

nuance recognizer



white paper ::

Bringing New Levels of Accuracy, Reliability, and Ease of Use to Speech-Based Self-Service Applications

NUANCE RECOGNIZER

In today's competitive business environment, customer care is a highly effective means of achieving differentiation and winning customer loyalty. As a result, speech applications are playing an increasingly pivotal role in defining the customer care experience. Companies are investing in speech-based self-service solutions to meet growing demand for convenient 24x7 access to information, transactions, and customer service. The most successful of these speech solutions balance automation with usability to deliver self-service interactions designed and managed from the customer's point of view. Relying heavily on speech recognition, these solutions enable more efficient phone-based interactions that result in improved customer satisfaction and significant cost savings.

Nuance Recognizer is the best-of-breed speech recognition software that increases the efficiency of self-service solutions to help companies deliver a consistently excellent customer experience. Built upon the combined expertise, research, and code bases of the two market-leading speech engine providers, it delivers the industry's highest recognition accuracy and enables natural, human-like conversations that drive more satisfying self-service interactions. Built-in load balancing, multitenancy, and centralized logging support more efficient use of computing resources. Centralized server management—with Operations, Administration and Management (OA&M) capabilities, reporting, and other analytical functions—simplifies deployment and maintenance of self-service solutions. Together these capabilities add up to unparalleled levels of accuracy, reliability, and ease of use that transform the way companies care for key constituents: customers, business partners, and even employees.

Nuance Recognizer helps drive higher business performance through more efficient self-service solutions that increase automation rates and improve customer satisfaction. It optimizes self-services interactions and enhances the customer experience by delivering outstanding value in three key areas:

- **Accuracy** for more efficient customer interactions
- **Reliability** for more efficient use of resources
- **Ease of use** for more efficient deployment and maintenance

accuracy

Frustrated by lengthy, circular self-service interactions that too often end in opt-outs to a live agent, customers are demanding more usable speech applications that make it easier to resolve issues on the first call. That's why accuracy is such a key concern when evaluating speech recognition software for self-service applications. In fact, in a May 2006 internal survey, Nuance's customers and partners identified accuracy as what they valued most from a speech solution provider.

Nuance Recognizer was specifically designed to offer the industry's most accurate speech recognition through natural language processing, noise robustness, and support for an unparalleled range of languages and dialects. In fact, Nuance Recognizer has been shown to deliver exceptional accuracy results compared to predecessors such as OpenSpeech® Recognizer 3 and Nuance 8.5. (See table 1.)

Nuance Recognizer, by the numbers

- 47** Staff-years invested in the design and development of Nuance Recognizer v9 from January 2006 through its release in April 2007
- 27** Average reduction in error rate compared to its recognition engine predecessors
- 10** Average years of speech experience of Nuance Recognizer team members
- 44** Global languages and dialects supported by Nuance Recognizer
- 250** Cumulative years of speech experience of the speech R&D team at Nuance
- 6** Acquired companies whose technology was integrated into Nuance Recognizer

Language	Achieved RERR% vs. OSR3	Achieved RERR% vs. Nuance 8.5
U.S. English	27%	26%
Australian English	35%	29%
UK English	15%	32%
German	33%	16%
Canadian French	27%	39%
French	14%	N/A
Spanish	45%	N/A
Indian English	27%	N/A

Table 1. Relative Error Rate Reduction (RERR) for Nuance Recognizer, from internal Nuance benchmark testing. Results represent averages across multiple recognition tasks such as digit strings, alphanumeric spellings, and itemlists such as stocks or city names.

better performance

Nuance Recognizer effectively ignores background noise, measurably improving speech clarity and raising accuracy rates dramatically for wireless, hands-free, and noisy environments. This allows today's multi-tasking and increasingly mobile customers to make service-related calls from virtually any location for maximum convenience. In addition, it *automatically* adapts to each specific deployment and adjusts acoustic models to local dialects, regional accents and telephony channel characteristics for improved recognition accuracy—without any additional effort—even if caller demographics shift over time.

advanced acoustic models

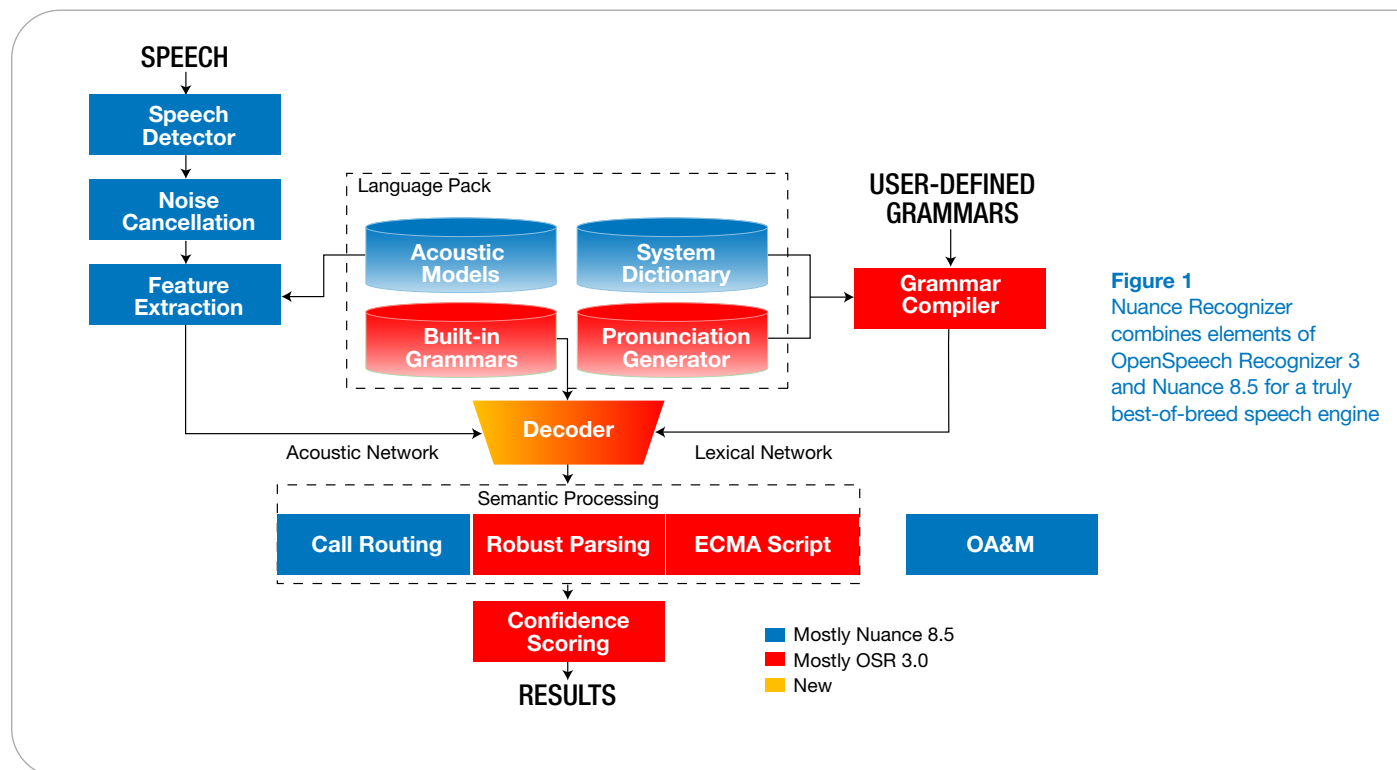
Background noise, including conversations, coughs, traffic, and mobile phone static all pose a key challenge in speech recognition. Nuance Recognizer, however, uses advanced acoustic models that are inherently robust to noisy input to deliver higher accuracy, fewer re-prompts, and superior barge-in. Thanks to strong connections within academia and established relationships with data collection consortia, Nuance has extremely broad and deep access to acoustic and linguistic data. As a result, Nuance Recognizer’s acoustic models have been trained on a vast range of real-world data, including noisy data from all sorts of environments, to deliver unmatched noise robustness.

endpointer detection

Nuance Recognizer is able to separate speech from background noise with remarkable precision due to its superior endpointing and speech detection algorithms. Improved endpointer detection enables the system to determine when speech started and ended—even in extremely noisy mobile environments—for more accurate transcription. In fact, Nuance Recognizer has been shown to exhibit exceptional accuracy compared to predecessors such as OpenSpeech Recognizer 3, delivering an average 45% improvement in detecting speech in noisy environments.

more natural conversations

Nuance Recognizer delivers the industry’s most accurate natural language capabilities to efficiently handle open input and to power mixed-initiative dialogs. These flexible, more intelligent conversations take into account information provided proactively by callers—for instance, a caller who answers several questions at once, or who talks over a prompt in order to correct the system. As a result, speech-based self-service applications can deliver better performance results—higher automation rates, reduced call length, and fewer hang-ups/opt-outs—and a more rewarding interaction experience for callers.



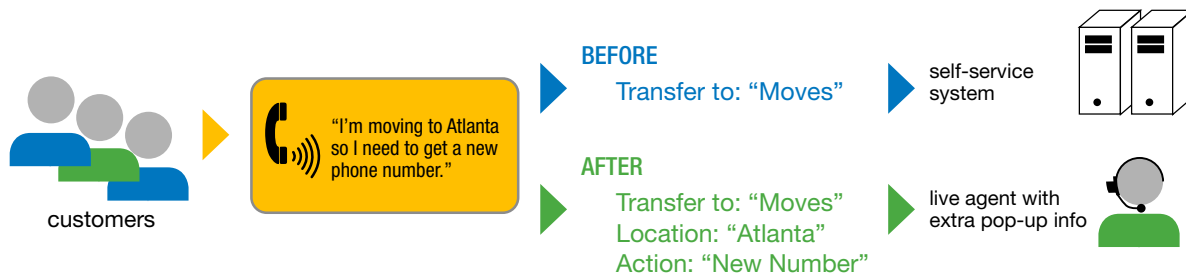


Figure 2 Improvements to SpeakFreely allow Nuance Care Solutions to identify multiple intentions from a single open-ended request

SpeakFreely™

Nuance Recognizer allows applications to understand an exceptionally broad range of words and phrases without requiring highly complex grammar rules that would be prohibitively time consuming and expensive to develop using traditional recognition techniques. This ability, known as SpeakFreely, uses statistical language models (SLMs) and statistical semantic models (SSMs) to determine what callers mean—not just what they say. Rather than requiring prompts with long lists of choices, Nuance Recognizer can offer open-ended questions (for instance, “How can I help you today?”) and handle responses with very accurate recognition results.

Nuance’s latest SpeakFreely breakthrough encourages more efficient customer interactions by reacting to multiple pieces of information in a caller’s unconstrained request. The result is faster handling of caller requests by reducing the back-and-forth between the system and the customer for additional information, as shown in figure 2.

multilingual input support

In order to compete more effectively, many companies are striving to provide global reach with local access. To that end, it’s important that a company’s self-service applications support multiple languages and dialects as appropriate to meet the needs of its customer base. Nuance Recognizer supports more than 44 languages and dialects spoken around the world. It can even recognize from several language grammars in parallel to address the needs of caller populations that speak multiple languages. As a result, Nuance Recognizer is the ideal engine to power speech self-service systems targeting customer bases that span wide geographic regions or those that are multilingual.

reliability

With customers ascribing value to a company based on their experiences across self-service channels, it's critical that companies implement speech self-service applications that offer consistently high availability and optimal performance. Nuance Recognizer meets companies' reliability requirements by providing a distributed, resource-managed system architecture that supports load balancing, scalability, and consolidated logging and monitoring.

resource management & load balancing

Even the most robust and well-designed speech system will enhance the customer experience only if it is available and operating at peak performance. Nuance Recognizer's architecture provides scalable operation while minimizing hardware requirements and costs. Taking advantage of the latest Media Resource Control Protocol (MRCP v2), it supports load balancing and fault tolerance across speech recognition and text-to-speech operations. Nuance Recognizer ensures efficient utilization of system resources for dramatic hardware savings compared to other systems.

The Nuance Resource Manager, which manages routing of client requests across distributed servers, uses a resource discovery protocol to automatically assign each incoming call to a server based on availability. The Resource Manager can automatically detect if a server in the cluster goes down and re-direct calls as appropriate to ensure fault-tolerant operation. In addition, the distributed, resource-managed architecture enables servers to be taken off line one at a time for maintenance without risk of service disruption. This way, companies can rest assured that their speech self-service systems are always accessible to customers and ready to handle even peak call volumes in the most efficient manner.

consolidated logging and monitoring

Nuance Recognizer's consolidated call logging and monitoring capabilities improve application management and tuning by providing for secure server side logging of information. Combining the one-file-per-call logging schema of Nuance 8.5 with the log content and format of OpenSpeech Recognizer 3, Nuance Recognizer makes call logs readily accessible to speech application developers, resulting in easier system deployment and optimization. For added security, logging can be disabled for sensitive data captured by Nuance Recognizer, such as credit card numbers. Thanks to multitenancy, multiple speech-based applications can share the same server and still be tracked separately for logging and reporting tasks. In addition, Nuance Recognizer APIs support precise, automatic log merge with OA&M capabilities to facilitate efficient speech system management and maintenance.

ease of use

When evaluating speech recognizers, it is important to look beyond feature/function alone. In order to deliver a superior customer experience, boost business performance, and provide strong return on investment, a speech system must also be easy to deploy and maintain over time. Nuance Recognizer responds to these requirements with support for current industry standards and availability on a comprehensive range of open platforms. In addition, Nuance Recognizer provides centralized server management with powerful, secure OA&M capabilities, including consolidated logging, monitoring, and reporting, that simplify efficient management and maintenance of a company's speech-based self-service systems.

Standards Support

Standards-based solutions offer the benefits of application portability and investment protection. Standards also promote market expansion by encouraging developers to write code that will likely operate on a broader range of platforms. Nuance not only embraces standards, but also takes a leadership role in establishing standards for creating speech-enabled applications, Text-To-Speech output, and communication protocols for speech on VoIP and Wide Area Networks.

Reflecting this commitment to open, standards-based solution, Nuance Recognizer follows the following speech standards:

- Speech Recognition Grammar Specification (SRGS)
- Semantic Interpretation for Speech Recognition (SISR)
- Natural Language Semantic Markup Language (NLSML)
- Extensible Multi-Modal Annotation (EMMA)
- Media Resource Control Protocol (MRCP, both v1 and v2)

centralized server management

The Nuance Management Station provides centralized server management with powerful, secure OA&M capabilities. These advanced performance management and reporting features enable system administrators and operators to more efficiently manage and maintain all aspects of their speech systems to ensure high service availability. The Nuance Management Station is centrally accessible from any location using a Web-based GUI console and collects logs from managed nodes running any service or process across the network. It also supports saving and sharing of role files to speed and ease network configuration.

In addition, the Nuance Management Station enables administrators and operators to:

- Set up the elements of the Nuance network
- Define clusters—groups of servers—according to service type and configure them together
- Start, stop and set service properties
- Monitor host and services vital signs, such as CPU utilization and available memory—on a single screen
- Manage data and generate reports
- Automatically send alerts and reports via email

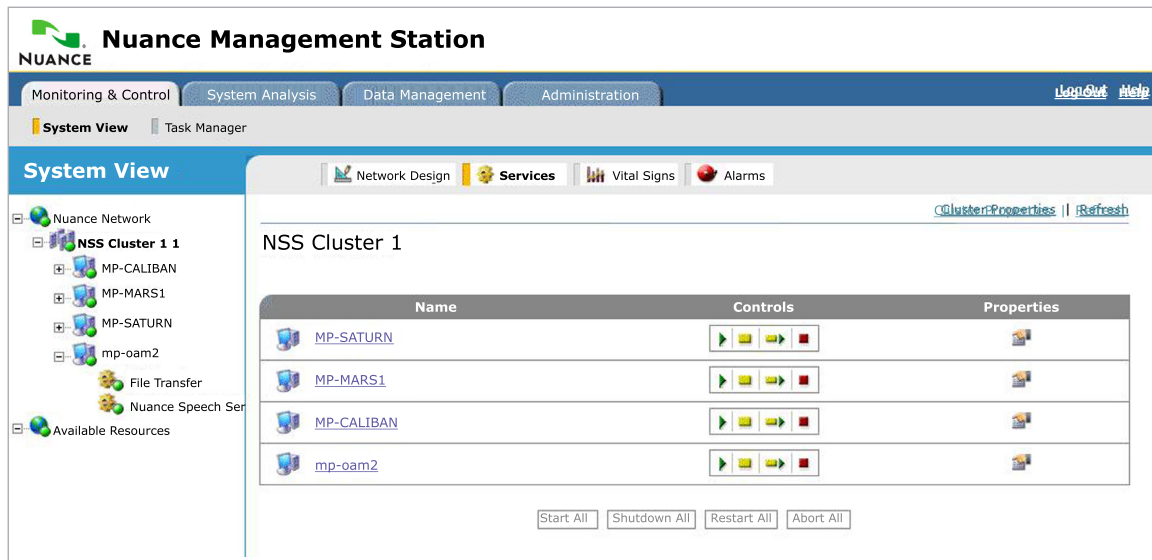


Figure 3 The Nuance Management Station provides centralized logging and monitoring capabilities to more efficiently maintain speech services across multiple server clusters.

reporting and analytics

Nuance Recognizer includes robust reporting capabilities for monitoring application performance and usage. Application reporting tracks metrics and gathers business intelligence needed to recommend broader tuning initiatives, while event logging and reporting accelerate and simplify troubleshooting and problem resolution. Sophisticated, yet easy-to-use reporting and analytics enable business managers to assess how well the system is delivering on key company objectives. This is accomplished using tools that provide automatic real-time measurement of performance against specific cost reduction or revenue generation metrics defined early in the speech application planning process. Some examples of reports include task completion rates, voice site hits, call volume, and task durations. With Nuance Application Reporting, a tuning tool included with Nuance Recognizer, application developers are able to effectively analyze and optimize voice application performance and usability.

CONCLUSION

With customer care emerging as *the* key differentiator in today's crowded markets, it's no longer sufficient for companies to automate customer service without considering the effects on callers. That's why more and more organizations are partnering with Nuance. Combining best-in-class technology with the expertise, experience, and practice gained through hundreds of speech deployments worldwide, Nuance Care Solutions are helping organizations to provide the most flexible, efficient, and personal caller interaction experience.

Nuance Recognizer is the speech recognition solution of choice for companies looking to achieve differentiation and win customer loyalty by delivering an exceptional customer experience. It combines the natural conversational capabilities of OpenSpeech® Recognizer and SpeakFreely™ with the administrative resources of Nuance 8.5 to deliver a best-of-breed speech recognizer that drives business performance by dramatically enhancing the efficiency of speech-based self-service interactions. Providing unprecedented accuracy, reliability, and ease of use, Nuance Recognizer is helping companies around the world to transform the way they care for customers, business partners and employees. Like all Nuance Care Solutions, it helps organizations enhance business efficiency, increase caller satisfaction, and achieve differentiation through superior customer care.

about Nuance Communications, Inc.

Nuance is the leading provider of speech solutions for businesses and consumers around the world. Its technologies, applications and services make the user experience more compelling by transforming the way people interact with information. Every day, millions of users and thousands of businesses experience Nuance's proven customer care solutions that empower companies to improve business performance, better satisfy customer demands and differentiate themselves through superior service. For more information, please www.nuance.com.

© 2007, Nuance Communications, Inc. All rights reserved. Nuance, the Nuance logo, and The experience speaks for itself, and HumanTouch are trademarks or registered trademarks of Nuance Communication, Inc. and/or its affiliates in the United States and/or other countries. All other trademarks are the properties of their respective companies.

WP1/07