered for any AI implementation:


- Be iterative. Sometimes it takes many cycles of review and improvement to get the most out of AI.
- Consistently reflect and self-criticise. It's easy to assume that a model does what we want because we built it, but the fact is, a model built on real-world data will reflect real-world biases, and it's important that we're aware of that.
- Be realistic. It's important to know now what ML can and can't do, and what the limits are. Trying to create General AI, or one that requires too much unexplainable context is doomed to failure. Keep it small, and in the realm of the ML you are able to use.
- Stand on the shoulders of those who came before you. Even if your problem hasn't been solved, it's possible someone else has already solved a similar problem and you can start where they left off.

THE VIEW FROM CX DECISION-MAKERS

Survey respondents were asked their opinion on how important various customer experience developments would be to their organisation in the next two years.

Perhaps the most striking finding was that the most important factor determining the future success of the customer experience programme was not technology-related, but rather a requirement for the continuing and strengthening executive commitment to improving customer experience, without which the multi-departmental CX initiatives could not hope to succeed.

Figure 23: Importance of CX developments in the next 2 years



All of the customer experience developments and investments offered as choices to survey respondents were seen as being important, but it is interesting to note that omnichannel and the application of artificial intelligence to self-service were not yet universally seen as being vital to the success of the customer experience programme, supporting the view that AI is still at the beginning of its use in the customer contact space.

THE FUTURE ROLE OF AI IN THE CONTACT CENTRE

There seems little doubt that the eventual overall roadmap of the contact centre industry will lead to significant levels of AI involvement in customer contact. In the foreseeable future, AI is likely to be used to improve self-service and assisted 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 do well.

Businesses' interactions with customers will become a highly polarised mixture of the automated and the personal touch. Moving a large proportion of interactions onto self-service reduce 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, offering personalised service without the need for a human agent in some cases.

We can expect to see personal technology applications seeking out the best deals on offer, or interacting with a business on behalf of customers 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.

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 people call people 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 experience that this is the most effective way of making sure their issue is resolved.

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: the rise of the robots will be slow, but inexorable.

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.

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.

We help the biggest and most successful vendors develop their contact centre strategies and talk to the right prospects. We have shown the UK government how the global contact centre industry will develop and change. We help contact centres compare themselves to their closest competitors so they can understand what they are doing well and what needs to improve.

If you have a question about your company's place in the contact centre industry, perhaps we can help you.

Email: info@contactbabel.com

Website: www.contactbabel.com

Telephone: +44 (0)191 271 5269

FURTHER CONTACTBABEL REPORTS

The 2021 UK Contact Centre Decision-Makers' Guide: results of the largest annual survey of UK contact centre operations. Free to download.

The 2021 UK Contact Centre HR & Operational Benchmarking Report: detailed information on salaries, attrition, absence, recruitment and performance benchmarks, costing £350 + VAT.

The Inner Circle Guides: detailed analyst reports on key technologies, including:

- Cloud-based Contact Centres
- Customer Interaction Analytics
- First-Contact Resolution
- Fraud Reduction and PCI DSS Compliance
- Omnichannel
- Outbound & Call Blending
- Remote Working
- Self-Service
- Voice of the Customer
- Workforce Optimisation.

All Inner Circle Guides are free to download. Further information and downloadable reports can be found at www.contactbabel.com.