Nexidia Technology and Architecture for Interaction Analytics

nexidia

Technology: Neural Phonetic Speech Analytics™

Nexidia has always invested in original speech technology research and invented the process of rapidly searching audio known as phonetic indexing. Today, Nexidia Interactions Analytics employs Neural Phonetic Speech Analytics,™ which combines the strengths of Automatic Speech Recognition (ASR) and phonetic indexing to yield highly accurate results. The outputs of Neural Phonetic Speech Analytics include word level transcripts and sentiment scores for early discovery and query building within Interaction Analytics, as well as phonetic indices that supports searching across 100% of interactions for deep dive root cause analysis.

Word level transcripts and sentiment data are delivered via a breakthrough technique that applies artificial neural networks to the investigation of customer interactions. Neural networks are computational designs that apply machine learning and pattern recognition to complex problems. Influenced by the structure of the human brain, artificial neural networks are well suited to computational problems with large numbers of variables, especially when those problems are of a probabilistic nature. Neural Phonetic Speech Analytics employs neural networks in new and innovative ways for the analysis of audio, expanding the depth, breadth, and accuracy of the results without driving computational costs beyond what is practical for any business. Neural Phonetic Speech Analytics supports customer-specific libraries that ultimately describe the extent of the transcription vocabulary. The technology utilizes two types of models: the acoustic "fingerprints" of words and phrases, along with language models. These are leveraged for early discovery or "narrative modeling," and provide the data to build taxonomies of structured audio searches or "topical modeling."

PHRASE RECOGNITION

The narrative models begin with an industry-specific lexicon based on real world conversational data and acoustic properties delivered by Nexidia. Using interactions and other text data (chats, emails, after-call surveys, social media, product and training documentation, CRM case notes, etc.) referencing the customer's environment, products, and services, Nexidia enhances the default vertical model with language and sentiment specific to the customer. This process utilizes natural language processing techniques to identify relevant terminology, while ignoring the words and phrases that are not useful to analysis, such as "thank you for calling" and "how can I help you?" The technology automatically adds new terms to the language model over time, adapting to new terminology as the business evolves. Because this process is automatic, there is no need for professional services to update the model periodically; however, customers can at any time add to these lists, "forcing" particular phrases to be included. This is especially useful when the customer knows of upcoming terminology, such as new products, that may as yet not be present in any interactions or data source.

Neural Phonetic Speech Analytics also provides input and suggestions to the building of structured queries, or topical modeling. The system guides the user through the process of using the information derived from the narrative model to quickly create the topical model (structured queries organized in taxonomies) that best supports the business objectives.

The system recommends additional terms to include, increasing the breadth of the query; it suggests terms to exclude in order to prevent false alarms; and it automatically dispositions the query, setting the best threshold possible. The end result is a better-structured query, one that recalls more true occurrences of the topic with fewer false alarms, with minimal effort on the part of the user. Customers armed with the knowledge of their business and corporate goals can focus on outcomes versus technology.

PHONETIC SEARCH

Most importantly, Neural Phonetic Speech Analytics generates a time-aligned phonetic index. This index is based on phonemes, the distinct sounds that make up language. Because phonemes are simply uttered sounds, the indexing is not affected by factors such as background noise, languages, dialects or speaking styles. Nexidia Interaction Analytics uses the structured queries derived from topic modeling to search across a phonetic index of 100% of customer interactions. This quantified approach calculates the statistics for root cause analysis, data mining, and metric-driven performance management.

Neural Phonetic Speech Analytics™ supports phrase recognition and sentiment detection to uncover emerging trends in the contact center.

The phonetic index can be searched directly on words or phrases or using special operators such as Boolean strings or time-based proximity to other content. Nexidia's proprietary search engine identifies and matches the phonetic equivalent of the search string and returns relevancy-ranked results. When words or business logic need to be changed, the system can rapidly re-query the index without having to re-process the audio. The result is a process that not only creates the truest representation of spoken audio, but also enables the fastest, most accurate access to the information contained within the audio files.

Architecture: Nexidia Search Grid™

As Neural Phonetic Speech Analytics is the technology that powers Nexidia Interaction Analytics, Nexidia Search Grid is the architecture that manages and operates the neural phonetic engine. The Search Grid is a highly scalable, distributable, parallel processing system for efficiently indexing audio (both phonetically and via transcription), storing and managing the resultant phonetic indices and transcription results, and then very quickly searching them.

BIG DATA METHODOLOGY

Based on MapReduce principles for parallel processing, the Search Grid consists of three major components: gateway node, data nodes, and compute nodes. The gateway node receives requests for work from a client application, in this case Nexidia Interaction Analytics. The most common work requests are for indexing new audio files and searching existing indices with search terms or structured queries. The gateway node "maps" the request; it determines which compute nodes should perform the computations and which data nodes should store or retrieve the indices, and then distributes the work. If the request involves indexing a new audio file, the designated compute node will retrieve the audio file from the specified media store and then index it, sending the resultant index file to the appropriate data node. If the request involves searching a set of existing indices, the data nodes storing the needed indices serve them up to the designated compute nodes, which then execute the searches. All of the results, such as index status or search results are then "reduced" by the gateway node and returned to the client application. What this means to the end user is an architecture with unsurpassed speed, greatly impacting the efficiency of the search functions of the application.

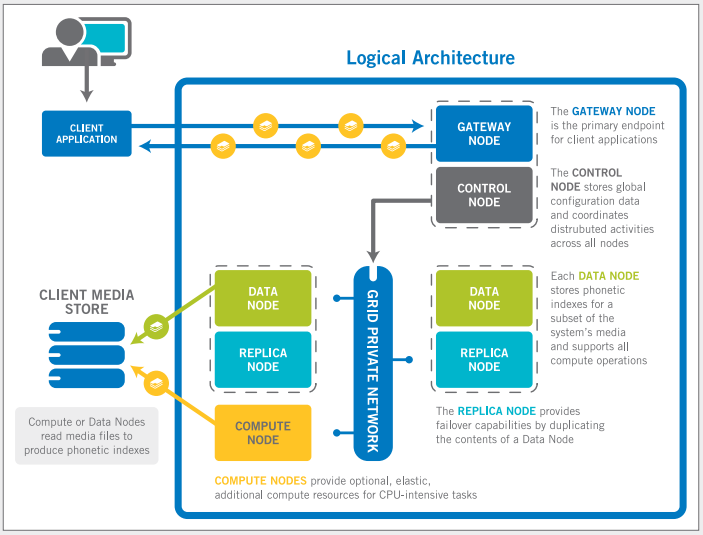

Nexidia Grid Logical Architecture

UNSURPASSED SCALABILITY

The parallel processing design of the Nexidia Search Grid makes it highly scalable horizontally.

The architecture of the Search Grid allows for data partitioning across any number of data nodes. Consequently, the Grid can grow to support any amount of audio indices simply by adding additional data nodes as needed. The Search Grid also supports elastically scaling the compute power. If a job requires more compute power than is already deployed, additional compute nodes may be added for the duration of the need, then removed. These additional compute nodes need only meet the prerequisites for the Search Grid software; they do not need to match the configuration of the other nodes of the system. For instance, the deployment team may deploy the core, permanent Search Grid servers directly to physical servers, but then deploy the elastic compute nodes via virtual machines. Nexidia Search Grid monitors the workload on every node and distributes new work accordingly, assigning more work to those nodes with more capacity.

Finally, the Neural Phonetic Speech Analytics processes hosted by the Nexidia Search Grid are designed to take maximum advantage of the available resources on any individual server. For instance, deployment teams can scale individual compute node servers vertically with the installation of additional graphical processing units (GPUs) in the server. When one or more GPUs are present, the Search Grid compute node on that server automatically shunts all the computations it can to the GPU(s), thus dramatically increasing the overall throughput of that compute node.

With the ability to scale both horizontally and vertically, Nexidia Search Grid optimizes hardware utilization by dynamically scaling to meet demand. Where other solutions require multiple instances of software in order to keep up with large implementations, Nexidia Search Grid manages with one instance of the system, providing unparalleled computational efficiency and scale.

VoiceProducts

VOICE • VIDEO • DATA

nexidia

8555 E. 32nd St. N. | Wichita, KS 67226
P: 800.466.1152 | www.voiceproducts.com | sales@voiceproducts.com