

Annual Report for Period:09/2007 - 08/2008

Submitted on: 08/31/2008

Principal Investigator: Picone, Joseph .

Award ID: 0414450

Organization: Mississippi State Univ

Submitted By:

Picone, Joseph - Principal Investigator

Title:

Nonlinear Statistical Modeling of Speech

Project Participants

Senior Personnel

Name: Picone, Joseph

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Lazarou, Georgios

Worked for more than 160 Hours: Yes

Contribution to Project:

Professor Lazarou has assumed supervision of the students while Dr. Picone is on a sabbatical.

Post-doc

Graduate Student

Name: Patil, Sanjay

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr. Patil is leading our research into Lyapunov exponent estimation and particle filtering.

Name: Raghavan, Sridhar

Worked for more than 160 Hours: Yes

Contribution to Project:

Mr. Raghavan is responsible for development of the baseline speaker recognition system that will be used to assess our research. He is also improving our Support Vector Machine and Relevance Vector Machine models that are used as alternative statistical estimators. These will provide important comparison points for the nonlinear estimators.

Name: Prasad, Saurabh

Worked for more than 160 Hours: No

Contribution to Project:

Mr. Prasad joined the project in Fall'2005, and is leading our research on Kalman filters and unscented Kalman filters. These are being used in our particle filter implementation.

Name: Pannuri, Madhulika

Worked for more than 160 Hours: No

Contribution to Project:

Ms. Pannuri joined the project in Fall'2005 as an entry-level MS student. She is currently providing programming support to the project as she develops her background in this research.

Name: Srinivasan, Sundararajan

Worked for more than 160 Hours: No

Contribution to Project:

Mr. Srinivasan joined the project in Fall'2005, and is developing more robust ways to estimate Lyapunov exponents, as well as providing programming support while he gets up to speed on the research.

Name: May, Daniel

Worked for more than 160 Hours: Yes

Contribution to Project:

Daniel May is now a graduate student pursuing his MS in Computer Engineering. He has been developing new statistical models for features based on nonlinear measures.

=====

Daniel May joined the project in Fall'2005 as a senior undergraduate. He is assisting the release of our software, including Java applets developed on an NSF REU extension to this project.

Name: Ma, Tao

Worked for more than 160 Hours: Yes

Contribution to Project:

Speech researcher - initially responsible for maintaining and upgrading our experimental infrastructure

Undergraduate Student

Name: Irwin, Ryan

Worked for more than 160 Hours: Yes

Contribution to Project:

Ryan Irwin joined the project in Summer'2005 as an undergraduate working on the NSF REU extension. He has extended our Java-based pattern recognition applet to include Kalman and particle filtering.

Name: Holland, Wesley

Worked for more than 160 Hours: Yes

Contribution to Project:

Wesley Holland joined the project in Summer'2005 as an undergraduate working on the NSF REU extension. He has developed software that allows our speech recognition system to support XML grammar formats. This includes software to manipulate and transform context-free and regular grammars.

Technician, Programmer

Other Participant

Research Experience for Undergraduates

Organizational Partners

Department of Defense

DoD is a partial sponsor of this work. I am also currently working on a sabbatical with the group that sponsored the work. We have a close working relationship with them and transfer software from MS State to DoD on an as-needed basis. We try to address DoD needs that arise that are relevant to the project. Two examples of such work are confidence measures and lattice rescoring. We delivered some Perl tools to DoD for these computations, and also released these as part of our public domain software.

Other Collaborators or Contacts

We support a number of collaborators through the distribution of our public domain software.

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)

Findings:

Convergence of Kalman filtering, particle filtering, and other such iterative algorithms is very sensitive to prior knowledge. This is typically why these techniques don't work well for robust speech processing applications.

Our initial SVM-based speaker recognition system provides a small improvement over the GMM baseline, a finding that is consistent with what others in the community have found.

Our attempts at adding several nonlinear features independently have not yet provided an improvement in performance. We have certified our implementations against previously published work, so our results conflict with previously published results. We continue to explore and analyze these experiments.

Demonstrated a small improvement in feature analysis using a combination of standard features and nonlinear features.

Demonstrated a small, but consistent improvement on a pilot database for a new acoustic modeling approach based on a probabilistic mixed autoregressive model.

Demonstrated a small, but consistent improvement in speech recognition performance using a linear dynamic model.

Training and Development:

Several of our students have improved their technical writing skills through publications and documentation related to this project.

All students are learning a strict software engineering process we use in our lab, and have increased their knowledge considerably.

All graduate students have received special training through a graduate level course in Natural Language Processing that we were able to offer in conjunction with this project.

Introduced the students to a Java-based interactive development environment, Eclipse, that is rapidly becoming an industry-standard environment.

Outreach Activities:

We have hosted several open houses for a local high school that specializes in mathematics and sciences: The Mississippi School for Mathematics and Science. It is one of the best math and science schools in the country and attracts the best students from all over Mississippi.

We also presented a lecture at the Math Club at the same institution. This presentation emphasized the value of math in the engineering disciplines.

Journal Publications

S. Prasad, S. Srinivasan, M. Pannuri, G. Lazarou and J. Picone, "Nonlinear Dynamical Invariants for Speech Recognition", International Conference on Spoken Language Processing, p. 2518, vol. 1, (2006). Published,

S. Prasad, S. Srinivasan and J. Picone, "Reconstructed Phase Space of a Vector Time Series", 14th European Signal Processing Conference, p. 1, vol. 1, (2006). Accepted, but declined trip because the lead author dropped out,

S. Srinivasan, S. Prasad, S. Patil, G. Lazarou and J. Picone, "Estimation of Lyapunov Spectra From a Time Series", Proceedings of IEEE SoutheastCon, p. 192, vol. 1, (2006). Published,

S. Patil, S. Srinivasan, S. Prasad, R. Irwin, G. Lazarou and J. Picone, "Sequential State-Space Filters for Speech Enhancement", Proceedings of IEEE SoutheastCon, p. 240, vol. 1, (2006). Published,

S. Srinivasan, T. Ma, D. May, G. Lazarou and J. Picone, "Nonlinear Mixture Autoregressive Hidden Markov Models For Speech Recognition", Proceedings of INTERSPEECH, p. , vol. , (2008). Accepted,

D. May, T. Ma, S. Srinivasan, G. Lazarou and J. Picone, "Continuous Speech Recognition Using Nonlinear Dynamic Invariants", Proceedings of INTERSPEECH, p. , vol. , (2008). Rejected, but available from our project web site,

T. Ma, S. Srinivasan, D. May, G. Lazarou and J. Picone, "Robust Speech Recognition Using Linear Dynamic Models", Proceedings of INTERSPEECH, p. , vol. , (2008). Rejected, but available from our project web site,

Books or Other One-time Publications

J. Picone, "Risk Minimization Approaches in Speech Recognition", (2006). Book, Published
Editor(s): Idea Group Inc., USA

Collection: Kernel Methods in bioengineering, communications, and signal processing

Bibliography: G. Camps-Valls, J.L. Rojo-Álvarez, and M. Martínez-Ramón, "Kernel Methods in bioengineering, communications, and signal processing," The Idea Group, Inc., Hershey Pennsylvania, USA

J. Picone, "Statistical Optimization in Speech Recognition", (2009). research monograph, Published
Bibliography: under development

D. May, "Nonlinear Dynamic Invariants For Continuous Speech Recognition", (2008). Thesis, Published

Bibliography: D. May, Nonlinear Dynamic Invariants For Continuous Speech Recognition, M.S. Thesis, Department of Electrical and Computer Engineering, Mississippi State University, May 2008.

S. Srinivasan, "Nonlinear Mixture Autoregressive Hidden Markov Models For Speech Recognition", (2009). Thesis, under development

Bibliography: MS State Dissertation

T. Ma, "Robust Speech Recognition Using Linear Dynamic Models", (2009). Thesis, under development

Bibliography: MS State dissertation

Web/Internet Site

URL(s):

http://www.ece.msstate.edu/research/isip/project/nsf_nonlinear

Description:

We always maintain a web site for every project we execute. This web site includes software, data, publications, etc., related to the project.

Other Specific Products

Product Type:

Teaching aids

Product Description:

A tutorial on particle filtering.

Sharing Information:

This work is disseminated via a URL:

http://www.cavs.msstate.edu/hse/ies/whats_new/2005_06/

Product Type:

Teaching aids

Product Description:

We have developed a number of tutorials on key core technologies for this project.

Sharing Information:

These tutorials are available at:

http://www.cavs.msstate.edu/hse/ies/projects/nsf_nonlinear/doc/

Product Type:**Software (or netware)****Product Description:**

We have augmented our pattern recognition applet to include three time series analysis techniques: linear prediction, Kalman filtering, and Particle filtering.

Sharing Information:

The applet is available at:

http://www.cavs.msstate.edu/hse/ies/projects/speech/software/demonstrations/applets/util/pattern_recognition/current/index.html

Product Type:**Software (or netware)****Product Description:**

We have developed baseline implementations of several techniques for estimating nonlinearities in a signal. These are provided in MATLAB and used as reference implementations for our C++ code.

Sharing Information:

This software is located at:

http://www.cavs.msstate.edu/hse/ies/projects/nsf_nonlinear/downloads/software/matlab/

Contributions**Contributions within Discipline:**

We have replicated previously published work on particle filtering, Kalman filtering, and Lyapunov exponent estimation. We have applied these techniques to more comprehensive speech databases as a first step in characterizing their performance on a large-scale application.

We have provided reference implementations of Lyapunov exponents, Correlation Dimensions, and the Embedding Dimension for use in speech recognition experiments.

We have evaluated these features on both speaker recognition and speech recognition tasks.

We demonstrated that a reference implementation of a linear dynamic model for acoustic modeling of speech recognition did not give an improvement in performance on recognition and verification tasks of scale.

We combine traditional MFCCs with nonlinear dynamic invariants and demonstrated a relative improvement of 11.1% on the clean evaluation set for Aurora 4. However, an average relative decrease of 7.6% was observed for the data sets containing noise.

We demonstrated a linear dynamic model that explicitly models the evolution of hidden states using an autoregressive process. This smoothed trajectory model provided a 4.9% relative improvement on the Aurora-4 clean evaluation set, and a 6.5% relative improvement on the noisy evaluation set.

A nonlinear mixture autoregressive model was used to model state output distributions (MAR-HMM) in an HMM-based speech recognition system, and was shown to provide superior performance to a comparable Gaussian mixture model-based HMMs (GMM-HMM) with lower complexity on two pilot classification tasks. This ongoing work is now being evaluated on a much larger-scale task.

Contributions to Other Disciplines:

We have made software available as part of our public domain system, and also released a number of tutorials on our web site.

Contributions to Human Resource Development:

We have introduced two undergraduate students to speech research. Both are showing great promise and plan to pursue graduate studies.

A third undergraduate will enter graduate school in Spring'2006 and continue working on this project.

Contributions to Resources for Research and Education:

A project web site has been developed and maintained.

Contributions Beyond Science and Engineering:

Special Requirements

Special reporting requirements: None

Change in Objectives or Scope: None

Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

Contributions: To Any Beyond Science and Engineering

09/30/07 — 08/31/08: RESEARCH ACTIVITIES

The primary goal of this project is to develop novel nonlinear modeling techniques for speech and speaker recognition systems. In previous years of this project, we have explored various classical nonlinear models with limited success. This year we focused on two approaches, a nonlinear mixture autoregressive model and a linear dynamic model. We also completed experiments on adding nonlinear features to the acoustic feature vector.

An overarching challenge in this project has been dealing with the nonstationarity of the speech signal. We have explored several models that, if given ample amounts of data, appear capable of modeling nonlinear behavior. The parameter estimates of these models converge over time scales of seconds. However, the speech signal varies on the order of 10 to 30 msec. Some phonemes crucial to good recognition and verification performance last only a few tens of milliseconds, on the order of a few frames of data. We have yet to find success at estimating parameters of these nonlinear models over such short durations of speech. In such cases, one only has a few hundred samples with which to compute the model parameters. Worse, coarticulation phenomena often impinge on these samples, making it very difficult to estimate phonemes with short duration. This is one of the great challenges of speech processing – estimating model parameters from a signal that is evolving on very short time scales. We even attempted to oversample the signal in an effort to increase the amount of data, but even this approach suffers from some well-known drawbacks.

One of the reasons we shifted the focus of the original proposal from recognition to verification was to provide a longer time epoch over which parameters could be estimated. Unfortunately, though in verification you have access to a long speech utterance, the fact that the signal model is continually evolving during this time makes it difficult to separate out speaker variations from phoneme variations. Nonlinear models attempting to directly estimate long-term suprasegmental parameters related to the speaker's identity also did not perform well.

Therefore, in the final year of this project we focused on approaches that blend well with the traditional, piecewise linear approximation that has been so successful in speech recognition. We also attempted to leverage as much of the hidden Markov model infrastructure as possible. In this report, we summarize findings of the three approaches mentioned above. One M.S. thesis is available that describes the feature extraction work in more detail; two PhD theses are under development that describe the other work in greater detail. All are available from our web site while under construction in an effort to disseminate the project information as quickly as possible.

A. Continuous Speech Recognition with Nonlinear Dynamic Invariants

This year we completed our work on nonlinear dynamic invariants, culminating with the publication of an M.S. thesis [1] and submissions to several conferences including INTERSPEECH 2008 [2][4]. We have previously reported our general approach was to combine traditional MFCCs with nonlinear dynamic invariants in an effort to produce a more robust feature vector for continuous speech recognition.

We continued our analysis of three standard dynamic invariants: Lyapunov exponents, fractal dimension, and Kolmogorov entropy. Lyapunov exponents associated with a trajectory provide a measure of the average rates of convergence and divergence of nearby trajectories. Fractal dimension is a measure that quantifies the number of degrees of freedom and the extent of self-similarity in the attractor's structure. Kolmogorov entropy defined over a state-space, measures the rate of information loss or gain over the trajectory. These measures search for a signature of chaos in the observed time series. Since these measures quantify the structure of the underlying nonlinear dynamic system, they are prime candidates for feature extraction of a signal with strong nonlinearities. The motivation behind studying

such invariants from a signal processing perspective is to capture the relevant nonlinear dynamic information from the time series – something that is ignored in conventional spectral based analysis.

Our preliminary experiments provided strong support that the addition of these nonlinear invariants the standard MFCC feature vector will improve the accuracy of speech recognition tasks. We next evaluated these features on two sets of continuous speech recognition experiments involving the Aurora 4 task (a variant of the WAJ 5K task).

The recognition results for the clean test set were very encouraging. Each of the MFCC/invariant feature combinations resulted in a significant recognition performance increases over the baseline MFCC experiments. Correlation entropy resulted in the largest relative improvement of 11.1%. The recognition results for the noisy test sets were less encouraging as each experiment resulted in a performance decrease (higher error rate) compared to the baseline. These results contradict our theory that the addition of invariants would result in a feature vector that is more robust to noisy conditions unseen in the training set. We also examined more closely some state-space filtering methods which we hoped would enhance the algorithms' robustness to noise. These also did not result in any improvements. This work is described in greater detail in [1].

B. MixAR Modeling of Speech Signals for Speaker Recognition

Gaussian mixture models are a very successful method for modeling the output distribution of a state in a hidden Markov model (HMM). However, this approach is limited by the assumption that the dynamics of speech features are linear and can be modeled with static features and their derivatives. In this work, a nonlinear mixture autoregressive model is used to model state output distributions (MAR-HMM). This model can handle both static and dynamic features.

The MAR-HMM model we have developed is a generalized version of [5] that has been extended to handle vector observations, so that we can operate on the speech feature vector stream rather than speech samples. One property of MAR that is of particular relevance is the ability of MAR to model nonlinearity in time series. Though the individual component AR processes are linear, the probabilistic mixing of these AR processes constitutes a nonlinear model. In a GMM, the distribution remains invariant to the past samples due to the static nature of the model. For MAR, the conditional distribution given past data varies with time. This model is capable of modeling both the conditional means and variances. Thus, MAR can model time series that evolve nonlinearly.

To better understand the efficacy of the MAR-HMM model, we evaluated its performance on two simple pattern recognition tasks. The first task represents data with known nonlinearities. The MAR, with just 2 components and 8 parameters, achieved 100% classification accuracy using only static features. The GMM approach using only static features was unable to do much better than a random guess strategy since the two classes have similar static marginal distribution. With the inclusion of delta coefficients, the GMM performance increases significantly, but even in this case it achieved only 85% accuracy with 28 parameters. Though delta features capture some amount of dynamic information in the features, it is still only a linear approximation, and we cannot capture their nonlinear evolution with just GMMs.

The second task was a simple phone classification task. To test the efficacy of MAR-HMMs in speech modeling, we made 16 kHz recordings of three distinct phones – “aa” (vowel), “m” (nasal), and “sh” (sibilant). For each phone and for each speaker, 35 recordings were made to serve as training database, while another 15 were reserved for testing. Silence was removed so that we could focus on the ability of the approach to model speech.

We evaluated the performances of 2-, 4-, 8-, and 16 mixture GMM-HMM and MAR-HMM with the 13 dimensional static MFCC features. For an equal number of parameters, MAR outperformed GMM significantly. For instance, MAR-HMM achieved a phone classification accuracy of 94.4% with only 320 parameters while a GMM system using 432 parameters could only achieve a 93.3%.

To determine whether MAR is more effective at exploiting dynamics than what GMM can achieve using dynamic features, we also performed another experiment with 39-dimensional features containing both static as well as velocity and acceleration coefficients. In this case, the results were not conclusive. While MAR-HMM showed an accuracy rate of 97.8% with 472 parameters, GMM-HMM attained only 96.7% accuracy with 632 parameters.

Unfortunately, the performance of MAR-HMM saturated with an increase in the number of parameters. For example, MAR-HMM at 1888 parameters achieved only 98.9% accuracy while GMM-HMM achieved 100% with 1264 parameters. We suspect that this could be due to the fact that our parameter estimation and likelihood computation procedures assume that the features are independent. It is well-known that the static MFCC features are uncorrelated (at least, theoretically), but obviously the delta features are correlated with the static ones. While this should also cause problems for GMM, the problem is more acute for MAR because in this case, unlike GMMs, we employ the past history explicitly.

This work is described in detail in a paper that will be presented at INTERSPEECH'2008 [6].

C. Linear Dynamic Models for Speech Recognition

Last year we proposed Linear Dynamic Models (LDMs) as an alternative to hidden Markov models (HMMs) for robust speech recognition in noisy environments. HMMs in speech recognition typically utilize a diagonal covariance matrix assumption in which correlations between feature vectors for adjacent frames are ignored. LDMs use a state space-like formulation that explicitly models the evolution of hidden states using an autoregressive process. This smoothed trajectory model allows the system to better track speech dynamics in noisy environments.

This year we refined this model and completed its evaluation on a complex speech recognition task – the Aurora 4 task. This task is attractive because it contains speech degraded by a variety of noise conditions (digitally) and it removes the language model problem by using a closed-set 5K language model. An HMM baseline system was used to generate alignments at the phone level. Each phone instance was treated as one segment. A total of 40 LDM phone models, one classifier per model, were used to cover the pronunciations. Each classifier was trained using the segmental features derived from 13-dimensional frame-level feature vectors comprised of 12 cepstral coefficients and absolute energy. The full training set has as many as 30k training examples per classifier. Each phone-level classifier was trained as a one-vs-all classifier. The classifiers were used to predict the probability of an acoustic segment.

The HMM system achieves up to 46.9% and 36.8% accuracy for the clean evaluation data and noisy evaluation data respectively. The LDM classifiers achieved superior performance to the HMM classifiers with a classification accuracy of 49.2% for the clean evaluation data and 39.2% for the noisy evaluation data. This represents a 4.9% relative and a 6.5% relative increase in performance over a comparable HMM system with 3-state models.

We are currently developing a HMM/LDM hybrid decoder architecture to model the frame correlation using LDMs as well as utilizing HMMs techniques for phone segment alignment. Preliminary experiments will be presented on the Alphadigits (AD) and Resource Management (RM) speech corpora.

This HMM/LDM hybrid decoder architecture will be a good evaluation of LDMs on continuous speech recognition tasks, and can be compared to other hybrid decoders we have developed that utilize other nonlinear statistical models (e.g., support vector machines and relevance vector machines).

D. Additional Comments

Progress in the final year of this project was slowed somewhat because the two PhD graduate students working on the project had opportunities to do internships. Mr. Sundar Srinivasan spent Summer'2008 at Motorola's research lab working on rapid adaptation techniques in speech recognition. Mr. Tao Mao spent the summer with Intel's speech research group working on methods for speeding up speech recognition on Intel processors. He will continue there through December 2008. These students were identified by these employers based on their work on this project. For each of them, it was their first chance to do an internship during their graduate studies, so the PIs felt it was worthwhile to let them pursue these opportunities. Each had a very productive summer.

As a result, we applied and received an extension of the project through December 2008. This will allow Mr. Srinivasan to present his work at INTERSPEECH, and to complete some experiments on larger-scale speech processing tasks. He is expected to complete his dissertation by August 2009.

Mr. Tao will complete his graduate studies as a graduate teaching assistant in the ECE Department at MS State. This will allow him to complete his dissertation on the work previously described. He is expected to graduate by December 2009.

The final report for this project will be delivered by the end of the 2008 calendar year, when we expect to have expended all of the funds.

E. References

- [1] D. May, *Nonlinear Dynamic Invariants For Continuous Speech Recognition*, M.S. Thesis, Department of Electrical and Computer Engineering, Mississippi State University, May 2008.
- [2] D. May, T. Ma, S. Srinivasan, G. Lazarou and J. Picone, "Continuous Speech Recognition Using Nonlinear Dynamic Invariants," submitted to INTERSPEECH, Brisbane, Australia, September 2008.
- [3] D. May, S. Srinivasan, T. Ma, G. Lazarou and J. Picone, "Continuous Speech Recognition Using Nonlinear Dynamic Invariants," submitted to the Association for Computational Linguistics: Human Language Technologies, Columbus, Ohio, USA, June 2008.
- [4] D. May, S. Srinivasan, T. Ma and J. Picone, "Continuous Speech Recognition Using Nonlinear Dynamical Invariants," submitted to the International Conference on Acoustics, Speech, and Signal Processing, Las Vegas, Nevada, USA, March 2008.
- [5] C.S. Wong and W.K. Li, "On a Mixture Autoregressive Model," *Journal of the Royal Statistical Society*, vol. 62, no. 1, pp. 95-115, February 2000.
- [6] S. Srinivasan, T. Ma, D. May, G. Lazarou and J. Picone, "Nonlinear Mixture Autoregressive Hidden Markov Models For Speech Recognition," to be presented at INTERSPEECH, Brisbane, Australia, September 2008.
- [7] T. Ma, S. Srinivasan, D. May, G. Lazarou and J. Picone, "Robust Speech Recognition Using Linear Dynamic Models," submitted to INTERSPEECH, Brisbane, Australia, September 2008.