



Enhancing Customer Support with AI

.....

5 crucial metrics to consider

Intro

The purpose of this article is to give leaders interested in AI for customer support a general guideline on how to best evaluate different solutions, and help decide whether the technology is right for their business. After reading this article, you should be able to understand what specific benefits your company might expect from the AI product you're considering, and properly assess whether or not it's worth moving forward to a deeper investigation of that product.

The following evaluation criteria provided in this article is based on the experience of AI leaders from Google, Amazon, and Facebook, with intimate knowledge of most major players out there.

Background

The world moves quick, and by 2020 it is predicted that virtual agents will be involved in 85% of all types of business-customer interactions, [saving businesses nearly \\$8 billion annually by 2022](#). That's a lot of hype. Let's focus today on separating the hype from reality. "Artificial Intelligence" and "Machine Learning" have become buzzwords, and it's critical to ensure your investment will be spent on solving real business needs—enhancing the customer experience, improving labor efficiency, and driving cost savings—in an intelligent and realistic manner.

To facilitate an easier decision-making process, we've laid out five capabilities every AI system should have, and what to look for when evaluating them for your business:

1. AI Training

2. Performance and Accuracy

3. Customization and integration

4. Self-learning

5. Return on Investment



1. AI Training

Before deciding on any solutions, the first thing to ask is how the AI training process looks like, just like you need to understand the upfront investments before getting the ball rolling for any project.

Do I need to manually teach the system? What does the AI training process look like?

The more knowledge that AI has about your business, the more effective it can be. The accumulation and processing of this knowledge, however, typically constitutes some form of investment. Most AI applications for support will tackle the training (and the cost involved) in one of two ways.

Some companies will provide you with the tools necessary to create a conversational agent of your own. This type of manual training imposed upon your team is akin to teaching a toddler a new language, and will have a lower barrier to entry than other options. The AI may already possess a simple understanding of natural language, but will require internal effort to ensure this can be applied towards your customers specific use cases. Considering the scope of customer expectations today, this approach typically requires several months of internal effort to produce a relatively rigid AI understanding of your customers intentions, draft a large set of rules for each scenario, and have the capability to respond to them accordingly. It also leads to a higher overhead cost in the long-term, as self-evolution will be all but nonexistent (refer to Point #4 below for more information).

Alternatively, it is possible to take a more data-driven approach. While the initial overhead may be higher, the long-term costs associated would be minimal, especially if the AI is capable of self-improvement. Advanced technologies will utilize existing support tickets or chat transcripts to train the AI, mimicking the way your human agents have handled cases in the past. This example is comparable to a teacher sharing textbooks with a student, enabling them to form an in-depth, but guided, understanding of the scope ahead. As a result, this usually leads to more effective solutions with a frictionless onboarding process and a highly tailored solution for your team.

2. Performance and Accuracy

Artificial intelligence is still a fast-growing field in both academia and industry, especially when it comes to natural language processing. For support virtual assistants, they are only as effective as their ability to break down complex language and execute complex tasks.

Is the AI based on the latest Deep Learning and NLU technology, or more traditional search or Rule-based systems?

There exists a great deal of “AI” solutions out there. Unfortunately, many companies still heavily rely on rule-based or keyword-based implementation, built in a brute-force manner by a large engineering team. This old technology leads to limited opportunities for automation and optimization.

To optimize performance and ensure a positive user experience, affirm that the technology utilizes deep learning and the latest advancements in Natural Language Understanding. It may also be useful to ask if they are leveraging in-house proprietary technologies, or wrapping up open source 3rd party libraries.

What is the typical system accuracy, and how is this figure calculated?

Proprietary AI technology is easier said than done. It's important to benchmark the vendor's system accuracy and understand how that accuracy is evaluated. One systematic way to measure accuracy is to sample AI involved conversations and categorize them into accurate ones (true positive and true negative), and inaccurate ones (false positive and false negative).

Does the core team have the right expertise? What proprietary technology does the company have in place?

What type of AI experience do the founders and core team possess? Have they accomplished anything significant in the space? Consider asking what type of proprietary technology the company has in place.

3. Customization and Integration

Be sure to understand the resource allocation that is necessary to fully harvest the benefits of an AI agent. What is the typical timeline and resources involved in customizing the solution and integrating into existing workflows and applications?

How is the solution integrated with your existing workflow?

AI-augmented workflows are one of the key drivers in modernizing customer support teams. Per a recent [PwC survey](#), 59% of executives believe that the lack of integration of new and existing technologies poses a barrier to achieving expected results from digital technology initiatives. Will your solution be able to ENHANCE the productivity of your team, while reducing workload? Make sure your “value-add” is not adding additional work to your team.

To achieve the highest performance, your platform should be able to integrate with a number of different channels, including but not limited to prior chat conversations, emails, Knowledge Bases, and CRM's or ticketing platforms. Enabling such integrations will ensure your AI can learn dynamically and quickly based on historical conversations between your customers and employees, thus ensuring a consistency with branding and company tone.

What types of use cases should my solution be ready to facilitate?

Your solution should be able to deliver meaningful messages, take direct orders from users, answer frequently asked questions (FAQs), access personal information, account statuses, previous actions and more. A system that analyses all these interactions will deliver a good user experience, and contextually-aware analytics.

4. Self-learning


Having your AI virtual agent trained and integrated with your workflows would be an exciting milestone to celebrate, right? Well, not quite yet... Just like your employees, the AI agent will need time to learn and expand on both its depth and coverage of the knowledge.

What needs to happen if I want my AI agent to learn more? What is the level of self-learning currently attained by the solution?

Some platforms will require you to manually teach basic skills: If a customer asks THIS, respond with THAT. This is a very time-consuming process, and unless you have a team of dedicated engineers who can monitor each input and output, this will more often than not lead to frustration and a very slow improvement rate.

Thankfully, with deep learning and data-driven semantic parsing, manual training would be mostly avoided. Self-learning will ensure increased resolution rates over time, and less overhead required from a team in the long-run.

Considering your future resource allocation and service ROI, this is one of the most crucial areas to consider. Does it feel like an ongoing maintenance burden for your team, or an opportunity to consistently achieve more?



5. Return on Investment

Finally, we must mention what is on every decision maker's mind when evaluating a new enterprise support platform: Return on Investment.

What is the time to market?

How long will it take to bring your solution to market? You should discuss ways to ensure a quick and reliable timeline when launching and refining a solution.

Most advanced enterprise solutions take 1-2 weeks to train and deploy a fully integrated AI-powered virtual agent.

How can it improve the key metrics?

Will the solution decrease average response times and time to resolution? What are the expected resolution rates of your agent, and how much workload will that actually deflect from your team? Will these metrics be backed by data?

What is the pricing model?

Pay-per-call: Cost is based on usage of the number of software API calls. A pay-per-call pricing will give you flexibility to pay-as-you-go when launching your AI agent, but may be hard to have a concrete ROI projection.

Pay-per-performance: Volume based pricing is simple to understand, but does not inherently represent confidence in a solution. If you can, look for performance-based pricing that guarantees you to attain certain performance metrics (KPMs) and ensure you could achieve a strong ROI.

Monthly license fees: A flat license fee will provide you with regular AI agent operation (and sometimes support). Make sure to evaluate expected resolution rates and productivity improvements prior to going with this model to ensure a reliable ROI.

In Conclusion

As you continue your search for an intelligent AI platform, make sure you find a model flexible to your needs. No matter what platform you choose to go with, make sure you're working with an experienced team who understands your company's needs, priorities, budget, and strategic goals. We at PerceptAI proudly differentiate ourselves by utilizing the latest data-driven technology to integrate seamlessly with existing platforms and enhance business operations within the enterprise.

If you've found this article to be useful, or have any feedback on more you'd like to see, let us know!

EVALUATION QUESTION	SCORE	ADDITIONAL INFO
AI Training		
Do I need to manually teach the system? Any coding, data labeling, rule designs required?		
What does the AI training process look like? Is the solution tailored to my company?		
Performance and Accuracy		
Is the AI based on the latest Deep Learning and NLU technology, or more traditional search or Rule systems?		
What is the typical system accuracy, and how is this figure calculated? 70%, 90%?		
Does the core team have the right expertise? What proprietary technology does the company have in place?		
Customization and integration		
How is the solution integrated with my existing workflow?		
What types of use cases should my solution be ready to facilitate?		
Self-learning		
What needs to happen if I want my AI agent to learn more? Is there any maintenance cost?		
What is the level of self-learning currently attained by the solution? Can it expand the knowledge coverage by itself?		
Return on Investment		
What is the time to market? Does it take weeks or months?		
How can it improve the key metrics? How are the key metrics backed by data?		
What is the pricing model? Is it Pay-per-call, Pay-per-performance, or Monthly license fees?		