

Mining Recorded Calls: Understanding the Value of the CallMiner approach

Introduction

CallMiner is a suite of client and server applications that enable call center managers to gain valuable insight into what was said in recorded calls. The applications allow access to information buried deep within recorded calls that, until to now, was uneconomical to extract. The CallMiner suite of applications has 3 principal components:

1. CallMiner CoreEngine is the server based tool that converts the recordings into standard SQL for later comparison and analysis. This non-proprietary database can be used by CallMiner and Customer applications directly.
2. TrendMiner is a web-based client application that allows users to quickly and easily perform ad-hoc queries and trend analysis on mined data.
3. Tireless Supervisor is a client and server based tool that uses patterns to classify and score calls based on user-defined criteria. Tireless Supervisor distinguishes itself in that very complex ideas and patterns are used to correctly match the seemingly infinite possible ways a customer or agent can express an idea. This is something that a simple word spotting/searching tool just cannot do reliably.

Business Value

- Ability to optimize call scoring by providing “look ahead” scoring. Scoring every call recorded with Tireless Supervisor, then showing those scores to supervisors allows them to pick the most important calls to manually score and evaluate.
- Ability to classify calls automatically and into far more categories than manual processes are capable of. Tireless Supervisor uses patterns to match a call type and put that call into one or more matching classifications. This removes the reliance on manual processes accomplishment and allows for much greater statistical validity (all calls are classified, not just a random and much smaller sample) This allows for easy analysis and trending of classes over time.
- Ability to judge a call based on its fit to a behavior. This is also referred to as the “likelihood” score. What is the likelihood that a particular call fits into a particular category? i.e. what is the likelihood that a call is a money laundering call?
- Ability to rapidly reduce a large call volume down to a manageable set based on some search criteria. Typically by discovering and tweaking search parameters in an ad hoc fashion on the fly. Using our TrendMiner tool, several thousand calls can be sifted almost immediately

to a more manageable set, and then listened to to quickly pick those there are most important.

Technical Value

- Works on 34 different audio formats including highly compressed wav, vox and others.
- Works with 22 different languages and dialects, languages can be added quickly
- Open database design allows for simple integration with other data driven and analysis tools
- Languages include: US English, UK English, German, Austrian German, Spanish, Spanish - Central and North American, Spanish - South American, Italian, French, Dutch, Flemish, Portuguese, Brazilian Portuguese, Catalan, Finnish, Swedish, Norwegian, Danish and Hungarian.
- Open extensibility architecture, allows new products to be added by 3rd parties without changing underlying CallMiner components

Differentiators between CallMiner and phonetic search engines.

Scaling and Growth.

CallMiner is capable of mining and dealing with very large volumes of calls, especially over a large block of time due to our small database growth.

The CallMiner database grows very slowly in comparison to most phonetic search engines. CallMiner grows at about 0.3MB/hour of audio mined while phonetic search engines proprietary database grows at almost 16MB/hour of audio.

This becomes an issue with 100% recorders that record and archive a large volume of data. For instance, at one customer, who estimates that they will mine about 1600 hours of audio a day:

Callminer	= 0.3mb/(hr of audio) * (1600 hrs/day) = 480 MB /
day (one backup CD per day @ 50 cents)	
Phonetic Search	= 16MB/(hr of audio) * (1600 hrs/day) = 175,000 MB /
day (three backup DLT tapes @ \$30/apiece)	

Phonetic search's disk space requirements are greater than 50 times that of Callminer's, resulting in much higher hidden costs that affect the overall Total Cost of Ownership of the solution.

When a customer wants to maintain a full year of searchable data(using the above example), CallMiner is only 175GB of data (a desktop hard drive), but a phonetic search engine is a whopping unmanageable 9.35 Terabytes of data (a server room of hard drives)

Conversion Speed.

CallMiner's speed needs to be examined in two areas, one is our ability to create the database, and the other is our ability to search and retrieve from the database. CallMiner continues to improve our speeds of creation; currently our production code produces 40 hours of audio per server per day. This number is linear to processor, which means on a Multi-processor box, we expect to see as much as 160 hours of audio per server per day (4 Way processor).

Search Speed.

CallMiner uses established SQL database technology to perform our searches, and because database technology is very mature, our search speeds are very very good. Most Phonetic search tools are very proud of the speed of their search engine, some proclaiming 30 hours of audio can be searched per second for a single term. CallMiner, using databases of 2.7M calls; over 95,000 hours of audio have shown a reliable search speed of 4800 hours of audio per second per term. Over 160 times faster than phonetic search tools. Also, consider the 30 hours per second per term, and compare that to the 584,000 hours a year that a Customer produces. A single search for phonetic search tools would take over 5 ½ hours, but only 2 minutes for CallMiner.

Accuracy.

Large Vocabulary Speech Recognizers take a beating on accuracy by Phonetic and Dialog speech search tools. The reason is simple. Dictation (LVSR) engines are not very accurate on what they normally output. In a lab, the ability to search phonetically versus the ability to find words in the output of a LVSR produces dramatic and quotable results in favor of phonetic engines. This concept breaks down when you start applying "Business Accuracy", a term CallMiner uses to describe how accurate the results are in relation to the customer's business needs. In a recent large scale trial with Continental Airlines, CallMiner accurately (98%+) identifies and classified over 3300 calls, quickly solving their real-world business problem. Real world accuracy where it counts - not in the lab, but at the customer. CallMiner achieves these kinds of results because we don't just use a single technology to solve this problem - we apply multiple speech recognition algorithms, natural speech linguistics, statistical methods and advanced pattern matching to solve the dubious problem of deciphering a caller's needs and wants. This, coupled with an intense understanding of the call center industry is what makes classification possible and important.

Conclusion

CallMiner combines several technologies to solve the business problem of converting recorded conversations into real business value. Using the CoreEngine to mine the calls, TrendMiner to provide Ad hoc searches and Tireless Supervisor to provide an automated method of call classification and scoring, CallMiner quickly generates and then sifts through information that is otherwise sitting useless on a shelf.