

AI for the Contact Center: 5 Key Considerations

Artificial intelligence (AI) is today's trending technology. Everybody's talking about it, and companies in every industry want to understand what it can mean for them.

The contact center is one place that is clearly ripe for AI-driven transformation.

It's no wonder that nearly every vendor selling into the contact center environment is touting their AI offering. It's not easy to evaluate AI solutions, especially when the words from different vendors are so similar and the "how it works" isn't visible. But, the depth of AI technology different products offer varies widely—and can have a significant impact on the real value you receive. As you explore how AI can serve your business, this brief can help you ask the right questions and sort "AI light" offerings from technology that can deliver transformative results.

Overview

These are five key topics to consider as you're evaluating AI offerings:

1. What is the general approach to AI?
2. How are the machine learning models built?
3. Is the data used to train the models relevant to the business?
4. How will the solution behave over time? What support is needed?
5. AI feels like a black box.
How can I understand the workings of this solution better?

1. What is the general approach to AI?

Contact centers have been going and growing for 40+ years now, so there's a lot of existing technology. Legacy contact center vendors are eager to include AI capabilities as part of their offerings—but, they're limited in how deep they can go. They're burdened with building onto their existing stacks. And, their primary expertise is not AI.

Newer vendors take various approaches in creating their offerings. Smaller companies are often just adapting low/no-code models or generic models from the open domain. They do this because they simply don't have the resources to build something from the ground up for the specific application they have in mind. Big cloud providers have resources to invest in developing state-of-the-art AI models, but the broad use of them means that they aren't particularly well suited for domain or industry-specific applications.

Very few are reimagining what's possible when you approach contact center challenges with AI as a foundational technology. If your aim is to go beyond incremental improvement and see transformative results, seek out vendors who are building AI native software specifically for the contact center.

Here are some questions to ask to help you understand your vendor's approach:

- **What is your primary business and area of expertise?**
AI expertise with a contact center focus will yield the best results.
- **What can you tell me about your research and data science teams?**
You're looking for a deep team committed to pushing the boundaries of AI and actively engaging in product development and ongoing evaluation.
- **Are your machine learning models built in house?**
If they're not, you are likely looking at an "AI light" offering that will only provide limited benefit.

2. How are the machine learning models built?

There are at least three major modeling techniques used to get computers to act in intelligent ways:

- **Software models that explicitly encode human knowledge and expertise (so-called heuristics or rule-based systems)**

In systems that are ‘rules based,’ performance will be limited by what the ‘experts’ already know and typically have no easy mechanism to update the rules that govern the system’s behavior/performance.

- **Software models that can learn their own rules from very large sets of training data (so-called unsupervised machine learning)**

True ML systems discover new patterns and formulate new rules using more and more data the longer the system is in production—which typically leads to improved performance.

- **Software models that learn their own rules from training data which are then modified by way of human expertise (so called ‘supervised learning’).**

Supervised learning systems have the ability to find far more details than we could have taught the system up front but also depend on humans to help set up the initial patterns and define the action to take. Typically these systems will find new patterns and we then have to manually look at them to help figure out what the new pattern actually IS (this is also called annotation) and what action should be taken.

A supervised learning example

Let’s take a look at an example of an automatic speech recognition (ASR) model used for speech-to-text transcription. Nearly all commercial ASR models are trained with human annotated data (supervised learning).

LibriSpeech, a speech corpus of 1,000 hours of transcribed audiobooks, is one of the most popular benchmark datasets for ASR R&D. Thanks to the dataset, ASR models can learn rules such as phoneme-grapheme mappings (e.g., “f” -> “f”, “ai” -> “igh”) and context to next word mappings (e.g., “How are” -> “you”). Note that the training data can barely encode all possible rules so the models researchers develop need to generalize similar rules given the ones included in the data.

Modern speech-to-text systems employ sophisticated neural networks to implicitly encode the rules. Systems such as Google-ASR and Facebook-ASR are built based on a transformer, an advanced sequential neural network architecture.

ASAPP built a best-in-class ASR system explicitly designed and trained for enterprise contact centers. It is based on a novel multi-stream convolutional neural network (CNN) model, which outperformed the Google-ASR and Facebook-ASR on the LibriSpeech dataset in 2020.

Here are some questions to ask to help you understand how your vendor builds their machine learning models:

- **Are you building your own models?**
Leaders in AI will be building their own models rather than using off-the-shelf tools.
- **What examples can you provide of model performance?**
You’re looking for insight into how model performance is measured and the associated results.

3. Is the data used to train the models relevant to the business?

The ability of machine learning models to perform well in your environment is highly dependent on the quality of the data used to train it. The quality of the data can be characterized by two aspects: 1. The relevance of the data 2. The size of the data.

About data relevance

Most of the public datasets for conversational AI research, such as ATIS and MultiWoZ, focus on human-machine interactions in order to build chatbot systems. As a result, the data is typically dissimilar to complex human-human conversations, and so less than ideal for training models for conversations that take place between agents and customers (whether by chat or phone.) Further, the lexicon is unlikely to be specific to your industry, and definitely not specific to your business.

About data size

Obtaining a large amount of training data can be expensive, especially considering the data annotations by domain experts required for supervised learning. Starting with domain-specific data and a good model, it is possible for a vendor to quickly scale its machine learning capability to different customer domains with less data. This problem is known as meta learning or few-shot learning. Popular approaches for solving the problem include prototypical networks and Model-Agnostic Meta-Learning (MAML).

Here are some questions to ask to help you understand what data is used to train your vendors machine learning models:

- **Is the core data set you're using from human-to-human conversations ?**
If not, you can expect lower performance from the model.
- **Can the model be tuned for my industry? For my business specifically?**
Industry-specific training will improve performance. Custom training on your data will make it even better.
- **How do you ensure that you have enough training data?**
You're looking for recognition of the challenge and thoughtful approach to addressing it.

4. How will the system behave over time? What support is needed?

Simple AI systems recognize patterns and suggest or take action based on a set of rules. These systems require administration to supervise all learning and to write all rules for behavior. They are limited by what the administrator already knows and their ability to imagine all the rules needed. Updates typically require recognition of the need for new learning, the initiative of the admin to instigate updates, and substantial manual work to update the system.

More advanced AI systems continuously self-learn, finding new patterns and adjusting suggestions and actions as they process new data. It is necessary to help set up the initial patterns, and define actions based findings (rules), but then the system learns on its own. It can highlight new patterns, trends, and anomalous behavior, suggest adjustments, and alert an admin to review.

Here are some questions to ask to help you understand how a system will behave over time and what is required to support it:

- **Is the system all rules-driven or is it self-learning?**
You want to understand what resources you'll need to allocate to keep your system current.
- **How do you get feedback on what's working and what's not?**
You want to understand how to assess how well the system is serving the needs of your team and what insight you'll get on any changes needed.
- **How can learning be applied across different customer channels?**
You want to understand if learnings from voice/calls and chat/messaging are completely siloed or can be used together to inform insights.

5. AI feels like a black box.

How can I understand the workings of this solution better?

As mentioned in the opening, there are a broad range of AI offerings for the contact center on the market today. Vendors who are using AI as a foundational technology are pushing the boundaries of what's possible with AI to deliver extraordinary results. While it all can feel a bit like magic, your vendor should be able to give you some understanding of what's behind the curtain.

Here are some questions to ask to help dive a little deeper into your vendor's AI:

- **Can I meet with one of your research scientists directly?**
You're looking for an easy "yes" followed by a meeting with a research scientist (ideally one that is actually responsible for the models you'll be relying on) that answers your questions and makes you feel knowledgeable.
- **How is this different than <another product I'm considering>?**
You're looking for more than a checklist of features here. You want to understand why the differences matter and how they'll impact your business.
- **Where can I read more about your AI?**
An active research team will be able to point you to blog posts, papers, peer-reviewed publications, and patents for a more in depth understanding.
- **How does this model grow / change as my business grows and changes?**
Adaptability is the hallmark of great machine learning. Look for responses that suggest minimal manual intervention.

About ASAPP

Since 2014 we've pushed the boundaries of AI to change what's possible in CX. Our software is delivering transformative results for some of the world's largest consumer brands. Companies using ASAPP are lowering costs, driving new revenue, and increasing both customer and agent satisfaction at the same time.

Ready to learn more?

Get in touch at ask@asapp.com