

## Revolutionizing Call Center Performance with their Voice AI Platform, OBSERVE.AI, is providing the ability to derive Intent, Sentiment and Emotion from the Voice Stream and Provide Recommendations to Agents in Real-Time



**Sharath Keshava**  
Co-Founder

**OBSERVE.AI**  
[www.observe.ai](http://www.observe.ai)

**Contact:**  
**Sharath Keshava**  
**(510) 458-7886**  
[sharath@observe.ai](mailto:sharath@observe.ai)

**Interview conducted by:**  
**Lynn Fosse, Senior Editor**  
**CEOCFO Magazine**

**CEOCFO: *Mr. Keshava, what is the concept behind Observe.AI?***

**Mr. Keshava:** Observe.AI is a Voice AI platform for call centers. We have been around for the last eighteen months. We started out as a Y Combinator

incubated start-up and raised a Series A round last May. We have a strong fifty member team based in San Francisco and Bangalore, India. The goal is to work with enterprises that have large support teams and help them get visibility into their voice calls, give them insights into what happens on all their calls and find opportunities to,

1. Train or coach an agent to manage the customer better by giving him & his supervisor visibility
2. Provide recommendations to an agent in real time to make them more effective in managing customers
3. Find opportunities to automate / deflect these calls by building conversational maps and voice bots

The ultimate goal is to optimize business processes for call centers to improve their end customer experience and make agents life better. That is Observe.AI for you.

**CEOCFO: *How does that work on a day to day basis? Would you give us an example of someone you are working with - how the engagement starts and what you find?***

**Mr. Keshava:** I think that every organization today, if you look at it from an Integration standpoint, will generally have a Contact Center Management Solution. They generally receive all these calls or make all these calls through their call center infrastructure. This is where Observe.AI usually gets connected, especially if we are monitoring these calls in real-time. As a call center receives/makes calls we will start receiving these calls and start processing. Then based on the analysis we do we start providing insights to the agent directly, or to a supervisor who is monitoring these agents in real time. There are scenarios where call centers do not have a real time infrastructure or not prepared or do not want real time guidance. In these scenarios we do post call analysis. Here we usually integrate with their recording infrastructure and at a pre-determined frequency receive these calls and start processing them. Post processing, we deliver the insights both to the agent and supervisor on their respective Observe.AI dashboards. of the customer. We process every single call received by a call center thus giving them full visibility on every single customer interaction.

**CEOCFO: *What are some of the typical things you look at in a call? Are there things that would surprise people you are able to review and that people may not recognize can be a game changer or make a difference for a company?***

**Mr. Keshava:** Enterprise support is an 800 pound gorilla. There are so many things that can be optimized in enterprise support and there are a few that we have done with some of the customers that we have worked with. One of the most common use cases and the easiest to set up is essentially quality monitoring. Quality monitoring, as a process today, is extremely manual. People usually have QA analysts whose primary role is to pick up a sample set of calls and listen to those calls manually. They generally have an Excel sheet or a CRM, where they have fifteen to twenty parameters that they are looking to evaluate a call on. Then while they are listening to the call manually they evaluate the calls from these fifteen to twenty parameters and then give that feedback to an agent. This process is completely manual and since they pick up a random set of calls the sample set is not exhaustive and at best they audit 0.5-1% of calls for each agent. One of the things that we are trying to do with our VoiceAI platform is to completely automate the entire process. We are analyzing every single call. Once we onboard a customer for a quality monitoring use case the quality analyst (QA) logs into our system and configures the same 15-20 parameters that he/she wants to measure in every call and we analyze each call for these 15-20 parameters and present the insights to the supervisor on his/her dashboard. This will take out the randomness that was associated with the process earlier. We have gone a step further, where a QA can use our system to directly give feedback to an agent. They do not have to do set up a one on one again to give feedback to an agent, since we are already analyzing all the calls. We have an agent interface within our system, where they can set thresholds that can automatically be triggered to an agent or they actually have an inbox where they can leave a comment to an agent for that specific instance of the call. Then the agent will receive that feedback automatically. That is the simplest use case for quality monitoring that we get used for.

**“Ability to understand speech is at the core of our technology stack. Our ability to understand audio signals and derive intent, sentiment, emotion and silences from the voice stream is the core technology differentiator.” - Sharath Keshava**

The second use case that many financial institutions use us for is the compliance use case. There are multiple things that should not be said on a call or should be said on a call. Take an example of a company selling an insurance policy. Once they sell an insurance policy they are supposed to read out about seven different disclosures on what this insurance policy entails. On every call, it is mandated by law, that all those seven things are said on a call prior to booking a sale. How do you know, in the thousands of calls that your agents make every month, if these disclosures are being read out? That is a very simple thing that can be tracked in our system. The second thing in any banking process is authentication. An agent will have to ask a few questions (say 3 questions) to authenticate a customer, such as, “What are the last four digits of your social security number, what is your date of birth, confirm your mailing address etc.” How do you know if an agent is asking those three questions on every single call and if the customer is giving the right answers to all these questions? That is another thing that you can configure on our system, so that you can track for compliance on all of your calls. The third thing, which is most important, is you have to take a consent to recording every call. In most of the states you have to tape an explicit consent where a customer actually says yes to recording. How do you know that your agent has captured that consent explicitly on every single call? These are very small nuances, but this is super important and critical for a large bank that probably gets thousands of calls every day and there is no way for them to proactively monitor for compliance. However, with our system they can do that because we are analyzing every single call. This is a second commonly used use case.

The third use case is primarily used by sales teams. This is where our entire concept of AI comes in. Take for example, a large ticketing website. They basically receive calls from customers who probably want to book a hotel or who want to book a holiday package. These calls have a higher transaction value. One of the common use cases of our product is while the call is happening can we tell the enterprise whether this lead was a hot lead, a warm lead or a cold lead. This lead classification system is completely driven by our NLP engine. Once we have identified a lead as a hot lead, a warm lead or a cold lead, then the enterprise can determine, say for example, I received a thousand calls today.

Observe.AI's system qualified one hundred of them as hot leads, but only twenty percent of them closed, so what happened with the remaining eighty leads.

1. Was there a problem with our agent? If that is the problem then we can go back and train the agent so that my conversion rates can go up.
2. Is there a problem with my availability? What if the customers are asking for a place in San Francisco and we do not have enough inventories on that. I can take this feedback to my inventory team and give them that.

These are all the insights that they can get from our system that they never had before. This basically covers both visibility and then the ability for an enterprise to help in increasing conversions.

**CEOCFO: *How do agents respond to the monitoring? In general, are they able to recognize that there is a good reason behind it?***

**Mr. Keshava:** If you look at the language on our website, the goal is actually make every agent a super-agent. Super agents would be a more marketing term, but the underlying message here is how we can make you better and how we can actually help you to do your job better. That is the whole reason of actually making this entire tool agent facing, because most of the other companies in this spectrum have essentially built tools for the supervisor, so for an agent they believe that, "This is big brother actually listening to my calls and complaining about everything that I am not doing right to my supervisor and all I get to hear is my supervisor ranting about how I am not doing the job right." However, what we have tried to do is actually make this entire tool agent facing, so that an agent has full visibility into every single metric that he is being tracked on.

If a supervisor goes to track and has about twenty things that he wants to see done right on a call, then the agent knows what these twenty things are that he is being tracked on. If he is not doing well on one or two things, even before the supervisor reaches out to him, he can actually go and look at the system, look at the things that he has not done right on a call or on a given day, and then try and improve that proactively, rather than wait to be reprimanded by a supervisor. Agents have taken it in the right spirit because everybody wants to do their job well. Until now, there was no one that could give them that feedback consistently and at scale and here we are giving them all the visibility on how they are essentially being evaluated on and this is what they actually like about our product. Initially any change is slightly harder. However, once people start realizing what the product can actually do for them on a personal basis, they have actually loved our tool. Therefore, we see huge adoption from the agent (call center representatives) community where they start using our tools very, very proactively. They are the agents who spend a good thirty to forty five minutes every day just reviewing all their calls on our system and trying to improve their performance.

**CEOCFO: *What is it about your technology that allows you to be so thorough and comprehensive when others cannot?***

**Mr. Keshava:** One core differentiator is our innate understandings of the enterprise support landscape. Before we even started building the product we actually spent a lot of time with large call centers. In fact, me and my cofounders spent many days/weeks in Manila taking people in the call center industry. We met a ton of people, both at an exec level and with people who would eventually end up using our products, as well as for understanding workflows in most of these large contact centers. What is the workflow in these support processes, what are the different opportunities that we can optimize using AI? I think that understanding helped us to be very empathetic in actually building up this product.

The second thing is that we started working with larger call centers from day one, so the use cases and the data sets we got access helped us to be enterprise-ready much faster than a typical Bay area start-up who would have smaller start-ups as their initial customers. To build a cutting-edge AI platform you need large data sets, you need a specific use cases and you need a large customer who will take a leap of faith and help you build a core infrastructure. If not, you will end up with inconsistent data sets, you will end up spreading yourself thin with smaller customers and the machine learning models you will build will never be accurate enough to solve for a use-case. One thing that I realized in the growth of Observe.AI in the last eighteen months is that many people have been burnt out by AI because setting up an AI system is a lot harder than people think, there are many false positives that come out of the system & there is no silver bullet. If you are saying an agent is not doing the right job then you have to back it up with data. When you are trying to show this to an agent and if an agent finds five things out of ten is incorrect he will lose trust and once they start losing confidence adoption goes down significantly. That is one thing that we have been cognizant of. In building our product we essentially picked one large customer and we worked with that customer for a good six to eight months before we made our product available to anybody else. That enabled us to bring to market a platform that was easy to configure & get started, had a well thought out interface, was easy for people to adapt, the default transcription was far superior than anyone else in the industry and the language models we had configured on audio had much higher accuracies on detecting sentiments, emotions, silences. All that helped in better adoption of the product.

**CEOCFO: *Are you able to help with the customer that is aggressive and really challenges the customer service rep or in ways that they are not prepared for? Are you able to detect some of the more outlying situations or at some point might you be able to help agents understand how to deal with a customer that is not acting the way they expect?***

**Mr. Keshava:** That is a good question. That is the second part of the product. The first part of the product is us trying to do post analysis and providing insights on why a customer called, did the agent do everything he was supposed to do, did the agent solve the customer problem and what was the overall customer experience on this call. The second part of the product is actually helping an agent in real-time. In this industry, the churn is so high; on an average every call center

reports between forty and forty five percent of attrition. Therefore, there is a huge turnover in this industry. Every customer support rep you will end speaking two will probably not have spent more than three to six months in the process and it is impossible for a rep to know everything about the product he/she is supporting in such a short timeframe. That is where the entire augmentation framework comes into the picture.

We are not just helping an agent “post” the call is done, we are also helping an agent while the call is happening, as to what the different kinds of answers he can actually give to the same question. For example, you call into a bank and you say, “My credit score is 740. What are the different platinum cards available for me?” One of the first things that most agents would do is say, “Can I put you on hold while I go and fetch that information?” Then he or she goes to the intranet or knowledgebase where they will go and start searching for information. You will probably have spent five minutes to get hold of this agent and then he puts you on hold for another five minutes. He is not doing it intentionally, but he does not know all the answers. That is where an AI system like ours can be helpful, where we can retrieve that information for the agent in real-time and popup those recommendations to an agent so that an agent does not have to go and search for this information and then from a customer standpoint you do not have to wait for that information to be provided to you.

This entire information retrieval process is something that we do by connecting with their data bases, by connecting with their knowledgebase. Since we are always on as a system, we are also looking at all the different types of call that other agents are receiving and how different agents are responding to the same kinds of questions and we can bring in all those best practices to another agent who is on the phone by building conversation maps. This is where we start helping agents in real-time.

**CEOCFO: What is your strategy for the next year or so? What is the focus for Observe.AI?**

**Mr. Keshava:** Every company in this domain talks about them doing something unique in terms of Voice AI, so it has become very crowded. What we are trying to do is build some very strong differentiators in this domain. This is our approach to AI and automation:

Ability to understand speech is at the core of our technology stack. Our ability to understand audio signals and derive intent, sentiment, emotion and silences from the voice stream is the core technology differentiator. In the contact center context this differentiator helps us be more accurate with our insights and this has a significant impact on helping our customers deliver better customer experiences. While most companies do call analysis once they convert a voice call into a text and then try to apply AI/NLP on top of it we are actually trying to do this directly on a voice stream. That is because the problem of trying to do analysis after this is converted into text is that you lose out the tonality quotient. A simple example like, “I am tired of waiting” verses “I am tired of waiting” with a different tone will look exactly the same once it is converted into text but will have very different interpretation when you hear the audio stream. In the first phrase I said I was tired of waiting, but I was not very frustrated about it. In the second phrase I actually showed a lot of frustration while saying it. This tonality is completely missing when you are trying to do the analysis once it is converted into text. We are actually doing all of this analysis on the audio stream itself. We have seen, with our customers, where our accuracies have gone up by double digits and we can really understand the context of what is happening in a conversation. That is a core differentiator that we are bringing to the table.

The second thing that you are seeing is that most automation efforts for companies fail because they tried to automate preemptively without understanding the process really well. One thing that we are doing is we always lead with our insights product where the goal is to understand the process, understand the call structure, call anatomy, and on the basis of our understanding we build something called, Conversation Maps. Once we build Conversation Maps, we will have a very clear understanding on opportunities available to optimize a business process. If we believe there is merit in optimizing the process by providing recommendations to the agent in real time we implement our augmentation product. If we believe that there is an opportunity to fully deflect the call by automating the process by building a VoiceBot we implement a conversational bot. This process has helped us win customers for the longer term and we have seen many more happy customers, because now automation is a well thought through approach and not a forcing function. The goal is process optimization and automation or augmentation is a means to achieve efficiency.

