AI USE CASE

Using AI technology to optimize call center outcomes

NetApp

Uncover trends, patterns, and insights in your call center data.

NetApp and SFL Scientific have combined their expertise to help enterprises address important AI use cases quickly. Call centers are an important part of enterprise operations in many industries. This paper describes the rapid implementation of a state-of-the-art deep learning model to detect sentiment in near-real-time during call center interactions, providing insight into the customer's state of mind, employee performance, and more. The general methodology described is applicable to a broad range of NLP and other AI challenges.

SFL SCIENTIFIC
DISCOVER MORE

The importance of call center agents has grown significantly as a result of the Covid-19 pandemic. Customer call centers are now a primary point of contact between many businesses and their customers. Today's agents are not just problem solvers and order takers. In many industries, they are significant contributors to sales.

Because of the number of calls these centers process, meaningful assessment of performance may be next to impossible without automation. Artificial intelligence (AI) is emerging as an innovative new tool in tracking the success of call center interactions.

Sentiment analysis uses natural language processing (NLP) to determine whether the sentiment expressed during a customer call is positive, negative, or neutral. Using this approach, your call center can take advantage of vast amounts of previously untapped data. For instance, you could use sentiment analysis to correlate customer sentiment with regard to specific brands or products, track overall customer satisfaction, or monitor the sentiment of individual customers.

The ability to capture sentiment in real time could enable agents (and their managers) to track the effectiveness of conversations in progress to better satisfy customer needs, with sentiment displayed in a simple graphical dashboard, as illustrated in Figure 1.

NetApp and SFL Scientific have developed an easy-to-implement AI pipeline that captures and displays the sentiment of call center conversations in real time. Our solution can be quickly deployed on premises, trained, and tailored to your specific requirements to provide a better customer experience and to gain greater insight from every call center interaction.

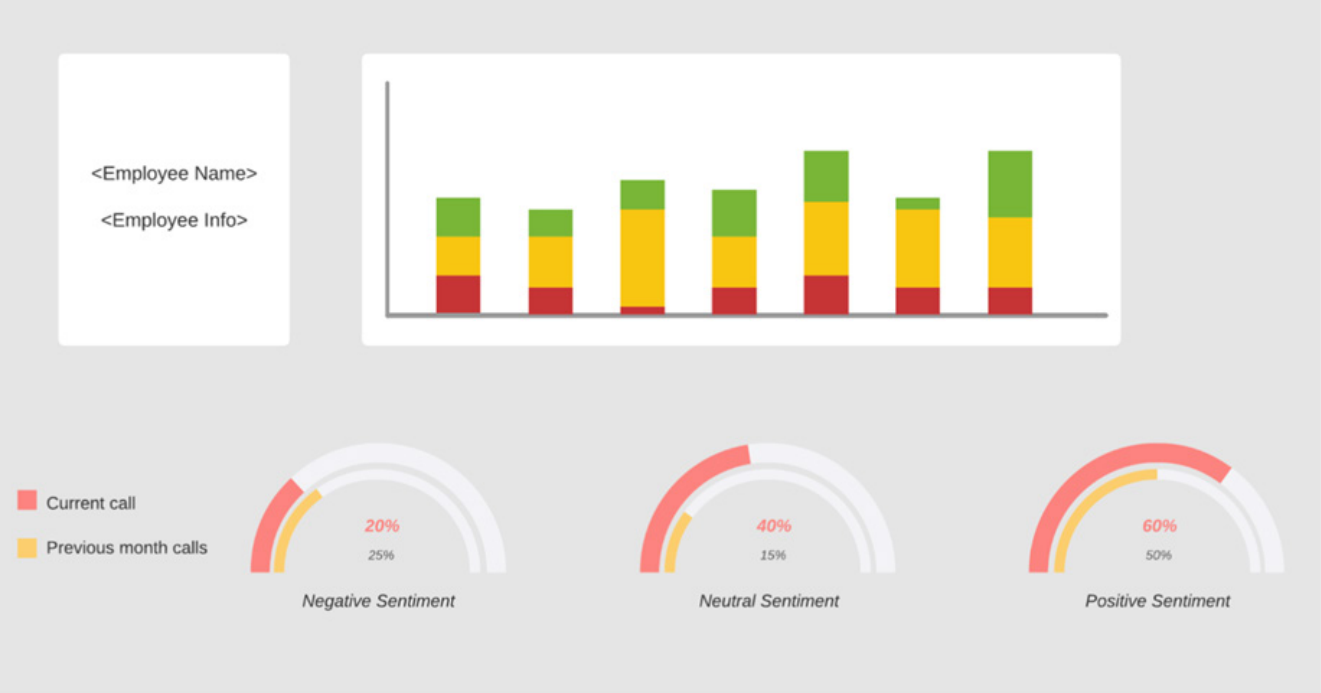

Figure 1) Sample interface dashboard showing current call sentiment and historical data.

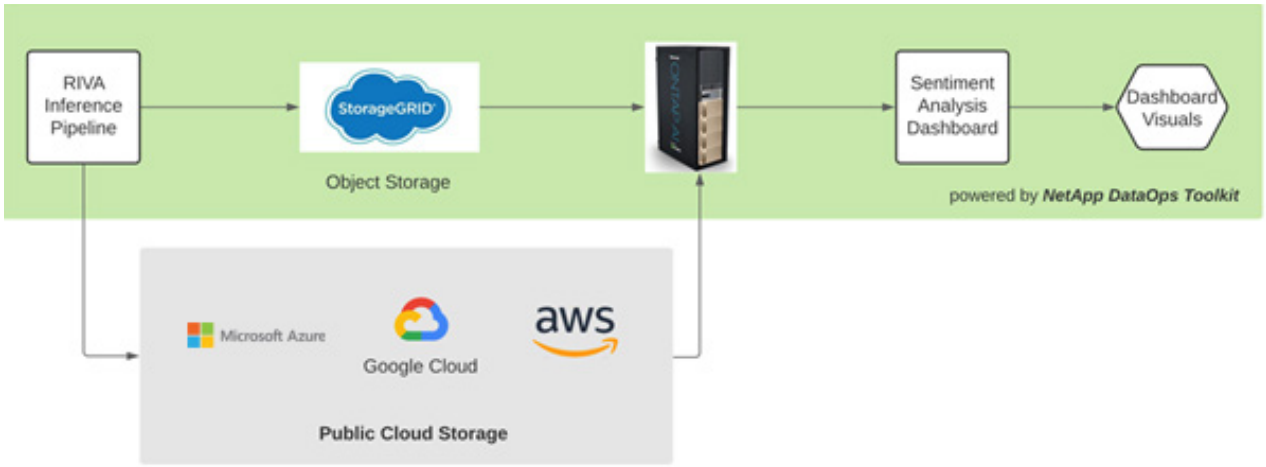

Figure 2) Object storage can be incorporated into the pipeline to accommodate distributed or cloud-based use cases and to provide a long-term archive of sentiment data for trend analysis.

An AI pipeline for call sentiment

Understanding customer sentiment is a great way to optimize customer interactions and improve customer satisfaction. Visual tools like the sentiment dashboard shown in Figure 1 can be used to guide the interactions between call center employees and customers, enabling your employees to steer conversations in a more positive direction. In addition, tracking the sentiment of both the customer and agent may make it possible to identify potential problem calls and intercede before it's too late.

By combining NetApp's hardware and data management tools with SFL Scientific's comprehensive data science and data engineering skills, we developed an end-to-end AI pipeline that is applicable to call center use cases. The pipeline uses a modular architecture, can be deployed quickly, and can be continuously refined to improve accuracy and satisfy the needs of your industry and business.

Every customer call is processed through a pipeline with pretrained speech recognition and sentiment modules to determine sentiment at the level of an individual sentence. Audio is ingested in real time and converted to text. Then each sentence is assigned a sentiment value of positive, negative, or neutral. Sentiment is tracked and scored, with results sent to the appropriate dashboards and stored for retraining and later analysis. The result is insight into call center operations, call center agents, and customer problems and concerns, as well as trends in product and business sentiment over time.

Pipeline components

The pipeline is built using NVIDIA software running on NetApp ONTAP® AI hardware to facilitate training and inferencing. NetApp ONTAP AI combines the capabilities of NVIDIA DGX and NetApp AFF to remove performance bottlenecks and accelerate AI workloads, while the NetApp DataOps Toolkit streamlines data storage and handling. Our pipeline ingests voice data in real time and passes it on for transcription and sentiment analysis.

The pipeline can optionally be architected to incorporate on-premises or cloud-based object storage to accommodate the needs of distributed and virtual call centers and for long-term archiving, as shown in Figure 2.

Audio transcription

Speech-to-text conversion is accomplished by using automated speech recognition (ASR). Each audio stream is preprocessed—including conversion to WAV format and separation into sentences—and passed to NVIDIA Riva, a GPU-accelerated software development kit (SDK) for NLP. The Riva server returns a text representation of each audio sentence and passes it on for sentiment analysis.

Sentiment analysis

NVIDIA Riva is also capable of classifying text in various ways. This project uses a pretrained module to detect sentiment, eliminating a significant amount of time-consuming training effort. The module classifies the sentiment of each sentence as neutral,

positive, or negative. Sentiment is sent as output to the appropriate dashboard, providing near-real-time insight. The end-to-end pipeline is illustrated in Figure 3.

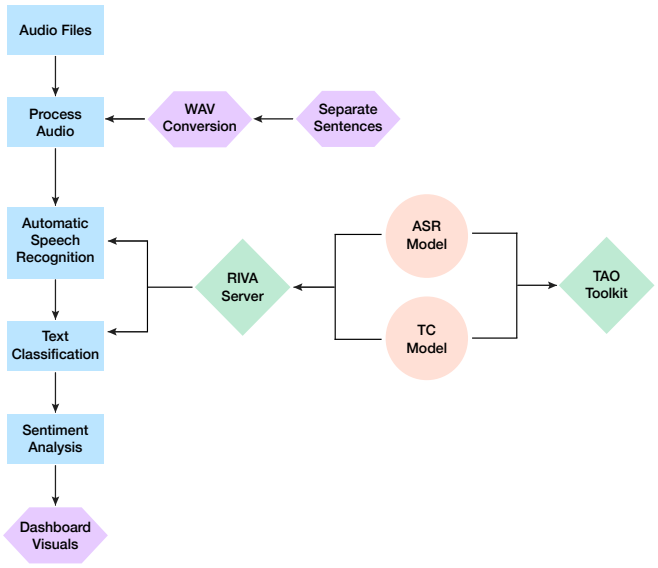

Figure 3) Pipeline to determine sentiment from call center audio.

Proof of concept

To assess the functionality and performance of this pipeline, we needed a real-world audio dataset that could be used to test both our audio transcription and sentiment modules. Because companies are understandably reluctant to share recorded call-center data, we settled on a dataset containing thousands of earnings call audio recordings and transcripts for S&P 500 companies. This dataset is a good stand-in for the thousands of calls that large call centers receive every day.

In testing, our pipeline is capable of processing hundreds of calls in parallel. Because this dataset includes transcriptions, it is ideal for training and fine-tuning the speech recognition and text classification models. Using this dataset, updated sentiment analysis appears on the output display less than 1 second after a sentence is ingested.

Transfer learning and fine tuning

The NVIDIA NGC catalog includes many pretrained models for use in the Riva SDK. These models provide a starting point for pipeline optimization that can be quickly fine tuned to meet specific call center needs by using transfer learning.

Transfer learning is the process of fine-tuning a previously trained neural network for a similar or new use case. NetApp and SFL Scientific started with pretrained deep learning models available with NVIDIA Riva. These models can be fine-tuned for your needs by training with a set of data that is specific to your industry and call center.

For instance, almost every company uses a variety of industry-specific jargon and company-specific product names. Transfer learning can allow the speech recognition model to accurately recognize and transcribe these terms, improving overall accuracy. This learning can be accomplished by using existing company audio and video files for which transcriptions exist, such as quarterly earnings calls.

Typically, our data scientists experiment with a range of model configurations and maximize model performance by using multiple model tuning experiments. They make adjustments during each iteration and select the model that performs the best.

Training and data management tools

The NetApp DataOps Toolkit and the NVIDIA Train, Adapt and Optimize (TAO) Toolkit are used in this solution to streamline the process of transfer learning and fine tuning. Using these tools to operate and maintain unified data pipelines reduces complexity, enables teams to efficiently manage data, and ensures traceability.

The NetApp AI Control Plane pairs machine learning operations (MLOps) with NetApp technology to simplify the management of AI data and facilitate experimentation. The NetApp DataOps Toolkit makes it easier to manage the large volumes of data that call centers produce.

Used separately or together, these tools significantly speed up AI projects, enabling you to quickly set up and clone the volumes needed for training, evaluate results, and iterate quickly, ensuring reproducibility, reliability, and compliance.

Retraining

You can periodically retrain your call center model by using the tools described in the previous section and text from a set of call transcripts. Transcripts are reviewed by managers to annotate and verify sentiment scoring. This annotated data is then fed through the training pipeline to keep your AI models up to date on your latest products, services, and terminology.

Sentiment analysis use cases

Our sentiment analysis pipeline can be extended to address a broad array of call center needs, such as the following.

- **More detailed sentiment metrics.** The simple dashboard illustrated in Figure 1 could be extended in a number of ways.
  - **Separate customer and call center agent sentiment.** Determine who's driving the overall sentiment for the call. The customer? The agent? Both?
  - **Analyze turn taking.** Is the customer or the agent dominating the call, and how does that relate to sentiment?
  - **Cutting in.** Is someone on the call repeatedly talking over or cutting off the other speaker?
- **Automatic escalation.** Intervene when a call turns negative.
  - If customer sentiment is below threshold for X period, trigger Angry Customer escalation.
  - If agent sentiment is below threshold for Y period, trigger Angry Agent escalation.
  - If customer and employee sentiment scores are below threshold for Z period, trigger Dispute escalation.
- **Trend monitoring.** Detect important trends that could affect call center success.
  - Is the sentiment of repeat callers improving or declining over time?
  - Is the overall sentiment of customers contacting the call center improving, declining, or staying the same over time?
  - Are call center personnel expressing more positive sentiment or less positive sentiment in interactions with customers over time (collectively and individually)?
  - Are there particular customers and call center personnel who should never be connected to each other?
- **Chatbot monitoring.** If a customer is interacting with a chat bot (text or voice) and becoming upset, transfer to a call center or live chat agent.

More AI opportunities

The methodology we used to create sentiment analysis tools for call center use can also be applied to a wide variety of enterprise AI use cases. For instance, most companies have a wealth of data available from social media feeds. What are people saying about your company, brand, product, or service? What are they saying about your competitors? Is there a trend in sentiment? What was the impact of that new product rollout or the latest financial report? Simple sentiment scoring can be augmented by using tools such as weighted approaches to scoring, voice stress analysis, tone analysis, speech rate, and other key sentiment indicators.

Looking beyond sentiment analysis, the combination of transfer learning, experimentation, iterative fine tuning, intelligent data management, and production deployment with regular retraining can be applied to a wide range of NLP and other AI use cases in your business. If you have the data, we can help you gain the insight you need. NetApp and SFL Scientific can help you get your AI project to production more quickly with fewer missteps.

About our partnership
The partnership between NetApp and SFL Scientific brings together SFL's proven data science and data engineering expertise and NetApp's industry-leading AI hardware and software.

Get started
To find out whether NetApp and SFL Scientific can help you accelerate your path to AI, visit netapp.com/ai and fill out the form to request a contact.

NetApp

Overcoming the challenges of building an AI architecture for machine learning (ML), deep learning (DL), and other AI workflows requires careful infrastructure planning. NetApp helps customers build AI pipelines that smooth the flow of data and deliver optimal results.

NetApp offers a complete portfolio of AI hardware and software solutions and services to get your data flowing from edge to core to cloud. netapp.com/ai

SFL SCIENTIFIC
DISCOVER MORE

SFL Scientific is a data science consulting and professional services company operating on the highest end of technology development.
We specialize in data engineering, machine learning, Artificial Intelligence, and big data, continuously consolidating a pool of the very best data science and analytics talent for our clients.
www.sflscientific.com

+1 877 263 8277

About NetApp
In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere. To learn more, visit www.netapp.com

NetApp