



Chatbase

Building Customer Service Virtual Agents in a Data-driven Way

A Chatbase White Paper

Table of contents

Executive Summary	2
Background	2
Causes and types of failure	3
Risks of the status-quo approach to building virtual agents	4
Benefits of the Chatbase data-driven approach to building virtual agents	5
Optimizing the experience	8
Conclusion	8
Footnotes	9
About Chatbase	10

Executive Summary

Today, deploying an *AI-powered customer service virtual agent* in the contact center is a visionary way to complement human agents (and reduce costs) by providing personalized, non-queued service 24/7. However, the status-quo approach to designing such solutions (i.e., intuition and brainstorming) is slow, based on guesswork, and just scratches the surface on functionality. In contrast, the data-driven approach taken by Chatbase Virtual Agent Modeling is faster (cutting weeks, months, or even years from development time), more accurate, and more comprehensive. Virtual Agent Modeling uses Google's machine learning and search capabilities to uncover critical insights across live-chat interactions for building the best voice or chat virtual agent for specific use cases. Post-deployment, continually analyzing and optimizing virtual agent performance with Chatbase Virtual Agent Analytics helps expand use cases and ensure good user experiences over time.

Background

Over the past decade, iconic companies like Google, Amazon, and Netflix have set a new standard in customer experience, arguably making it more differentiating than price or product. As McKinsey & Co. explained in a 2017 report, "Increasingly, customers expect from all players the same kind of immediacy, personalization, and convenience that they receive from leading practitioners such as Google and Amazon."¹ Customer service is perhaps the most important part in that experience, with 7 in 10 U.S. consumers saying they've spent more money with a company that delivers great service.²

However, many companies struggle to meet that standard in the face of rising call and live-chat volumes in the contact center. Even with the availability of interactive voice response (IVR) systems, or phone trees, and other traditional self-service options, large companies in interaction-heavy industries employ as many as tens of thousands of live agents, at high turnover rates, with the goal of delighting customers around the clock. After those massive investments, 7 in 10 businesses still list improving customer experience as their top priority.³

For companies today, deploying an *AI-powered customer service virtual agent* is a visionary and viable way to augment human agents, and reduce costs, by providing personalized, non-queued service 24/7. According to research from Accenture, 80% of live-chat sessions can be automated with a well-designed virtual agent,⁴ and perhaps for that reason, IDC

predicts that by 2021, the 2nd largest amount of business budget spent on AI will be for building customer service virtual agents.⁵

Several established contact center software vendors are addressing this opportunity by offering their own solutions, or partnering with an AI leader like Google Cloud.

Whichever development approach is involved, however, companies building AI-powered customer service virtual agents have to be

Eighty percent of live-chat sessions can be automated with a well-designed virtual agent

careful to create an experience that meets user expectations. The stakes are especially high, because few customers will tolerate more than one bad service experience; e.g., more than half of Americans have scrapped a planned purchase or transaction because of bad service.⁶

Causes and types of failure

Many builders of virtual agents rely on the natural language understanding (NLU) in their conversational experience platform to enable extended, open-ended conversations.

Unfortunately, that approach usually results in failure; NLU technology, while impressive and advancing rapidly, is not yet up to the task of recognizing every phrasing variation per supported customer question, or anticipating new and relevant customer questions. As a result, “UMM” errors abound, causing bad user experiences that frustrate users.

- ***Unsupported requests***: Requests that the virtual agent was not designed to handle. Result: a fallback response (e.g. “Sorry, I didn’t get that”)
- ***Misunderstood requests***: Requests for which the virtual agent misclassified the user intent (e.g., “insurance options?” triggering an <add-insurance> intent instead of <view-insurance-plans>). Result: a response to a slightly different request than the one intended, often with no easy options for repairing the conversation
- ***Missed requests***: Requests for which the virtual agent identified the correct intent but failed to recognize certain phrases or words (e.g. “cancel insurance” is recognized but “remove insurance” isn’t). Result: a fallback and missed opportunity to service the customer

Unfortunately, as you’ll learn below, the status-quo approach to building a virtual agent does little to reduce the risk of these “UMM” errors.

Risks of the status-quo approach to building virtual agents

For companies building an AI-powered customer service virtual agent, the first critical task is knowing which one to build. In other words, it's important to know in advance which high-priority use cases the virtual agent will need to handle and all the variations in phrasing involved, how many conversation turns are involved in those requests, and how top-performing live agents respond. The ultimate goal is a virtual agent that is an expert in the use cases for which it is designed, and for which use cases can expand over time.

It's important to know which requests the virtual agent will need to handle and all the variations in phrasing involved

Typically, this discovery process is manual and inherently limited in scope: for each use case, stakeholders and subject-matter experts get into a meeting room over weeks to document common customer requests, best practices for handling them, and so on. In other words, discovery is based on informed guesswork and anecdotes. For those teams with access to chat or call logs the process is more scientific, but still rarely touches more than a small sample of data.

That approach has three main drawbacks:

- **It misses edge cases**, particularly for large contact centers with millions of interactions per month or year. Examples can include related questions, confirmations, objections, special requests, modifications, and context switches.
- **It misses trends**, because a manual process cannot keep up with changing products, campaigns, and seasons, nor with evolving user interests and needs.
- **It's slow**, often taking weeks to completely document a single intent and its associated phrase variations.

The end result is a virtual agent that can handle just a small percentage of possible conversations for likely just a handful of use cases. Post-deployment, those gaps lead to "UMM" errors and bad user experiences that are toxic to a brand.

Benefits of the Chatbase data-driven approach to building a virtual agent

The biggest weakness of the manual discovery process is that it leaves most of the business value of customer interaction data off the table. This data contains everything a company could ever want to know about which requests customers make, how they make them, and how the best live agents respond.

Thus the far more efficient, and less risky, alternative to manual discovery is an

automated, *data-driven* approach that uncovers hidden insights across all customer interactions – for example, to find every possible way that customers request to delay a payment, even in cases they didn't initially know how to articulate the question.

The far more efficient, and less risky, alternative to manual discovery is an automated, *data-driven* approach

Chatbase Virtual Agent Modeling uses Google's machine learning and search capabilities to implement this capability. It analyzes thousands of live-chat transcripts, categorizing customer issues into "drivers" (see Figure 2) and then digging deeper to find specific intents per driver (see Figure 3). For complex intents it finds, Chatbase models simplified yet rich flows (see Figure 3) developers can use to build a voice or chat virtual agent that handles up to 99% of interactions, responds helpfully to more supplemental questions, and knows exactly when to do a hand-off to a live agent. In addition, a semantic search tool (based on technology used by Google Search) helps find the potentially thousands of phrase variations (aka training phrases) per intent. When this analysis is complete, developers can export results to their virtual agent via Dialogflow or other options.

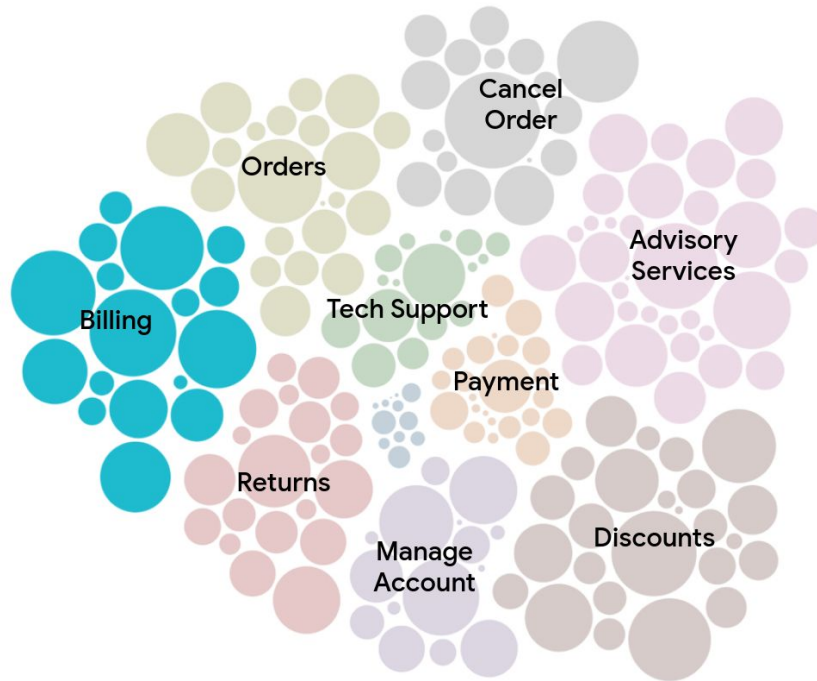


Figure 1. Virtual Agent Modeling categorizes all requests into high-level drivers (use cases).

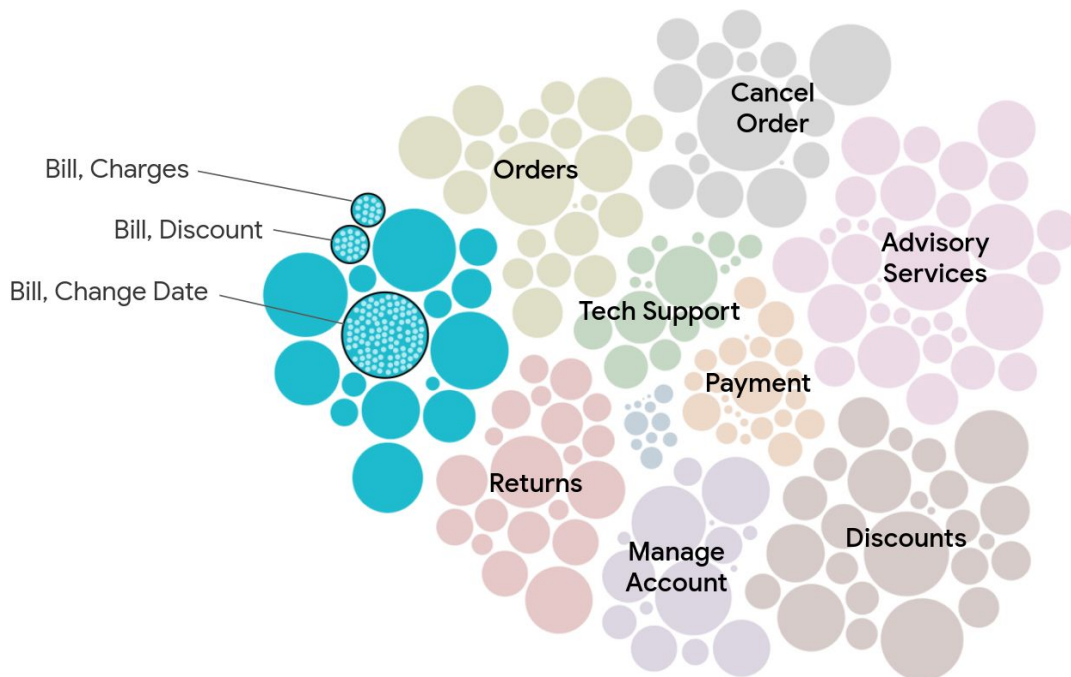


Figure 2. For each driver, Virtual Agent Modeling digs deeper to find popular intents and their training phrases (for improved virtual agent accuracy)

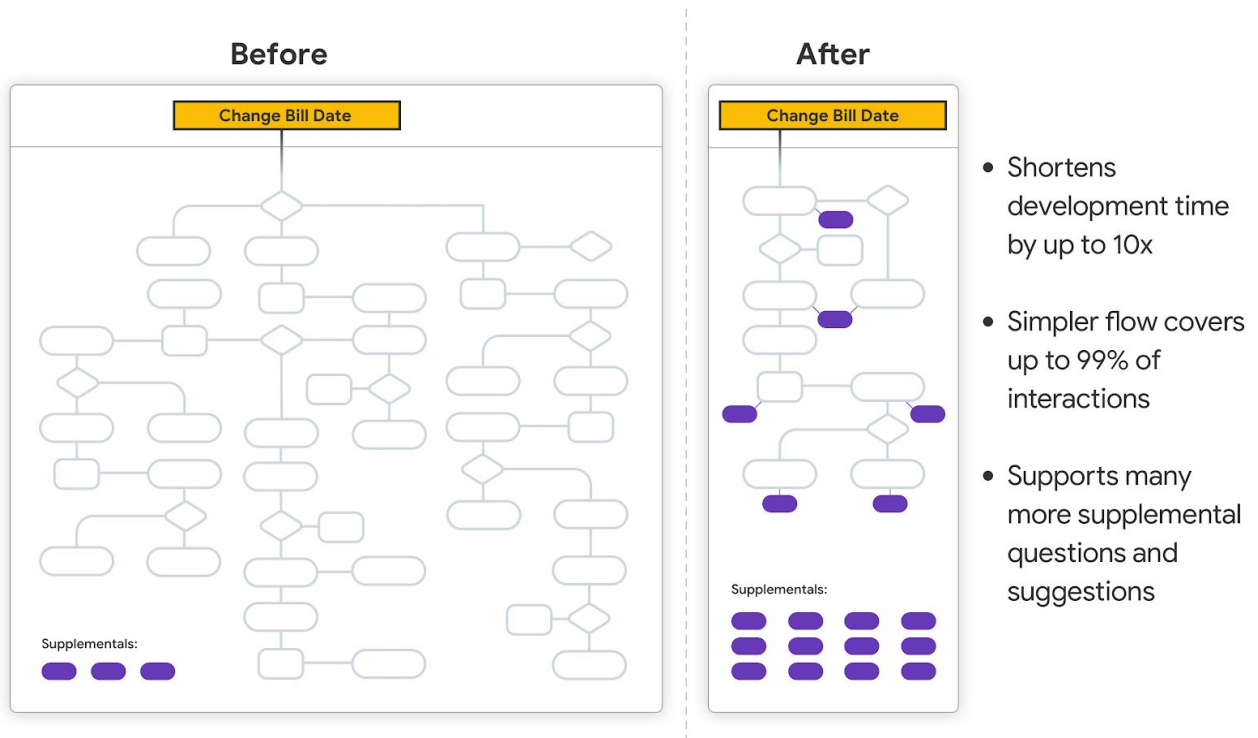


Figure 3. Virtual Agent Modeling creates simplified yet rich flows for building a voice or chat virtual agent that handles up to 99% of interactions.

Benefits for builders of virtual agents include:

- **More automation opportunities.** Virtual agent modeling captures edge cases and hidden trends comprehensively, finding far more opportunities for a virtual agent to add value in a given use case than the manual method.
- **Improved customer satisfaction.** Having up to thousands of training phrases per intent makes it much easier to build a virtual agent with a low rate of missed and misunderstood requests, reducing customer frustration.
- **Shorter development time.** Google’s machine learning algorithms can analyze logs at a massive scale -- in fact, the more data involved, the better the results. The analysis fleshes out conversation paths and training phrases per intent in days, versus a manual process that can take weeks or months.

For one large (Fortune 100) enhancing its customer service virtual agent, these capabilities have been invaluable: In just under two weeks, Chatbase Virtual Agent Modeling discovered hundreds of intents, each mapped to a top-level driver, and thousands of training phrases

for many of these intents. The company anticipates that this new approach to design will speed up the development process by over 10x – furthermore, it expects to nearly double the number of interactions its virtual agent can handle, from 53% of them covered to 92%.

Optimizing the experience

Post-deployment, the ability to uncover functionality gaps via message-level analytics (not just tracking of aggregate KPIs) is also important. [Chatbase Virtual Agent Analytics](#) complements [Chatbase Virtual Agent Modeling](#) by providing, in addition to tracking of KPIs, tools for analyzing conversation flow, handled and not handled messages, plus conversion and retention efficiency. In addition, using machine learning, it suggests optimizations for fixing broken user experiences (i.e. misunderstood and missed customer requests).

Developers can then improve virtual agent performance over time, and expand the intents coverage in their virtual agent to handle new use cases.

Conclusion

Although more and more companies consider AI-powered customer service virtual agents as strategic, the status-quo approach to their design usually fails to deliver a solution that meets that standard. Instead, the Chatbase data-driven modeling approach helps developers build a virtual agent that is engineered to precisely handle specific, real-world use cases. Over time, the number of use cases covered can expand as the model is improved.

[Chatbase Virtual Agent Modeling](#) and [Chatbase Virtual Agent Analytics](#) are bookend products that make that goal possible. To schedule a personal demo, visit chatbase.com.

Footnotes

1. *Customer Experience: New capabilities, new audiences, new opportunities*. McKinsey & Co., 2017
2. American Express 2017 Global Customer Service Barometer
3. *A Customer-Obsessed Operating Model Demands A Close Partnership With Your CI*. Forrester Research, 2016
4. *Chatbots in Customer Service*. Accenture Digital, 2016
5. *Worldwide Semiannual Cognitive Artificial Intelligence Systems Spending Guide*. IDC, 2018
6. American Express 2017 Global Customer Service Barometer

About Chatbase

Chatbase (chatbase.com) is a cloud service for more easily building, analyzing, and optimizing AI-powered customer service virtual agents. Powered by Google machine learning and search capabilities, the Chatbase approach is faster, more accurate, and more comprehensive than manual methods, helping contact centers more quickly deploy and improve a virtual agent that delights customers. Chatbase is brought to you by Area 120, an incubator operated by Google.

Chatbase

Area 120 @ Google

2 Harrison St.

San Francisco, CA 94105