AUTOMATED QUALITY MONITORING FOR CALL CENTERS USING SPEECH AND NLP TECHNOLOGIES

G. Zweig, O. Siohan, G. Saon, B. Ramabhadran, D. Povey, L. Mangu and B. Kingsbury

IBM T.J. Watson Research Center, Yorktown Heights, NY 10598

ABSTRACT

This paper describes an automated system for assigning quality scores to recorded call center conversations. The system combines speech recognition, pattern matching, and maximum entropy classification to rank calls according to their measured quality. Calls at both ends of the spectrum are flagged as "interesting" and made available for further human monitoring. In this process, the ASR transcript is used to answer a set of standard quality control questions such as "did the agent use courteous words and phrases," and to generate a question-based score. This is interpolated with the probability of a call being "bad," as determined by maximum entropy operating on a set of ASR-derived features such as "maximum silence length" and the occurrence of selected n-gram word sequences. The system is trained on a set of calls with associated manual evaluation forms. We present precision and recall results from IBM's North American Help Desk indicating that for a given amount of listening effort, this system triples the number of bad calls that are identified, over the current policy of randomly sampling calls. The application that will be demonstrated is a research prototype that was built in conjunction with IBM's North American call centers.

1. INTRODUCTION

Every day, tens of millions of help-desk calls are recorded at call centers around the world. As part of a typical call center operation a random sample of these calls is normally re-played to human monitors who score the calls with respect to a variety of quality related questions, e.g.

- Was the account successfully identified by the agent?
- Did the agent request error codes/messages to help determine the problem?
- Was the problem resolved?
- Did the agent maintain appropriate tone, pitch, volume and pace?

This process suffers from a number of important problems: first, the monitoring at least doubles the cost of each call (first an operator is paid to take it, then a monitor to evaluate it). This causes the second problem, which is that therefore only a very small sample of calls, e.g. a fraction of a percent, is typically evaluated. The third problem arises from the fact that most calls are ordinary and uninteresting; with random sampling, the human monitors spend most of their time listening to uninteresting calls.

This work describes an automated quality-monitoring system that addresses these problems. Automatic speech recognition is used to transcribe 100% of the calls coming in to a call center, and default quality scores are assigned based on features such as key-words, key-phrases, the number and type of hesitations, and the average silence durations. The default score is used to rank the calls from worst-to-best, and this sorted list is made available to the human evaluators, who can thus spend their time listening only to calls for which there is some a-priori reason to expect that there is something interesting.

The automatic quality-monitoring problem is interesting in part because of the variability in how hard it is to answer the questions. Some questions, for example, "Did the agent use courteous words and phrases?" are relatively straightforward to answer by looking for key words and phrases. Others, however, require essentially human-level knowledge to answer; for example one company's monitors are asked to answer the question "Did the agent take ownership of the problem?" Our work focuses on calls from IBM's North American call centers, where there is a set of 31 questions that are used to evaluate call-quality. Because of the high degree of variability found in these calls, we have investigated two approaches:

1. Use a partial score based only on the subset of questions that can be reliably answered.

2. Use a maximum entropy classifier to map directly from ASR-generated features to the probability that a call is bad (defined as belonging to the bottom 20% of calls).

We have found that both approaches are workable, and we present final results based on an interpolation between the two scores. These results indicate that for a fixed amount of listening effort, the number of bad calls that are identified approximately triples with our call-ranking approach. Surprisingly, while there has been significant previous scholarly research in automated call-routing and classification in the call center , e.g. [1, 2, 3, 4, 5], there has been much less in automated quality monitoring per se.

2. ASR FOR CALL CENTER TRANSCRIPTION

2.1. Data

The speech recognition systems were trained on approximately 300 hours of 6kHz, mono audio data collected at one of the IBM call centers located in Raleigh, NC. The audio was manually transcribed and speaker turns were explicitly marked in the word transcriptions but not the corresponding times. In order to detect speaker changes in the training data, we did a forced-alignment of the data and chopped it at speaker boundaries. The test set consists of 50 calls with 113 speakers totaling about 3 hours of speech.

2.2. Speaker Independent System

The raw acoustic features used for segmentation and recognition are perceptual linear prediction (PLP) features. The features are

Segmentation/clustering	Adaptation	WER
Manual	Off-line	30.2%
Manual	Incremental	31.3%
Manual	No Adaptation	35.9%
Automatic	Off-line	33.0%
Automatic	Incremental	35.1%

Table 1. ASR results depending on segmentation/clustering and adaptation type.

Accuracy	Top 20%	Bottom 20%
Random	20%	20%
QA	41%	30%

Table 2. Accuracy for the Question Answering system.

Accuracy	Top 20%	Bottom 20%
Random	20%	20%
ME	49%	36%

Table 3. Accuracy for the Maximum Entropy system.

Accuracy	Top 20%	Bottom 20%
Random	20%	20%
ME + QA	53%	44%

Table 4. Accuracy for the combined system.

mean-normalized 40-dimensional LDA+MLLT features. The SI acoustic model consists of 50K Gaussians trained with MPE and uses a quinphone cross-word acoustic context. The techniques are the same as those described in [6].

2.3. Incremental Speaker Adaptation

In the context of speaker-adaptive training, we use two forms of feature-space normalization: vocal tract length normalization (VTLN) and feature-space MLLR (fMLLR, also known as constrained MLLR) to produce canonical acoustic models in which some of the non-linguistic sources of speech variability have been reduced. To this canonical feature space, we then apply a discriminatively trained transform called fMPE [7]. The speaker adapted recognition model is trained in this resulting feature space using MPE.

We distinguish between two forms of adaptation: off-line and incremental adaptation. For the former, the transformations are computed per conversation-side using the full output of a speaker independent system. For the latter, the transformations are updated incrementally using the decoded output of the speaker adapted system up to the current time. The speaker adaptive transforms are then applied to the future sentences. The advantage of incremental adaptation is that it only requires a single decoding pass (as opposed to two passes for off-line adaptation) resulting in a decoding process which is twice as fast. In Table 1, we compare the performance of the two approaches. Most of the gain of full offline adaptation is retained in the incremental version.

2.3.1. Segmentation and Speaker Clustering

We use an HMM-based segmentation procedure for segmenting the audio into speech and non-speech prior to decoding. The reason is that we want to eliminate the non-speech segments in order to reduce the computational load during recognition. The speech segments are clustered together in order to identify segments coming from the same speaker which is crucial for speaker adaptation. The clustering is done via k-means, each segment being modeled by a single diagonal covariance Gaussian. The metric is given by the symmetric K-L divergence between two Gaussians. The impact of the automatic segmentation and clustering on the error rate is indicated in Table 1.

3. CALL RANKING

3.1. Question Answering

This section presents automated techniques for evaluating call quality. These techniques were developed using a training/development set of 676 calls with associated manually generated quality evaluations. The test set consists of 195 calls.

The quality of the service provided by the help-desk representatives is commonly assessed by having human monitors listen to a random sample of the calls and then fill in evaluation forms. The form for IBM's North American Help Desk contains 31 questions. A subset of the questions can be answered easily using automatic methods, among those the ones that check that the agent followed the guidelines e.g.

- Did the agent follow the appropriate closing script?
- Did the agent identify herself to the customer?

But some of the questions require human-level knowledge of the world to answer, e.g.

- Did the agent ask pertinent questions to gain clarity of the problem?
- Were all available resources used to solve the problem?

We were able to answer 21 out of the 31 questions using pattern matching techniques. For example, if the question is "Did the agent follow the appropriate closing script?", we search for "THANK YOU FOR CALLING", "ANYTHING ELSE" and "SERVICE REQUEST". Any of these is a good partial match for the full script, "Thank you for calling, is there anything else I can help you with before closing this service request?" Based on the answer to each of the 21 questions, we compute a score for each call and use it to rank them. We label a call in the test set as being *bad/good* if it has been placed in the bottom/top 20% by human evaluators. We report the accuracy of our scoring system on the test set by computing the number of *bad* calls that occur in the bottom 20% of our sorted list and the number of *good* calls found in the top 20% of our list. The accuracy numbers can be found in Table 2.

3.2. Maximum Entropy Ranking

Another alternative for scoring calls is to find arbitrary features in the speech recognition output that correlate with the outcome of a call being in the bottom 20% or not. The goal is to estimate the probability of a call being *bad* based on features extracted from the automatic transcription. To achieve this we build a maximum
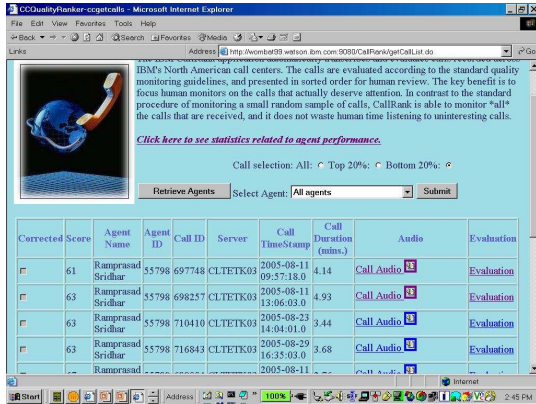
Fig. 1. Display of selected calls.

Fig. 2. Interface to listen to audio and update the evaluation form.

entropy based system which is trained on a set of calls with associated transcriptions and manual evaluations. The following equation is used to determine the score of a call C using a set of N predefined features:

$$P(class/C) = \frac{1}{Z}exp(\sum_{i=1}^{N} \lambda_i f_i(class, C)) \qquad (1)$$

where $class \in \{bad, not - bad\}$, Z is a normalizing factor, $f_i()$ are indicator functions and $\{\lambda_i\}_{\{i=1,N\}}$ are the parameters of the model estimated via iterative scaling [8].

Due to the fact that our training set contained under 700 calls, we used a hand-guided method for defining features. Specifically, we generated a list of VIP phrases as candidate features, e.g. "THANK YOU FOR CALLING", and "HELP YOU". We also created a pool of generic ASR features, e.g. "number of hesitations", "total silence duration", and "longest silence duration". A decision tree was then used to select the most relevant features and the threshold associated with each feature. The final set of features contained 5 generic features and 25 VIP phrases. If we take a look at the weights learned for different features, we can see that if a call has many hesitations and long silences then most likely the call is *bad*.

We use $P(bad|C)$ as shown in Equation 1 to rank all the calls. Table 3 shows the accuracy of this system for the bottom and top 20% of the test calls.

At this point we have two scoring mechanisms for each call: one that relies on answering a fixed number of evaluation questions and a more global one that looks across the entire call for hints. These two scores are both between 0 and 1, and therefore can be interpolated to generate one unique score. After optimizing the interpolation weights on a held-out set we obtained a slightly higher weight (0.6) for the maximum entropy model. It can be seen in Table 4 that the accuracy of the combined system is greater that the accuracy of each individual system, suggesting the complementarity of the two initial systems.

4. END-TO-END SYSTEM PERFORMANCE

4.1. Application

This section describes the user interface of the automated quality monitoring application. As explained in Section 1, the evalua-
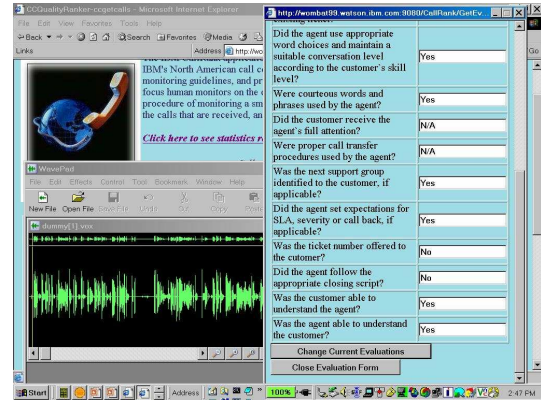
tor scores calls with respect to a set of quality-related questions after listening to the calls. To aid this process, the user interface provides an efficient mechanism for the human evaluator to select calls, e.g.

- All calls from a specific agent sorted by score
- The top 20% or the bottom 20% of the calls from a specific agent ranked by score
- The top 20% or the bottom 20% of all calls from all agents

The automated quality monitoring user interface is a J2EE web application that is supported by back-end databases and content management systems [1] The displayed list of calls provides a link to the audio, the automatically filled evaluation form, the overall score for this call, the agent's name, server location, call id, date and duration of the call (see Figure 1). This interface now gives the agent the ability to listen to interesting calls and update the answers in the evaluation form if necessary (audio and evaluation form illustrated in 2). In addition, this interface provides the evaluator with the ability to view summary statistics (average score) and additional information about the quality of the calls. The overall system is designed to automatically download calls from multiple locations on a daily-basis, transcribe and index them, thereby making them available to the supervisors for monitoring. Calls spanning a month are available at any given time for monitoring purposes.
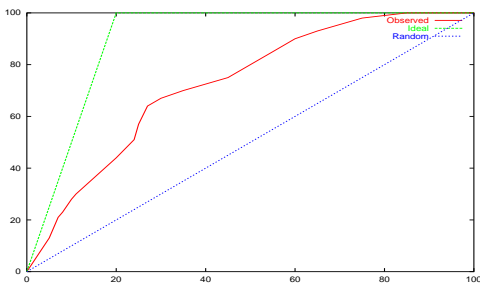
4.2. Precision and Recall

This section presents precision and recall numbers for the identification of "bad" calls. The test set consists of 195 calls that were manually evaluated by call center personnel. Based on these manual scores, the calls were ordered by quality, and the bottom 20% were deemed to be "bad." To retrieve calls for monitoring, we sort the calls based on the automatically assigned quality score and return the worst. In our summary figures, precision and recall are plotted as a function of the number of calls that are selected for monitoring. This is important because in reality only a small number of calls can receive human attention. Precision is the ratio

---

[1]In our case, the backend consists of DB2 and IBM's Websphere Information Integrator for Content and the application is hosted on Websphere 5.1.)
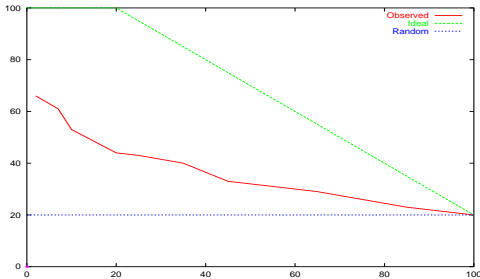
Fig. 3. Precision for the bottom 20% of the calls as a function of the number of calls retrieved.

Fig. 4. Recall for the bottom 20% of the calls.

Fig. 5. Ratio of bad calls found with QTM to Random selection as a function of the number of bad calls retrieved.

Fig. 6. Scatter plot of Human vs. Computer Rank.

of bad calls retrieved to the total number of calls monitored, and recall is the ratio of the number of bad calls retrieved to the total number of bad calls in the test set. Three curves are shown in each plot: the actually observed performance, performance of random selection, and oracle or ideal performance. Oracle performance shows what would happen if a perfect automatic ordering of the calls was achieved.

Figure 3 shows precision performance. We see that in the monitoring regime where only a small fraction of the calls are monitored, we achieve over 60% precision. (Further, if 20% of the calls are monitored, we still attain over 40% precision.)

Figure 4 shows the recall performance. In the regime of low-volume monitoring, the recall is midway between what could be achieved with an oracle, and the performance of random-selection.
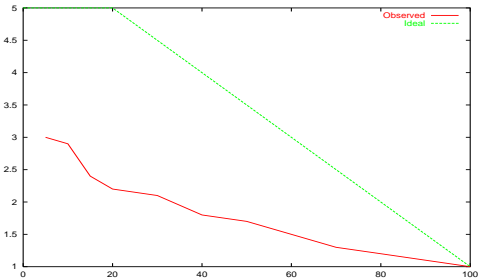
Figure 5 shows the ratio of the number of bad calls found with our automated ranking to the number found with random selection. This indicates that in the low-monitoring regime, our automated technique triples efficiency.
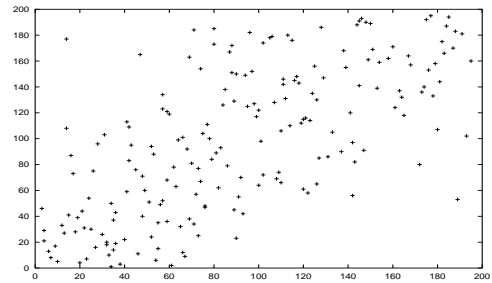
4.3. Human vs. Computer Rankings

As a final measure of performance, in Figure 6 we present a scatterplot comparing human to computer rankings. We do not have calls that are scored by two humans, so we cannot present a human-human scatterplot for comparison.

5. CONCLUSION

This paper has presented an automated system for quality monitoring in the call center. We propose a combination of maximum-entropy classification based on ASR-derived features, and question answering based on simple pattern-matching. The system can either be used to replace human monitors, or to make them more

efficient. Our results show that we can triple the efficiency of human monitors in the sense of identifying three times as many bad calls for the same amount of listening effort.

6. REFERENCES

[1] J. Chu-Carroll and B. Carpenter, "Vector-based natural language call routing," *Computational Linguistics*, 1999.

[2] P. Haffner, G. Tur, and J. Wright, "Optimizing svms for complex call classification," 2003.

[3] M. Tang, B. Pellom, and K. Hacioglu, "Call-type classification and unsupervised training for the call center domain," in *ARSU-2003*, 2003.

[4] D. Hakkani-Tur, G. Tur, M. Rahim, and G. Riccardi, "Unsupervised and active learning in automatic speech recognition for call classification," in *ICASSP-04*, 2004.

[5] C. Wu, J. Kuo, E.E. Jan, V. Goel, and D. Lubensky, "Improving end-to-end performance of call classification through data confusion reduction and model tolerance enhancement," in *Interspeech-05*, 2005.

[6] H. Soltau, B. Kingsbury, L. Mangu, D. Povey, G. Saon, and G. Zweig, "The ibm 2004 conversational telephony system for rich transcription," in *Eurospeech-2005*, 2005.

[7] D. Povey, B. Kingsbury, L. Mangu, G. Saon, H. Soltau, and G. Zweig, "fMPE: Discriminatively trained features for speech recognition," in *ICASSP-2005*, 2004.

[8] A. Berger, S. Della Pietra, and V. Della Pietra, "A maximum entropy approach to natural language processing," *Computational Linguistics*, vol. 22, no. 1, 1996.