

Smartpaper

Conversational A.I. for the Enterprise



Watson

IBM

Table of Contents:

Conversational A.I. and Corporate Strategy.....Page 3

1. Do you need a chatbot?
2. Developing Your Ideas: The Use-Case
3. Defining the Purpose
4. Target User
5. Viewpoint
6. Tone and Personality
7. The Conversation Team
8. Build your Chatbot

Conversational Design.....Page 6

1. Question Collection
and Understanding Utterances
2. Ground Truth
3. Basics of Dialog
4. The Conversational Flow
5. Designing Responses

User Experience.....Page 9

1. Deployment Channels
2. Optimization
3. Memory and Personalization
4. Machine Learning

Conversational A.I. and Corporate Strategy

The state of chatbots has quickly evolved from a fad to an essential part of corporate strategy. No longer a nascent technology, chatbots have matured into solutions that enterprises across industries are taking seriously. But how do we distinguish a simple chatbot from human-like, AI-driven conversational solutions that we now see businesses adopting? In this three part series, we will examine what it means to deploy “Conversational A.I.” for the enterprise.

Whether to provide better customer service, drive employee engagement, or improve other solution areas that chatbots play a role in, Conversational AI allows you to build robust agents that carry out conversations that communicate knowledge, answer questions, or walk you through a process. Designing a conversation to sufficiently meet consumer needs and return real business value requires a nuanced strategy and in-depth considerations. Once implemented, these conversations become more than just another way to interact with your brand—they become a source of data to understand what your customers are looking for and how they are engaging with your organization. Ultimately, it provides you with new insights and ways to optimize the customer experience.


Conversational AI is therefore much more than just a simple chatbot. However, for the purpose of concision and clarity, we will use the term “chatbot” as an all-encompassing category of cognitive applications that use natural language understanding to enable conversations between man and machine.

Chatbots harness the power of human conversation to exponentially enhance the expertise, agency and impact of your organization on the world. Let’s learn how to design one for your business.

1. Do you need a chatbot?

Before deciding on whether to implement a chatbot into your operations, first determine if there is a need. Making sure a chatbot is the right course to solving your business challenge is unequivocally the primary step. Are your employees stuck on the phone answering sales or support questions over and over again, hindering employee productivity and performance? A bot can help automate these common workflows so your employees can focus on what matters. Are users struggling to find information on your website? A chatbot can guide

them to the right information quickly and accurately. Do you want customers to apply for a mortgage? You can teach your chatbot all the steps needed to underwrite your policy, and ask questions along the way. The experience you are creating for your clients is paramount, so before moving quickly to development, figure out the problem areas that you, your users and employees are struggling with, and if a chatbot is the best fit to meet everyone’s needs.



2. Developing Your Ideas: The Use-Case

Ideating your use-case is the next logical stage. There is no concrete process in discovering the ideal use-case, but [design thinking principles](#) are often used to define this stage—to generate ideas and then prioritize them. Here are some basic practices to approach the initial ideation step:

Step 1. Define objectives

Outline what the overall objectives are for your business and for your users. Differentiating the advantages for both you and the end user is necessary to evaluate the impact of a chatbot. We suggest brainstorming potential ideas while keeping in mind your limitations and pinpointing untapped opportunities.

Step 2. Find your gaps and opportunities

Use analytics and accessible core metrics such as highest volume of calls, topics with highest number of requests, and customer satisfaction pain points to analyze weaknesses, prove the need for your solution, and evaluate return on investment.

Step 3. Lay out realistic goals

Though the potential to create may seem overwhelming, keep your initial scope narrow. If you try and solve too much, too soon, you won’t end up solving anything at all. Take banking as an example. Start with a chatbot that solves credit card issues for customers, get that right, and then expand to new domains

over time. From the onset, you should provide users with a firm understanding of how the chatbot can help them, while also effectively managing expectations.

3. Defining the Purpose

You have a proposal for a use-case. Now you need to define the *purpose* of the chatbot: What *problem* are you trying to solve and *for whom*? Though the use-case and the purpose appear similar, it is important to distinguish between the two. The *purpose* is linked inextricably to how you are positioning your chatbot in conversational interactions: The way you engage with your users.

The concept of a purpose can be difficult to grasp at first, as it's the part of designing a chatbot that is farthest removed from hands-on-keyboard developing. However, conversational interactions should always have a clear purpose — a well-defined job to do (as the corpus of knowledge grows, your purpose may grow too). Here are a few examples:

- **Encourage** the user to complete an online form or process
- **Persuade** the user to purchase a product or service
- **Prequalify** the user for something
- **Engage** the user and generate interest in something
- **Educate** or inform the user
- **Entertain** the users

Explicitly defining the purpose will help you achieve a coherent solution that meets your needs and the needs of your end-users. To support your purpose, you must carefully design key moments in the conversational interaction—proactively engaging users at the right time with key messages and questions. This involves preliminary research by gathering key information about users' interests, concerns and behaviors. It involves a combination of business analysis and strategy, and an understanding of the human factors involved in a human-computer interaction. There are multiple, interrelated elements to consider, and some of the concepts are nuanced.

4. Target User

Like in any conversation, there are two parties; In this case, it's your chatbot and the end-user. The human interacting with your chatbot needs to be listened to, and their expectations must be met. Build out a persona for the target users as you would in product strategy or a marketing campaign, so that when you build out dialog, everything aligns back to that

user. It is critical to have an audience clearly in mind when determining how your chatbot will respond.

You may have a specific and well-defined audience, like high-wealth investment banking clients or university students, or you may have a relatively broad audience, like all the customers of a particular bank. Generally, the more specific the users, the stronger, more distinctive the tone and personality of the solution can be. Be sure to pinpoint end-user pain points, and the methods of communication they tend to be receptive to. Be sure to consider future audiences for the solution as it evolves—not just the audience for the current iteration of the solution.

5. Viewpoint

Armed with a purpose and a persona, carefully identifying the viewpoint of your chatbot will support your business objectives and provide a consistent user experience. The goal is to figure out whom the solution speaks *on behalf of*.

Decide how closely aligned your organization is to the viewpoint. If the viewpoint is that of an employee, exhibiting explicit sales-like behavior will make sense. However, if the viewpoint is somewhat independent of your organization—completely aligned with the user—then overt selling behavior is likely to feel jarring. The viewpoint, purpose and use-case all work in concert to direct the conversational flow and dialog structures. It is essential to be as specific as possible in these areas as you work through your conversational design, while keeping in mind future changes, to not hamper or limit future iterations and conversations. Here are key questions to consider when defining the viewpoint:

- Who is the solution acting on behalf of?
- What role should the solution play in the relationship between your organization and your end users?
- What relationship should the solution have with the end user (to achieve the purpose of the solution)?

6. Tone and Personality

The tone and personality is the “voice” your chatbot speaks in. The tone and personality should *reflect* the solution's viewpoint; if the viewpoint is that of an employee, the tone and personality should closely reflect the ideal staff member. If the viewpoint is that of a partially or fully independent advocate or guide for the end user, then the tone and personality should not be that of an enforcer or authority figure.

The tone and personality used by your solution can help or hinder achieving the purpose. If the purpose of the chatbot is to engage users proactively and encourage them to take an action, a friendly, informal tone will be more successful than a formal or authoritative tone. If the purpose is to provide information or guidance about a weighty or serious topic, a casual, chatty tone may undermine credibility of the information provided.

Be careful not to confuse tone and personality with viewpoint – *who* the solution speaks on behalf of. Tone and personality should be consistent with the viewpoint, but the viewpoint alone does not determine what tone and personality a solution should have.



7. The Conversation Team

Creating a vision and building out a chatbot requires a number of crucial skills. Using intents, entities and dialog as your tools, designing a conversation for customer support and engagement depends on coordination with individuals across functions—to not only help record and keep track of the project, but also to include all areas of expertise that might inform the customer experience (e.g., content, marketing, product, etc.).

In the early stages, teams are small (two or three members) and roles may not have strong delineation. Usually, chatbots begin as experimental and with little, if any, budget, so there may not be a clear determination if the project will move forward past the initial phase. Therefore, team members must be prepared to take on work outside their usual scope. As the vision becomes tangible and is on its path to approval, roles should ideally become more focused. The following roles are suggestions and may not directly reflect the makeup of your team:

Line-of-Business (LOB) Leader

- Pitches the project for funds, or to sell to external clients.
- Partners with developers during the initial build and is a key stakeholder/advocate who must be regularly updated on demonstrated value.

Project Lead

- Given project oversight, including reporting performance and progress to the LOB leader.
- Manages team resources and/or becomes a product manager.

Conversational Strategist

- Defines the positioning, personality, and workflows for the virtual agent through analysis of end-user persona and existing business processes.
- Crafts customer conversations that result in good end-user interactions for customers.
- Collaborates and works cross-functionally to build user-friendly experiences.

Data Scientist or Engineer

- Analyzes historical user data, especially any natural language communication the particular end-user group has had in the same domain.
- Using the knowledge of what users ask, they create the machine learning model for intent recognition, and builds the dialog process workflows in which have been designed by the Conversational Strategist.

Developers

- There are multiple types of developers that will be useful for building a chatbot: front-end, back-end, and low-code, for instance, ideally with some direct experience or early interest/training in AI or conversational systems.
- Create a customer user interface for users to interact with the bot.
- Responsible for any integration to company backend to supply the conversation with any supplemental context it may need to personalize responses and perform automations.
- Work closely with conversational strategist to build out the solution.

8. Build Your Chatbot

Once you've developed a clear but multi-faceted approach to building your chatbot, the next phase is development. In the implementation, approaching the journey is just as important as the journey itself. Educating yourself on the

basics of conversational design and its impact on your business is essential to formulating a successful plan of action. This preparation will allow you to communicate, learn, and monitor against a standard, allowing your business to build a customer-ready and successful project.

Delving into the nuances of building a chatbot for business requires patience, preparation, and collaboration—but with a readied plan and a clear understanding of the basic tenets, your development team will thank you.

Conversational Design

Conversational Design is a new area of design that has evolved alongside the proliferation of chatbot technology. Just as User Experience (UX) design is important in building a great traditional interface, Conversational Design is important for creating great chatbots.

When designing a conversational system, there are key concepts and foundational elements that lay the groundwork for development. These concepts will reappear as your team works through the mechanics of building your chatbot. In future phases of development, these will prove integral to understanding, iterating and improving your solution.

As with standard UX design, the first thing to understand is who the user is and what they want to achieve. Refer back to Part 1 to inform your conversational design. As a leader of the team's strategy, you can significantly improve collaboration by understanding how to conceptually design a conversation.

1. Question Collection and Understanding Utterances

Chatbots must be trained on the end users' most frequently asked questions – not on the questions we think the system should answer. Question collection is the process of collecting questions that represent the ideal user scenario. A question is a type of utterance, any input a user provides when prompted. An utterance can be a sentence, a question, nonsense—anything understandable as words.

You should perform careful research on the types of questions your end-users will ask. Ultimately, the content of these questions is what you will need to train your system on. The purpose of training is to improve your chatbot's ability to accurately detect intents from a user's utterance. There

are different approaches you can take to researching end-user utterances, which may include public crowdsourcing, surveying targeted individuals, reviewing past customer support logs, or conducting online forums.

It is best to use real examples of user utterances, however, this may not be possible initially. If that is the case, create empathy maps and try to get into your user's head when first choosing with questions to include. Later on, as you improve the bot with testing and real usage, you can include real utterances from recorded chat logs.

A good chatbot framework will use Natural Language Understanding to understand user utterances based on examples. This way, you don't have to provide every possible utterance, just enough to capture the various ways a user can say it.

These are just a few examples of ways to collect and iterate on existing data that will influence questions to train your chatbot with. Keep the following question set pit-falls in mind during question collection:

- **Too many out-of-domain questions:** This is a warning sign for a domain or use case not being defined correctly, or end-user perception of a solution being different than expected.
- **Insufficient number of questions:** This will almost certainly become a problem down the road, when building and testing the solution. Insufficient questions will lead to insufficient training data, or reuse of questions for multiple purposes.
- **Team selected questions to be included in test sets:** You have no way of assessing which questions your solution will learn from. Selecting questions based on complexity or frequency is particularly bad, as it will lead to inadequate training or an over-fit system (i.e., a system being better trained on certain questions).
- **Misspelled questions were corrected:** Train the system on incorrectly spelled questions just as frequently as on correct ones. Correcting spelling errors will likely lead to reduced system quality. However, take into account the possibility the end-user's device will run a spell check!

2. Ground Truth

Once you know what users are going to ask, and, equally as important, how they ask it, you can start building your ground


truth, the clustering of questions into intents. The success of your solution depends on training the system with well-defined intents formed by clustering representative questions: establishing ground truth. Having a set of representative questions is the critical success factor.

There are two things to focus on here: intents and entities. An intent is the action someone wants to take, and an entity is typically the object he or she wants to act upon.

Intent: An intent represents the purpose of a user's input, such as a question about business locations or a bill payment. You define an intent for each type of user request you want your application to support. The questions will allow you to determine the scope of knowledge that the system must be trained on and which entities, if any, will be important. Collect a minimum of 10 representative questions for each unique intent.

Entity: An entity represents a term or object that is relevant to your intents and that provides a specific context for an intent. For example, an entity might represent a city where the user wants to find a business location, or the amount of a bill payment. Entities provide exact matches for keywords and synonyms, and can also be used as system event triggers using predefined event keywords sent from external applications.

For example: “Where is the pool” and “where is the restaurant” would both have an intent of #location (intents are denoted by the hashtag/pound symbol) and individual entities of @pool and @restaurant (entities use the @symbol).



3. Basics of Dialog

Once you have your ground truth, it's time to begin designing dialog and developing the application layer. This is more of

a standard software development exercise, and is all about getting the 'brain' developed in the last couple steps to be able to interact with your users where you want them to.

Designing the dialog is the way to customize the exact words, phrases, questions and answers your chatbot engages with your user. It determines the behavior of your solution in its interactions, attributing to your overall strategy.

Response: What the chatbot returns to the end-user based on the intents and entities it recognizes in inputs. Not all answers are text, some are actions.

Exit Strategy: While users may ask very common questions with answers that require precision and are highly tailored to the dialog flow (often called short-tail responses), users often ask questions that are outside the predicated conversation. In this case, your solution may use content from a range of different sources to provide responses to users' questions (long-tail responses). This provides your chatbot with the capacity to answer specific questions that you may not have prepared for or designed in your dialog. Content can be drawn from a knowledge base: web search, websites, manuals and official documents to name a few.

If the question is too complex and out of scope, handing the user off to a live agent is currently the most successful way to continue a fluid, human-like conversation. This process is referred to as Hybrid AI—since it combines the power of chatbot technology with human expertise and understanding.

Small talk and chitchat: Personality is often reflected through small talk or chitchat. Given the relevant variables, making sure the chatbot engages in the appropriate type of small talk is significant. It may seem frivolous, given the complexity of your business challenge, but it's paramount for a friendly user experience. It makes the chatbot seem more human-like, engaging, intelligent, and even fun.

Proactivity and reactivity: Proactivity is the degree to which your chatbot proactively engages and guides the user. Reactivity, on the other hand, is when the chatbot sits back and waits for the user to ask a question. Generally bots are reactive, that is they respond to requests/questions from the user. However, a bot can also take the initiative and be proactive. This is useful for educational bots that are guiding the user through a task, or where there is a clear task that users want to do. In the case of a banking chatbot, if users mostly want to check their balance, you can prompt them to do so, without them even asking.

Another way to be proactive is to fetch information before the user even asks for it. For example, a weather bot may fetch and display the forecast for a user based on their location before they even ask for it, as that will most likely be their request.

4. The Conversational Flow

When designing dialog, it's important to consider the conversation flow. What is the necessary sequence of steps to follow during a task? There are four key processes to think about while building a flow.

I. Outcomes

In order to build the flow, you need to understand what the particular outcome or task of that specific flow is, in alignment with the chatbot's overall objective. Based on a user's intent, what are they looking for and where do we want to lead them? When choosing specific outcomes, make sure to understand the limitations of the information your chatbot can provide in order to succeed in the places your chatbot is most confident. Here are a few examples:

- Make a reservation
- Cancel a credit card
- Sign up for an account
- Purchase an order
- Get delivery information

Based on what the user wants, what you want from the conversation, and what the bot actually knows, we can figure out how to achieve the ultimate outcome, i.e. we can think about the steps needed to complete the task. This can include gathering information from the user, getting information from somewhere else (e.g. by calling a service), and checks or confirmations needed. Write down the steps you think are needed to complete this task—they don't necessarily have to be in the right order at this stage.

II. Beginning the Conversation

It is important to consider how the flow begins. First, introduce the chatbot's personality and give the user an impression of its behavior. Then, give the user some information about what the bot can do.

III. Ending a Flow

Think about where the flow ends—this could be in many places, not just the step where the task is completed. You may have steps where you decide not to allow the user to complete the

task or to direct them elsewhere. However, most importantly, make sure you don't have any accidental dead-ends.

IV. Guidance

It is important to provide guidance to the user to help them complete the task. In addition to guidance in the introduction, you should provide guidance at key points in the flow, particularly where the user has to make a choice or provide necessary information. If the user asks for help, then the bot should provide understandably some guidance. Allow questions to be asked in context and then in return, provide contextual answers. One way to direct the conversation is by being very explicit about what information is needed and avoiding open-ended questions.

A common form of guidance is providing the user with canned options. Canned options are options displayed for the user to select. These are provided to users in place of requiring their typed responses. These can be very useful and can help to streamline a particular flow. However, it's important not to use them too often, as they can stop a natural conversation and make it feel mechanical. Some common examples of when to use them include:

- To introduce the chatbot's main features
- To push users through clear common paths
- To deal with an off-topic utterance that may cause a transaction to fail
- To control the scope of utterances

Canned responses need not be exhaustive. You can still provide the text box to type a response, but even just being able to see some example responses can help to guide the user.

To prevent unnecessary confusion, create fallbacks for your bot, i.e. what to do when the bot doesn't understand where the conversation is going. Some of these may coincide with the flow endings, but generally you will need to use fallbacks to guide the users back on topic and achieve the desired outcome.

5. Designing Responses

While you are building your flow, you are not required to come up with the exact wording of responses, as long as you write enough so that the user choices are clear. Only once your flow is complete, you can think more closely about the actual text. When writing responses you should think about how a human might write in a chat. Keep responses natural, concise, and

break them up as a human would when messaging another human. Here are a few items to consider while designing responses:

I. Provide Variation

It is good practice to provide some variations in responses particularly for the fallback responses where the bot may not understand the utterance. Variations do not need to be wildly dissimilar—these differences can be subtle—but they prevent a repetitive or “robotic” feel.

II. Restating Intent

Repeating back utterances to clarify intent can be useful to confirm that you have understood the user before you move them onto the next step. Anticipating questions can also steer the conversation if the user is not sure where to begin. Make sure to keep users informed of their progress, particularly when they complete a task.

III. Clarifying Answers

It is often worth giving an answer even if you are unsure if it is what your user is looking for. Give your best guess and ask for confirmation that you provided a suitable response. For example, if the user asks for their account balance, give the current balance of their most requested or last requested account, and then check that’s the account they meant, rather than inserting another step in the flow to ask which account.

IV. Adjusting Tone

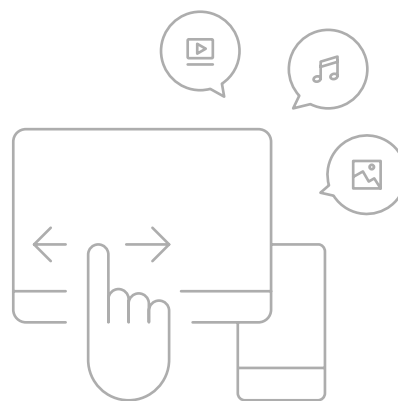
You can adjust responses based on the tone of the users input. There are services available that analyze the tone of text. You may want to adjust a response if the user seems angry—responding with a different tone might reduce friction and ease the conversation.

V. Be Honest

Responses are a good opportunity to implement personality and chitchat, however, one thing to keep in mind is to be honest about the fact that the user is interacting with a chatbot. This does not mean you need to state this upfront, however, pretending to be a human can feel disingenuous. Users are more forgiving of your chatbot’s mistakes if you are clear that it’s a chatbot.

VI. Messaging Shaping

Message shaping is understanding how a user will respond to what you write. Understanding the person interacting with your chatbots helps to develop ways of controlling their expectations, reduce technical complexity of the conversational flows and improve customer satisfaction even in a negative situation.



The User Experience

While largely dictated by the flows and responses, the user experience (UX) also has other important design features that require close consideration. Deployment channels, design aesthetic, and user interface are a few pieces that can affect the performance of your chatbot.

Depending on where the chatbot is deployed, UX options may be limited. Deploying to a messaging platform means you have less control over the surrounding space, whereas if it is integrated to your own website you may have free rein for design.

Where you do have control, consider augmenting the experience with additional information such as images, videos, multimedia, and other design elements. Where you do not control the surrounding space, media should be included within the chat itself.

Additionally, there are many cases where a digital experience is faster and simpler for a user. With chatbots, you can include interactive objects within the chat interface. Non-text enhancements can include things like buttons, carousels, cards, and web-views to name a few. For instance, when asking a user a yes/no question, it can be simpler for the user to simply tap a button, than to type a word. This leaves much less room for error on the user behalf. In a shopping scenario, carousels are a good opportunity for users to scroll through options to see what they are buying—a user can simply swipe around the carousel until they find what they are looking for.

To augment the experience, also consider implementing other cognitive APIs or supplemental services. Do you need it integrated with a camera for visual recognition activities? Would you like your chatbot to analyze the tone of your users? Do you need speech-to-text so that a user can interact with

their voice? Integrating other APIs will elevate your chatbot's ability to converse naturally and smoothly. Deciding which other cognitive elements to utilize on top of your chatbot depends mainly on the use case and brand voice.

1. Deployment Channels

At this point, your chatbot can anticipate what you expect your customers to ask and you are ready to put it in front of real users. There are a number of places you can station your chatbot to maintain an omni-channel experience. Common places include:

- Messenger apps
- Mobile apps
- Email
- Social Media channels
- Chat widget on your website
- SMS channels
- IVR system
- Voice command application
- Website Interface

Consider your target user and use case when deciding what channel to deploy in. Are these exclusively customers of your business? How will they discover your bot? Where do they typically like to ask questions and get support? Your chatbot should be located where your users are looking to get their questions answered.

2. Optimization

The process of building your chatbot does not end at deployment. In order to satisfy users and maintain productive conversations, you need to optimize. This stage starts once the chatbot is deployed to production and is engaging with end users. Think about what data can be stored to improve conversations and personalize the end-user experience. This data is used not only to complete tasks and avoid asking a user for the same information repeatedly, but also to improve future responses in a conversation and even future conversations altogether.

It is useful to get feedback from users to enable you to improve your chatbot based on real user interactions. There are a number of ways to do this, ranging from thumbs up/down buttons to directly asking the user to tracking the user's tone. One of the most important places to get feedback is after

completing a task. The user has had a complete experience within the bot and has accomplished something, so they could have some interesting feedback. If the feedback is good, it could also be a chance to move them on to a new task or even upsell.

You can adjust conversation flows based on the stored information. Chatbots save chat histories, analyze conversations, and recommend opportunities for improvement in a continuous learning cycle. All activity is monitored—when a chatbot is unable to answer a question, frequent utterances, times for live agent handoff, etc.

3. Memory and Personalization

The chatbot should remember that they have spoken to you before and provide a different introduction in subsequent interactions, even if it's just "Welcome back" vs "Welcome." It is also important for the chatbot to remember the position in a flow. If a user gets half-way through a task, then leaves and comes back later, it should be able to pick up where they left off. You may want to consider including some reminder text, so that when the user comes back to the conversation, they know how far they got in the conversation.

Building out preferences for the user can enhance the overall experience—creating a sense of personalized conversation. For example, taking the banking chatbot, if it discovers that the user is not interested in upgrading their account, then don't keep asking them. Another example is to make use of user history. If a particular user checks their balance 99% of the times they visit, why not streamline the flow and make it easier for them to get their balance? Simple things like preselecting or reordering things can also help improve the user's journey.

4. Machine Learning

In order to develop a continuous cycle of improvement, you will want to closely monitor all conversational activity in order to engage in supervised machine learning. Review all the conversations and add new content so that next time these questions come around, your bot knows how to understand them and respond appropriately. Your chatbot logs will most likely show two areas of enhancement:

- **Content missed during initial training:** With a supervised feedback loop, chatbots can start to point out things it needs

help with or specific patterns, but still wants you as the expert to tell him what to do with it. This is called supervised learning, and is the most appropriate method for creating a consumer facing chatbot.



– **Content outside of the scope of the initial use case, but still valuable to your customers:** As we discussed, you want to set yourself up for success by starting with a narrower domain so that your users know what to ask about, and Watson has a better chance of already having learned about it. Over time, you can add more and more domains, and the best way to do this is based on user feedback. Customers will naturally start asking about other topics and this is the best way to determine where to expand to next.

What's Next?

By now, it should be evident that Conversational AI can play an important role in your company's success. Breaking down the process of creating Conversational AI allows you to make smaller successive wins, as opposed to tackling your pain points all at once. There are three overarching missions that you must execute on in order to achieve those wins: defining where chatbots fit into your corporate strategy, building a conversational design, and optimizing the user experience.

With IBM's Watson Conversation Service, you have the power to build out Conversational AI for your enterprise, addressing specific challenges for both your customers and employees. Businesses across all industries are already attaining immense value with their chatbots; the time is now.

© Copyright IBM Corporation 2017

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
July 2017

IBM, the IBM logo, ibm.com, and Watson are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/us/en/copytrade.shtml>

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The information in this document is provided "as is" without any warranty, express or implied, including without any warranties of merchantability, fitness for a particular purpose and any warranty or condition of non-infringement.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Statement of Good Security Practices: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed or misappropriated or can result in damage to or misuse of your systems, including to attack others. No IT system or product should be considered completely secure and no single product or security measure can be completely effective in preventing improper access. IBM systems and products are designed to be part of a comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. **IBM does not warrant that systems and products are immune from the malicious or illegal conduct of any party.**

