



# Avaya Conversational Intelligence™: A Real-Time System for Spoken Language Understanding in Human-Human Call Center Conversations

Jan Mizgajski<sup>1,3</sup>, Adrian Szymczak<sup>1</sup>, Robert Głowski<sup>1</sup>, Piotr Szymański<sup>1,2</sup>, Piotr Żelasko<sup>1,4</sup>,  
Łukasz Augustyniak<sup>1,2</sup>, Mikołaj Morzy<sup>1,3</sup>, Yishay Carmiel<sup>1</sup>, Jeff Hodson<sup>1</sup>, Łukasz Wójciak<sup>1</sup>,  
Daniel Smoczyk<sup>1</sup>, Adam Wróbel<sup>1</sup>, Bartosz Borowik<sup>1</sup>, Adam Artajew<sup>1</sup>, Marcin Baran<sup>1</sup>,  
Cezary Kwiatkowski<sup>1</sup>, Marzena Żyła-Hoppe<sup>1</sup>

<sup>1</sup>Avaya, Santa Clara, California, United States

<sup>2</sup>Wrocław University of Science and Technology, Wrocław, Poland

<sup>3</sup>Poznan University of Technology, Poznań, Poland

<sup>4</sup>AGH University of Science and Technology, Kraków, Poland

jimizgajski@avaya.com

## Abstract

Avaya Conversational Intelligence™ (ACI) is an end-to-end, cloud-based solution for real-time Spoken Language Understanding for call centers. It combines large vocabulary, real-time speech recognition, transcript refinement, and entity and intent recognition in order to convert live audio into a rich, actionable stream of structured events. These events can be further leveraged with a business rules engine, thus serving as a foundation for real-time supervision and assistance applications. After the ingestion, calls are enriched with unsupervised keyword extraction, abstractive summarization, and business-defined attributes, enabling offline use cases, such as business intelligence, topic mining, full-text search, quality assurance, and agent training. ACI comes with a pretrained, configurable library of hundreds of intents and a robust intent training environment that allows for efficient, cost-effective creation and customization of customer-specific intents.

**Index Terms:** speech recognition, Spoken Language Understanding, speech analytics, real-time systems, business intelligence

## 1. Introduction

Call center conversations are a valuable, but still underutilized asset for organizations. The majority of speech recognition products focus on offline use cases, such as speech analytics, quality assurance, or agent training. This means that business process changes, potentially identified during the analysis, have to be enforced outside of the application. With real-time conversation understanding, organizations can recognize the problem and act during the call, enabling such use cases as real-time supervision and coaching, agent assistants and robotic process automation, or live compliance and brand risk monitoring. We present an end-to-end system for real-time call center conversation understanding and several use cases built on top of it.

## 2. Overview of Features and Capabilities

### 2.1. Real-Time Capabilities

Avaya Conversational Intelligence™ has the following real-time capabilities.

**Large Vocabulary Speech Recognition**, which is based on Kaldi [1] and our proprietary models, is capable of transcribing

thousands of concurrent calls with very low phrase latency, providing high accuracy and computing efficiency.

**Transcript Refinement** - the transcript is further refined with automatic: punctuation (prediction of question marks, periods, and commas) [2], truecasing (recovering the proper casing of recognized words), readability turns (a word sorting algorithm to improve the readability of overlapping utterances from multiple speakers).

**Intent and Entity Recognition** - our Spoken Language Understanding engine comes with a fast and easy to train intent and entity recognition accompanied by a library of hundreds of pretrained intents. With capabilities such as: fuzzy matching, support for long intents (7+ words), recovery from speech recognition errors, and entity parsing (matching entities as parts of intents and parsing them into a machine-readable format), it is able to capture complex intents with high precision and recall. Entities are divided into two robust categories: *system entities* (numbers, amounts of money, dates, duration, quantities, volumes, percentages, spelling, names, surnames, locations) and *user-defined entities* – placeholders, such as BRAND or PRODUCT, that can be filled in by the customer, thus lowering the barrier of entry in both customizing existing intents and defining new ones.

**Business Rules Engine** - business rules model sequences and aggregates over the stream of events, such as intents, entities, or call events, to allow organizations to define complex triggers used to power such use cases as real-time supervision or agent assistance.

### 2.2. Post-Call Capabilities

The real-time capabilities are complemented with post-call analysis.

**Unsupervised Keyphrase Extraction** - we run a proprietary unsupervised keyphrase extraction algorithm to extract additional context for analytics to better service out-of-ordinary cases that are not covered by the configured set of expected intents.

**Call Summary** - calls are summarized using natural language generation based on all extracted information, such as intents, entities, and keyphrases.

### 2.3. Batch Capabilities

For customers that wish to work with recordings rather than live audio, we offer speaker diarization and classification (customer, call center agent), as well as all real-time and post-call capabilities in batch processing mode.

### 2.4. Aggregation, Search, and Business Intelligence

Based on the data extracted from multiple conversations, we offer:

- information retrieval with aggregation and filtering on multiple criteria, such as full-text search, matched intents and entities, call metadata (agents, projects, etc.), customer-supplied metrics and labels
- hierarchical topic mining
- visualizations and dashboards

### 2.5. Built-In Intent Training System

To allow fast training and customization, our product comes with built-in training including:

- mass call annotation with phrase grammars
- automatic, self-adapting synonym recommendations
- support for entities as parts of phrases
- instant verification on historical data sets
- easy-to-use UI, optimized for rapid intent definition
- feedback loops and false-positive training

## 3. Key Applications

In this section we present some of the key applications of Avaya Conversational Intelligence™.

### 3.1. Sentinel

Sentinel is a dashboard of live calls showing the most recent intents. Calls are assigned risk scores based on customer-defined business rules that can be triggered by various events in the call (intents, entities, etc.). High-risk calls are highlighted and enriched by additional context. A supervisor can start watching the transcription and events of the problematic call and take corrective action.

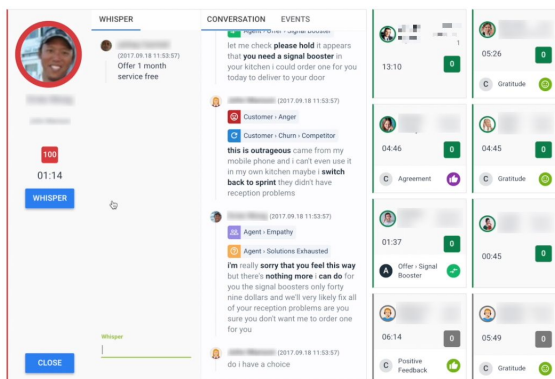


Figure 1: Sentinel application, showing live supervision of a problematic call.

### 3.2. Explorer

With Explorer, our customers can query, aggregate, filter, and visualize all historical calls. Each query can be saved and turned into a dashboard widget for future reference or comparison with multiple queries. If a particular call proves to be interesting, the customer can inspect it in detail along with its recording, transcription, and all extracted data.

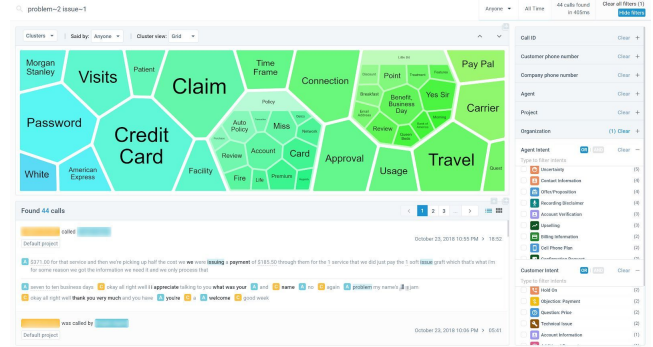


Figure 2: The Explorer app showing results for a query with topic mining visualization and some available filters.




Figure 3: Call Details panel with a fully transcribed, refined, and summarized call.

### 3.3. Streaming API

Streaming API enables our customers to subscribe to the real-time stream of transcriptions, intents, entities, and other derivative events to build their own applications, such as agent assistants, robotic process automation, or custom live-supervision dashboards.

## 4. References

- [1] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The kaldi speech recognition toolkit," 2011, iEEE Catalog No.: CFP11SRW-USB. [Online]. Available: <http://infoscience.epfl.ch/record/192584>
- [2] P. Zelasko, P. Szymanski, J. Mizgajski, A. Szymczak, Y. Carmiel, and N. Dehak, "Punctuation prediction model for conversational speech," in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018.*, 2018, pp. 2633–2637. [Online]. Available: <https://doi.org/10.21437/Interspeech.2018-1096>