Information Research - Vol. 30 No. iConf (2025)

Empowering customer service with generative AI: enhancing agent performance while navigating challenges

Charles Costa and Souvick Ghosh
DOI: https://doi.org/10.47989/ir30iConf47566

Abstract

Introduction. As large language models (LLMs), such as GPTs, become more intelligent, a key area of exploration is how these technologies can improve the customer experience. Contrary to common belief, many consumers, including Gen Z, prefer human-provided customer service, illustrating the importance of human-AI collaboration in the space.

Method. By leveraging the author's real-world knowledge of enterprise knowledge management and customer service delivery, we reviewed numerous literature about AI, knowledge management, and service design and synthesised practical insights for industry professionals to build a successful AI strategy.

Analysis. We examined the gap between academic research on generative AI and how such solutions can be deployed in real-world enterprise contact centers.

Results. We outline how companies are adopting generative AI, offer strategies for training AI models to assist with customer issues, discuss AI-powered localization, and share tips for facilitating human-AI collaboration.

Conclusion. This research found that companies will require humans for effective customer service delivery because there is no complete substitute for human judgment in mission-critical settings. Despite this, companies can still see significant performance and customer satisfaction improvements by focusing on human-AI collaboration, and empowering human agents to focus on specialized tasks while the AI addresses rudimentary ones.

Introduction

The discussion surrounding whether AI agents will replace human workers has been particularly concentrated on the customer service industry. Considering that customers spend around 14 billion hours annually engaging with customer service representatives and the turnover rate in customer service call centers is approximately 60% per year, costing $10,000 to $21,000 per agent, it is logical for businesses to look for cost-saving measures in this domain (Buesing, et al., 2024).

Although it is understandable to assume that the increased efficiency of AI agents will drive massive reductions in human staffing within contact centers, there are many indicators that contradict this view. Consider the fact that 71% of Gen Z customers believe that phone calls are the quickest and most effective way to resolve customer service matters. Furthermore, Gen Z consumers are 35-40% more likely to call a business than millennials (Blackader, 2024). Despite the benefits of generative AI, the technology isn't a solution for all customer service challenges because LLMs don't provide the empathy and connection provided by human agents in resolving complex matters. Even with automation, companies still need to recruit talented customer service agents to assist with matters that are too complex for automated resolutions (Blackader et al., n.d.).

Research goals and contributions

The phrase '*trust but verify*' is relevant to any business that intends to incorporate generative AI into their customer service divisions. To justify the investment in AI technology, customer service personnel must have the confidence that the information presented to them is accurate, and customer service leaders must be sure that technologies are improving efficiency (Leidner & Mousavi, 2024). The goal of this paper is to apply knowledge from academia to overall customer experience knowledge gained from the author's industry experience by answering the following questions.

- **RQ 1:** what tactics can companies deploy to effectively integrate generative AI technology into their customer service functions on a worldwide scale?
- **RQ 2:** how can generative AI systems be best designed to fit the needs of human customer service agents?

The need for effective human and AI agent collaboration

First, it is important to clarify that this document is specifically focused on companies that integrate AI into an agent desktop, which is considered a solution that allows customer service agents to access internal knowledge bases, product information, and account details from a single tool (NICE, n.d.). To understand how AI can be utilized by customer service agents, one must consider the relevant information science theories. The bounded rationality model is relevant because it states that people make suboptimal decisions due to a combination of (Korzynski et al., 2023):

- time constraints.
- an imperfect understanding of information; and
- human cognitive constraints.

Because customer service agents are frequently rated based on the average handling time (the amount of time spent on resolving a customer contact), human agents must rush through pages of information and make judgment calls based on the customer's input, the information displayed on their agent desktop (e.g., account status, previous notes from agents, purchase history, etc.) and information from the knowledge base (Korzynski et al., 2023).

Even the most talented agent cannot read the customer's mind and will always have a different interpretation of the issue than the customer. Generative AI is a valuable solution to overcome the previously mentioned issues because models such as ChatGPT (https://openai.com/chatgpt/overview/) enable agents to ask questions in a natural language (e.g., 'Why is this customer's account suspended') and receive relevant responses (Korzynski et al., 2023). Because agents are not forced to use internal jargon to query internal databases, the generative AI experience helps to bring parity to how the agent and the customer understand the situation.

AI agents offer significant benefits for self-service customer-facing assistance, as illustrated by SAP Concur (https://www.concur.com/), which realized a 30% decrease in case creation after adopting these technologies. Generative AI-powered search solutions not only improved the ability of customers to express their needs, but AI models could also synthesize responses from various sources. This eliminated a significant friction point: the need for users to search multiple documents to fully answer their question (Garst, 2024).

Consider the hypothetical example of users receiving assistance when using an online website to buy/sell items. When users require assistance, they can access the chatbot through the help pages. Since a company can have a variety of users with different needs (buyers, business sellers, consumer sellers, etc.), the AI agent should consider the customer's input along with the customer's account information to determine the most relevant answers to the questions. If the customer prefers to interact with a human or if the matter is too complex for AI agents, there should be a seamless handoff to the human agent.

When the human agent enters the conversation, they should receive an AI-generated synthesis of the previous discussion along with relevant account information, facilitating a more personalized experience. Figure 1 presents the interaction scenario between the customer, the AI model, and the customer service agent.

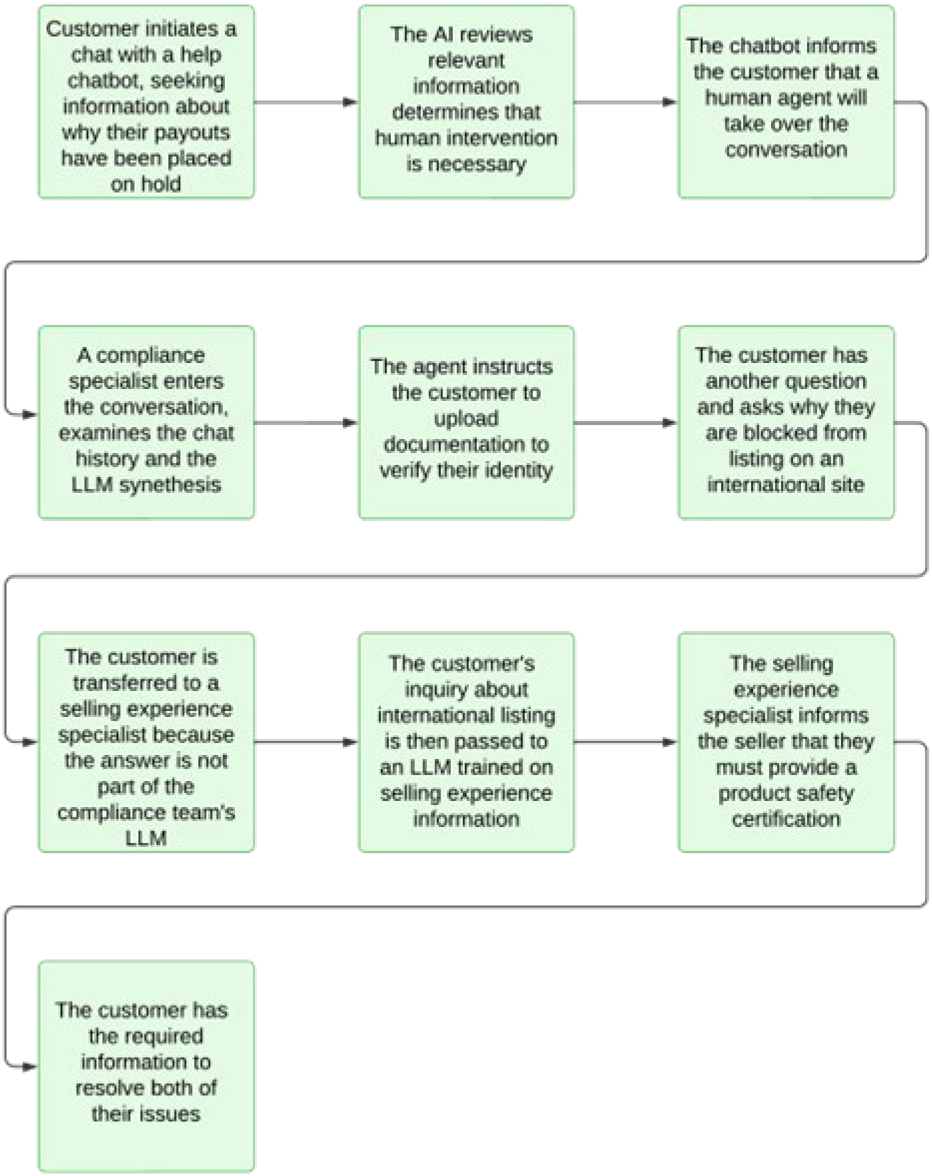

Figure 1. Customer-AI-customer service agent interaction

Klarna (https://www.klarna.com/us/) is one of several companies that have seen success with AI-powered chatbots for customer support. In the first month of operation, the AI assistant handled 2.3 million conversations (two-thirds of Klarna's support chat volume), contact handling times dropped from 11 minutes to less than two minutes, and the number of repeat inquiries decreased by 25%. The company estimated these improvements would result in $40 million in efficiency savings. It is important to remember that one-third of support contacts still require a human agent, meaning the automated system enables humans to focus on specialized tasks (Klarna, 2024) (OpenAI, 2024).

Although generative AI technologies can significantly improve the customer experience, they must be customized to fit the needs of the organization. One of the risks related to AI that companies face is hallucinations, defined for this work as 'a plausible but false or misleading response generated by an artificial intelligence algorithm' (Merriam-Webster, Incorporated, 2024). One of the most egregious situations that illustrates the importance of proper training of the AI model is the deployment of

Microsoft's AI bot Tay, where the chatbot repeated user verbatim phrases and processed information without considering the context of the message (Ars Technica Staff, 2016).

Another notable example occurred in 2023 when a Chevrolet dealership added a chatbot to their website to assist with customer service. Because the dealership did not adjust the chatbot logic to handle only tasks related to providing customer service, users were able to prompt the bot to write Python scripts, respond to provocative questions, and even agree to sell a 2024 Chevy Tahoe for $1 (Notopoulos, 2023).

Enterprise organizations can protect themselves from rogue AI models by defining clear use cases, evaluating the sensitive information that the model can access, and establishing performance metrics for ongoing monitoring (Baquero et al., 2020). Some enterprise knowledge management methods are already conducive to AI model training, such as fine-grained access controls, and audit trails.

Localization: meeting the needs of global customers

Given that enterprises frequently have a global reach, localization is an important capability for any AI model used in a customer service setting. For example, if an online business is trying to do business in the United States, Canada, France, Germany, and Spain, the company would need to have at least five versions of all relevant customer service content. Localization can be challenging because while buyers and sellers must follow the applicable laws where they reside, if they engage in global commerce, they will also have to follow the rules and regulations of the countries within which they are doing business. Consider the following example:

- A business in the United States attempts to sell products in Spain but the eCommerce software blocks their listings because they did not provide an Extended Producer Responsibility (EPR) number.
- A business in China offers products through a UK website because they want to accept payments in GBP. Although they do not accept returns for purchases made on the US site, if a UK buyer wants to make a return, the business will have to honour UK commerce regulations.

Although multilingual pre-trained models trained on hundreds of languages have shown promise in natural language generation (NLG), performance frequently suffers because each language must compete for the same computing resources. When there are not sufficient computing resources for each language, it may fall victim to the 'source language hallucination' problem, where information with the correct meaning is provided in the wrong language. 'The curse of multilinguality' can also apply, which is when a model's performance decreases as the number of pre-training languages increases, or if there is limited information for an individual locale (Piccinno et al., 2023).

Enterprise knowledge management teams often must manage hundreds of thousands, if not millions, of content assets, which can make AI-powered localization difficult if the enterprise has specialized LLMs trained on small subsets of such information. The mmT5 model works to overcome the challenges mentioned above by segmenting language-specific information from general materials (Piccinno et al., 2023). By strategically freezing locale-specific information during fine-tuning, the model can virtually eliminate the source language hallucination problem.

Despite mmT5 being relatively new, the underlying concept of documenting localization nuances in a dedicated hub has been done by many, including the author. The table below is an example of a localization based on the author's experiences at a large e-commerce retailer.

	US	UK	Germany	Australia
Supported payment currencies	United States Dollar (USD)	British Pound Sterling (GBP)	Euro (EUR)	Australian Dollar (AUD)
Supported shipping carriers	FedEx, United States Parcel Service (USPS), UPS (United Parcel Service)	DHL, UPS, FedEx, Royal Mail	DHL, UPS	Australia Post, Fastway, Sendle
Bank processing times	1-3 business days	0-4 business days	1-2 business days	1-3 business days

Table 1. Product nuances across a hypothetical online shop

Information tokenization is another technique to improve localization. Consider emails that are used to educate users about global rules and regulations. A company may have dozens of email templates to inform users about the actions they should take to be successful on the site. Although much of the information can be hard-coded, the portions of content that vary across countries can contain blocks of content that were repurposed from help pages. By reusing content across the organization, knowledge management teams within enterprises can develop more accurate impact assessments for policy/product changes, and the teams only need to write once for all channels.

While it is understandable to believe AI is not capable of understanding linguistic nuances, AI-powered localization is delivering results that come close to traditional machine translation. Expedia (https://www.expediagroup.com/home/default.aspx) uses GPT to provide personalized travel suggestions in multiple languages, with a quality comparable to that of conventional machine translation systems. VMware (https://www.vmware.com/) is another company that is exploring AI in localization by enabling content managers to verify errors in source texts (the initial draft content intended for localization) (Savenkov, 2023). Continuing the theme of companies successfully embracing AI, Warner Brothers Discovery (https://www.wbd.com/) decreased expenses by 50% and the time required to caption content by 80% after embracing AI technologies (Google Cloud, 2024).

Integrating AI into the customer service delivery process

Customer service delivery is often the product of multiple humans working together to solve tasks. Recent developments in AI systems have changed customer service delivery by integrating an automated agent into the mix. Collaborative intelligence (CI) systems that facilitate cooperation deliver the most value to enterprises because they leverage the strengths of both humans and AI models. The core capabilities of such systems include (Blaurock et al., 2024):

- **Engagement:** requesting human agent input during the decision-making process;
- **Transparency:** displaying the logic used to make a decision;
- **Process control:** enabling human agents to adjust parameters for decision making
- **Outcome control:** requiring the human agent to make the final decision on how to proceed; and
- **Reciprocal strength enhancement:** empowering human customer service agents to focus on their strengths (e.g., providing an emphatic response) through efficient information retrieval.

Although generative AI systems are helpful for synthesizing information, the customer service agent still must evaluate the key elements of the customer interaction. For example, the agent must confirm details such as 'Whether the contact is related to buying or selling activities?' or 'What is the exact transaction the customer refers to?' and make general judgment calls based on the nature of the situation at hand.

Generative AI in an agent desktop or any other customer service solution is not a system that instantly diagnoses the issue and offers a solution. Rather, the system is a helper that asks the right questions to get the human agent the most relevant information. Similar to guided logic tools, where agents evaluate refund requests or purchase disputes using predefined logic (ServiceNow, 2023). The main difference between LLMs and guided logic tools is that LLM outputs are tailored to the current task, while guided logic tools must have the exact pathways and all possibilities programmed in advance, which is not only a more difficult task but lacks flexibility.

Examining material hurdles to fully automating customer service

Notwithstanding the promise of AI technology, a lack of trust in technology will ensure humans are required for certain customer service tasks in the foreseeable future. For example, the Digital Services Act (DSA) is a robust law that requires certain online service providers to take extensive efforts to identify 'illegal online content' and take efforts to prevent it from being shared (Apostle et al., n.d.). The law's authors also show concerns with AI automation in multiple parts of the law:

> *…providers concerned should, for example, take reasonable measures to ensure that, where automated tools are used to conduct [content moderation activities], the relevant technology is sufficiently reliable to limit to the maximum extent possible the rate of errors* (European Parliament, 2022, para. 26).
>
> *Providers of online platforms shall ensure that the decisions, [related to user appeals of content removal], are taken under the supervision of appropriately qualified staff, and not solely on the basis of automated means* (European Parliament, 2022, Art. 20, para. 6).

Though the details of DSA case appeals and enforcement are beyond this paper's scope, it's crucial to note that legally, service providers must have human staff manually process appeals for eligible adverse actions. This alone illustrates the need for humans to remain involved in mission critical tasks/situations.

Conclusion

Despite widespread predictions that AI models would spell the end for human customer service agents, both technological and societal factors have shown otherwise. Generative AI models cannot provide the empathy of humans during service calls, and the technical constraints make it impractical for AI models to handle every customer query.

Organizations can mitigate AI risk by building on existing knowledge management best practice and facilitating collaboration between their human agents and AI systems.

In this paper, we examined the gap between academic research on generative AI and how such solutions can be deployed in enterprise contact centers. We outlined how companies are adopting generative AI, offered strategies for training AI models to assist with customer issues, discussed AI-powered localization, and shared tips for facilitating human/AI collaboration.

About the authors

Charles Costa is currently pursuing a Master of Library and Information Science degree from San Jose State University in San Jose, California. They also work full-time as a senior content strategist who focuses on financial services content and automated policy emails for an eCommerce marketplace.

His research interests are in human-computer interaction in technical communication and enterprise knowledge management. He can be contacted at charles@charlescosta.net.

Dr. Souvick Ghosh is an Assistant Professor in the School of Information at San Jose State University, California, USA. He received his Ph.D. from Rutgers University, and his research focuses on human-centred, ethical, and conversational AI. He is particularly passionate about the ethical application of AI and its conversational capabilities within complex sociotechnical systems. He can be contacted at souvick.ghosh@sjsu.edu.

References

Apostle J., Schröder C., Yavorsky S., Egan, H., & Kawkabani, R. (n.d.). Europe's Digital Services Act (DSA) is approved: What you need to know. Retrieved January 4, 2025, from https://www.orrick.com/en/insights/2022/10/europes-digital-services-act-dsa-is-approved-what-you-need-to-know

Ars Technica Staff. (2016, March 26). Tay, the neo-Nazi millennial chatbot, gets autopsied. Ars Technica. https://arstechnica.com/information-technology/2016/03/tay-the-neo-nazi-millennial-chatbot-gets-autopsied/

Baquero, J. A., Burkhardt, R., Govindarajan, A., & Wallace T. (2020, August 13). Derisking AI by design: How to build risk management into AI development. https://www.mckinsey.com/capabilities/quantumblack/our-insights/derisking-ai-by-design-how-to-build-risk-management-into-ai-development

Blackader, B., Luchtenberg, D., & Buesing, E. (n.d.). The evolution of customer care: AI and the Gen Z effect [Broadcast]. Retrieved July 31, 2024, from https://www.mckinsey.com/capabilities/operations/our-insights/the-evolution-of-customer-care-ai-and-the-gen-z-effect

Blackader, B. (2024, April 12). Why your kids aren't calling you, but they are calling their bank. https://www.mckinsey.com/capabilities/operations/our-insights/operations-blog/why-your-kids-arent-calling-you-but-they-are-calling-their-bank

Blaurock, M., Büttgen, M., & Schepers, J. (2024). Designing collaborative intelligence systems for employee-AI service co-production. Journal of Service Research. https://doi.org/10.1177/10946705241238751

Buesing, E. Gupta, V., Higgins, S., & Jacobson, R. (2020, June 22). Customer care: The future talent factory. https://www.mckinsey.com/capabilities/operations/our-insights/customer-care-the-future-talent-factory

Digital Services Act, PE/30/2022/REV/1 (2022). http://data.europa.eu/eli/reg/2022/2065/oj

Garst, A. (2024, October 10). Driving transformation with AI: How SAP is leveraging genai for support excellence. https://www.coveo.com/blog/sap-concur-generative-ai/

Google Cloud (Director). (2024, September 24). Warner Brothers Discovery: AI-assisted captioning on Google Cloud [Video recording] https://www.youtube.com/watch?v=D8F55oFJjRQ

Klarna. (2024, February 27). Klarna AI assistant handles two-thirds of customer service chats in its first month | Klarna International. https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/

Korzynski, P., Mazurek, G., Altmann, A., Ejdys, J., Kazlauskaite, R., Paliszkiewicz, J., Wach, K., & Ziemba, E. (2023). Generative artificial intelligence as a new context for management theories: analysis of ChatGPT. Central European Management Journal, 31(1), 3–13. https://doi.org/10.1108/CEMJ-02-2023-0091

Merriam-Webster, Incorporated. (2024). Hallucination definition & meaning—Merriam-Webster. In Merriam-Webster. Merriam-Webster, Incorporated. https://www.merriam-webster.com/dictionary/hallucination

NICE. (n.d.). What is an agent desktop? Retrieved September 3, 2024, from https://www.nice.com/glossary/what-is-an-agent-desktop

Notopoulos, K. (2023, December 18). A car dealership added an AI chatbot to its site. Then all hell broke loose. Business Insider. https://www.businessinsider.com/car-dealership-chevrolet-chatbot-chatgpt-pranks-chevy-2023-12

OpenAI. (n.d.). Klarna's AI Assistant Does the Work of 700 Full-Time Agents | OpenAI. Retrieved June 29, 2024, from https://openai.com/index/klarna/

Pfeiffer, J., Piccinno, F., Nicosia, M., Wang, X., Reid, M., & Ruder, S. (2023). mmT5: Modular multilingual pre-training solves source language hallucinations (No. arXiv:2305.14224). arXiv. https://doi.org/10.48550/arXiv.2305.14224

Savenkov, K. (2023, June 28). Real-world genai applications: global enterprises using ai for enhanced localization programs » Intento. Intento. https://inten.to/blog/real-world-genai-applications-global-enterprises-using-ai-for-enhanced-localization-programs/

Service Now. (2023, August 2). Configuring guided decisions. https://www.servicenow.com/docs/bundle/vancouver-customer-service-management/page/product/customer-service-management/concept/setting-up-guided-decisions.html