

The 9 Principles of Conversational AI

HOW TO ACHIEVE PERSONAL INTERACTIONS WITH YOUR CUSTOMERS AT SCALE



Georgian

The use of messaging apps, speech-based assistants and chatbots to automate communication and create personalized customer experiences at scale.



Hundreds of millions of people use Facebook Messenger, Kik, WhatsApp and other messaging platforms to communicate with their friends and family every day. Millions more are experimenting with speech-based assistants like Amazon Alexa and Google Home. As a result, messaging and speech-based platforms are rapidly displacing traditional web and mobile apps to become the new medium for interactive conversations. If you're a business that's important because it will change every aspect of when, where and how you engage and communicate with your customers.

We refer to the use of messaging apps, speech-based assistants and chatbots to automate communication and create personalized customer experiences at scale as conversational AI. Conversational AI applications enable long-running interactions with customers via text or voice using the most intuitive interface available: natural language. Conversational interactions are driven by words, whether in full sentences or in a menu. And, unlike social media, they can support engaging, two-way interactions with private audiences. When combined with automation and artificial intelligence (AI), these interactions can connect humans and machines through virtual assistants and chatbots.

Yet the real power of conversational AI lies in its ability to simultaneously carry out highly personalized interactions with large numbers of individual customers. Conversational AI can fundamentally transform an organization, providing more ways of communicating with customers, while facilitating stronger interactions and greater engagement.

Consumers are already using conversational AI platforms in place of email, phone calls and face-to-face communication to talk with family and friends. Younger people, in particular, tend to prefer messaging over other forms of communication. Not only that, social events are being organized via messaging apps, traditional texting is giving way to more functional and secure messaging applications, and early adopting consumers are using voice interfaces to make purchases.

Today's consumers expect to be able to communicate with businesses on their favorite chat platforms, using text and voice, just as they currently do with friends and family. They want to do so because it's easier, less intrusive and quicker than using other communication channels. For businesses like yours, that means you have to be able to hold seamless, synchronous conversations with consumers across whatever channel they happen to be using at the time, no matter where they are. Importantly, those conversations could be a short one-off request or response, or part of a longer-running customer engagement.

As conversational AI becomes a business imperative, you need to be thinking about how to best prepare your organization for it.



Today's consumers expect to be able to communicate with businesses on their favorite chat platforms, using text and voice, just as they currently do with friends and family.

How to Use This White Paper

This white paper introduces the principles of conversational AI. In it, we also provide a framework for applying them to your organization, a maturity model that you can use to measure organizational progress toward conversational AI and a hypothetical case study to help put the ideas in context.



A Framework for Conversational AI



PURPOSE

Begin by defining the purpose of conversational interactions in your business. Identify where conversational AI will make the biggest impact in your business. Ensure you understand the conversational platforms your users are already on. Document the application programming interfaces (APIs) and services your application will require to create a complete conversational experience.



DESIGN

Recognize that creating conversational applications is different from traditional software development. Conversational design fundamentally is about language and human communication, personas that bring a brand to life, and incorporating the right level of automation.



DATA

Conversational AI applications rely on data, both content and context, to engage users. Content is the information exchanged during natural dialogue with customers while context enables conversational applications to anticipate user needs.



INTEGRITY

Finally, consider the integrity of the application with respect to user data. Set user expectations appropriately around what conversational information you are recording, how it will be used, and how you will protect that information.



Maturity Levels

LEVEL ONE

CAPABILITY AND INGREDIENTS

Your organization has intrinsic company characteristics, assets or strategic intent that give you an advantage for executing on conversational AI. You only have a basic understanding of priorities and opportunities, and there are few or no specialist skills within your organization.

LEVEL TWO

READINESS AND PROCESS ADVANTAGE

You have formalized company standards, practices, procedures and skills that give you an advantage for execution. You have partially documented your approach and priorities, and have some specialist conversational AI skills within your organization.

LEVEL THREE

EXECUTION ADVANTAGE

You're exploiting advantages with mature processes that capture outcomes. You're also using data and insights to drive continuous improvement. Last, but not least, you have complete documentation of your approach, including roadmaps for implementation, and a complete set of appropriate skills to support the implementation of the thesis within your organization.

Case Study—ServiceMatch

ServiceMatch is a matchmaking service for consumers and local service providers. It enables consumers to enter a query, identify a vendor such as a plumber, painter or house cleaner, and then set up a meeting. ServiceMatch then follows up to help consumers resolve any outstanding issues, close out open tickets, and rate their overall experience.

ServiceMatch has determined that its existing web and mobile applications aren't easy to use and that the updates it provides to customers and service providers aren't always timely. ServiceMatch has also noticed that its customers have begun to respond positively to text message reminders that it has started sending out. In some cases, customers are responding to texts as if they expect to engage in a conversation with the company. ServiceMatch now believes that conversation-based interaction is an opportunity to improve communication overall and ultimately drive deeper engagement with all parties.





PRINCIPLE 1

Identify where you'll have the greatest conversational impact

Customer expectations are changing. Anticipate opportunities to meet their unstated needs.

LEVEL ONE

Opportunities identified in an ad-hoc manner.

LEVEL TWO

Business processes partially documented and prioritized.

LEVEL THREE

All business processes prioritized, roadmap developed.

Messaging and voice-based conversational interfaces can add value to businesses by creating engaging user experiences that will reach wider audiences, increase loyalty and drive revenue.

It's important to know where to start by understanding which conversational interactions deliver the greatest value. The most impactful applications are those that engage large audiences through intuitive, personalized interactions and automate business processes to achieve greater operational efficiency. For example, you might be able to expand your potential audience by exposing users to product capabilities via messaging channels, increase user engagement and improve retention through a natural language user interface, or scale up your customer service function while reducing costs by using chatbots.

Identify conversational opportunities

Start the prioritization process by mapping the breadth of how your customers, your customers' customers, your employees and your business partners interact with your business and products through web, mobile, social, email and other channels.

Consider high-friction product interactions that could benefit from context-aware dialogues, such as data entry workflows or interactive product help. Focus on areas where conversations are already happening, such as in the customer care center or throughout the sales process. Include interactions that are primarily informational (responding to questions) and transactional (creating a business request to order a product, or alter a service).

Prioritize opportunities and develop a roadmap

Once you have identified the opportunities for conversational AI, assign a relative value to each opportunity based on the criteria that are currently of greatest strategic importance to your business. If your business is in a growth stage, for example, then opportunities that expand your customer base through outbound messaging and social sharing features may outrank those that increase retention of current customers through personalized engagement.

Use these priorities to develop and maintain a roadmap that will deliver the highest value.



Key Takeaways

1

Identify all possible business processes, both informational and transactional.

2

Prioritize processes that help you meet your strategic business needs.

3

Develop a roadmap for conversational AI that will drive high customer satisfaction and company value.

ServiceMatch

ServiceMatch ran a discovery workshop and mapped out a wide range of potential business processes. The company identified the following high-priority candidate processes for conversational AI:

- New-customer onboarding
- Processing service requests from repeat customers
- Outbound customer contact for customer satisfaction surveys
- All of the above for service providers

The company's strategic business priorities are growth and customer engagement. The ServiceMatch team scored each potential business process based on criteria that support these high-level priorities. As a result of that work, the team decided to initially focus on processing service requests from repeat customers.



PRINCIPLE 2

Understand your audience

Not all customers are created equal, nor should your channels or interactions.

LEVEL ONE

Basic understanding of the platforms the audience is using.

LEVEL TWO

Research and testing used to develop a complete understanding of audience preferences.

LEVEL THREE

Interaction data used to refine understanding of audience's preferences.

The best place to meet your users is where they are. Engaging users on the channels and platforms that they're already using can eliminate the friction and churn caused by requiring them to install and switch to new applications. It also allows you to tap into the larger network of friends and contacts that your users have already established.

Know your audience

Once you have selected business processes to deliver through conversational interfaces, identify current and potential user segments as well as their preferred conversational channels and platforms. Users within the target audience can be segmented by demographics, product engagement levels, the messaging or voice platforms they already use, or the value that they bring to your business.

In addition, you should understand how the different user segments typically interact with conversational interfaces, be it via text or voice. For instance, your audience might consist of millions of millennials who regularly use Kik to share their experiences with friends via messages that include text, emoji and stickers, or they might be a group of baby boomers who primarily communicate about family events via traditional text messages.

When identifying your target audience, use your knowledge of your current audience as a point of comparison to understand user differences and

expectations. In addition, seek to supplement that knowledge by conducting surveys, interviewing existing and prospective customers, and leveraging third-party research. Based on the knowledge you gain about your audience, choose the best channel and platform to deliver the selected business function to your audience.

Choose your channel carefully

It's possible that your customers' preferred mediums might not be optimized for the specific business process that you would like to deliver through conversational interactions. For example, if your current audience uses WhatsApp, the lack of chatbot support on this platform will limit your ability to automate and scale interactions. Similarly, if your target audience uses Skype, the platform's limited security features may make it difficult to develop a conversational banking experience.

In cases like these, it may be best to design the conversational experience for channels that meet your functional and technical requirements, despite the need to recruit and migrate users to them. While these types of limitations will likely diminish over time, you may need to consider bringing users into an alternate channel or modifying the business process that you're delivering.

Continue to evolve your understanding of your target audiences as your user segments shift. Expand the channels and platforms that you are using to deliver the business function accordingly.



Key Takeaways

1

Identify your current and potential audience.

2

Learn what they need, what they do and how best to engage them, including the right platform, channel and device to use.

ServiceMatch

Initially, ServiceMatch provided a pure web interface for its services and then quickly added support for a mobile browser interface. Feedback from customers on the mobile browser experience has been negative. At the same time, ServiceMatch discovered that the majority of its customers are active on Facebook.

Rather than developing a native mobile application, ServiceMatch felt that delivering an experience on Facebook Messenger was critical. This channel also provided an opportunity for customers to rate service providers and share their experience with friends and contacts, thereby allowing ServiceMatch and its vendors to strengthen their brand and expand their reach.



PRINCIPLE 3

Build complete experiences

Ensure that you have all of the right connections to the systems needed for an end-to-end conversation.

LEVEL ONE

Some back-end service integrations accessible.

LEVEL TWO

All service integrations required for end-to-end conversations identified, most are accessible.

LEVEL THREE

Optimized integrations for complete end-to-end conversational experiences.

Any time you deliver a business process through a conversational interface, your users should be able to complete the desired task from beginning to end without needing to manage multiple applications, switch channels or interrupt the process flow. For example, if you're developing a sales chatbot to increase in-store revenue, the interaction should allow the customer to not only discover a new product and ask questions about it, but also locate a store that has the product in stock. Plus, it should provide the option to order the product for in-store pickup or home delivery and accept payments, all within a single conversational interface.

Document the entire process

Document all of the individual steps in the selected process, along with any system access dependencies required to complete the task end to end. Your goal should be to obtain a complete view of how the business process currently operates and to identify all requirements for a full conversational AI interaction.



Key Takeaways

1 Map out the services needed to deliver the end-to-end business process.

2 Ensure that you have access to the full range of required API services to build complete experiences.

ServiceMatch

Having selected inbound customer service requests from existing customers to be the business process to implement as a conversational application, the team has mapped out each of the individual steps within the business process. They are:

- New-feature promotion (promoting new platform functionality)
- Service request (new or repeat)
- Service provider selection
- Scheduling
- Service delivery (human labor)
- Confirmation of service completion
- Rating (of service, and of client)
- Opportunity identification and promotion of related services
- Ongoing nurturing/nudging of the client

ServiceMatch has also identified a number of dependencies on APIs, service integrations and data sources required to deliver an end-to-end experience for registered users via chat. These include:

- User authentication
- User transaction history
- User-to-service provider matching model
- Scheduling API and calendar integration
- Outreach platform, contact history and appointment reminder system
- Rating system
- Opportunity management model and opportunity management system



PRINCIPLE 4

Content enables conversations

Businesses need to capture, curate and create content to deliver the best possible set of conversational interactions.

The best source of content to use in a conversational context is the data you already collect from existing customer conversations through email, phone, or searches they have performed on your website or support knowledge base. Use that collected content as the basis for enabling conversations. Identify gaps and quality issues in content, and develop a plan to address them.

Capture conversational data

Capturing data first involves identifying all reliable existing sources of useful content and creating processes for extracting, combining and storing it in a format that makes it usable for messaging, chatbot and voice assistant applications. This process may include recording and transcribing customer support calls, merging knowledge bases and linking questions with their corresponding responses. Content can also originate from existing product interactions in natural language form, such as alerts, notifications or recommendations sent to users. Third-party data, such as social media conversations about products, may also provide content that can be used to understand user intents.

LEVEL ONE

Examples of traditional customer interactions captured over time.

LEVEL TWO

Process in place to capture and integrate all relevant customer interactions.

LEVEL THREE

Content is refined and expanded based on measurements of relevance, quality and completeness.

Curate and filter

You will also need to curate, sample and filter any content that you gather to ensure that it's representative of the types of interactions that will be modeled in the conversational application. Doing so will likely require some level of analysis, such as determining which interactions yield successful vs. unsuccessful outcomes and filtering content for topical relevancy and accuracy. This step is also important for identifying and remediating potential sources of bias in the conversational models that you train.

Fill in the content holes

Not every conversational interaction will have existing conversational content that can be readily used. In cases where content is missing or unavailable, it can be created by engaging script writers and editors. Over time, machine learning and natural language generation may partially or fully automate the process.



Keep improving

As soon as conversational experiences go live with users and data is being generated through real interactions, you can tune the methods that you're using for content management.

Whenever possible, capture and analyze interaction data in real time to identify new content (e.g., user verbatims and system responses) that can be integrated into the system for future interactions. For example, if a customer asks a question about an unknown product, the system should recognize that the product isn't included in its current vocabulary, add the product to its existing content repositories and training materials, and develop an appropriate response for a similar future request. Similarly, if a customer responds negatively to a specific system response, or the outcome of an interaction is unsuccessful, the response should be flagged or modified to improve the system.

Key Takeaways

1

Capture, curate and create all the content necessary to carry on the most engaging conversations.

2

Enrich your content with interaction data to improve conversations over time.

ServiceMatch

The ServiceMatch team has created an inventory of all customer interactions that they have archived, including customer emails, reviews and social media posts. They have also processed, transcribed and stored customer voice interactions from their call center. They now plan to use this content to predict customer intents and model future conversational interactions. For each additional channel that ServiceMatch adds, the team plans to capture all of the content that's generated.

ServiceMatch has also hired contract copywriters and editors to fill in a number of identified gaps in the dialogue and to repurpose the raw call center content into usable dialogue for a conversational application.



PRINCIPLE 5

Capture user context

Being knowledgeable about the context of each interaction and user will improve user satisfaction.

LEVEL ONE

Basic contextual data for customer interactions captured from internal sources.

LEVEL TWO

Broad contextual data from both internal and third-party sources captured and integrated.

LEVEL THREE

Contextual data expanded by learning from every customer interaction.

Context improves experience

The best conversations are those in which participants know something about each other. Context about who specific users are, what their preferences are, and what they've already said and done will help personalize interactions, reduce friction and streamline the conversational flow. A simple example of using context is remembering a user's name from one conversation to the next. Other examples include knowing a customer's location or preferred time of day to be contacted.

Conversational context is rich

There are several types of context that can help enrich conversations:

Demographic. Attributes of the user that don't typically change very often, if at all, such as gender, date of birth, occupation, education, home location and marital status, among others. These attributes can be collected once and usually can safely be assumed to be accurate for some time.

Interests and preferences. Characteristics that describe the user and influence his or her general activities, such as being a nature lover, a sports fan, a traveler, an athlete or a foodie. Like demographic information, these types of contextual information are also relatively static.

Activities. The things that the user is doing now, such as buying a car, planning a trip, looking for a restaurant or simply transferring money, and the immediate circumstances, including the current location and weather, among others.

Mood. Whether hurried, relaxed, frustrated, angry or experiencing any number of other feelings, the user's mood may influence his or her communication style and the conversational flow.

Transactional. The user's history of past transactions and the system data to support those transactions, e.g., bank balances, past requests, or open status records.

These different aspects of context can be collected over varied time frames, can only be derived in real time in some cases, and could have a short or long shelf life. Understand how to collect and curate all contextual information, maintain it over time and use it to inform and enrich your conversations.

The best conversational applications will develop a comprehensive and persistent understanding of their users through contextual data. That, in turn, will enable highly personalized interactions that you can use to measurably improve outcomes.



Key Takeaways

1

Record rich contextual data about each interaction.

2

Consider location, user preferences, demographic information, transaction history, conversational outcomes, user state and even the local weather.

3

Monitor context and continuously improve your understanding of the user and the conversational context to improve the quality and effectiveness of future interactions.

ServiceMatch

By breaking down the steps in its customer interactions, ServiceMatch determines that each step has different objectives and context. For example:

Registration. Registration needs to be simple. ServiceMatch can offer a Facebook login as a single sign-on that captures and reuses customer profile information (persistent context) and supports the selection of Facebook as the primary interaction channel.

Vendor search. ServiceMatch uses machine learning to discover and understand different market segments (e.g., plumbers) and the nature of the service providers who offer services on the platform. They use this context to improve the matching and recommendation engine that drives chat interactions with customers.

Scheduling. ServiceMatch processes service requests for a specific vendor by using context from calendars, locations and dates.





PRINCIPLE 6

Design for conversation

Give people credit for their intelligence. Don't try to be cute or clever; instead, take advantage of tools and technologies to deliver a truly valued experience.

Conversational design involves understanding and applying a range of linguistic and communication conventions to improve conversational interactions, from menu-driven flows to natural language dialogues. Some techniques are simple and effective, such as recognizing verbatim user intents or remembering to greet a user at the beginning of a session. Others are more complex, such as understanding dialogue structure or determining the user's emotional state.

Understand intent

Conversational applications need to understand the user's intent (what the user needs or wants to accomplish). Simplify the task by limiting the focus of the application or by using structured menu-like interactions for user input. More sophisticated approaches will use natural language processing to enable the automatic identification of the user's intent, conversational topic, named entities referred to within the text, and even speaker sentiment.

LEVEL ONE

Limited set of user intents and responses, support scripted conversations.

LEVEL TWO

Wide range of user intents and responses, use some conversational conventions and natural language.

LEVEL THREE

Set of intents and responses continuously expanded, extensive use of natural language and conversational conventions.

Respond appropriately

When generating responses, most conversational applications will use a retrieval model in which a finite collection of predefined possible responses are used to simulate dynamic behavior. Only the most sophisticated applications will attempt to generate truly dynamic responses; this might require the subject domain to be narrow. Avoid responses from retrieval systems that seem artificial and responses from dynamic systems that seem inappropriate or irrelevant.

Be perceptive

The best conversationalists know when to listen and when to speak. They understand what's being said right now, and are able to recall what's been said in the past. They can also respond with the appropriate words and ideas within an appropriate time frame, apply the right level of formality for the given situation, and even recover when there's a misunderstanding.

Being able to do all these things, and more, demonstrates what we refer to as conversational intelligence. The best conversational applications demonstrate the highest possible levels of conversational intelligence and carry on natural, engaging dialogue with their users.



Learn from every conversation

Define success for each conversational process, then measure and optimize user interactions to improve performance. Record what contributed to the success, or failure, of every interaction and use the data to improve performance. Monitor conversational flow, order of interaction, use of humor, available contextual information and speed of the user response.

Key Takeaways

1

Linguistics and human communication inform conversational design.

2

Interactions can range from highly structured button and menu sequences to free-form natural language.

ServiceMatch

The initial phase of ServiceMatch's conversational interface will be limited to button-driven interactions and a narrow set of user intents and flows that are focused on finding and selecting a provider and booking a service call. Though simple from a user interface perspective, the experience will still deliver rich functionality to the user via robust back-end integrations. The second phase of the interface is planned to go beyond buttons and menus and add richer natural language support to handle surveys and service provider feedback by applying natural language processing (NLP) to user verbatims.

The user experience (UX) team at ServiceMatch has added new skills to support the work: a linguist and a copywriter. Their focus will be on creating the most effective outbound notifications and prompts, and to determine the best interaction flows. Software developers are working alongside the expanded UX team, testing the use of NLP for understanding user intents. The developers are also tasked with building a retrieval-based response system for the second phase. That system will support predefined, language-based responses being matched to user intents in real time.



PRINCIPLE 7

Personify your business

Conversational AI extends the way customers view a company's brand and image. Make sure it enhances their view.

By using your brand to inform the design of a conversational AI application, you can bring your brand values to life. Persona design should therefore start with an understanding of your organization's brand strategy and core values. Doing so will help your team design personas that are appropriate, recognizable, consistent, engaging, relevant and memorable.

Develop the right persona

Your conversational application's persona should also take into account the nature of the use case for the application and the types of users it will interact with. Will it be an expert, an evangelist, a motivator or a comedian? For example, the persona of a financial services chatbot responsible for transferring money between accounts for small business owners might use formal and precise language and an older personality (e.g., representing a middle-aged bank manager) and a formal name. By contrast, a chatbot for the same bank aimed at marketing savings accounts to millennials over text messages might use humor or colloquialisms to be more engaging, and have a more youthful personality and a more casual or shorter name.

You should also consider how the persona will be kept consistent with other outward-facing communications including your website, customer emails and even human interactions. It's important to think about how

LEVEL ONE

Some brand values and characteristics identified.

LEVEL TWO

Persona designed to display key traits to support brand strategy.

LEVEL THREE

Brand impact of persona measured and monitored to drive improvements.

your chatbot might change as your brand changes, so alignment with the brand strategy should be an ongoing process. If there are news, events or public relations issues that need to be addressed, it may be necessary to consider how they will be dealt with through the chatbot, if at all.

Be recognizable and unique

The extent to which an organization can create a unique persona for its brand will vary between conversational AI platforms. For example, voice frameworks such as Amazon Echo, Google Assistant and Apple Siri provide a highly recognizable voice for their respective services. The sound of these voice assistants is part of the brand that third-party brand applications must operate within. However, even though these frameworks provide a branded voice persona, they also support the creation of unique functionality and responses, including pre-recorded prompts.

Make use of customization opportunities to reflect your brand to the greatest degree possible by carefully designing unique traits into your custom content.



Stay aligned

Finally, as conversational AI applications learn and evolve, you should monitor performance to ensure that you maintain brand alignment. This is especially true if your responses to the user are being dynamically generated. In the simplest case, that might involve human review of a sample of the chatbot interaction logs or voice responses. In more advanced scenarios, it would likely involve automated monitoring and analysis of sentiment, word use and interaction outcomes using conversational analytics tools.

Key Takeaways

1

Conversational AI allows you to build a persona that represents your brand through conversational interactions.

2

Understand the values and characteristics that you would like to convey and develop a conversational brand strategy.

ServiceMatch

ServiceMatch's brand — what it wants to be known for — is about making it easier to find and use the best service providers. Interactions with a conversational application for ServiceMatch should be approachable, friendly and efficient, using plain language that is easy for both customers and service providers to understand. Humor may be another device used to make the application more approachable.

Due to the nature of the business, there are situations in which one or both parties (the consumer and/or the service provider) are unsatisfied. ServiceMatch will pay particular attention to the language it uses for handling such situations, such as a consumer failing to make an appointment time. If the situation can't easily be resolved in an automated fashion, the application will then quickly escalate to a human operator who has been trained appropriately to respond.



PRINCIPLE 8

Bot or not?

Make sure that the right person or technology is taking part in the conversation.

The increasing use of AI to power messaging and voice chat will allow you to use conversational AI to create unique customer conversations at scale. However, not every conversational AI interaction will support full automation, nor does every interaction require it. Instead, many interactions will begin or end with human-to-human live chat.

Over time, automation will increasingly be used to scale conversational interactions across multiple users and platforms and take on an increasing number of repetitive tasks. Humans will therefore be able to focus on more complex tasks requiring judgment and take on the responsibilities of training and evaluating the performance of conversational systems. In the meantime, there is a wide range of choices for how to approach automation and how quickly it should be integrated into conversational processes.

Is automation right for you?

There are many benefits to automating your conversational AI applications. These include low-cost scalability, consistency, deep personalization and fast recall of huge volumes of data. Yet full automation may either be too difficult or too expensive to achieve. Understand how much automation is right for your particular situation. Then focus on automating the processes that you can, while recognizing opportunities to seamlessly hand off conversations and tasks between bots and humans where appropriate.

LEVEL ONE

Strengths and limitations of automated and human-in-the-loop conversations understood.

LEVEL TWO

Design process identifies opportunities for automating conversations.

LEVEL THREE

Processes automated and optimized through seamless collaboration between humans and machines.

Automation has its limits

In many cases, humans may be better suited to handle outliers or infrequent edge cases that would otherwise incur very high upfront training costs to automate. For instance, it might be more cost-effective to train and maintain a small team of customer service representatives to handle unusual change requests for multi-city plane tickets than try to train a chatbot to manage the full range of possible challenges that could arise in the process.

Working together

Humans and machines should work together seamlessly to meet user requirements. Where access is unavailable for certain functions, it may be reasonable to hand off the conversation to a human in the loop as a temporary workaround. When including humans in the workflow, it's important to manage user expectations.

For example, if an automated chatbot is unable to solve a customer support request, the interaction may be escalated to a human chat operator within the same conversation with minimal time delay. The chatbot should then take over the conversation once again when the human operator has finished assisting.

The best conversational applications will use a combination of automation and humans-in-the-loop with a seamless and bidirectional handoff between the bot and the human.



Key Takeaways

1

Automation is a continuum, and humans are almost always needed in the loop, either as trainers, evaluators or both.

2

Humans can serve as an on-ramp to automation.

3

Define a strategy for finding the right balance, pace and integration points for automation, and treat it as a design element.

4

Work to achieve seamless handoffs between humans and machines.

ServiceMatch

At first, human ServiceMatch agents will continue to engage customers directly, with the agent-customer interactions being analyzed. A subset of those interactions, including brokering an appointment time between the two parties, are being prioritized for suitability to be automated using a chatbot.

In addition, ServiceMatch has analyzed customer query logs on its website (including completed and abandoned queries) for service providers and determined the optimal moment of outreach for a proactive “how may I help you” interaction.

As part of the onboarding process, ServiceMatch explicitly asks new customers for their preferred channel and method for reminders and updates. Customers can choose to pay for human phone call updates or free automated text/SMS or Facebook Messenger updates.

A key point of escalation identified during the chatbot design process is when either the customer or the service provider is dissatisfied with the interaction and a human is required to mediate a solution. ServiceMatch is paying particular attention to detecting negative sentiment and intent, and evaluating ways to smoothly hand off to a human operator in those situations.



PRINCIPLE 9

Keep private conversations private

Recognize that conversations are inherently more intimate. Users may have higher expectations of privacy than other interaction channels.

Conversational interactions are intimate and will likely come with high expectations of privacy from users. Look for opportunities to implement stronger conversational privacy than the market norm to create competitive barriers. Be explicit with users about how you manage conversational data and set expectations with your users about the nature of the conversation. Explain how they are going to benefit, and what you'll need to do with their data to deliver those benefits.

Privacy is part of design

Make security and privacy part of your conversational design process, and build on platforms that enable secure communications. Avoid recording conversational data unnecessarily, and avoid using private data outside of its intended context or beyond its intended purpose. Be aware that users may overshare and provide more information than is necessary or expected.

Develop a conversational policy

You should have a conversational privacy policy and proactively make security a part of your conversational design process. At a minimum, you should have a blanket privacy policy that covers the use of conversational data. A better approach is to have an explicit policy that addresses the specific risks, requirements and opportunities conversational data

LEVEL ONE

Blanket privacy policy provided that covers the use of conversational data.

LEVEL TWO

Explicit conversational privacy policy implemented as part of design process.

LEVEL THREE

Security controls proactively implemented and improved to deliver measurable value to users.

If you discover a new use for interaction data, ensure that it's either already covered in the conversational privacy policy, or inform your users before taking action with it. Be careful about how you use interaction data to train conversational systems, especially if you incorporate user verbatims into your responses. For example, always make sure that it isn't possible to recognize individuals in responses learned from user behavior.

Language is more personal

Web and mobile applications primarily deal with structured user interaction and input data. It is therefore possible to locate personally identifiable information (PII) so that it can be handled separately or deleted. However, conversational interactions can include significant amounts of unstructured data from the user that contain PII. This makes it more difficult to identify which data may have more stringent privacy requirements.

Furthermore, the recordings of conversational data and content from the user can be used to derive insights such as the emotional state of the user. It is therefore necessary to take a much broader view of which data does and does not constitute PII in conversational applications.



Key Takeaways

1

Interactions with conversational interfaces come with greater intimacy and trust than web- and app-based interactions.

2

Set expectations with users explicitly about how conversational data will be used.

3

Explain how users benefit from analysis of their conversational data to build trust.

4

Develop a conversational privacy policy and update it to reflect changes in your product and the user experience.

ServiceMatch

ServiceMatch has adapted its privacy policy to include conversational clauses explaining how conversational data is collected, protected and used. In addition to being on the website, the first time that users log in to the application, they are given the opportunity to opt in to the policy, and the conversational elements of the policy are called out directly.



Conclusion

Conversational AI presents an opportunity for companies to create natural, engaging and deeply personalized experiences that are highly scalable. To capture maximum value from conversational AI, organizations should take a methodical approach to prioritize and implement conversational opportunities. The principles of conversational AI provide a framework for getting started and will help you avoid the most common mistakes.

Those mistakes include focusing on the wrong part of the business, delivering a poor user experience due to lack of contextual information about the user, and failing to apply good conversational design to your bot or voice assistant.

As an organization, you can focus your efforts in the right place by using a selection and prioritization process that identifies where the impact will be greatest. Competing ideas for where to use conversational AI can then be evaluated based on their ability to drive the greatest improvement in key metrics such as customer engagement, revenue or cost.

The best chatbot (or voice assistant) user experiences are those where the applications seem to know what users want, or at the very least don't require users to repeat themselves.

The goal is to have as much context about the user as possible before the interaction takes place. For a customer-facing conversational application that may include previous transactions and support issues, their preferred language and dialect, their name and location, who they are connected to on the social graph and how long they have been a user or customer. The more context your application has, the less work the user has to do, and the better the experience.

Conversational design is new, and for many organizations, it will be difficult. It signifies a fundamental shift away from the visual design principles that have dominated application development for the past 30 years and toward designing with language and conversational flow in mind. Even where a chatbot uses a menu-based interaction that the user selects, rather than free-flowing natural language, designing the user experience will focus on words and language and the natural flow of dialogue.

The best organizations will apply well-understood linguistic techniques to design superior language-based interactions, as well as sophisticated chatbot and voice assistant applications. The challenge for organizations is to recognize that different skills are required and to build teams that are capable of understanding and applying good conversational design.

Conversational AI Maturity Model

PRINCIPLE	LEVEL 1	LEVEL 2	LEVEL 3
1. Identify where you'll have the greatest conversational impact	Opportunities identified in an ad-hoc manner.	Business processes partially documented and prioritized.	All business processes prioritized, roadmap developed.
2. Understand your audience	Basic understanding of the platforms the audience is using.	Research and testing used to develop a complete understanding of audience preferences.	Interaction data used to refine understanding of audience's preferences.
3. Build complete experiences	Some back-end service integrations accessible.	All service integrations required for end-to-end conversations identified, most are accessible.	Optimized integrations for complete end-to-end conversational experiences.
4. Content enables conversations	Examples of traditional customer interactions captured over time.	Process in place to capture and integrate all relevant customer interactions.	Content is refined and expanded based on measurements of relevance, quality and completeness.
5. Capture user context	Basic contextual data for customer interactions captured from internal sources.	Broad contextual data from both internal and third-party sources captured and integrated.	Contextual data expanded by learning from every customer interaction.



PRINCIPLE	LEVEL 1	LEVEL 2	LEVEL 3
6. Design for conversation	Limited set of user intents and responses, support scripted conversations.	Wide range of user intents and responses, use some conversational conventions and natural language.	Set of intents and responses continuously expanded, extensive use of natural language and conversational conventions.
7. Personify your business	Some brand values and characteristics identified.	Persona designed to display key traits to support brand strategy.	Brand impact of persona measured and monitored to drive improvements.
8. Bot or not?	Strengths and limitations of automated and human-in-the-loop conversations understood.	Design process identifies opportunities for automating conversations.	Processes automated and optimized through seamless collaboration between humans and machines.
9. Keep private conversations private	Blanket privacy policy provided that covers the use of conversational data.	Explicit conversational privacy policy implemented as part of design process.	Security controls proactively implemented and improved to deliver measurable value to users.



Georgian

At Georgian, we're building a platform to provide a better experience of growth capital to CEOs and their teams.

Georgian's platform is designed to identify and accelerate the best growth-stage software companies, taking an intelligent, data-driven approach to solving the key challenges CEOs face as they grow their businesses.

We invest in high growth companies across North America that harness the power of data in a trustworthy way. Based in Toronto, Georgian's team brings together software entrepreneurs, machine learning experts, experienced operators and investment professionals.

info@georgian.io
georgian.io