

LA-UR- 98-4312

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE: MAXIMUM-LIKELIHOOD CONTINUUM IMPROVING MATCHING
(MALCOLM): AN ALTERNATIVE TO HMMS

RECEIVED

AUG 18 1999

OSTI

AUTHOR(S): David A. Nix, John E. Hogden

SUBMITTED TO: Neural Information Processing Systems
Denver, CO
December, 1998

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos

Los Alamos National Laboratory
Los Alamos New Mexico 87545

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Maximum-Likelihood Continuity Mapping (MALCOM): An Alternative to HMMs

David A. Nix
dnix@lanl.gov
Computer Research and Applications
CIC-3, MS P990
Los Alamos National Laboratory
Los Alamos, NM 87545

John E. Hogden
hogden@lanl.gov
Computer Research and Applications
CIC-3, MS B265
Los Alamos National Laboratory
Los Alamos, NM 87545

Abstract

We describe Maximum-Likelihood Continuity Mapping (MALCOM) as an alternative to hidden Markov models (HMMs) for processing sequence data such as speech. While HMMs have a discrete "hidden" space constrained by a fixed finite-automata architecture, MALCOM has a continuous hidden space (a *continuity map*) that is constrained only by a smoothness requirement on paths through the space. MALCOM fits into the same probabilistic framework for speech recognition as HMMs, but it represents a far more realistic model of the speech production process. We support this claim by generating continuity maps for three speakers and using the resulting MALCOM paths to predict measured speech articulator data. The correlations between the MALCOM paths (obtained from only the speech acoustics) and the actual articulator movements average 0.77 on an independent test set not used to train MALCOM nor the predictor. On average, this unsupervised model achieves 92% of performance obtained using the corresponding supervised method.

1 INTRODUCTION

Hidden Markov models (HMMs) are generally considered to be the state of the art in speech recognition (e.g., Young, 1996). The strengths of the HMM framework include a rich mathematical foundation, powerful training and recognition algorithms for large speech corpora, and a probabilistic framework that can incorporate statistical phonology and syntax (Morgan & Boulard, 1995). However, HMMs are known to be a poor model of the speech production process. While speech production is a continuous, temporally dependent process, HMMs treat speech production as a discrete, finite-state system where the current state depends only on the immediately preceding state. Furthermore, while HMMs are designed to capture temporal information as state transition probabilities, recent results by Boulard *et al.*, (1996) demonstrate that when the transition probabilities are replaced by constant values, recognition results do not significantly deteriorate. This indicates that the conventional HMM framework does not adequately capture temporal information, and complex modifications are required to do so (such as Boulard *et al.*'s conditioning the transition probabilities on the acoustics as well as the state).

Given these deficiencies, it is desirable to consider alternatives to the HMM approach that maintain its strengths while improving upon its weaknesses. This paper describes an alternative model called Maximum-Likelihood Continuity Mapping (MALCOM). We first review a general statistical framework for speech recognition so that we can compare and contrast the HMM and MALCOM formulations. Then we consider what the abstract hidden state represents in MALCOM, and we demonstrate empirically that in contrast to HMMs, the paths through MALCOM's hidden space are related to the movements of the speech production articulators.

2 A GENERAL FRAMEWORK FOR SPEECH RECOGNITION

Consider an unknown speech waveform that is converted by a front-end signal-processing module into a sequence of acoustic vectors \mathbf{X} . Given a space of possible utterances, W , the task of speech recognition is to return the most likely utterance W^* given the observed acoustic sequence \mathbf{X} . Using Bayes' rule this corresponds to

$$W^* = \operatorname{argmax}_W P(W|\mathbf{X}) = \operatorname{argmax}_W \frac{P(\mathbf{X}|W)P(W)}{P(\mathbf{X})}. \quad (1)$$

In recognition, $P(\mathbf{X})$ is typically ignored, and the posterior $P(W|\mathbf{X})$ is estimated as the product of the prior probability of the word sequence, $P(W)$, and the probability that the observed acoustics were generated by the word sequence, $P(\mathbf{X}|W)$. The prior $P(W)$ is estimated by a *language model*, while the production probability $P(\mathbf{X}|W)$ is estimated by an *acoustic model*. In continuous speech recognition, the product of these terms must be maximized over W ; however, in this paper, we will restrict our attention to the form of the acoustical model only.

Every candidate utterance W corresponds to a sequence of word/phone models \mathcal{M}_w such that $P(\mathbf{X}|W) = P(\mathbf{X}|\mathcal{M}_w)$, and each \mathcal{M}_w considers all possible paths \mathbf{Y} through some "hidden" space. Thus, for each candidate utterance, we must calculate

$$P(\mathbf{X}|\mathcal{M}_w) = \int_{\mathbf{Y}} P(\mathbf{X}|\mathbf{Y}, \mathcal{M}_w)P(\mathbf{Y}|\mathcal{M}_w)d\mathbf{Y}. \quad (2)$$

2.1 HIDDEN MARKOV MODELS

Since HMMs are finite-state machines with a given fixed architecture, the path \mathbf{Y} through the hidden space corresponds to series of discrete states, simplifying the integral of Eq. (2) to a sum. However, to avoid computing the contribution of all possible paths, the "Viterbi" approximation—considering only the path that maximizes Eq. (2)—is frequently used without much loss in recognition performance (Morgan & Bourlard, 1995). Thus,

$$P(\mathbf{X}|\mathcal{M}_w) \approx \operatorname{argmax}_{\mathbf{Y}} P(\mathbf{X}|\mathbf{Y}, \mathcal{M}_w)P(\mathbf{Y}|\mathcal{M}_w). \quad (3)$$

The first term corresponds to the product of the emission probabilities of the acoustics given the state sequence and is typically estimated by mixtures of high-dimensional Gaussian densities. The second term corresponds to the product of the state transition probabilities. However, since this second term is what can be made constant without degrading performance, the modeling power of the conventional HMM must reside in the first term.

Training the HMM system involves estimating both the emission and the transition probabilities from real speech data. The Baum-Welch/forward-backward algorithm (e.g., Morgan & Scofield, 1994) is the standard computationally efficient algorithm for iteratively estimating these distributions.

2.2 MAXIMUM-LIKELIHOOD CONTINUITY MAPPING (MALCOM)

In contrast to HMMs, the multi-dimensional MALCOM hidden space is continuous—there are an infinite number states and paths through them. While the HMM is constrained by a fixed architecture, MALCOM is constrained by the notion of continuity of the hidden path. That is, the path must be smooth and continuous: it may not carry any energy above a given cutoff frequency. Unlike the discrete path in an HMM, the smooth hidden path in MALCOM attempts to emulate the motion of the speech articulators in what we call a *continuity map* (CM).

Unless we know how to evaluate the integral of Eq. (2) (which we currently do not), we must also make the Viterbi approximation and approximate $P(\mathbf{X}|\mathcal{M}_w)$ by considering only the single path that maximizes the likelihood of the acoustics \mathbf{X} given the utterance model \mathcal{M}_w , resulting in Eq. (3) once again. Analogously, the first term, $P(\mathbf{X}|\mathbf{Y}, \mathcal{M}_w)$, corresponds to the acoustic generation probability given the hidden path, and the second term corresponds to the probability of the hidden path given the utterance model. This paper focuses on the first term, since this is the term that produces conventional HMM performance.¹

Common to all \mathcal{M}_w is a set of N probability density functions (pdfs) Φ that define the CM hidden space, modeling the likelihood of \mathbf{Y} given \mathbf{X} for an N -code vector quantization (VQ) of the acoustic space. Since these pdfs model the low-(e.g., 6)-dimensional CM space instead of the high-(e.g., 39)-dimensional acoustic space, MALCOM requires many fewer parameters to be estimated than the corresponding HMM.

3 THE MALCOM ALGORITHM

We now turn to developing an algorithm to estimate both the CM pdfs Φ and the corresponding paths \mathbf{Y} that together maximize the likelihood of a given time series of acoustics, $\mathcal{L} = P(\mathbf{X}|\mathbf{Y}, \Phi)$. This is an extension of the algorithm first proposed by Hogden (1995). Starting with random (but smooth) \mathbf{Y} , the MALCOM training algorithm calculates a CM by iterating between the following two EM-like steps: (1) Given \mathbf{Y} , reestimate Φ to maximize \mathcal{L} ; and (2) Given Φ , reestimate smooth paths \mathbf{Y} to maximize \mathcal{L} .

3.1 COST FUNCTION

To define a cost function for \mathcal{L} , we make two dependence claims and one independence assumption. First we claim that \mathbf{y}_t depends on \mathbf{y}_{t-1} . This says that the path configuration at t depends on the configuration just prior to t , an expression of the continuity constraint described above. We make another natural claim that \mathbf{x}_t depends on \mathbf{y}_t , that the path configuration at time t influences the corresponding acoustics. However, we do make the conditional independence assumption that

$$\mathcal{L} = P(\mathbf{X}|\mathbf{Y}, \Phi) = \prod_{t=1}^n P(\mathbf{x}_t|\mathbf{y}_t, \Phi). \quad (4)$$

Note that Eq. (4) does not assume that each \mathbf{x}_t is independent of \mathbf{x}_{t-1} (as is often assumed in data modeling); it only assumes that the *conditioning* of \mathbf{x}_t on \mathbf{y}_t is independent from $t-1$ to t . Since \mathbf{x}_t depends on \mathbf{y}_t , \mathbf{y}_t depends on \mathbf{y}_{t-1} , and \mathbf{x}_{t-1} depends on \mathbf{y}_{t-1} , \mathbf{x}_t is *not* assumed to be independent of \mathbf{x}_{t-1} .

¹However, we are currently developing a model of $P(\mathbf{Y}|\mathcal{M}_w)$ to replace the corresponding (and useless) term in the conventional HMM formulation as well (Hogden *et al.*, 1998).

With a log transformation and an invocation of Bayes' rule, we obtain the MALCOM cost function:

$$\ln \mathcal{L} = \sum_{t=1}^n [\ln P(\mathbf{y}_t | \mathbf{x}_t, \Phi) + \ln P(\mathbf{x}_t) - \ln P(\mathbf{y}_t | \Phi)]. \quad (5)$$

We model each $P(\mathbf{y}_t | \mathbf{x}_t, \Phi)$ by a probability density function (pdf) $p[\mathbf{y}_t | \mathbf{x}_t, \Phi_j(\mathbf{x}_t)]$, where the particular model Φ_j depends on which of the N VQ codes \mathbf{x}_t is assigned to. Here we use a simple multi-dimensional Gaussian for each pdf.² Next, we need an estimate of $P(\mathbf{y}_t | \Phi)$, which can be obtained by summing over all VQ partitions: $P(\mathbf{y}_t | \Phi) \approx \sum_{j=1}^N p(\mathbf{y}_t | \mathbf{x}_j, \Phi_j) P(\mathbf{x}_j)$. We estimate $P(\mathbf{x}_j)$ by calculating the relative frequency of each acoustic code in the VQ codebook.

3.2 PDF ESTIMATION

For step (1) of training, we use gradient-based optimization to reestimate the means of the Gaussian pdfs for each acoustic partition, where the gradient of Eq.(5) with respect to the mean of pdf i is

$$\nabla_{\mu_i} \ln \mathcal{L} = \sum_{t=1}^n \left\{ \Sigma_i^{-1} (\mathbf{y}_t - \mu_i) - \frac{\sum_{j=1}^N p[\mathbf{y}_t | \mathbf{x}_j, \Phi(\mathbf{x}_j)] P(\mathbf{x}_j) \Sigma_j^{-1} (\mathbf{y}_t - \mu_j)}{\sum_{j=1}^N p[\mathbf{y}_t | \mathbf{x}_j, \Phi(\mathbf{x}_j)] P(\mathbf{x}_j)} \right\} \quad (6)$$

where Σ is the covariance matrix for each pdf. For the results in this paper, we use a common radially symmetric covariance matrix for all pdfs that is reestimated after each path optimization step.³ In doing the optimization, we employ the following algorithm:

1. Make an initial guess of each μ_i as the means of the path configurations corresponding to the observed acoustics $\mathbf{X} \in x_i$.
2. Construct $\nabla_{\mu} \ln \mathcal{L}$ by considering Eq. (6) over all N acoustic partitions.
3. Determine a search direction for the optimization using, for example, conjugate gradients and perform a line search along this direction (Press *et al.*, 1988).
4. Repeat steps [2]–[3] until convergence.

To avoid potential degenerate solutions and since the scale of the continuity map is meaningless (only its topological arrangement matters), after each pdf optimization, the N pdf means are orthogonalized and scaled to zero mean, unit variance before new paths are estimated.

3.3 PATH ESTIMATION

For step (2) of training, we use gradient-based optimization to reestimate \mathbf{Y} , where the gradient of the log likelihood function with respect to \mathbf{y}_t is given by

$$\nabla_{\mathbf{y}_t} \ln \mathcal{L}_t = \frac{\nabla_{\mathbf{y}_t} p[\mathbf{y}_t | \mathbf{x}_t, \Phi(\mathbf{x}_t)]}{p[\mathbf{y}_t | \mathbf{x}_t, \Phi(\mathbf{x}_t)]} - \frac{\nabla_{\mathbf{y}_t} \sum_{j=1}^N p[\mathbf{y}_t | \mathbf{x}_j, \Phi(\mathbf{x}_j)] P(\mathbf{x}_j)}{\sum_{j=1}^N p[\mathbf{y}_t | \mathbf{x}_j, \Phi(\mathbf{x}_j)] P(\mathbf{x}_j)}. \quad (7)$$

In doing the optimization, we employ the following gradient-based algorithm:

1. Make an initial guess of the path \mathbf{Y}^0 as the means of the pdfs corresponding to the observed acoustic sequence \mathbf{X} .

²We are currently exploring the use of multi-modal mixtures of Gaussians to represent the pdfs for sounds such as stop consonants for which the inverse map from acoustics to articulation may not be unique.

³However, we are currently exploring the effects of individual and diagonal covariance matrices.

2. Low pass filter \mathbf{Y}^0 .
3. Construct $\nabla_{\mathbf{Y}} \ln \mathcal{L}$ by considering Eq. (7) over all t .
4. Determine a search direction for the optimization using, for example, conjugate gradients (Press *et al.*, 1988).
5. Low-pass filter this search direction using the same filter as in step [2].
6. Perform a line search along the filtered direction (Press *et al.*, 1988).
7. Repeat steps [3]–[6] until convergence.

Since neither the line search direction nor the initial estimate \mathbf{Y}^0 contains energy above the cutoff frequency of the low-pass filter, their linear addition—the next estimate of \mathbf{Y} —will not contain energy above the cutoff frequency either. Thus, steps [2] and [5] implement the desired smoothness constraint.

4 MALCOM PATHS AND SPEECH PRODUCTION: DATA AND METHOD

To demonstrate our claim that MALCOM paths are topologically related to articulator motions, we construct a regression predictor from \mathbf{Y} to measured articulator data using the training data and test the quality of this predictor on an independent test set.

Our speech corpus consists of data from two male and one female native speakers of German. This data was obtained from Dr. Igor Zlokarnik and recorded at the Technical University of Munich, Germany using electro-magnetic articulography (EMA) (Perkell *et al.*, 1992). Each speaker's articulatory measurements and acoustics were recorded for the same 108 sentences, where each sentence was about 4 seconds long.

The acoustics were recorded using a room-placed microphone and sampled using 16-bit resolution at 16 kHz. Prior to receiving the data from Munich, the data were resampled at 11025 Hz. To represent the acoustic signal in compact vector time-series, we used 256-sample (23.2 msec) Hamming-windowed frames, with a new frame starting every 5.8 msec (75% overlap). We transform each frame into a 13th-order LPC-cepstral coefficient vector \mathbf{a}_t (12 cepstral features plus log gain—see Morgan & Scofield, 1994). A full acoustical feature vector \mathbf{x}_t consists of a window of seven frames such that \mathbf{x}_t is made up of the frames $\{\mathbf{a}_{t-6}, \mathbf{a}_{t-4}, \mathbf{a}_{t-2}, \mathbf{a}_t, \mathbf{a}_{t+2}, \mathbf{a}_{t+4}, \mathbf{a}_{t+6}\}$. To VQ the acoustic space we used the classical k -means algorithm (e.g., Bishop, 1996). We used 512 codes to model the vowel data, and 256 codes each to model the stop consonants, the fricatives, the nasals, and the liquids (1536 codes total).⁴

The articulatory data were recorded using EMA and consist of the (x, y) coordinates of 4 coils along the tongue and the y -coordinates of coils on the jaw and lower lip. Figure 1 illustrates the approximate location of each coil. The data were originally sampled at 250 Hz, but we resampled the data to 172.26 Hz to match one articulatory sample for each 75%-overlapping acoustic frame of 256 samples. The articulatory data were subsequently low-pass filtered at 15 Hz to remove measurement noise.

Sentences 1–90 were used as a training set, and sentences 91–108 were withheld for evaluation. A separate CM was generated for each speaker using the training data. We used an 8 Hz cutoff frequency because the measured articulatory data had very little energy above 8 Hz, and a 6-dimensional continuity map was used because the first six principal components capture 99% of the variance of the corresponding articulator data (Nix, 1998).

Since the third term in Eq. (5) is computationally complex, we approximated Eq. (5) by only its first term (the second term is constant during training) until $\ln \mathcal{L}$, calculated at the

⁴This acoustic representation and VQ scheme were determined to work well for modeling real articulator data (Nix, 1998), so they were used here as well.




Figure 1: Approximate positions of EMA coils for speech articulation measurements.

end of each iteration using all terms, started to decrease. At this point we started using both the first and third terms of Eq. (5). For each training iteration, convergence was defined as when the maximum movement of a path or a centroid is $< 10^{-4}$. Final convergence was defined as when the correlation of the paths from one iteration to another was > 0.99 in all dimensions. This usually took about 30 iterations.

To demonstrate that the MALCOM hidden paths capture information related to articulation, we used the same training set to estimate a non-linear regression function from \mathbf{Y} to the corresponding measured articulator data. If MALCOM paths had no information in common with speech articulation, we would expect this regression to fail. We used an ensemble of 10 single-hidden-layer, 32-hidden unit, multi-layer perceptrons trained on different 2/3-training, 1/3-early stopping partitions of the training set, where the results of the ensemble on the test set were averaged (e.g., Bishop, 1995). A linear regression produced results approximately 10% worse than those we report here.

To compare the performance of MALCOM at capturing articulatory information to the corresponding supervised method—where articulatory data is available for training as well as evaluation—we used the same training set to map from the same acoustic VQ codes to the measured articulatory data. This involves only the pdf optimization step of MALCOM, since the paths are fixed as the articulator measurements. The resulting pdfs are then used in the path optimization step to determine paths for the test data acoustics. We can then measure what fraction of this supervised performance that MALCOM attains.

5 RESULTS AND CONCLUSIONS

The results of this regression on the test set are listed in Table 1. On average, the MALCOM paths had a correlation of 0.77 with the actual articulator data, compared to 0.84 for the comparable supervised method. Thus, MALCOM was able to achieve over 90% of the performance of the corresponding supervised method. MALCOM captured considerable articulatory information from the speech signal *using only the speech acoustics*.

Given that (1) MALCOM fits into the same probabilistic framework for speech recognition as HMMs and (2) MALCOM's hidden paths capture considerable information about the speech production process, we believe that MALCOM is a viable alternative to the HMM for speech processing tasks. Our current work emphasizes developing a word model to complete the MALCOM formulation and test a full speech recognition system. Furthermore, MALCOM is applicable to any other task to which HMMs can be applied, including

Table 1: Results of using MALCOM to infer articulatory trajectories for independent test data. For "MALCOM" and the corresponding "Supervised" method, the table gives correlations between predicted and actual articulatory trajectories for the measurements illustrated in Figure 1. "Fraction" lists the fraction of the supervised performance that MALCOM attained.

Speaker M1

Method	T1x	T1y	T2x	T2y	T3x	T3y	T4x	T4y	LLy	LJy	AVE
MALCOM	0.81	0.84	0.81	0.88	0.75	0.83	0.75	0.66	0.77	0.76	0.79
Supervised	0.85	0.88	0.86	0.92	0.84	0.90	0.85	0.80	0.82	0.78	0.85
Fraction	0.95	0.95	0.94	0.96	0.89	0.92	0.88	0.83	0.94	0.97	0.92

Speaker M2

Method	T1x	T1y	T2x	T2y	T3x	T3y	T4x	T4y	LLy	LJy	AVE
MALCOM	0.74	0.83	0.76	0.88	0.79	0.84	0.86	0.69	0.73	0.79	0.79
Supervised	0.80	0.89	0.81	0.93	0.85	0.90	0.91	0.81	0.82	0.83	0.86
Fraction	0.93	0.93	0.94	0.95	0.93	0.93	0.95	0.85	0.89	0.95	0.92

Speaker F1

Method	T1x	T1y	T2x	T2y	T3x	T3y	T4x	T4y	LLy	LJy	AVE
MALCOM	0.66	0.82	0.75	0.84	0.72	0.73	0.77	0.47	0.77	0.76	0.73
Supervised	0.73	0.88	0.80	0.88	0.79	0.82	0.84	0.62	0.83	0.79	0.80
Fraction	0.90	0.93	0.94	0.95	0.91	0.89	0.92	0.76	0.93	0.96	0.91

fraud detection (Hogden, 1997) and text processing.

Acknowledgments

We would like to acknowledge James Howse and Mike Mozer for helpful comments on this manuscript. This work was performed under the auspices of the U.S. Department of Energy.

References

- Bishop, C.M. (1995). *Neural Networks for Pattern Recognition*, NY: Oxford University Press, Inc.
- Bourlard, H. Konig, Y., & Morgan, N. (1996). "A training algorithm for statistical sequence recognition with applications to transition-based speech recognition," *IEEE Signal Processing Letters*, 3(7), 203-205.
- Hogden, J. (1995). "Improving on hidden Markov models: an articulatorily constrained, maximum-likelihood approach to speech recognition and speech coding," Los Alamos National Laboratory Technical Report, LA-UR-96-3945.
- Hogden, J. (1997). "Maximum likelihood continuity mapping for fraud detection," Los Alamos National Laboratory Technical Report, LA-UR-97-992.
- Hogden, J., Nix, D.A., Gracco, V., & Rubin, P. (1998). "Stochastic word models for articulatorily constrained speech recognition and synthesis," submitted to Acous. Soc. of Am. Conf., 1998.
- Morgan, N. & Bourlard, H.A. (1995). "Neural Networks for Statistical Recognition of Continuous Speech," *Proceedings of the IEEE*, 83(5), 742-770.
- Morgan, D.P., & Scofield, C.L. (1992). *Neural Networks and Speech Processing*, Boston, MA: Kluwer Academic Publishers.
- Nix, D.A. (1998). *Probabilistic methods for inferring vocal-tract articulation from speech acoustics*, Ph.D. Dissertation, U. of CO at Boulder, Dept. of Computer Science, in preparation.
- Perkell, J.S., Cohen, M.H., Svirsky, M.A., Matthies, M.L., Garabieta, I., & Jackson, M.T.T. (1992). "Electromagnetic midsagittal articulometer systems for transducing speech articulatory movements," *Journal of the Acoustical Society of America*, 92(6), 3078-3096.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., & Flannery, B.P. (1988). *Numerical Recipes in C* Cambridge University Press.
- Young, S.J. (1996). "A review of large-vocabulary continuous speech recognition," *IEEE Signal Processing Magazine*, September, 45-57.