



Revolutionizing customer experiences with generative AI

Customer experience is the cornerstone of business success in a customer-centric market environment.

Generative AI is coming in full force to disrupt how customers and employees interact with companies. It is helping to plan vacations, answer customer support questions, and guide users through onboarding.

Learn to harness the full capabilities of generative AI to enhance your customers' experiences.

How to enhance customer experiences with generative AI

Resolve issues quickly

Customer support can quickly retrieve documentation applicable to a customer service request, without knowing exactly what the search needs to be. Get relevant information quicker, saving employees' and customers' time.

Get answers faster

Rather than getting a list of results, customers can get that one answer they're looking for. Need the cheapest, round-trip flight from London to NYC next Friday? Get one flight that fits that criteria without being bothered by other results.

Empower your team with the right information

Deploy advanced chatbots and virtual assistants that can find relevant, timely information. Your team doesn't want their search to result in a PDF they then need to search through. They want an excerpt from that PDF that tells them the exact answer they were looking for.



Three benefits of generative AI



Increase efficiency

Delivering the right information, faster with process automation.



Personalize experiences

Presenting personalized answers and recommendations.



Accelerate competitive differentiation

Providing innovative experiences and timely responses.

Ready to take the next steps?

Empower your customers with real-time, relevant answers using Search AI — generative AI's computational intelligence combined with search technology's precision. With Search AI, you can securely bridge your customer information and events data with LLMs to generate answers that are up-to-date, accurate, relevant, and specific to your customer experience issues.

Get the steps to implement generative AI through Search AI at your organization.

[Find out more](#)

AI and generative AI cheat sheet

Artificial Intelligence

Artificial intelligence (AI):

The ability of machines to perform tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and decision-making.

Artificial intelligence for IT operations (AIOps):

The application of AI, machine learning (ML), and analytics to improve the day-to-day operational work for IT operations teams.

Deep learning:

A subfield of neural networks that has many layers, allowing it to learn significantly more complex relationships than other machine learning algorithms.

Machine learning (ML):

A branch of AI that focuses on the use of data and algorithms to imitate the way humans learn, gradually improving accuracy over time. One way they do this is with neural networks that utilize interconnected nodes in a layered structure that resembles the human brain.

Natural language processing (NLP):

A subfield of artificial intelligence that focuses on enabling machines to understand, interpret, and generate human language.

Neural networks:

A type of machine learning algorithm that consists of interconnected layers of nodes that process and transmit information. It is inspired by the structure and function of the human brain.

Generative AI

Generative AI:

A branch of AI centered around computer models capable of generating original content that mimics human creativity. By leveraging the power of large language models, neural networks, and ML, generative AI models are trained to learn the underlying structures, relationships, and patterns to produce new and unique outputs like images, video, code, and more.

Large language model (LLM):

A deep learning algorithm that can perform a variety of natural language processing (NLP) tasks.

Prompting:

A prompt is an instruction given to an LLM. Few-shot prompting teaches the model to predict outputs through the use of examples.

Retrieval augmented generation (RAG):

A framework that enables users to “feed” an LLM private or proprietary, external data so it has the most up-to-date information.

Hallucinations:

When an LLM produces a false or nonsensical output or one that does not match the user’s intent. Because large language models are not search engines or databases — they only predict the next syntactically correct word or phrase — they can appear to produce results that are factually incorrect or contradictory, especially if the data set they are trained on contains contradictory information.

