



AI Matters: An Inside Look at How CallMiner Powers Business Performance Improvements

Learn why CallMiner was named
the Leader in The Forrester Wave[®] :
Conversation Intelligence for
Customer Service, Q3 2023



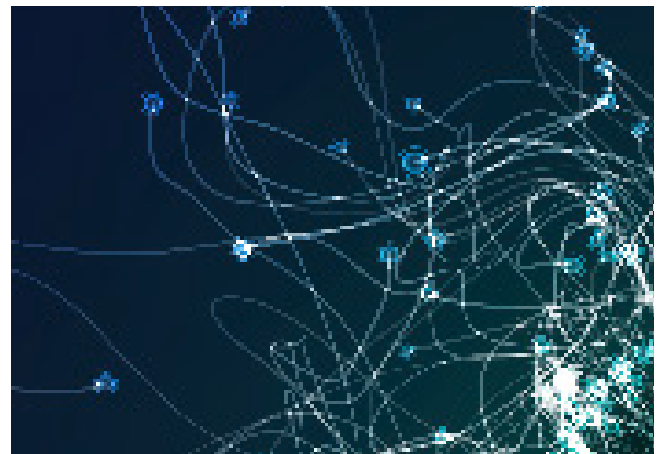
AI Matters: An Inside Look at How CallMiner Powers Business Performance Improvements

Customers interact with brands, products, and services on more channels than ever before. As the number of channels has increased, so have expectations for the customer experience (CX).

According to a [report from Salesforce](#), 66% of customers expect companies to meet their unique needs. When contacting a company, 83% of customers anticipate that they will connect with someone immediately. And despite wanting consistency, more than half of customers describe their service interactions as fragmented.

With these sky-high expectations comes a greater responsibility for organisations to not just respond quickly, but truly understand the context behind why customers interact with brands. Using that context, brands can reduce friction, make critical business changes and improve the overall CX. But, given the wealth of conversations happening at every touchpoint, how can organisations separate the signal from the noise?

Many teams manually review customer interactions or rely on surveys to collect solicited feedback. Both of these approaches are flawed. Manual review can be costly, and focuses only on a small subset of interactions. Likewise, the average survey response rate is only between **5 and 15%**. People who tend to answer survey questions usually are at opposite ends of extremes: extremely satisfied or extremely disappointed. This leaves a majority of customer interactions in the dark, with high margin of error and potential bias in the resulting analysis.



Instead, conversation intelligence from CallMiner can help organisations analyse 100% of customer interactions across every channel. This level of analysis helps brands extract key trends about their customers and personalise the experience to them, regardless of the channel they're using. Conversation intelligence can also analyse survey data to develop a complete picture of the voice of the customer (VoC) to influence CX, and more.

This technology uses artificial intelligence (AI) to extract meaningful insights from unstructured data to drive action, such as improving how teams interact with one another, as well as customers, prospects and partners. CallMiner's technology empowers organisations to understand conversations both in real time and post-interaction, giving teams the flexibility to achieve a variety of business performance improvement goals.

CallMiner uses both customer interaction data and metadata around those interactions, including what channel, who reached out, date and time, interaction ID, if they've been a customer or patient before, and more. Both structured and unstructured data can be mined for insights. Many teams leverage this data by extracting it into existing data lakes to paint a full picture of each experience – from both the business and customer perspective. As a result, organisations can get a better understanding of what's going on across their customer base, because they gain a complete view of every interaction in an omnichannel environment.

According to a [press release from Gartner](#), "By 2025, 60% of organisations with voice of the customer (VoC) programs will supplement traditional surveys by analysing voice and text interactions with customers... Business and CX leaders will continue to seek ways to diminish reliance on traditional surveys." Gartner continues on to say, "Most customer service organisations today are deploying customer surveys, but customer surveys are perceived to bring less value than other methods of VoC collection."

These other methods include capturing insight from emails, messages, posts and calls. To do this properly, these organisations need AI and machine learning (ML) technology in their corner.

Let's learn more about how this technology works within conversation intelligence.


Understanding how AI and ML works in language

Many companies claim to use AI and ML behind the scenes. As a result, it can be hard to distinguish exactly how this technology is used and why. Let's look at [McKinsey's definition](#) of each of these terms to provide more clarity on what we mean by AI and ML

- "AI is typically defined as the ability of a machine to perform cognitive functions we associate with human minds, such as perceiving, reasoning, learning, and problem solving. Examples of technologies that enable AI to solve business problems are robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning."

- "ML algorithms detect patterns and learn how to make predictions and recommendations by processing data and experiences, rather than by receiving explicit programming instruction. The algorithms also adapt in response to new data and experiences to improve efficacy over time. Most recent advances in AI have been achieved by applying ML to very large data sets.."

Specifically within CX, the AI used within conversation intelligence technology distills relevant customer insights from massive datasets. What's more, using this technology can mitigate some of the bias that humans have. Instead of projecting back answers organisations may want to hear, AI uncovers what's actually happening with customers, prospects, and more. We'll cover more on that later, but for now, let's look at how conversation intelligence evolved to what it is today.



Evolution of AI in CX

As a pioneer in conversation intelligence, CallMiner has leveraged cutting-edge research in AI and ML in its products over the last 20 years. The latest state-of-the-art (SOTA) advancements in AI are always changing. CallMiner has demonstrated the ability to not only innovate and keep pace with the ever-evolving AI landscape, but also to drive the future of conversation intelligence with its own research team.

A key differentiator for CallMiner is its laser-focus on outcomes, rather than locking into particular technology preference. The company also has the resources and architecture to change with the times, while newer entrants in the space may not have the same vision and flexibility to keep pace with SOTA AI as it shifts. Since AI builds upon previous advancements in the field, CallMiner can lean on its experience to achieve research milestones faster. Based on years of expertise, the company focuses on developing the best possible solutions to meet customers' desired outcomes.

Why start with speech?

Unlike many AI products that started with text analytics, CallMiner built its product on the foundations of speech analytics. CallMiner's analytics use semantic building blocks which account for the unique characteristics of unscripted speech. This approach enables CallMiner's analytics to be more nuanced and capture context better than text analytics.

Text analytics often depends on the rules and conventions of formal language, exhibited in the ways emails, documents, and even chats are structured. To contrast, transcribed spoken conversations are less rigid than text. These require a deeper level of analytics to understand context and meaning, while dealing with multiple speakers and recognition errors. CallMiner's experience and advanced algorithmic approach make the company successful within the complex digital world of omnichannel communications.

Let's take a look at how AI has evolved in this field over the years.

ERA 1 - Machine transcription with human analysis

Speech recognition AI is used to transcribe conversations to text. It has a long history dating back to the 1950s, when Bell Labs built the **first documented speech recogniser** called "Audrey," which recognised strings of digits with pauses in between. The field really took off in the late 1990s to early 2000s, fueled by major advances in mathematics and computing power.

One drawback of machine transcription is that humans still need to interpret each transcript manually to gather insights. This involves a lot of effort for humans, and isn't practical for high volumes of interactions. However, there's still value in this technology when there's a need to dive deeply into individual interactions. Many workflows still require these in-depth explorations, particularly when it comes to mentoring and training employees on the front-lines of customer interactions. Completely removing human-in-the-loop analysis is impractical when deriving conversation intelligence.

ERA 2 - Word spotting

Word spotting, otherwise known as keyword spotting, uses AI to look for the presence or absence of certain words. Some of the first research in this field happened in the **late 1980s** and early 1990s. Typically, this technology is used for sentiment analysis. For example, a word spotting algorithm could see if the words "awesome" or "terrible" are present within a call transcript, aiding with the practice of sentiment analysis.

Although helpful, word spotting still needed to evolve significantly to increase its utility for conversation intelligence. First, the technology relies on perfect transcript accuracy. Even the best transcribers aren't 100% accurate. Second, many words have multiple meanings. While the word "rejected" in the context of a credit application does not inherently imply a negative sentiment, "rejected" in other contexts could be considered negative.

Language is far more complicated than the individual words being said. Instead, meaning is derived from how words interact. A good example is when a positive word like “good” is used sarcastically in a negative context like “good grief.” Even so, word spotting is still useful in many contexts where industry or company-specific words must be detected. For example, competitor mentions might trigger certain automated events within a robotic process automation (RPA) system, taking much of the lift off of individual agents.

ERA 3 - Rules-based approaches

In this era, pioneers in natural language processing (NLP) began capturing the complexity of human language in conversation. Rules-based approaches became less about detecting individual words and more about understanding how the words interact. Early **statistical models in the 1980s and 1990s** were some of the first automated rule-based systems for language that didn't rely on labor intensive, hand-written rules. Rules were an important evolution in AI, because they were able to capture the context of when something was said (metadata), what specifically was said (semantics) and how something was said (acoustics). Rules also allowed for the filtering of scenarios. Using an example from above, a rules-based system would filter out instances of the terminology “rejected” from a negative customer sentiment analysis when it's used in association with “credit application.” Since credit applications are either accepted or not accepted, the “rejected” terminology wouldn't necessarily imply a customer complaint, as it might in other scenarios. Even so, rules-based technology needed to evolve further, since it's difficult to capture every way something could be said within a finite set of rules. Rules also capture bias in how people think something should be said, rather than letting the data unveil what is actually happening. However, the value still remains in situations where there are only a few ways in which something can be said (or if things can only be said in a specific way). Good examples include legal disclosures, compliance mandates, and certain agent scripts.

ERA 4 - Machine learning

Machine learning was another AI breakthrough for conversation intelligence. Instead of describing data by creating rules (as was done in rules-based approaches described above), ML uses advanced algorithms to make decisions based on big data.

For example, instead of predetermining the topics that should be in a conversation and writing rules to capture them, CallMiner uses unsupervised learning to analyse and cluster any subset of interactions into explorable and expandable topics. Unsupervised learning uses algorithms to discover patterns in unlabeled datasets without human oversight. As another example, instead of having to define all the ways a dental patient can say “braces,” CallMiner uses ML and NLP to provide organisations with related words and phrases that are being used similarly in real conversations (such as “aligners” or “those plastic straighteners”). In the more recent months of this era, CallMiner has stayed cutting edge by focusing on explainable AI. There is a misconception that ML is a “black box,” and that it's impossible to know why a model is making the decisions it does. Through explainable AI, every decision a model makes is interpretable and can be easily scrutinized to find out the reasons for its selections. Merely predicting events is no longer enough—machine learning can be used to uncover actionable features of an interaction and make measurable, trackable changes.

ERA 5 - Generative AI

Generative AI is a key technological innovation that is helping shape the future of conversation intelligence. It is designed to create content such as images, text, audio, and more. Generative AI, a category of AI systems that includes Large Language Models (LLM), is capable of understanding context, generating coherent paragraphs, answering questions, providing explanations, and even completing sentences.

CallMiner leverages a number of generative models and LLMs in the platform. Generative models are used in conjunction with other AI models (ML clustering, NER, rules-based machine learning, etc.) and features such as Vector DB/Search to enrich analytics and insights, foster ease of use across the platform, and drive business results for a range of use cases including for CX, operational, and coaching use cases. Appropriate guardrails and security features are used to limit risk of hallucination (i.e. making up or giving the wrong answer) and data security concerns. CallMiner's generative AI capabilities generate net new data from current data. An example of this is CallMiner's contact summarisation, which understands customer conversations and writes a human-like paragraph explaining what happened during the conversation. This capability offers summary consistency and lessens after call work for agents – reducing costs.

Using conversation intelligence for feature extraction

With CallMiner, it's much simpler to prepare data for use in a prediction model. Conversation intelligence platforms can help extract the most important features from a conversation to use as input for a prediction model.

For example, customers are building prediction models through CallMiner by creating features for "sale or no sale" models, predicting net promoter scores (NPS), and developing churn models to determine the likelihood of a customer leaving.

How CallMiner uses AI to help organisations drive business value from customer conversations

Every conversation intelligence solution on the market uses speech recognition to provide a transcript based on what customers say. But, once an organisation gets the transcript, what do they do with that information? That's what separates CallMiner from the rest of the competition.

Most tools focus on transcription accuracy, relying on rudimentary word spotting to derive insights from a conversation. To contrast, CallMiner uses ML on a sophisticated workflow to determine conversation accuracy. This approach allows organisations to identify interactions or specific events that happen in a conversation. For example, word spotting may misidentify a certain word a customer says, such as "eyeliner" instead of "aligner." CallMiner's technology identifies context of the conversation and can identify the appropriate substitution, understanding the company is dental provider.

With CallMiner, each conversation starts as unstructured data, regardless of its channel of origin. From there, the data is moved through a thoughtfully organised workflow to structure it and provide robust insights that fit an organisation's business needs. CallMiner uses AI to achieve this process and drive organisation-wide value. Here are some of the key benefits of this approach.

- 1. Accelerate adoption and ease of insights:** CallMiner uses unsupervised learning algorithms and AI tooling to help analysts easily build and customise their own models and provide out-of-the box analytics for quick, successful, and comprehensive insights. While our graphical user interface (GUI) is optimised for intuitive workflows, our robust API allows for easy integration of all outputs into existing data lakes, lakehouses, and pipelines for analytics.
- 2. Uncover trends:** CallMiner uses clustering algorithms and ML paired with outcomes to help uncover leading indicators in an organisation's data, surfacing trends teams might miss on their own. Using a proprietary AI algorithm and correlation analysis, these indicators can then be tracked over time (or even retroactively) in order to monitor how initiatives can affect the direction of trends.
- 3. Identify a suite of emotion:** Like other competing solutions, CallMiner reveals sentiment in customer conversations. But CallMiner's AI takes that process a step further with the ability to identify emotions within a conversation, such as surfacing a moment when a customer is stressed. These emotion markers were created with supervised machine learning techniques and rigorous fine-tuning of unsupervised outputs. CallMiner does this by using ML, processing unstructured text, creating comprehensive annotation schemas, and identifying acoustic measures.
- 4. Categorise interactions:** CallMiner uses AI to identify and categorise conversations (for example, to identify at-risk situations for compliance, customer satisfaction, and even quality of care in a healthcare setting). CallMiner also analyses and categorises interactions to identify and provide the next-best-action at the point of service. Interactions are categorised using many different types of AI and ML, including supervised machine learning and transformer-based models, acoustic-derived metadata, outputs from NLP pipelines, and customisable decision tree-based algorithms.
- 5. Gain flexible ROI (or business value):** Some competing vendors focus their solution to deliver specific (and limited) kinds of ROI. CallMiner's AI is flexibly built to uncover a wide range of business value. Some of CallMiner's customers have realised 10+ ROI types from the actionable insights in their analysed data. CallMiner achieves this by maximising big data to fine tune models and pipelines to new use cases as they are identified. Think of CallMiner's AI as reusable building blocks that can be leveraged across use cases, minimising the lift for each new initiative.



Here's how specific CallMiner AI features are used to achieve the benefits described above.

	State-of-the-art Speech recognition systems	NLP Pipelines	Unsupervised Learning and Clustering	Supervised Machine Learning and Transformer Models	Proprietary AI algorithms	Acoustic feature extraction and AI	LLM and generative models
Accelerate Adoption and Ease of Insights	X		X	X	X	X	X
Uncover Trends	X	X	X		X		X
Full Suite of Emotions	X	X	X		X	X	X
Categorise interactions	X	X		X	X	X	
Flexibility in ROI	X	X	X	X	X	X	X

Not all AI and ML is made equal

With vendors making so many AI claims, it can be difficult to tell which technology will support the best possible business outcomes. Often the truth lies in the size of the dataset upon which an AI system is trained and tested. For building and testing AI models, size matters.

In other words, AI becomes more effective with larger datasets. That's why larger AI organisations like Microsoft, IBM, OpenAI and Google have an advantage when it comes to building large, generic services like speech recognition, facial analysis, information retrieval, language transformer models, and predictive modeling.

CallMiner, too, has been tried and tested with big data. In conversation intelligence specifically, technology providers need to support both depth and breadth. That means supporting sheer volume around unique phrases or topics, as well as multiple verticals, languages, and sources (like calls, chats, email) to adequately test new models and techniques.

Why? More training data reduces the risk of "overfitting" a model. That's when a model adjusts its outcomes based on a one-off bias in the data. One example of a bias is a conclusion that customers named Monique always buy more on Tuesdays. Each of these mistakes can potentially result in lost revenue for a customer.

When you're using big data, even a 1% increase in model performance could mean a better experience for thousands of customers. In some scenarios, accuracy matters more than others. For example, in medical imagery, if a model is designed to detect cancerous cells, it's incredibly important to fine-tune models for low error rates. In other situations like advertising, the impact of an incorrect prediction is much lower.

To contrast, the breadth of data removes one-off AI biases and improves accuracy. CallMiner models have been tested on multiple use cases and verticals across a sample of 100+ research partners, as well as large datasets with upwards of 250 billion words.



What's more, a high availability of data reduces the average time required to build a model by more than 50-75%. That means there's no delay or lag in starting a prediction model, since the data is already collected and waiting to be engineered by the AI team.

For example, in the contact centre world, many companies want to build a churn prediction model to understand whether someone is likely to cancel. To make this type of model accurate, you would need a million customers' worth of data – and all of them would have to churn. How long does it take to lose a million customers? Most companies wouldn't want to find out the answer to that question.

Rather than starting from scratch, CallMiner has been mining and managing massive volumes of this type of data for years. Both the longevity and the bandwidth (or volume) of data matter. That means companies don't have to wait to start their predictions or rely on smaller (and as a result potentially limited) datasets.



The future of conversation intelligence

As pioneers in the field of conversation intelligence, CallMiner has been at the forefront of driving value and pushing innovation during each of the AI eras described above – from machine transcription to today's SOTA ML. CallMiner's diverse team of data scientists, physicists, neuroscientists, mathematicians, linguists and engineers capitalise on the newest, most relevant breakthroughs in AI. They quickly apply these techniques to customer conversation data and integrate them into the CallMiner platform. And, with dedicated personnel focused on Responsible AI, CallMiner ensures its products are built to the highest ethical standards, delivering technology that's both unbiased and fair.

What is Responsible AI?

Whether you want to call it Responsible AI, ethics in AI, or detecting bias in AI – there's a major need to ensure that machine learning doesn't add to or mirror the injustices of our world. In fact, ML and AI can even help us counteract and fight these injustices.

At CallMiner, Responsible AI is not simply a checklist that one researcher completes one time throughout the research process. It cannot even be delegated to a single team. Responsible AI is a company-wide effort that succeeds through the diversity and dedication of those working to achieve it. It is a process that must be evaluated by a diverse set of minds and backgrounds in order to account for a larger universe of perspectives and experiences. This process must also be repeated throughout the life of an AI-driven tool, from the time the idea is conceived, through its development and deployment, and consistently as users begin to explore and understand it.

The CallMiner framework for Responsible AI is based around five foundational ideas.

Find injustice – The goal of Responsible AI is to recognise and address the injustices of our world, which manifest themselves within our data or are reflected by our models.

Embrace the unknown – No one person will have the answer to any given problem, and there are some questions that we come across that have yet to be answered anywhere in the field.

Strive for transparency – Few algorithms are biased by nature. We must monitor the training and application of these algorithms to guarantee fair and ethical AI.

Adapt as knowledge grows – Responsible AI is an imperfect, incomplete, and relatively new aspect of the field, and any approach must be able to adapt to the new challenges and injustices that appear every day.

Work without ego – The first step to Responsible AI is well-documented self-evaluation and meaningful and sometimes uncomfortable conversations that help us work towards a brighter future.

To read more about CallMiner's approach Responsible AI, [check out this blog series.](#)

So what's next for CallMiner as we continue to research and apply the latest in SOTA AI?

Here are a few of the trends that the team is currently investigating.

Improving and generalising AI and ML for low resource languages.

AI typically performs well on high resource languages like English, Spanish and Chinese. The future of AI will focus on providing the same AI tools for languages with fewer resources. Currently, CallMiner supports an expanding catalogue of lower resource languages like Catalan, as well as regional dialects of high resource languages such as English, Spanish, French and Chinese.

Helping organisations shift toward prescriptive outcomes

AI has been trending toward outcome-based ML – or helping organisations discover how to make changes rather than simply delivering predictions. This trend will continue as ML expands into the world of “what-ifs.” For example, organisations may ask, “How will my outcomes change if I update this procedure?” or “Will changing my screening process improve NPS?”. CallMiner’s goal is to have AI understand the correlations and causations that happen while introducing both real and hypothetical changes.

Leveraging large language models like GPT-3, making big data even more crucial.

Ever since the GPT series of transformer models was introduced by OpenAI, the industry’s fascination with large language models has continued to skyrocket as quickly as the datasets it takes to train them. This trend will continue in this current era, increasing the need for big data models like CallMiner’s that generalise across use cases and ROIs.

Developing models that improve the customer experience (CX)

While the pandemic has highlighted the need to build a better CX across every channel, the shift toward forging stronger customer connections isn’t going away. In the future, there will be a continued desire for ML models that improve CX. These models demand explainable AI, so teams can understand the “why” behind the outcomes they get from their algorithms. To that end, CallMiner will be adding more custom predictions and outcome models to the product, aimed at delivering value and ROI even faster. In addition, CallMiner will use AI and ML to uncover topics and insights within unstructured data. That process will improve the product’s ability to find relevant and important topics from conversations, with little to no human effort.

Tapping into AI to drive transformational growth

With nearly 20 years of experience in the contact centre and beyond, CallMiner has pioneered the use of AI to drive measurable CX improvements. Using the team’s extensive knowledge of AI and ML, CallMiner has helped hundreds of organisations better understand their customers and drive transformational growth. For CallMiner, it’s all about using the right AI technology to drive specific business outcomes, goals and operational challenges – rather than simply using AI for AI’s sake.

Through cutting edge research and applied AI, CallMiner continues to invest in the best possible conversation intelligence technology to achieve customers’ desired outcomes. When evaluating a conversation intelligence solution, organisations should ask what it can do to meet specific business needs or goals. That’s where CallMiner’s unparalleled track record across a wide variety of industries makes all the difference.

“CallMiner”, “Illuminate”, “Eureka”, “Eureka!”, “Eureka Analyze”, “Eureka Coach”, “Eureka Alert”, “Eureka Redact”, “Eureka API”, “MyEureka”, “MyEureka”, “EurekaLive”, the “CallMiner Eureka” logo, the “CallMiner MyEureka” logo, the “CallMiner EurekaLive” logo, “Listen to your customers. Improve your business.”, “Feedback Is A Gift”, “Listen”, “Engagement Optimization”, the “Engagement Optimization” logo, and “EO” are trademarks or registered trademarks of CallMiner, Inc. in the United States and foreign jurisdictions. Other product names mentioned herein may be the trademarks of their respective owners.

About CallMiner

CallMiner is the global leader in conversation intelligence to drive business performance improvement. Powered by artificial intelligence and machine learning, CallMiner delivers the industry's most comprehensive platform to analyse omnichannel customer interactions at scale, allowing organisations to interpret sentiment and identify patterns to reveal deep understanding from every conversation. By connecting the dots between insights and action, CallMiner enables companies to identify areas of opportunity to drive business improvement, growth and transformational change more effectively than ever before. CallMiner is trusted by the world's leading organisations across retail, financial services, healthcare and insurance, travel and hospitality, and more.



©2022 by CallMiner. All rights reserved.

Corporate Headquarters

200 West Street
Waltham, MA 02451

Sales

sales@callminer.com

Social

facebook.com/callminer_inc

twitter.com/callminer

linkedin.com/company/callminer