How conversational AI is enabling the experience economy

Received (in revised form): 12th January, 2022

Laetitia Cailleteau
Managing Director, Accenture, UK

Laetitia Cailleteau is a managing director at Accenture and leads the conversational AI domain globally, driving innovation, sales and delivery across the globe for multiple industries and clients for more than 22 years. She has authored several academic publications and holds patents in the conversational AI domain. Laetitia founded Accenture's Innovation Studio in London, a space for co-creation with clients to realise the potential of artificial intelligence, augmented reality, Internet of Things, using user-centred approaches. She is a reserve member of the Artificial Intelligence High Level Group at the European Commission and also a board member of OASIS Open Foundation in Europe.

Accenture, 30 Fenchurch St, London, EC3M 3BD, UK
E-mail: laetitia.cailleteau@accenture.com

Emer Gilmartin
Engineer and Computational Linguist, ADAPT Centre, Trinity College Dublin, Republic of Ireland

Dr Emer Gilmartin is an engineer and computational linguist. She researches human–machine interaction at the ADAPT Centre, Trinity College Dublin, with particular emphasis on making conversational AI more natural and efficient. She has published widely in the field. She also builds spoken dialogue systems for education, health and commercial purposes, and consults to industry on all aspects of natural language interfaces.

Address: ADAPT Centre, Trinity College Dublin, Republic of Ireland
E-mail: gilmare@tcd.ie

Fuad Hendricks
Senior Manager, Accenture, UK

Fuad Hendricks is part of Accenture's Data&AI Europe group, where he leads the delivery of cross-industry investments focused on differentiated transformation strategies, offerings and services to rapidly deliver value to clients where data is at its core.

Accenture, 30 Fenchurch St, London, EC3M 3BD, UK
E-mail: f.hendricks@accenture.com

Abstract The use of conversational AI (CAI), already experiencing strong growth, has accelerated throughout the COVID-19 pandemic. Natural or human language interfaces enable engaging, satisfying and efficient experiences across platforms and sectors. This paper briefly introduces CAI, provides a snapshot of where the technology is, with examples of recent use cases. It then explores the growing area of ethical considerations, best practice and legislation for CAIs, and maps human rights concerns to the technological landscape. It concludes with pointers towards useful, ethically sound applications where users can collaborate easily and flexibly with CAIs.

KEYWORDS: conversational AI, user experience, ethics, responsible AI, digital literacy, platform architecture

INTRODUCTION

Conversational AI (CAI) promises to let us use our natural language to communicate with services and systems facilitated by machines. Until very recently, interacting with computers relied on specialist knowledge and was only available on desktop computers. Recent advances in natural language processing and exponential growth in computing power have brought a wave of artificial agents or natural language interfaces (chatbots, virtual assistants, dialogue systems) which allow users to talk (or text) systems and access services through smart speakers and devices in their homes or the mobile messaging channels which have become familiar to large sections of the population.

CAI allows organisations to extend their services to people who do not have computers to access online services. This opens services to a wider group of users, delivering the same or similar experiences without requiring the requisite knowledge of the more digital-savvy. Depending where users are identified on the digital spectrum, a good mix of channel engagement can assist customer interaction. For example, telephony may still account for large volumes for managing customer needs, whereas more digitally literate, may want to engage through social messaging channels via smartphones. Using natural language interfaces helps bridge the gap between talking to a person and interacting efficiently with a machine. Thus, CAIs can help close the digital divide, encouraging inclusion across a wider spectrum of users. Organisations can also conserve resources by automating repeatable tasks through CAIs.

In addition, CAI helps companies leverage organisational data to extract meaningful insights and derive business value. AI helps unearth these data points that enhance the client experience by using the right data at the right time. For example, when users call into a contact centre, we can anticipate the typical questions that could be posed based on previous engagements with the client, such as correspondence, any call to action, etc., and be ready to respond or prompt as required. Other tools such as sentiment analysis can help systems respond appropriately and manage users more effectively. These new technologies can open ethical and legal questions, however, with issues of trust predominating. There is a need for clear guidance on the responsible and effective use of CAIs.

Below we briefly describe what CAI is, why it works, how it can be responsibly used, and provide guidelines for effective implementation of these technologies.

WHAT IS CAI?

People talk. We structure and experience our lives through communication with each other using natural language: spoken or written, formal or informal. Our social and pragmatic goals are met through conversation or dialogue — from casual chats about the weather and discussions with friends and family, through commercial transactions and consultation with experts to formal legal and ritual processes. We have built our societies, technologies and industries through our linguistic abilities.

Although we increasingly interact with machines and systems, these encounters have not always resembled human–human communication. Most human–machine interaction has relied on rigid physical or graphical interfaces — from buttons and dials on physical machines to coding languages and later menus and graphical user interfaces (GUIs) on our computers. These interactions require human users to learn and conform to the machine's communicative paradigms. As the range of applications of digital technologies grows, traditional interfaces become increasingly rigid, and interactions which could be accomplished easily using natural language become more complex to implement using GUIs.

Talking to machines as we would to another person is not a new idea but

requires a number of tools. Natural language processing (NLP) — the technology of understanding and generating natural language — has been under active development since the mid-20th century. Text-based and spoken dialogue systems — the larger group of technologies involving natural language communication — have been developed and deployed with real users for specific use cases in education, military training, healthcare and medical training, and transport. These systems have ranged from simple text and phone interfaces to embodied conversational agents (ECAs), capable of understanding and producing speech, vocal gestures such as laughter or feedback, facial expression and bodily gesture, posture and eye gaze. In the commercial landscape, interactive voice response (IVR) systems in telephony became common through the 1990s and 2000s, often based on VoiceXML. The usability of these systems was hampered, however, by the lack of computing power needed to understand the enormous variability of natural language, and they were limited to rigid menu-like dialogue structures, such as the familiar 'Press 1 for …' type dialogue encountered in IVR systems.

Dramatic advances in computing power and mobile technology since the mid-2000s have driven exponential improvements in NLP, and particularly in natural language understanding (NLU). These advances, combined with people's massive adoption of messaging services as a communications platform of choice, have resulted in an explosion of popularity for simple text and speech-enabled dialogue systems or CAIs, embedded in messaging channels. The COVID-19 pandemic has further fuelled growth in the sector, as the usage of and familiarity with messaging and video meeting platforms has grown in all sectors of the population, for social, business and educational purposes.

CAIs are dialogue systems or natural language interfaces which can map input from a user to an intent — an understanding of how the system can meet a user's needs. This intent is then mapped to an action the system can take to fulfil the intent. Any natural language interface needs to:

- Capture what a user is saying — whether voice or text;
- Decode the message to words in the case of voice through automatic speech recognition (ASR);
- Translate this message to machine-understandable form (NLU);
- Manage the process of getting the appropriate response to the user's request; this could be as simple as looking up an answer to a question under frequently asked questions (FAQs) or as complex as performing a number of tasks via application programming interfaces (APIs) to remote systems — dialogue management (DM);
- Formulate a human-understandable message back to the user — natural language generation (NLG);
- Render that message in text, speech, graphics, etc.

These basic elements can be augmented in various ways to offer a richer experience to the user and better understanding by the system (see Figure 1). CAIs are often used in situations similar to smart speakers (Alexa, Google Home), where a user gives a single command, eg asks a question, and the system fulfils the intent invoked in a single turn or very short exchange. The NLU does much of the work in these interactions by disambiguating the user intent from the list of intents available in the CAI.

Most commercial CAI platforms (CAIPs) also allow for 'slot-filling' dialogues, where the task involves getting several pieces of information from a user in order to fulfil an intent — as in a commercial transaction where the user must supply, eg dates, places and other features. Slot-filling architectures allow flexibility in that the system can use

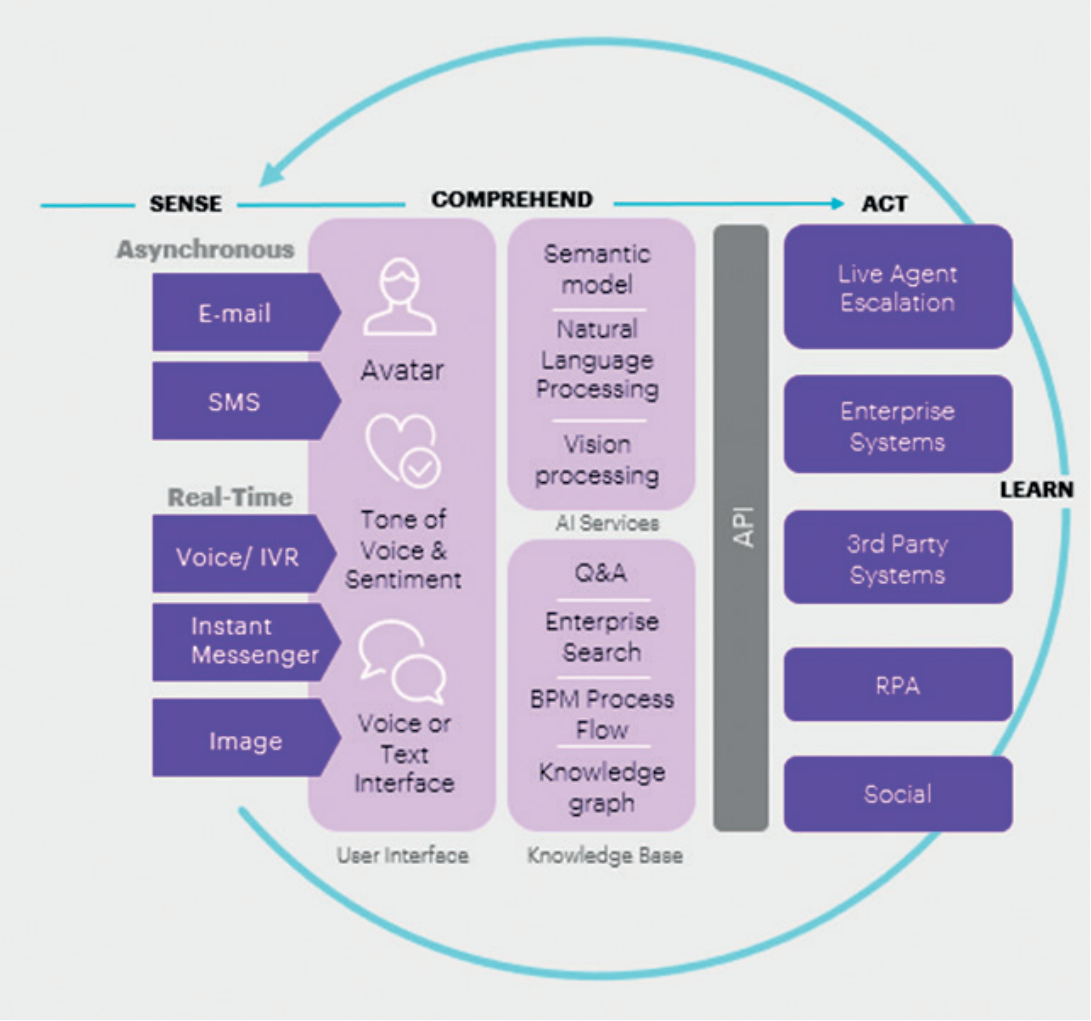

Figure 1: CAI elements

information from the user to fill several slots and also dynamically ask for missing items, rather than rigidly ask for one item at a time. As an example, in a flight booking system, the user can specify some or all of the necessary information (destination, date, type of flight) in the initial request and the system will automatically ask for any missing information rather than rigidly asking for each piece individually. Other features include automation of error/disambiguation routines, small talk, and the facility to hold one task in order to fulfil another proposed by the user and then resume the original task.

WHY CAI?

These newer, more flexible conversational interfaces allow users to interact with a system using speech or natural text, accomplishing tasks through linguistic communication rather than a preset series of menu items. Such an interface is more intuitive and efficient for human–machine interaction, particularly on mobile devices using communication channels (messaging apps) already in use for human–human interaction. This 'conversational' paradigm offers three main advantages:

- The first advantage is that by using conversational interfaces, the user does

not have to learn the system's way of communicating but rather interacts with the system in a natural fashion, where the system acts like a human interlocutor or agent;

- The second advantage is that most CAIs are deployed through messaging channels, including MS Teams, WhatsApp and Slack, allowing users to call or invoke an agent with a simple message, rather than switching out to open a downloaded and installed mobile or desktop application. This means that CAIs can do the work of such applications, with the advantage of the application 'living in the messaging channel';
- The third advantage is more subtle and likely to become important as systems become more efficient and NLP and spoken language processing (SLP) improve. In a sufficiently human-like chatbot or dialogue system, the user relates to the system as another entity or agent rather than as a machine interface. This 'agent' versus 'interface' distinction is the driver of user engagement with the system and thus with the services mediated by the system.[1] If users see the system as a relational agent, with human-like behaviour and motivation, their attitude and actions may follow deep-seated social norms of politeness and involvement, and attachments may form. This treatment by the user of the system as an agent greatly increases engagement and retention of users in contrast to systems which are regarded by users strictly as software interfaces or tools.[2] Hard-coded and ancient human social norms are thus leveraged to build a relationship, either short-term for the duration of the interaction or longer-term over several interactions, in which the system takes on a persona and becomes part of the user's business or social community, with obvious benefits in terms of user engagement and retention.

In the broader business context, CAIs can provide solutions and savings anywhere a predictable or routine interaction is needed. CAIs can be used to replace information points such as FAQs with more flexible question-and-answer (QA) systems providing more tailored answers and understanding more varied questions, and can carry out simple commercial transactions such as bookings and sales or internal functions such as human resources (HR) procedures and training.

WHAT CAN CAIs DO?

With better NLU, modern CAIs can better understand user input and match to particular responses, tasks or intents. Systems can handle under-informing and over-informing by the user, automatically generating further queries to help accomplish the task in hand. Although the probability of correctly understanding input has risen greatly, well-designed CAIs can make timely and graceful handoff to a human operator when necessary. Systems can also handle a number of different tasks, including holding an ongoing task when a new task is introduced and fulfilled and resuming it seamlessly. Current CAIPs easily meet well-defined use cases which involve short interactions. These include such as customer relationship management (CRM), FAQ and information delivery to users, simple sales transactions, and are particularly suited to information, marketing, branding and conversational commerce applications. More complex applications can be implemented with careful and proficient design, but such applications remain mostly limited to graph-based or slot-filling dialogue flows. A popular application of CAIs is incorporation in new or existing IVR systems, adding better NLU to make IVR dialogues more flexible and conversational.

Much of the original development on the dialogue architectures used in CAIs was at a time when telecommunications were strictly

unimodal — either voice or text, telephone or telex/fax —and GUIs were just emerging. Thus early systems were based on exchange of linguistic information only, although people regularly combined speech and shared visual information to perform tasks; some of us will remember the days when we talked to a travel agent in their office while looking together at a screen displaying various options.

Modern smartphone and messaging channels are multimodal, allowing audio and visual information to be displayed in the same application. This allows CAIs to display information in a range of formats, using a mixture of text, graphics, audio and video — all of which can easily be incorporated into the rich media, card-based messaging formats available in messaging platforms. This allows for efficient communication of different types of information using the most suitable modality. Additionally, there are places in an interaction where user input needs to be very concise and clear — confirming a transaction or choosing among alternatives, for example — where rich media elements such as buttons, lists and carousels can be very useful and efficient.

The benefits of CAIs to the bottom line are many (see Figure 2) and there are also benefits to several organisational areas in addition to those resulting from better and faster service to customers.

In terms of sales and service, CAIs can:

- Increase first-time resolution rate while decreasing resolution time on calls;
- Drive continuous quality improvement through improved pace and consistency on interactions;
- Save costs, not only in terms of lower labour costs but also in reduction in need for physical space;
- Improve cross-selling or lead generation effectiveness, as bots can be linked to allow the virtual agent to propose extra services during an interaction. For example, if a customer is arranging a mortgage, the agent can offer services such as movers or insurance;
- Tracking of customer engagement or satisfaction on parts or all of the conversation through immediate feedback during the call or through metrics such as rates of call-back, abandonment or escalation to human agents.

The use of CAIs can also increase employee satisfaction and effectiveness:

- Employees are released from boring or routine tasks to work on more complex and satisfying interactions with customers;
- CAIs can be used to aid employees during calls, where assistant agents can provide information, suggest next best actions and provide an estimate of how a call is going.

Other benefits include:

- Extended hours of service, multilingual service;
- Easy adaptation of services to new tasks, markets and languages without the need to train new staff;
- Elasticity of provision — the offering can meet surges in demand without the need for extra staff;
- Traceability and searchability — calls, conversations are transcribed and searchable for particular queries but also for service improvement purposes.

CAIs can also further responsible business practices:

- Services can be made more inclusive, as CAIs can be designed to best serve different demographics. In addition to providing services in a number of languages, CAIs can be adapted for communication with more or less technical users, different cultures and communication styles and differing literacy levels. Systems can switch communication style depending on the user profile or cues gathered during an

interaction. Such personalisation has a long history in dialogue technology research. For example, the long-running Let's Go experimental voice transport information system from Carnegie Mellon University was adapted to better communicate with non-native speakers, seniors and users with varying degrees of knowledge of the city,[3] while health CAIs from Tim Bickmore's Relational Agents lab have been shown to be very effective with users with low levels of health literacy.[4] Systems can also tailor conversations based on user history;

- CAIs can also be more environmentally friendly. As an example, in measured implementation for a Nordic client, it was found that a CAI system was emitting 1,000 times less CO2 than a traditional call centre.

In summary, CAI is important to chief experience officers (CxOs) as it becomes a keyway for brands to interact with customers or citizens as we go in the Fourth Industrial Revolution. This is as revolutionary as websites in the Internet era.

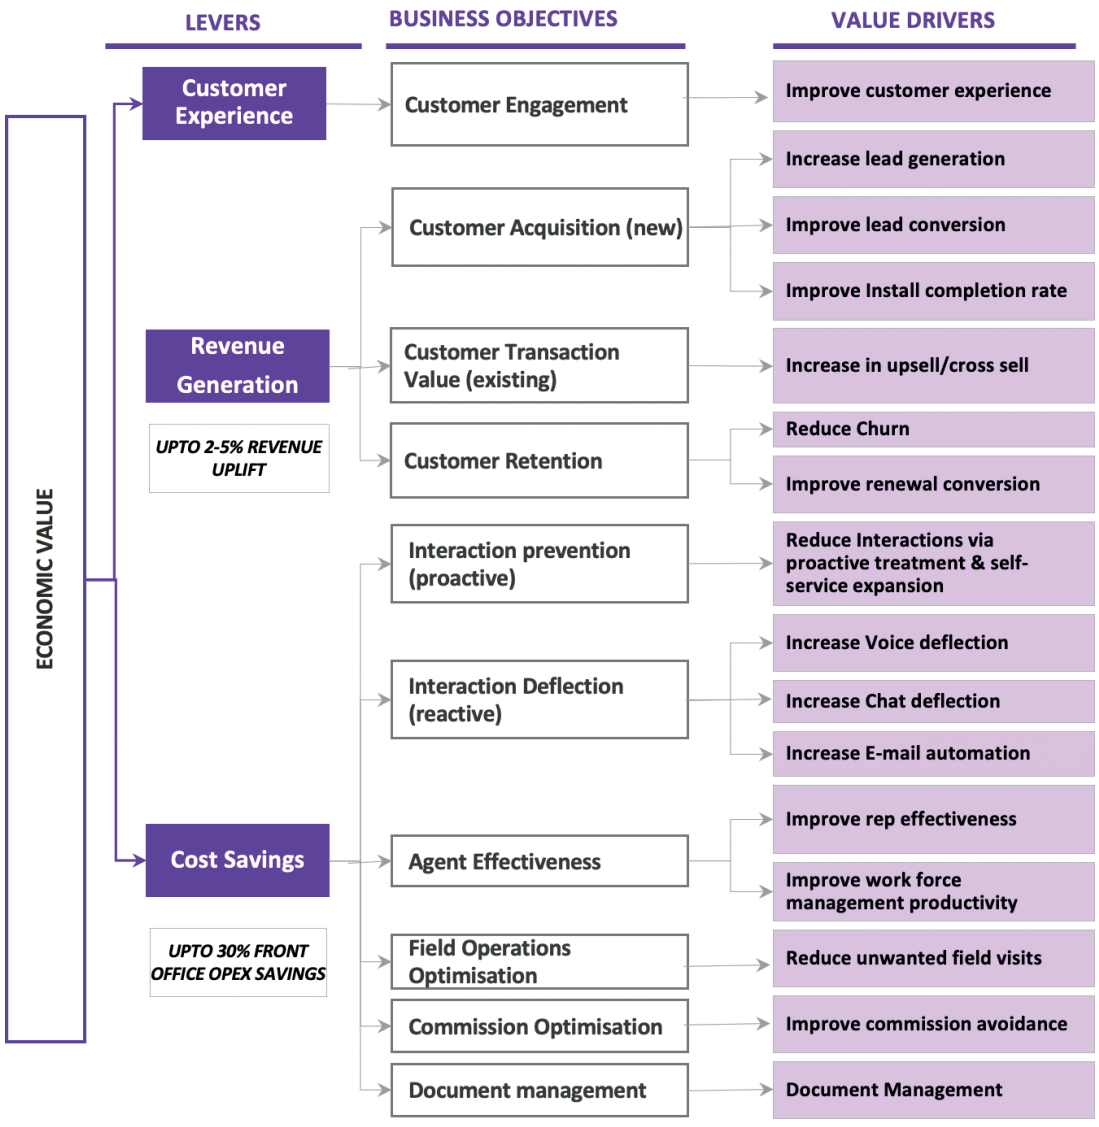

Figure 2: Value delivery through implementing CAI

CAI IS HERE TO STAY

In 2011, Gartner Summits projected that more than 85 per cent of customer interactions would be managed without a human by 2020, while a 2016 Techemergence survey of 35 AI start-up founders/executives on future trends in AI found that chatbots were expected to be the number one consumer application of AI over the next five years.[5] Some scepticism and concerns were subsequently raised, however, about over-hyping and inflated expectations of chatbots, which led to a dip in growth in early 2017, although interest has since picked up again. Gartner now estimates that there are more than 1,500 conversational platform vendors worldwide, with massive variation in quality and depth of offering.[6]

The CAI industry is built on AI and NLP services, recently made available on the cloud. Although there are some firms using custom NLP, most CAIPs use cloud-based services provided by large multinationals, including Google, Amazon (Lex), Microsoft (LUIS) and IBM. Although the big players do provide CAI development platforms, a number of smaller companies have developed richly featured low-code or no-code visual design solutions, often templated for specific use cases. These can often be extended with more expert programming (at the software development kit [SDK] and API level) to implement more complex work and conversation flows.

CAI technology development is moving quickly, with several features on the horizon (see Figure 3). While relatively simple CAIs have been very successful commercially, it should be noted that many of the currently emerging CAI capabilities, particularly around sentiment analysis, multimodal communication (eg gesture and voice biometrics), have already been developed as research projects and for bespoke applications; an interesting example is the virtual human technology produced by the University of Southern California's Institute for Creative Technology.[7] Such applications can require more complex dialogue management than that currently used in the CAI sphere, but much of the groundwork has been done.

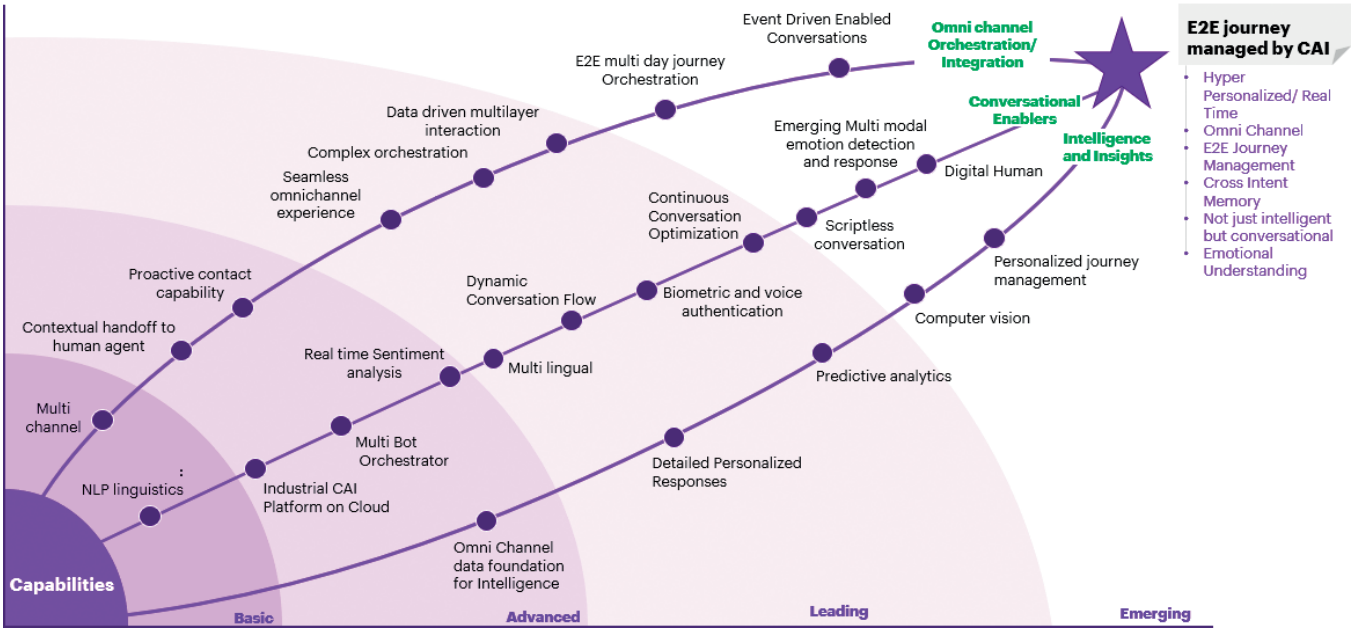

Figure 3: CAI solution maturing curve

CAI IN ACTION

The examples below demonstrate how CAI has helped to reduce costs and enhance customer experience in three different sectors.

Case study A: Large Spanish insurance company

A Spanish insurance company combined IVR and CAI for user identity verification and document management-related enquiries, thus solving for initial customer touch points and avoiding any delay in the verification process.

The agent handled verification of the user's identity before transferring them to a human agent who would take care of their enquiries. In addition, to address document management-related enquiries, the user was offered, via the IVR, to be assisted by a chatbot that could be accessed via a link sent by SMS to the user's mobile devices.

What the solution offered:

- Natural language capabilities — by training the virtual agent to interpret and interact with customers in natural language, the agent was able to engage in a natural human-type conversation with the customer. The system is built on Microsoft Bot Framework using Microsoft LUIS for NLU;
- Designed to secure and mask the customer information before and after making the identification process;
- Seamless integration with other channels such as a web portal or a mobile application;
- Document management-related enquiries to handle text-based interactions.

By implementing the solution, the virtual agent could handle 900 customer interactions daily, with over 70 per cent of users choosing to address their document management-related enquiries through the chatbot.

Case study B: Irish public services organisation

The public services organisation wanted to deflect calls from the contact centre with a voice-enabled virtual agent integrated into the telephony system. The call centre currently deals with approximately 55,000 calls per year about tax clearance. Consumers using the service often have questions encompassing how to apply, application status and application queries. The company developed a customer-facing CAI, accessed via telephony, targeted at automatically handling a portion of these customer calls.

The solution:

- Used an AI-powered virtual agent, leveraging best-in-class AI services for NLP, speech-to-text and text-to-speech. The virtual agent could be easily configured to switch between different AI vendors through use of Accenture's conversational AI technology-agnostic platform;
- Integrated with existing IVR systems using Accenture's patented voice adapter. The virtual agent was accessible to callers through the existing Cisco IVR system to allow seamless call routing between virtual and live agents.

The organisation was able to drive efficiencies by implementing a service that could handle tax clearance queries. Approximately 30 per cent of calls were handled by a combination of human and bot. Based on details passed over by the IVR system, the bot could offer proactive suggestions, helping to streamline the resolution of calls across six customer segments, handling 15 call reasons and manage over 55,000 calls previously handled by call centre agents.[8]

Case study C: Large US telco

The company wanted to broaden its customer care offerings by increasing

self-service capabilities in order to manage costs and drive the end-to-end experience without human intervention. The solution would leverage conversational AI to decrease call volumes and provide unique, pleasant, differentiated experiences.

The CAI achieved:

- A 75 per cent automation rate (ie customer completes transaction with the bot with no human intervention);
- Five million annual call deflections, with 380,000 calls deflected monthly;
- Satisfaction rating of 75 per cent, an increase of over 200 per cent from previous level.

Using the service, the company was able to save US$18m in operational costs over six months, with a projected US$150m saving over three years.

Work on use cases such as those described above has exposed a number of challenges in the introduction of CAIs and given rise to valuable lessons in best practice. These are described in the following section.

CAI AND ETHICS

All good relationships are based on trust. Nobody likes to feel duped. The technologies behind conversational AI — NLP, sentiment analysis, facial recognition, and so on — are now so powerful it is becoming hard for people to tell whether they are speaking with a human or an algorithm. CAIs and the organisations that deploy them need to be open and honest about their use, or risk losing customer trust and confidence.

There is also a whole spectrum of risks relating to gender and racial stereotyping, data privacy and potentially manipulative behaviour that companies need to be alert to. Recently, concerns have been raised very publicly around algorithmic bias, where pre-existing societal biases and inequalities can be reflected in the data used to train some algorithms — and, by extension, CAIs. Although there is significant work being done to resolve these issues, companies do need to be vigilant. A key point to understand is that these ethical risks are not due to technical limitations, but tend to arise as the result of decisions taken at the design stage, where the potential downstream impact may not always be fully appreciated.

Balancing the engagement with humans and machines: Responsible business

The relationship between humans and machines is not new. Machines have been around for a long time, from industrial applications to the phones we carry in our pockets. They have become part of our everyday experiences. As AI goes mainstream, it creates a new paradigm of human–machine engagement with attendant questions of how these new relationships should be managed. As human–machine affinity shifts, we need to strike a balance by implementing the technology responsibly. In particular, two dimensions should be considered:

- Autonomy of decision-making capabilities based on algorithms;
- Wider societal effect of AI technologies creating anxieties when people spend more time engaging with machines.

As technology has evolved, so too has our interaction and engagement or affinity with machines. The ‘interlocutor model’, or the way a human user sees a machine they interact with, is transitioning rapidly. As CAIs are integrated into our daily lives, our perception of them moves from dumb servants to something approaching a companion. We have noticed various degrees in this affinity and acceptance of these relationships and recognise the need to model them. We apply the metaphor of the focus and impact on an individual’s human rights as outlined in the European Convention on Human Rights. We have

developed this approach and ideas to best craft how conversational AI can provide the right experiences and value levers for users. Below we focus on three of these fundamental rights (summarised in Table 1) to illustrate our approach.

The right to equality and freedom

Looking human and the mimicry of human features and characteristics to better engage end users creates a case of commonality; however, this must be done in such a way that the solution does not infringe the right to equality and freedom. Poor or unconsidered choice of characteristics such as gender, age and ethnicity for a CAI could result in a sense of discrimination, stereotyping or reinforced bias.

The right to privacy

Safeguards are needed to guarantee against intrusive profiling. Behavioural and transactional data can be collected and machine learning (ML) methods used to make inferences about the user's emotional state or intent. This can be very useful in informing a CAI how best to communicate. As these methods give machines ever-enhanced understanding of humans, however, the danger of infringing privacy becomes increasingly relevant, and users may feel uncomfortable about how their data is managed. Careful data management, clearly informing users and adherence to best practices ensures the right levels of data protection is established.

The right to freedom of opinion, choice or thoughts

When CAI design creates human-like behaviours through simulation and stimulation of emotions to better engage with users, it must be made very clear that the CAI is a machine, to avoid exertion of empathetic persuasion or decision guidance that leaves a user feeling coerced into directed thinking.

CAI AND SOCIETY: REGULATION OF AI

AI, as we know, is an industry growing at breakneck speed. Company adoption of AI is becoming mainstream in its use of the technology. The use of large datasets required to derive maximum benefits from the technology does present inherent risks in the management and governance of datasets, especially relating to personally identifiable information (PII). Exposing data of this type can introduce many societal risks or negative consequences and fundamentally threaten a person's rights. To prevent and avoid threats of this nature, regulation must be introduced, and laws designed to regulate the use of data. We have observed that regulations may not have kept pace with the development of the technology, and thus there is a lag in policing of how companies make use of this technology.

Policy makers have realised the importance of regulations in this space — organisations such as the Organisation for Economic Cooperation and Development (OECD) (OECD AI Principles), Institute of Electrical and Electronics Engineers (IEEE) (definition

Table 1: Responsible use of CAI

	Looking human	Understanding the human	Behaving like a human
Description	The mimicry of human features and characteristics to better engage end users	Collecting behavioural and transactional data to infer the user emotional state or intent	The simulation and stimulation of emotions and behaviours to better engage with them
Effects	• Discrimination • Stereotyping • Reinforced bias	• Private data collection • Emotional interpretation and inferences	• Empathetic persuasion • Decision guidance
Impact	The right to equality and freedom	The right to privacy	The right to freedom of opinion, choice or thoughts

of standards for AI) and the European Union (EU) (defining AI regulation). The pandemic has accelerated the digitalisation of organisations and has acted as a catalyst for quicker progression of laws in this space. In June 2021, the European Commission (EC) released its proposal for the Artificial Intelligence Act,[9] legislation for a coordinated European approach on the human and ethical implications of AI. As a next step, organisations should understand the implications and prepare their systems to support compliance with these regulations.

As a result, AI chatbots have been described as limited and minimal-risk AI systems, ie AI systems with specific transparency obligations: when using AI systems such as CAIs, users should be aware that they are interacting with a machine so that they can take an informed decision to continue or step back. Further policy developments are expected in this space.

The European Court's Schrems II judgment was written into law in September 2021 by the EU as a new way of dealing with personal data outside of Europe. This has implications for call centres in particular. As a result of Schrems II, many call centre operations that have been outsourced will have to come back into Europe because customers' personal data cannot be transferred or viewed outside of the EU. Moving operations back into home countries will have an impact on operational costs, and may also lead to increased levels of automation to cope with the number of customer interactions coming through. The impact to organisations is material to organisational processes, technical capabilities and operations.

DESIGNING AND IMPLEMENTING GOOD CAIs

Planning interaction and engagement with end users, whether in-house or customers, are crucial steps for any CAI solution. A bad experience can change customers' view of a brand forever, while positive experiences can build loyalty and drive growth and become a massive enabler of the experience economy.

Below we outline important considerations in the process of designing and implementing CAIs well, many of which are encapsulated in Accenture's golden rules for CAI development (see Figure 4).

Figure 4: Golden rules for CAI

There are several CAI platforms on the market which organisations can use to implement a conversation AI solution. There are fundamental constructs to consider when deciding on the solution:

- Will it provide the right capabilities to meet business objectives? Can new conversation flows be easily designed and implemented when requirements change?;
- What is the complexity in setting up the platform? Can it integrate within existing environments and technologies?;
- Getting the right skills — in addition to technical architecture, engaging and easily understandable conversation design is essential. Conversational language (voice and text) is very different to more formal written language. A good conversation designer will craft system prompts and dosing of information to meet human affordances.

Define clear objectives and metrics

Use cases identified should be aligned to the organisation's goals and objectives. This creates opportunities to scale function and allow for enterprise-wide adoption. When aligning goals, value architects and designers should create the required metrics to reflect how the function will be measured and tracked.

Reinvent user experiences

The virtual assistant is in essence an extension of existing business capabilities in terms of task goals; however, the experience should be designed from the ground up. It is important to exploit the multimodality offered by modern mobile devices. For example, voice, text and visuals can be combined in an experience to boost its efficiency — a booking application can use voice or text to elicit user input while also displaying relevant information as a table or form. It is also vital to design conversations to match human affordances. A typical 'rookie' mistake is to try to directly take content from, for example, a form, FAQ or web app, and reproduce it in a CAI. This can result in unwieldy CAIs with overlong unnatural prompts and rigid dialogue structure. Good conversational design will ensure naturalness and ease of use, leading to enthusiastic uptake by users rather than the frustration often seen with older systems.

It will be advantageous to initially limit the scope of interactions to frequently used functions and design features from these. This approach will help to identify other integration opportunities — for example, connecting with APIs or connecting to existing infrastructure and services.

Find the right platform

Finding a platform able to scale and integrate seamlessly with existing and new services will provide best value in function to the organisation. Due diligence will aid understanding how a candidate platform can meet the varying requirements of the organisation, which will greatly help in choosing the right one.

The CAI vendor ecosystem is particularly crowded and constantly evolving, with many emerging players in the market. There are strengths and weaknesses for every vendor, and currently no single provider has emerged as a clear leader.

The landscape can be roughly divided into two groups:

- *AI platform services and open-source*: Plug-and-play platforms (commercial tools have pay-as-you-go pricing) offer flexibility and scalability but with additional integration work needed. They often require significant upfront training in order to build a robust virtual agent. Example vendors are Microsoft, Amazon, IBM and Google, all of which offer CAI design and implementation as an adjunct to their cloud NLP computing services. There are

also open-source toolkits from companies such as RASA;
- *CAI solution vendors*: These vendors typically offer a whole suite of tools with specialised functionality and often come with an upfront or recurring licence fee. Usually considered for specialised needs, as a platform or an addition to other frameworks. These solutions often offer templates for particular business areas or use cases and can either build and host bots or offer low-code or no-code solutions. Such services are provided by companies including Nuance, Amelia, Artificial Solutions and Kore.ai. Most of these solutions rely on cloud services for their NLP and many of the smaller companies are being acquired by major players.

A key decision to be made is whether the overall solution architecture should be based on custom development leveraging AI cloud services, virtual agent commercial frameworks, open-source or commercial platforms, or a combination of these options. There are pros and cons to each approach, often depending on the particular use case envisaged, and careful consideration should be given to this choice.

Given the speed at which the market evolves, evaluating the tools available with criteria specific to the organisation and project requirements is advisable. All options should be carefully considered with reference to stakeholders' needs and should consider all relevant information. Figure 5 shows possible factors to consider in the choice of a solution.

Design with data privacy first

This is a key concept for developing conversational AI solutions. As the solution becomes an extension of business capabilities, it introduces additional failure points for the business — and as it uses data extensively, the risk becomes higher due to the nature and type of data processed. Virtual assistants usually process personal data and, as stated earlier, compliance with regulations is paramount.

Integration with broader enterprise IT

The CAI solution should have capabilities that allows it to seamlessly integrate within existing environments, but also extend to applications within the AI ecosystem to help enhance functionality. This means the platform can retrieve customer information from internal CRM systems, or link into contact centre technologies.

Does your organisation have the right skills?

Getting the right skills and teams can be a complex exercise. In addition to technical architecture, engineering and data skills, it is important to mobilise conversation designers and domain small to medium enterprises (SMEs) to create engaging

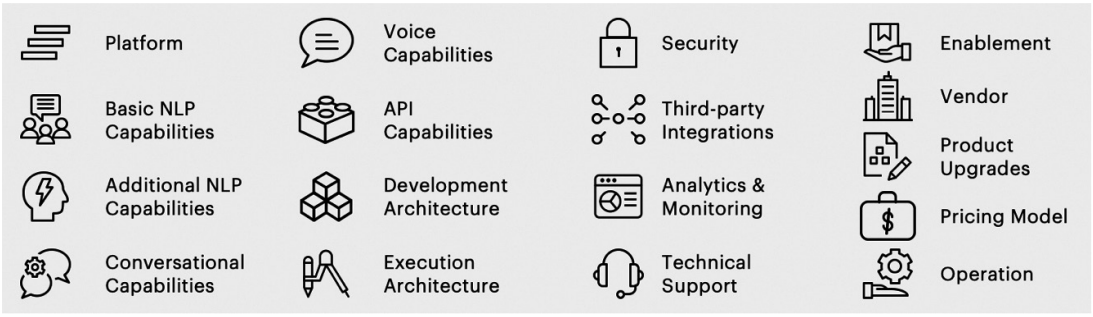

Figure 5: Factors for consideration when choosing a CAI solution

and easily understandable conversations. Conversational language (voice and text) is very different to more formal written language. A good conversation designer will craft system prompts and dosing of information to meet human affordances. Often, internal skills do not exist and may require training or even assistance from outside to help implement the solution. As CAI is becoming more mainstream, the designer market is starting to pivot to this new form of design.

Legal and marketing skills may also be required, especially for new service. Thinking about how the bot looks, its tone of voice, its style, gender — for voice or text bots — is important. Bots ultimately represent an extension of the brand or organisation. The legal element is to establish the approach on how to set an ethical governance and ensure compliance, making sure the right reflections and discussions are happening at the right time with the right stakeholders, that recording and regular reviews are organised, for example.

On the technical side, it is vital to ensure that the platform selected does not require specialised knowledge for the technology, but rather widely used skills that can be transferred to your organisation. Not only do you need to plan for the right team, but also, once deployed live, you should think about the effect on your current operations.

Exploiting organisational data to enhance the experience

Nowadays, everything is data-first, and data is the foundation for organisations to make decisions on how businesses operate. Conversational AI helps exploit organisational data to extract meaningful insights and derive business value. AI helps unearth these data points that enhance the client experience by using the right data at the right time and demonstrating why the service creates stickiness for customers.

CONCLUSIONS

CAI is coming of age and ready to deliver the future of digital experience. It enables strategic, personal, high-touch access to services. It is creating new ways of reaching people, making services available in different ways. It enhances the way organisations engage, creating deeper opportunities to explore and exploit data to enhance or create new services and products. It also drives operational efficiencies, opening up various channels across a range of client segments, ie voice, social messaging channels and mobile. All of these make the business case for utilising CAI increasingly palatable for organisations.

Advances in technology, particularly in NLU and cloud capabilities, allow for simpler development, deployment, integration and scaling of CAIs. CAIPs are actively working to extend the dialogue management capabilities of CAIs to more complex interactions. Many business processes can be fully or partially implemented in CAIs, saving resources and diverting human operators to higher-value tasks. CAI also brings ethical issues into the spotlight, with questions of trust particularly, leading to ethical and legislative efforts to safely manage the growing relationship between humans and artificial agents.

Our hope is that these powerful technologies will help bridge the digital divide, enabling people to communicate more naturally and easily with artificial agents, and allowing companies to use CAIs' ability to provide assistance 24/7 in multiple languages to increase access and take-up of their services.

References

1. Larsson, S. (2005), 'Dialogue systems: Simulations or interfaces', Proceedings of the 9th workshop on the semantics and pragmatics of dialogue, Dialor'05, available at http://dialor05.loria.fr/Papers/06-larsson_final.pdf (accessed 19th January, 2022).
2. Bickmore, T. W., Caruso, L., Clough-Gorr, K. and Heeren, T. (2005), '"It's just like you talk to a friend" relational agents for older adults', *Interacting with Computers*, Vol. 17, No. 6, pp. 711–735.

3. Raux, A., Bohus, D., Langner, B., Black, A. W. and Eskenazi, M. (2006), 'Doing research on a deployed spoken dialogue system: One year of Let's Go! experience', *Interspeech*, pp. 65–68, available at http://www.cs.cmu.edu/~dbohus/docs/letsgo_interspeech2006.pdf (accessed 19th January, 2022).
4. Bickmore, T. W., Pfeifer, L. M. and Jack, B. W. (2009), 'Taking the time to care: Empowering low health literacy hospital patients with virtual nurse agents', Proceedings of the SIGCHI conference on human factors in computing systems, ACM, pp. 1265–1274.
5. Suthar, S. (2017), 'Chatbots Are Pumping Up The Sales Process', Chatbots Magazine, available at https://chatbotsmagazine.com/chatbots-are-pumping-up-the-sales-process-72195dddba7e (accessed 19th January, 2022).
6. Revang, M., Mullen, A. and Baker, V. (2020), 'Architecture of Conversational AI Platforms', Gartner.
7. Rizzo, A., Lucas, G., Gratch, J., Stratou, G., Morency, L. P., Chavez, K., Shilling, R. and Scherer, S. (2016), 'Automatic Behavior Analysis During a Clinical Interview with a Virtual Human', *Studies in Health Technology and Informatics*, Vol. 220, p. 316.
8. View more on this case study: Accenture (February 2019), 'Improving Taxpayer Experience: Revenue's Voicebot', YouTube, available at https://www.youtube.com/watch?v=k9KUCq3Fdi8 (accessed 19th January, 2022).
9. European Commission (2021), 'Shaping Europe's digital future: Regulatory framework on AI', available at https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai (accessed 19th January, 2022).