



1005 Hamilton Court
Menlo Park, CA 94025
Phone 650.847.0000
Fax 650.847.7979
www.nuance.com

White Paper

The Nuance Project Method

A Guide to Speech Application Development

A large, light gray decorative graphic consisting of two overlapping, rounded rectangular shapes. The top shape is on the left and the bottom shape is on the right, with their corners overlapping in the center.

Contents

Introduction	3
Speech Application Overview	4
The Nuance Project Method	5
Phase 1: Requirements Analysis	6
Phase 2: Design	7
Phase 3: Implementation	8
Phase 4: Test	9
Phase 5: Tune & Monitor	10
Methodology Summary	12
Application Enablers and Tools	13
Glossary of Terms	14
About Nuance	15

The Nuance Project Method

A Guide to Speech Application Development

Introduction

Spoken conversation is one of the most natural and effective ways for people to access information and conduct business. There are more than 850 million wireline phones and 450 million wireless handsets in use today, compared with only 270 million Internet-capable PCs.¹ With the increased affordability of mobile service, the growth of wireless handsets is estimated to top one billion by 2003.² This makes telephones by far the most ubiquitous network device. Since voice recognition servers are accessible from any telephone, they are the perfect devices for offering services to the mass market.

With the rapid advances in computer processing and speech algorithms, telephony-based voice recognition systems can now “understand” natural spoken conversation. Since the mid-1990s, many prominent blue-chip companies have operated V-Commerce (voice-driven e-commerce) systems to deliver services such as flight information and real-time stock trading using both voice recognition and voice authentication. Nuance continues to lead the charge of voice interface service development into the 21st century, working with customers and partners to expand the Voice Web through top-quality, voice-enabled solutions.

This paper will:

- Explain what a speech application is, including the components and supporting infrastructure.
- Provide a clear and complete methodology for creating voice-enabled applications.
- Describe application enablers and development tools that Nuance offers for rapid deployment of speech applications.
- Describe the suite of speech application and voice interface services that Nuance and/or Nuance partners can provide.

¹ Yankee Group, IDC 1999

² Yankee Group, 2000

Speech Application Overview

As companies and organizations investigate the possibility of using a speech recognition or voice authentication system to reduce costs, improve customer service, and/or generate additional revenue, they need to know what steps are involved in making it a reality and a success. Introducing speech applications into a company should be approached as any new system development project; in a well-planned, consistent fashion and with a thorough understanding of both the strategic and tactical components. There are, however, certain aspects of such projects that are unique to voice-enabled applications. Specifically, designing the application so that the Voice User Interface (VUI) meets the needs and expectations of the target audience. This is very different from designing a GUI or a touch-tone application. However, with the right technology, process and implementation team, designing a VUI can be straightforward and simple.

Anatomy of a Speech Application

There are a number of key components required to implement a complete voice-enabled solution. They include the speech application platform, hardware and software, Nuance core speech software and servers, the telephony interface, application servers, the host functionality and database, and the system integration of all of these components. For maximum flexibility, Nuance's core speech software is available on all leading operating systems and on all market leading IVR and telephony platforms. **Figure 1** below is a simplified illustration of the key components of a speech application.

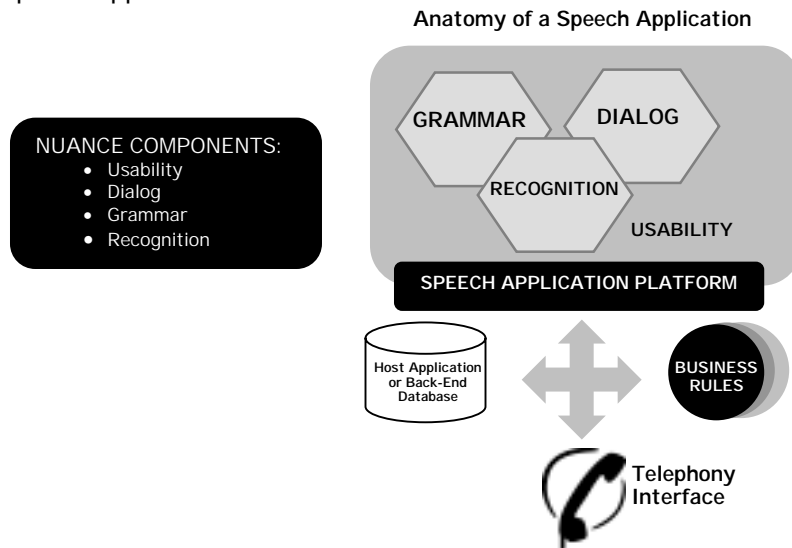


Figure 1. Anatomy of a Speech Application

The top half of **Figure 1** represents Nuance's software and application components, that are seamlessly integrated with the speech application platform. The Nuance core components include the speech recognition, the dialog, (a set of recorded voice prompts, responses, and directed call flows) and the grammars (defined sets of valid caller responses). The Nuance recognizer processes what is said, compares it to all acceptable possibilities in the grammar, and returns a result to the application code for appropriate action. By focusing on the needs of the user throughout speech application design and development, usability ties together the dialog, grammars and recognition to create a high-quality, voice-driven user experience.

Nuance speech applications can be deployed in several ways, and on many different speech application platforms. The application platform is comprised of Digital Signal Processing (DSP) boards or can use a Voice over IP (VoIP) software-only stack. The platform typically contains an application generation environment for creating a caller application using business rules to route and respond to calls. The platform can range in scope from a PC with DSP boards installed, to a large multi-server system, depending on a customer's expected call volumes and support needs.

The speech application platform loads grammars, plays audio prompts and routes the callers' utterances to the Nuance software for recognition. It can also access host systems and databases to retrieve content or customer data as needed. The telephony interface allows users to reach the application through the Public Switched Telephone Network (PSTN) or data network, comprising all available phone channels/networks including, landlines in the home/office, wireless cell phones, and connections to VoIP.

Regardless of the deployment method selected, the implementation of a voice-enabled application requires a proven process for planning, designing and deploying these applications. The remainder of this paper will focus on Nuance's speech application design and development methodology. It is this approach that enables Nuance to work with partners and customers to provide optimal voice interface solutions.

The Nuance Project Method

Nuance employs a proven methodology for planning, designing, and deploying successful speech applications and offers a full suite of speech application services. Together, they seek to exceed customer business goals, satisfy and surpass application requirements, and most importantly, optimize the user experience by providing a compelling voice interface to the end user.

Nuance's standard speech application development methodology is captured in **Figure 2**. The process has five key phases, each of which are explained in more detail in the remainder of this document. In every phase, Nuance's expertise in speech recognition technology and voice user interface design can be leveraged. Core competencies in each phase complement each other and ensure full preparation for moving from phase to phase during the project lifecycle.

By gaining an early understanding of the project requirements, Nuance can create an interaction that will be compelling and useful for both the customer and the target audience. Throughout the design and implementation stages, Nuance collects speech data and conducts usability studies to ensure the service being created is optimized to meet the customer's goals. To make sure high quality control standards (set by Nuance and the customer) are being met, thorough testing is conducted on each component of the application as well as on the fully integrated application. Finally, the system is tuned and monitored during both the pilot and production phases, in order to optimize performance, recognition accuracy, and the user experience for the target audience.

Nuance Project Method
Speech Application Design and Development

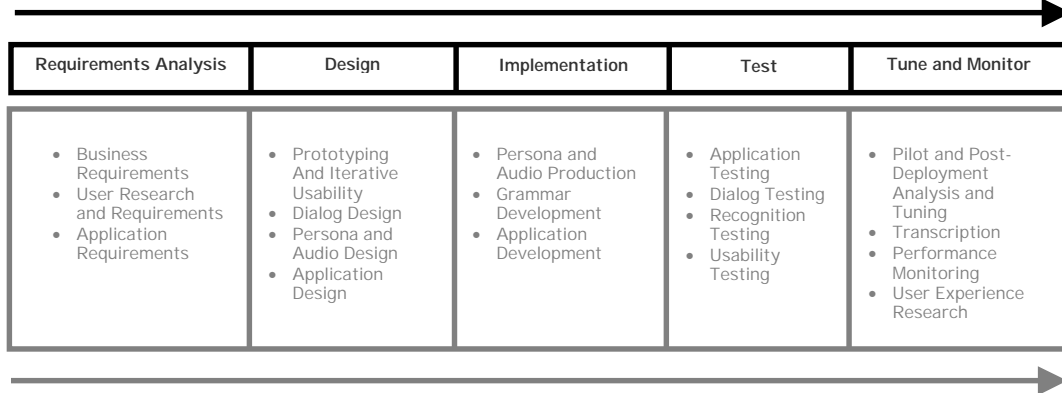


Figure 2. Nuance Project Method

Phase 1: Requirements Analysis

A thorough understanding of the needs and goals of a project is a strong predictor of its success. Therefore, Nuance gathers requirements in a regimented fashion. The business needs are assessed first, which then drive user and application requirements. A high-level understanding of the critical success criteria, expected users, system functionality, grammar and recognition complexity, and project risks are captured early on, allowing Nuance to accurately determine the scope of work.

The requirements analysis phase can be broken into three distinct components:

- Business Requirements
- User Research and Requirements
- Application Requirements

Business Requirements – Determining the business requirements of the application is the first step in the requirements analysis. This includes understanding the key reasons for building the service, how the proposed service will generate revenue, the branding and differentiating factors, and the metrics that will be used to judge overall success.

User Research and Requirements – Design considerations are captured based on user groups, usage patterns, usage environment, desired persona and usability metrics. The perceived needs of the user, how the user will interact with the system, and the reasons for calling and continuing to use the system are addressed in this phase. Nuance may have existing research data that can be leveraged, or new studies can be conducted to generate additional data if required. This results in a detailed description of who will be using the system including identification of primary and secondary user groups.

Nuance conducts a complete assessment of user demands on the service in order to design a compelling application that meets caller needs and matches their expectations. For example, anticipated usage patterns of the speech application, both as an individual interface, and in relationship to other communication methods used (e.g. GUIs, handheld devices, telephones, computers, and the Internet) are considered. The service usage environment and the implications of that environment on design are explored. Based on resulting goals for the user experience, Nuance identifies key usability metrics against which the service will be evaluated during design and upon deployment.

Application Requirements – Nuance also captures design guidelines based on recognition complexity and naturalness, grammars, database integration, and existing application platforms. The scope of the application, the use of additional speech technology such as Nuance Verifier™ (voice authentication), or Nuance Vocalizer (text-to-speech) are determined. Integration to back-end databases which include all types of caller data that must be captured by the application are also clarified at the outset of the project.

A partner or customer will first work with a Nuance sales engineer to compile these high-level requirements into a *Project Requirements Document (PRD)*. In most cases, the PRD provides Nuance with enough information to rapidly generate an accurate quote estimate. This quote, along with a description of the services to be rendered is presented to the partner or customer for review and approval in a *Services Proposal* with an accompanying *Statement of Work (SOW)*. Once a signed SOW is received, Nuance can allocate the necessary resources to kick off the project and initiate the design phase.

Phase 2: Design

The design phase focuses on planning the user interface by defining the dialogs, prompts and grammars as outlined in the Requirements Analysis phase. A key focus of design is to flesh out the overall application framework, including the introductory behavior of the application, menu states, and universal functionality of help and error responses. Design may be done as a custom project or it may be expedited through the use of one of Nuance's pre-packaged designs, built for specific markets where Nuance has developed strong domain expertise. These pre-packaged designs are discussed in detail later in this White Paper, in the section "Application Enablers and Tools."

The design phase can be broken into three distinct components:

- Dialog Design
- Persona Design
- Application Design

Dialog Design – This step involves first producing high-level *call flows*, then moves into detailed dialog design. The detailed design process includes writing a description of every call state and how each state transitions to another. It requires a framework for specifying call flow structure, state-specific prompts and natural language results, sample phrases for each state's grammar, help and error recovery, universal behaviors, and sample dialogs. To capture all this detail, Nuance creates a *Dialog Design Specification* that must be approved by the customer prior to moving forward with development.

Persona Design – Nuance research shows that any time a caller interacts with a speech recognition system, the caller perceives the system to have a personality. This is true whether or not the system was specifically designed to have one. It is, therefore, important to design a persona that is both appropriate and agreeable. Use of an appropriate persona facilitates comprehension, ease of use, and a positive user experience.

Persona design and dialog design are closely linked. In order to create a consistent and enjoyable user experience, the persona must infuse the entire dialog design process. Early in this stage, the persona designer can create one or more *persona definition* document(s): fictional autobiographies for potential characters of the

system. Nuance can also provide a number of off-the-shelf, standard personas to meet a variety of application needs in a cost-effective manner. Next, based on the initial call flow provided by the dialog designer, the persona designer creates a sample interaction between a caller and the application. These interactions may then be recorded using a variety of voices and with a variety of personas. The resulting *audio samples* are shared with the customer and potential users in order to determine an ideal match for the system. Once the customer selects a persona, a focused persona design effort begins, in which all baseline prompts are edited to evoke that persona and ensure clarity and comfort in the interaction. These prompts are then incorporated into the dialog design specification for customer approval.

Other non-speech audio is also considered at this stage such as audio logos for branding, latency covering sounds, and landmarks to instantly identify place and functionality by way of effects or music.

Application Design— Building on the details captured in the Requirements Analysis phase, Nuance works with the customer and/or partner to finalize and document how the various components of the application will be seamlessly integrated. This includes planning for the installation of the recognition software, storage of the application, access to the recorded prompts and other audio files, connecting the application to the platform and any required DTMF backup, and providing secure and reliable read/write access to backend databases or data feeds. Any required business rules that go beyond the dialog design are also defined (such as the ability to detect a first time caller for different prompting), and the estimated size/complexity of the necessary application grammars are scoped.

Throughout the design phase, *iterative usability* and rapid *prototyping* create a feedback loop to optimize the interaction of dialog design and persona prior to a full-scale implementation. Early and incremental usability testing is recommended on every project; it speeds time to market and ensures a design that creates a compelling user experience.

Phase 3: Implementation

In the implementation phase, the application is built, prompts are recorded and the recognition package is completed. These tasks can occur in parallel, since the dependencies and interfaces between the tasks (such as naming of the states, prompts and grammars) were addressed in the design phase, and documented in the Dialog Design Specification and persona definition documents. The three key components of implementation are:

- Persona & Audio Production
- Grammar Development
- Application Development

Persona & Audio Production – Nuance provides a team of specialists in voice talent coaching, prompt recording, and audio environment development. This team ensures that the project delivers on goals for the persona and overall audio experience, and that prompts are recorded in a fashion that facilitates clarity of communication, user comfort and user control. Specifically, the Nuance team:

- Identifies and engages specific voice talent that fits the persona
- Approves and certifies studios for prompt recording to ensure the highest standards of voice recording for the specific application

- Produces the final prompts on-site or in Nuance certified facilities around the world
- Produces and combines non-speech audio with prompts

Grammar Development – For each part of the dialog, the goal of grammar development is to anticipate what a user might say. The grammar, in conjunction with the prompts, is developed in a manner that encourages the caller to say something that will be understood. Nuance leverages its vast library of existing grammars associated with different types of applications, to define acceptable user utterances and their associated natural-language interpretation. Nuance will add to or edit those grammars, or create custom grammars, as required for the application. Nuance also offers a *Grammar Update Service* for information that needs to be constantly maintained such as names of US equities, mutual funds, and international airport/airline information.

Application Development – This step involves writing out the core speech application using any language supported by Nuance API's (including C++, Java with SpeechObjects, or VoiceXML). As with design, development may be accelerated with the use of application enablers created by Nuance to address specific vertical markets (e.g. brokerage, travel) or horizontal functionality (e.g. order management, customer support). These are explained further in the section of this white paper entitled "Application Enablers and Tools." If external data sources are required by the application (e.g. stock quote feeds, audio feeds, customer databases or a host), Nuance can use API specification documents to assist and facilitate integration tasks. Nuance will also build custom audio and backend interfaces, when necessary.

Prior to moving on to an in-depth testing stage, *unit tests* are performed on each separate component of the application. This involves the testing of each component in the design (i.e., dialog, persona, and application), grammar testing, and testing of each application module. The creator of each component (Nuance, a Nuance partner, or the customer) is responsible for conducting the respective unit test. Upon completion of unit testing, the application components are integrated and prepared for full application testing in the next phase.

Phase 4: Test

Once the application is developed and all individual components have been tested, Nuance works with the partner and/or customer to adequately test the full application leading up to the pilot. A quality assurance test plan is created to manage these efforts.

Three primary types of tests are conducted:

- Application Test
- Recognition Test
- Usability Test

Application Testing – Application tests are used to verify that the application is functional and all hardware and databases can be accessed. These tests include:

Dialog Traversal Test – consists of creating and executing a series of test cases that cover all the possible paths through the dialog. This test determines whether the right prompts are played, whether each state in the call flow is reached under all of the proper conditions, and whether all universal, error, and help behaviors are operational.

System QA Test – this testing stage involves rigorous, system-wide validation of the overall integration of the application with the platform hardware and APIs, as well as external interfaces such as databases, live data feeds and telephony.

System Load Test – Nuance uses load testing technology that simulates a high in-bound call volume to the new system. This test ensures that expected caller capacity can be handled, and that proper load balancing occurs across the system.

Recognition Test – After the dialog traversal test is completed, and any required modifications are made to the application to ensure compliance with the dialog specification, Nuance tests the recognizer to determine that it has been correctly integrated and is performing well. Nuance also develops test scripts for callers to generate utterances by talking to the application. These utterances are then transcribed by Nuance and scoring tools are used to determine the accuracy of the recognition results when compared with the transcription output.

Evaluative Usability Tests — While usability is also conducted early in the design process, it is often helpful at this stage to validate the performance of an application against the metrics laid out by the customer in the requirements analysis phase.

Phase 5: Tune & Monitor

Tuning is critical to ensuring that the system functions effectively and remains optimized for the user group as the service rolls out geographically and as users become more experienced. It is only through capturing, transcribing and analyzing live caller utterances that the application can be tuned for optimal performance. Nuance performs tuning on the application during the pilot(s) and, optionally, during post-deployment.

Regardless of when they are performed, tuning efforts focus on three principal areas: grammar, recognition, and dialog tuning.

- In *grammar tuning*, analysis and corrective actions increase the grammar coverage while removing unnecessary words, thereby resulting in overall improved system response and accuracy.
- *Recognition tuning* focuses on conducting experiments to determine the optimal search, confidence, and acoustic parameters of Nuance's software to maximize system accuracy and speed.
- *Dialog tuning* ensures that the prompts correspond appropriately to the grammars and that the dialog flow meets user expectations.

User experience and *live caller research* can also contribute valuable feedback from real users during the tuning process. This final phase in the Nuance Project Method, is executed in three stages:

- Pilot Analysis and Tuning
- Post-Deployment Analysis and Tuning
- Performance Monitoring and User Experience Research

Pilot Analysis and Tuning – In order to optimize the performance of the system before a full deployment, the application is made available through a pilot to a limited

audience of users. Nuance transcribes collected utterance data to analyze and determine performance metrics such as in-grammar accuracy, out-of-grammar rates and error distribution. The results are communicated to the customer or partner in a *performance report*. Experiments are run with modified grammars and/or parameters until recognition performance is optimized for the data. Then, a *revised grammar package* and/or *parameter recommendations* are released to the partner or customer for use in updating the system.

A true pilot is an iterative process. To achieve optimum application performance, Nuance generally recommends three cycles of data-collection, reporting and modification for each pilot. For some applications, Nuance may recommend conducting two pilots – the first with internal or "friendly" users and the second with a larger set of callers selected from the target end-user population.

Once pilot analysis and tuning is complete, and the customer is comfortable with the application performance, it can be moved into production and deployed to the full user base. This deployment may be coupled with a targeted communication to new users, or a wider media campaign, to introduce the new service.

Nuance also re-evaluates the application hardware requirements, which are heavily dependent on the size of the grammars and complexity of the dialog. Although this is estimated at the very beginning of the project, sizing the hardware based on real call data ensures the application is provisioned correctly, and will be ready to accept the volume of calls expected on the system.

Post-Deployment Analysis and Tuning – The same iterative process used during the pilots can be used to identify additional optimizations based on data collected from the full caller population. This data can also serve to track changes in usage due to increased caller familiarity or changes in the caller population. Revised grammar packages and parameter recommendations are provided by Nuance during this stage, in the same fashion as during the pilot.

Performance Monitoring and User Experience Research – Customers may schedule performance monitoring at regular intervals to check that the system is still maintaining peak performance. Adding additional callers often introduces new variables that may also need to be factored into a system's design. Nuance recommends two research techniques for gaining immediate insight from actual system users.

- *Live Caller Research* is designed to collect usability interviews from callers immediately following their interaction with a particular application. This ensures that callers remember enough of their interaction to provide detailed feedback on the application.
- *User Experience Research* also allows interaction with live users, but is designed to get feedback over time. This research technique is valuable for frequently used applications, where the design should accommodate the differences between frequent and infrequent callers, and potentially encourage particular behavior.

These techniques provide crucial insight for fine-tuning the dialog before an initial or subsequent release of the application.

Because tuning is extremely important for a speech application, Nuance has developed tools to facilitate this process. For more information, refer to the "Application Enablers and Tools" section of this white paper.

Methodology Summary

With industry leading knowledge of natural speech recognition applications, Nuance assists partners and customers at any stage of the speech application design and development process.

The following table summarizes the inputs and deliverables for each stage within the Nuance Project Method.

Project Phase	Inputs	Deliverables	Completion Criteria
Requirements Analysis	<ul style="list-style-type: none"> Requirements Documents 	<ul style="list-style-type: none"> Services Proposal and Statement of Work (SOW) 	Customer returns signed SOW and provides a corresponding purchase order number
Design	<ul style="list-style-type: none"> Services Proposal and SOW 	<ul style="list-style-type: none"> Dialog Design Spec. Persona Definition Voice talent selection Audio Environment System architecture 	Customer signs off on Dialog Design Specification and Persona/audio goals
Implementation	<ul style="list-style-type: none"> Design Outputs 	<ul style="list-style-type: none"> Prompt & Audio production Application code Recognition Package 	The "owner" (Nuance, partner or customer) of each component of the application conducts a unit test and ensures standards are met
Test	<ul style="list-style-type: none"> QA test plan Integrated application 	<ul style="list-style-type: none"> Results of integration, stress, dialog traversal and recognition tests 	Nuance and partner ensure all application test standards are met
Tune & Monitor	<ul style="list-style-type: none"> Application ready for Pilot Analysis and Tuning 	<ul style="list-style-type: none"> Performance reports Recognition Packages Live user research 	The application performance meets or exceeds pilot metrics and customer expectations.

Application Enablers and Tools

Nuance offers industry-specific application enablers and a suite of developer tools that enable the creation of sophisticated voice user interfaces while significantly reducing the development time and cost.

Application Enablers

Nuance Application Enablers are based on open-standards, with reusable speech application components that can be used independently and/or with VoiceXML to develop applications. They leverage the industry's best practices and can run on multiple telephony platforms. Application enablers include:

Order Management Suite (OMS) – Voice enables caller identification, shopping cart functionality, checkout (with shipping and payment options), credit card capture and verification, and order status.

Customer Support Suite (CSS) – Voice enables caller identification, frequently asked questions (FAQs), submission of new service requests, tracking of existing requests, and service bulletins.

The OMS and CSS suites include an XML layer for easy integration to existing e-commerce and CRM servers. Furthermore, the VUI is driven entirely by dynamic content supplied by the enterprise server – updates to this content don't require any updates to the speech application.

Broker Suite – Voice enables stock quote, mutual fund and trading applications. This bundle includes a design specification for retrieving quotes, holdings information, order status, and placing orders with qualifiers and durations. US equities grammars, and application grammars (including navigation, share amount, share price, etc.) are included, and mutual funds can be added to any Broker Suite application.

Voice Portal Suite – Powering Voice Web applications, these modules enable voice browsing of content for news, stocks, sports, horoscopes and weather. The modules can be implemented individually or combined together. Users can personalize their content by setting preferences within each module. The Voice Portal Suite can also be combined with voice dialing from a personal address book. To assist partners and customers, Nuance provides implementation and mentoring services around the Voice Portal Suite.

Nuance can also assist partners and customers with customization of these application enablers, or create new application enablers to meet specific application needs (unified messaging, auto-attendant functions, name and address capture, 401(k) account management, etc).

Grammar Products and Services

Nuance has a number of large, pre-built grammars for common application types (e.g. travel, brokerage) that can be leveraged to speed application development. In addition, Nuance can develop custom grammars for specific applications. Nuance also offers a comprehensive update service for keeping equity grammars current,

delivering optimal accuracy and usability for stock quotes and stock trading applications. Throughout the grammar development and update cycle,

Nuance follows a proven process to ensure coverage of expected caller responses, resolve ambiguities, provide accurate pronunciation models for all included words, and minimize errors.

Nuance Developer Tools

Nuance Developer Tools support the entire voice interface development life cycle. The tools employ intuitive graphical interfaces with “drag and drop” commands for ease of use. The developer tools use an open architecture that supports multiple hardware and software environments for maximum flexibility.

Nuance Grammar Builder – enables developers to rapidly create, review, test, manage, and deploy grammars. It also delivers over 100 commonly used grammars for rapid integration into speech applications.

Nuance V-Optimizer™ – helps developers understand system performance and quickly respond to changing user needs. It is used to analyze and tune piloted and deployed applications by providing access to performance data, system optimization features based on actual user behavior, and graphical report generation for system monitoring.

Nuance V-Builder™ – is the first complete graphical tool for building VoiceXML applications. It enables speech and web developers to rapidly prototype, develop, and test voice interfaces to web site and enterprise applications.

Glossary of Terms

In order to facilitate better understanding of the components of the Nuance Project Method, industry terms are defined below:

Barge-in – The ability of a caller to interrupt a recorded prompt or message to give a voice command.

Dialog – The interaction between the speech system and the caller. Includes the call flow (organization of the conversation), the prompts (everything the system can say) and the grammars (everything the user can say).

DTMF – Dual Tone Multi-Frequency, also known as touch-tone or keypad entries.

Grammar – The set of caller utterances allowed at any given time in the dialog.

IVR Unit – Interactive Voice Response Unit, a computer system that evokes and responds to a user's touch-tone or voice commands.

NLP – Natural Language Processing, the ability of a speech recognition system to translate many different utterances with the same meaning into a single result.

Persona – The consistent character of the system as reflected by voice, language use, and the audio environment. A persona is typically defined and or created for each application and is based on the target audience for that application.

Prompt – A speech system's audio output. Prompts are used for asking questions, presenting options, giving direction, or conveying information.

Speech Recognition System – A system that interacts with a caller to elicit speech, understands the speaker's meaning, and performs some pre-defined action based on that understanding.

SpeechObjects – Open, standards-based reusable speech application components that can be used independently or with VoiceXML to create robust speech applications.

Usability – The extent to which the voice application or service can be used by specified callers with effectiveness, efficiency and satisfaction. Also, the practice of studying a system's usability.

Voiceprint – A unique audio pattern that is created by the sound of a person's voice. Like a thumbprint, this can be used to identify individuals. It can replace the need for a user ID or PIN when trying to access secure areas of a voice application.

Voice Authentication – Also known as speaker verification. The process of matching a caller's voice to a pre-recorded voiceprint, providing secure access to systems and services through voiceprint matching.

VoiceXML – An extensible markup language that is an emerging standard for the VoiceWeb.

VUI – Voice User Interface, a broad term for the experience a caller is presented with when using the system. It is analogous to a GUI (graphic user interface) for a website. It includes the dialog components as well as the persona, and other audio files used to build the application.

WOZ – Wizard of Oz, a form of usability testing that simulates the flow of a voice application, but during which the caller is speaking with a human (i.e. – “the man behind the curtain”) rather than an implemented system.

About Nuance

Nuance is the leader in Voice Web software--speech recognition, voice authentication, text-to-speech and voice-browsing products that make the information and services of enterprises, telecommunications networks and the Internet accessible from any telephone. Every day, millions of people interact with Nuance systems at companies like BeVocal, Merrill Lynch, Nomura Securities and Tellme Networks. Nuance is headquartered in Menlo Park, Calif., with global sales offices and partners supporting solutions in multiple languages around the world. For more information about Nuance, visit www.nuance.com or call 888/NUANCE-8.

©2001 Nuance Communications Inc. Nuance, Nuance Verifier, Nuance Voyager, Nuance V-Builder, Nuance SpeechObjects and V-Commerce are trademarks of Nuance Communications Inc. All other trademarks are the property of their respective holders.