CONTINUOUS SPEECH RECOGNITION USING MULTILAYER PERCEPTRONS WITH HIDDEN MARKOV MODELS

N. Morgan[†] **& H. Bourlard**[†,‡]

[†]International Computer Science Institute
1947 Center Street, Suite 600
Berkeley, CA 94704, USA

[‡]Philips Research Laboratory Brussels,
Av. Van Becelaere 2, Box 8,
B-1170 Brussels, Belgium

ABSTRACT

We are developing a phoneme based, speaker-dependent continuous speech recognition system embedding a Multilayer Perceptron (MLP) (i.e., a feedforward Artificial Neural Network), into a Hidden Markov Model (HMM) approach. Artificial neural networks such as the MLP have recently been applied to a number of subproblems in speech recognition [1][2][3]. Using the interpolative capabilities of the MLP, statistical pattern classification can be performed over an undersampled pattern space [5] without many restrictive simplifying assumptions (such as the independence of input features). By using contextual information from a sliding window on the input frames, we have been able to improve frame or phoneme classification performance over the corresponding performance for simple Maximum Likelihood probabilities, or even Maximum a Posteriori (MAP) probabilities which are estimated without the benefit of context. We have also learned how to use an MLP to generate HMM emission probabilities for continuous speech recognition of a 1000 word vocabulary using a vocabulary independent (phonetic) training. Performance for a simple discrete density HMM system appears to be somewhat better when MLP methods are used to estimate the probabilities.

INTRODUCTION

We have performed a number of experiments with a 1000-word vocabulary continuous speech recognition task. Our frame classification results [6] are consistent with other research showing the capabilities of MLPs trained with back-propagation-styled learning schemes for the recognition of voiced-unvoiced speech segments [7], isolated phonemes [2], [3], [4], or of isolated words [8]. These results indicate that "neural network" approaches can, for some problems, perform pattern classification at least as well as traditional HMM approaches. However, this is not particularly mysterious. When traditional statistical assumptions (distribution, independence of multiple features, etc.) are not valid, systems which do not rely on these assumptions can work better (as discussed in [5]). Furthermore, networks provide an easy way to incorporate multiple sources of evidence (multiple features, contextual windows, etc.) without restrictive assumptions.

However, it is not so easy to improve the recognition of words in continuous speech by the use of an MLP. For instance, while it has been shown that the outputs of a feedforward network can be used as emission probabilities in an HMM [6], the corresponding word recognition performance can be very poor. This is true even when the same network demonstrates extremely good performance at the frame or phoneme levels. We have developed a hybrid MLP-HMM algorithm which (for a preliminary experiment) appears to exceed performance of the same HMM system using standard statistical approaches to estimate the emission probabilities. This was only possible after the original algorithm was modified in ways that did not necessarily maximize the frame recognition performance for the training set. We will describe these modifications below, along with experimental results.

METHODS

As shown by both theoretical [10] and experimental [9] results, MLP output values may be considered to be estimates of MAP probabilities for pattern classification. Either these, or some other related quantity (such as the output normalized by the prior probability of the corresponding class) may be used in a Viterbi search to determine the best time-warped succession of states (speech sounds) to explain the observed speech measurements. This hybrid approach (MLP to estimate probabilities, HMM to incorporate them to recognize continuous speech as a succession of words) has the potential of exploiting the interpolating capabilities of MLPs while using a Dynamic Time Warping (DTW) procedure to capture the dynamics of speech.

However, to achieve good performance at the word level, the following modifications of this basic scheme were necessary:

(1) MLP training methods - a new cross-validation training algorithm was designed in which the stopping criterion was based on performance for an independent validation set [11]. In other words, training was stopped when performance on a second set of data began going down, and not when training error leveled off. This greatly improved generalization, which could be further tested on a third independent test set.

(2) probability estimation from the MLP outputs - In the original scheme [10], MLP outputs were used as MAP probabilities for the HMM directly. While this helped frame performance, it hurt word performance. This may have been due at least partly to a mismatch between the relative frequency of phonemes in the training sets and test (word recognition) sets. Division by the prior class probabilities as estimated from the training set removed this effect of the priors on the DTW. This led to a small decrease in frame classification performance, but a large (sometimes 10 - 20%) improvement in word recognition rates (see Table II and accompanying description).

(3) word transition costs for the underlying HMM - word transition penalties had to be increased for larger contextual windows to avoid a large number of insertions. This was shown to be equivalent to keeping the same word transition cost but scaling the log probabilities down by a number which reflected the dependence of neighboring frames. A reasonable value for this can be determined from recognition on a small number of sentences (e.g., 50), choosing a value which results in insertions at most

equal to the number of deletions.

(4) segmentation of training data - much as with HMM systems, an iterative procedure was required to time align the training labels in a manner that was statistically consistent with the recognition methods used. In our most recent experiments, we segmented the data using an iterative Viterbi alignment starting from a segmentation based on average phoneme durations, and terminated at the segmentation which led to the best performance on an independent test set. For one of our speakers, we had available a more accurate frame labeling (produced by an automatic but more complex procedure [13]) to use as a start point for the iteration, which led to even better performance. These techniques improved our concomitant word recognition to be better than we could achieve either with the earlier MLP-HMM technique or the pure HMM technique (using, for both cases, HMMs with a single distribution for each phoneme, and a single vector-quantized feature).

EXPERIMENTAL APPROACH

We have been using a speaker-dependent German database (available from our collaboration with Philips) called SPICOS [12], The speech had been sampled at a rate of 16 kHz, and 30 points of smoothed, "mel-scaled" logarithmic spectra (over bands from 200 to 6400 Hz) were calculated every 10-ms from a 512-point FFT over a 25-ms window. For our experiments, the mel spectrum and the energy were vector-quantized to pointers into a single speaker-dependent table of prototypes.

Two independent sets of vocabularies for training and test are used. The training data-set consists of two sessions of 100 German sentences per speaker. These sentences are representative of the phoneme distribution in the German language and include 2430 phonemes in each session. The two sessions of 100 sentences are phonetically segmented on the basis of 50 phonemes, using a fully automated procedure [13]. The test set consists of one session of 200 sentences per speaker. The recognition vocabulary contains 918 words (including the "silence" word) and the overlap between training and recognition is 51 words. Most of the latter are articles, prepositions and other structural words. Thus, the training and test are essentially vocabulary-independent. Initial tests used sentences from a single male speaker. The final algorithms were tested on an additional male and female speaker.

The acoustic vectors were coded on the basis of 132 prototype vectors by a simple binary input layer with only one bit 'on'. Multiple frames were used as input to provide context to the network. In the experiments reported here, the context was 9 frames. When used, a hidden layer was varied from 20 to 200 units, while the size of the output layer was kept fixed at 50 units, corresponding to the 50 phonemes to be recognized. The input field contained $9\times132=1188$ units, and the total number of possible inputs was thus equal to 132^9. There were 26767 training patterns (from the first training session of 100 sentences) and 26702 independent test patterns (from the second training session of 100 sentences). Of course, this represented only a very small fraction of the possible inputs, and generalization was thus potentially difficult. Training was done by the classical "error-back propagation" algorithm, starting by minimizing an entropy criterion, and then the standard least-mean-square error criterion. In each iteration, the complete training set was presented, and the parameters were updated after each training pattern. To avoid overtraining of the MLP, improvement on the test set was checked after each iteration. If the classification rate on the test set was decreasing, the adaptation parameter of the gradient procedure was decreased, otherwise it was kept constant. After several reductions of learning rate, performance on the test set ceased to improve and training was stopped. In another experiment this approach was systematized by splitting the data in three parts: one for the training, one for the test and a third one absolutely independent of the training procedure for validation. No significant difference was observed between classification rates for the test and validation data.

The important idea in this procedure was that we stopped iterating by any one particular criterion when that criterion was leading to no new test set performance. This appeared to ameliorate the effects of over-fitting that had been observed in our earlier experiments, and greatly improved classification for frames of continuous speech.

The output layer of the MLP was evaluated for each frame, and (after division by the prior probability of each phoneme) was used as the emission probability in a discrete HMM system. In this system, each phoneme was modeled with a single conditional density, repeated $D/2$ times, where D was a prior estimate of the duration of the phoneme. Only selfloops and sequential transitions were permitted. A Viterbi decoding was then used for recognition of the first hundred sentences of the test session (on which word entrance penalties were optimized), and our best results were validated by a further recognition on the second hundred sentences of the test set. Note that this same simplified HMM was used for both the ML reference system (estimating probabilities directly from relative frequencies) and the MLP system, and that the same input features were used for both.

FRAME CLASSIFICATION RESULTS

Table I shows the frame classification performance for 5, 20, 50, and 200 hidden units. The peak at 20 hidden units for test set performance, in contrast to the continued improvement in training set performance, can be clearly seen. However, this effect is mild given the wide range in network size; using 10 times the number of weights as in the "peak" case only causes a degradation of 3.1%. We attribute this to the generalization-based training procedure in which training was halted when test set performance started to degrade.

Table I also shows results obtained with Maximum Likelihood (ML) and Maximum a Posteriori (MAP) estimates. In those cases, it is not possible to use contextual information, because the number of parameters to be learned would be 50×132^9 for the 9 frames of context. Therefore, the input field was restricted to a single frame. The number of parameters for these two last classifiers was then $50 \times 132 = 6600$, a small enough number of parameters to be estimated by relative frequencies of occurrence. This restriction explains why the MAP (also called Bayes) classifier, which is inherently optimal for a given pattern classification problem, is shown here as yielding a lower performance than the potentially suboptimal MLP.

WORD RECOGNITION RESULTS

Table II shows the recognition rate (100% – error rate, where errors includes insertions, deletions, and substitutions) for the first 100 sentences of the test session. All runs except the last were done with the 20 hidden units in the MLP, as suggested by the results above. Note the significant positive effect of division of the MLP outputs, which are trained to approximate MAP probabilities, by estimates of the prior probabilities for each class (denoted "MLP/priors" in Table II). Not shown here are the earlier improvements required to reach this level of performance, which were primarily the modifications to the learning algorithm described above. Additionally, word transition probabilities were optimized for both the Maximum Likelihood and MLP style HMMs. This led to a word exit probability of 10^{-8} for the ML and for 1-frame MLP's, and 10^{-14} for an MLP with 9 frames of context. After these adjustments, performance was essentially the same for the two approaches. Performance on the last hundred sentence of the test session (shown in the last column of Table II) validated that the two systems generalized equivalently despite these tunings.

An initial time alignment of the phonetic transcription with the data (for this speaker) had previously been calculated using a program incorporating speech-specific knowledge [13]. This labeling had been used for the targets of the frame-based training described above. We then used this alignment as a "bootstrap" segmentation for an iterative Viterbi procedure, much as is done in conventional HMM systems. As with the MLP training, the data was divided into a training and cross-validation set, and the segmentation corresponding to the best validation set frame classification rate was used for later training. For both cross-validation procedures, we switched to a training set of 150 sentences (two repetitions of 75 sentences) and a cross-validation set of 50 sentences (two repetitions of 25 each). Finally, since the best performance in Table II was achieved using no hidden layer, we continued our experiments using this simpler network, which also required only a simple training procedure (entropy error criterion only). Table III shows this performance for the full 200 recognition sentences (test + validation sets for Table II).

Finally, we duplicated the resulting analysis steps for two other speakers from the same data base. In this case, we used the first 50 recognition sentences to optimize the word entrance penalties separately for each method and speaker. Boostrap segmentations were unavailable for these speakers, so we labeled each training set (from the original male plus a male and a female speaker) using a Viterbi iteration initialized from a time-alignment based on a simple estimate of average phoneme duration. This reduced all of the recognition scores, underlining the necessity of a good start point for the Viterbi iteration. However, as can be seen from the Table IV results (measured over the full 200 recognition sentences), the MLP-based methods appear to consistently offer at least some measurable improvement over the simpler estimation technique. In particular, the performance for the two systems differed significantly ($p < .001$) for two out of three speakers, as well as for a multispeaker comparison over the three speakers (in each case using a normal approximation to a binomial distribution for the null hypothesis).

DISCUSSION

These results (all obtained with no language model, i.e., with a perplexity of 918 for a 918 word vocabulary) show some of the improvement for MLPs over conventional HMMs which one might expect from the frame level results (Table I). MLPs can sometimes make better frame level discriminations than simple statistical classifiers, because they can easily incorporate multiple sources of evidence (multiple frames, multiple features), which is difficult to do in HMMs without major simplifying assumptions. In general, the relation between the MLP and ML word recognition is more complex, because of interdependence over time of the input features. Part of the difficulty with good recognition may be due to our choice of discrete, vector-quantized features, for which no metric is defined over the prototype space. Despite these limitations, it now appears that the probabilities estimated by MLPs may offer improved word recognition through the incorporation of context in the estimation of emission probabilities. Furthermore, our new result shows the effectiveness of Viterbi segmentation in labeling training data for an MLP. This result appears to remove a major handicap of MLP use, the requirement for hand-labeled speech.

Interestingly, our best results were obtained using an MLP with no hidden layer. This suggests that, for the case of a single VQ feature, a single Perceptron model is rich enough for the probabilistic estimation. This network can also be trained more easily than networks with one or more hidden layers, particularly when an entropy criterion is used. Furthermore, although one of us has shown [10] that the MLP training can be embedded in the Viterbi segmentation process, these encouraging results were obtained using simple discrete ML estimation only within each iteration.

The features we have been using were chosen for their effectiveness in HMM systems, and different combinations may prove to be better for MLP inputs. In particular, we would expect that feature combinations which have not been vector-quantized should have more useful dependencies (both within-frame and over time) which the MLP may be able to learn and exploit. We intend to explore this possibility in the coming year.

Finally, is the MLP simply accomplishing a nice interpolation for the joint density estimates we seek? Perhaps so, and other smoothing techniques (kernal estimators, for instance), may work as well. Nonetheless, MLP approaches appear to offer a reasonable way of incorporating information from multiple sources of phonetic evidence into a continuous speech recognizer.

ACKNOWLEDGMENTS

Support from the International Computer Science Institute (ICSI) and Philips Research for this work is gratefully acknowledged. Chuck Wooters of ICSI and UCB provided much-needed assistance, and Xavier Aubert of Philips put together our Spicos materials.

Table I — Frame Classification
speaker m003

number of hidden units	% correct training	test
5	62.8	54.2
20	75.7	62.7
50	73.7	60.6
200	86.7	59.6
ML	45.9	44.8
MAP	53.8	53.0

Table II — Word Recognition - 1
speaker m003

system method	size of context	% correct test	validation
MLP	1	27.3	
MLP/priors	1	49.7	
MLP	9	40.9	
MLP/priors	9	51.9	52.2
ML	1	52.6	52.5
MLP/priors (0 hidden)	9	53.3	

Table III — Word Recognition - 2
speaker m003

method	context	test
MLP/priors (0 hidden)	9	65.3
ML	1	56.9

Table IV — Word Recognition - 3
3 speakers

speaker	MLE	MLP
m003	54.4	59.7
m001	47.4	51.9
w010	54.2	54.3

REFERENCES

[1] Franzini, M., Witbrock, M., and Lee, K.. "A Connectionist Approach to Speech Recognition", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 1, pp. 425-428, Glasgow, Scotland, 1989.

[2] Watrous, R.L. & Shastri, L.. "Learning phonetic features using connectionist networks: an experiment in speech recognition", *Proceedings of the First Intl. Conference on Neural Networks*, IV-381-388, San Diego, CA, 1987

[3] Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., and Lang, K.. "Phoneme Recognition: Neural Networks vs. Hidden Markov Models", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 1, pp. 107-110, New York, 1988.

[4] Makino, S., Kawabata, T., and Kido, K.. "Recognition of consonants based on the Perceptron Model", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 2, pp. 738-741, Boston, Mass., 1983

[5] Niles, L., Silverman, H., Tajchman, G., and Bush, M.. "How Limited Training Data Can Allow a Neural Network Classifier to Outperform an 'Optimal' Statistical Classifier", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 1, pp. 17-20, Glasgow, Scotland, 1989

[6] Bourlard, H., Morgan, N., and Wellekens, C.J.. "Statistical Inference in Multilayer Perceptrons and Hidden Markov Models with Applications in Continuous Speech Recognition", to appear in *Neuro Computing, Algorithms, and Applications*, NATO ASI Series, 1989

[7] Gevins, A., and Morgan, N.. "Ignorance-Based Systems", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 3, 39A5.1-39A5.4, San Diego, 1984

[8] Peeling, S.M., Moore, R.K.. "Experiments in Isolated Digit Recognition Using the Multi-Layer Perceptron", Royal Speech and Radar Establishment, Technical Report 4073, Malvern, Worcester, 1988

[9] Bourlard, H., and Morgan, N.. "Merging Multilayer Perceptrons and Hidden Markov Models: Some Experiments in Continuous Speech Recognition" International Computer Science Institute TR-89-033, 1989

[10] Bourlard, H., and Wellekens, C.J.. "Speech Pattern Discrimination and Multilayer Perceptrons", in *Computer, Speech, and Language*, vol. 3, pp.1-19, 1989

[11] Morgan, N., and Bourlard, H.. "Generalization and Parameter Estimation in Feedforward Nets: Some Experiments", International Computer Science Institute TR-89-017, 1989, also to be published in *Advances in Neural Information Processing Systems II*, Morgan Kaufmann, 1990

[12] Ney, H., and Noll, A.. "Phoneme Modeling Using Continuous Mixture Densities", *Proc. IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, Vol. 1, pp. 437-440, New York, 1988.

[13] Aubert, X.. "Supervised Segmentation with Application to Speech Recognition", in *Proc. Eur. Conf. Speech Technology*, Edinburgh, 1987, p.161-164