SVMS, SCORE-SPACES AND MAXIMUM MARGIN STATISTICAL MODELS

M.J.F. Gales and M. Layton

Engineering Department, Cambridge University, Trumpington St., Cambridge, CB2 1PZ U.K.
Email: {mjfg,ml362}@eng.cam.ac.uk

ABSTRACT

There has been significant interest in developing new forms of acoustic model, in particular models which allow additional dependencies to be represented than allowed within a standard hidden Markov model (HMM). This paper discusses one such class of models, augmented statistical models. Here a locally exponential approximation is made about some point on a base distribution. This allows additional dependencies within the data to be modelled than are represented in the base distribution. Augmented models based on Gaussian mixture models (GMMs) and HMMs are briefly described. These augmented models are then related to generative kernels, one approach used for allowing support vector machines (SVMs) to be applied to variable length data. The training of augmented statistical models within an SVM, generative kernel, framework is then discussed. This may be viewed as using maximum margin training to estimate statistical models. Augmented Gaussian mixture models are then evaluated using rescoring on a large vocabulary speech recognition task.

1. INTRODUCTION

There have been a wide-range of acoustic models applied to the speech recognition task. These range from the standard hidden Markov model (HMMs), to segmental models [1], switching linear dynamical systems (LDSs) [2], buried Markov models (BMMs) [3] and mixed memory models [4]. Many of these models can be viewed as state-space models and graphical models [2]. The underlying aspect of all these models is how to appropriately model the dependencies (and complexities) of the speech signal. For example, forms of dependencies include observation independence given the current state, as in an HMM, and independence given a continuous latent state-space variable, as in a LDS. The fundamental questions that must be answered when looking at these models is which latent variables should be included, what dependencies should be modelled, and how the distributions of the observations are altered by the dependencies. To the authors' knowledge, there are as yet no approaches that allow all these questions to be answered.

In this paper a structured approach is described to obtain the statistics that determine the dependencies to be modelled in the observation sequence. The approach adopted is to use a base statistical model, and then for each point on that distribution to construct a local exponential model. The base statistical model determines

the latent variables; the sufficient statistics are determined using higher order derivatives of the log-likelihood. This is similar to a constrained exponential model [5]. However here all the parameters of the model, including the parameters of the base distribution, may be trained. This will be referred to as an augmented statistical model [6]

Using this form of model, the number of model parameters rapidly increases. Though with sufficient training data, large numbers of model parameters can be trained, the use of robust training criteria that allow for good generalisation are useful. In this work maximum margin training, as used to train support vector machines (SVMs), is used. The paper is organised as follows. The next section describes the general theory of augmented statistical models and the forms that they take for Gaussian mixture models (GMMs) and HMMs. SVMs and generative kernels are then described. This is followed by a description of how maximum margin training can be used to train augmented models. The issues of applying these models to large vocabulary systems is then described followed by results on an large vocabulary task.

2. AUGMENTED STATISTICAL MODELS

2.1. The Exponential Family

Many standard forms of statistical model are based on the exponential family. The general form for the exponential family with parameters $\boldsymbol{\alpha}$ can be expressed as

$$p(\mathbf{o}; \boldsymbol{\alpha}) = \frac{1}{\tau} h(\mathbf{o}) \exp\left(\boldsymbol{\alpha}' T(\mathbf{o})\right) \qquad (1)$$

where $h(\mathbf{o})$ is the *reference distribution*, $\boldsymbol{\alpha}$ are the *natural parameters*, τ is the normalisation term (a function of both the reference distribution and the natural parameters) and the function $T(\mathbf{o})$ is a *sufficient statistic*. There are a number of standard examples, including the exponential distribution and Gaussian (Normal) distribution. The reason for the interest in members of the exponential family is that the sufficient statistics are of a fixed dimensionality and that conjugate priors can be defined simply for all members of this family. When dynamic data is being considered, such that each example is a series of observations, $\mathbf{O} = \mathbf{o}_1, \ldots, \mathbf{o}_T$, the range of possible statistics becomes very large. Dependencies between observations as in BMMs [3], as well as within the feature vector, can now be modelled.

An interesting subset of the set of all possible members of the exponential family is the *constrained exponential family* [5]. Here rather than allowing any form of statistics, a local exponential approximation to the reference distribution is used as the statistical model, where the local approximation replicates some of the properties of the reference distribution. In this paper a slightly more

general form of statistical model than the constrained exponential family is used. In addition to the values of the local exponential model, the reference distribution parameters may be learnt from the data. These models are referred to as *augmented statistical models* [6].

2.2. Augmented Statistical Models

Augmented statistical models are an attractive approach to building class-conditional probability density functions, since they yield a mathematically consistent formulation to add higher order dependencies into the model. First a base statistical model, $\check{p}(\mathbf{o}; \boldsymbol{\lambda})$, is defined. A member of the exponential family that locally approximates this base distribution for a particular set of parameters $\boldsymbol{\lambda}$ is then used as the final statistical model. The general form of augmented statistical model for a base statistical model can be expressed as

$$
\begin{aligned}
p(\mathbf{o}; \boldsymbol{\lambda}, \boldsymbol{\alpha}) &= \frac{1}{\tau}\check{p}(\mathbf{o}; \boldsymbol{\lambda}) \exp\left(\boldsymbol{\alpha}'\left\{\boldsymbol{\nabla}_{\lambda}^{(\rho)} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda}))\right\}\right) \\
&= \frac{1}{\tau}\overline{p}(\mathbf{o}; \boldsymbol{\lambda}, \boldsymbol{\alpha})
\end{aligned} \quad (2)
$$

where $\boldsymbol{\nabla}_{\lambda}^{(\rho)} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda}))$ is the vector form of all the derivatives[1] up to order ρ,

$$
\boldsymbol{\nabla}_{\lambda}^{(\rho)} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda})) = \begin{bmatrix} \boldsymbol{\nabla}_{\lambda} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda})) \\ \frac{1}{2!}\mathrm{vec}\left(\boldsymbol{\nabla}_{\lambda}^{2} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda}))\right) \\ \vdots \\ \frac{1}{\rho!}\mathrm{vec}\left(\boldsymbol{\nabla}_{\lambda}^{\rho} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda}))\right) \end{bmatrix} \quad (3)
$$

τ is the appropriate normalisation term, thus

$$
\tau = \int_{\mathcal{R}^d} \overline{p}(\mathbf{o}; \boldsymbol{\lambda}, \boldsymbol{\alpha}) d\mathbf{o} \quad (4)
$$

For this augmented model to be a valid PDF, the integration must be bounded. Additionally, it must satisfy $p(\mathbf{o}; \boldsymbol{\lambda}, \boldsymbol{\alpha}) \geq 0$.

If the base statistical model is itself a member of the exponential family, the augmented statistical model will also be a member of the exponential family, though not necessarily of the same form as the base distribution. This is not true for situations where the base statistical model is not a member of the exponential family, for example the Gaussian mixture model discussed in section 2.3.

It is interesting to contrast the nature of the dependencies that are incorporated into the augmented model compared to the base model. Independence assumptions in the base statistical model are reflected in the independence assumptions in the augmented model. However this is not the case for the conditional independence assumptions. For example, for the augmented GMM (A-GMM) discussed below, the observations are not conditionally independent given the base component that generated them – the posterior, $P(n|\mathbf{o}; \boldsymbol{\lambda})$ in equation 7, causes the likelihood to be a function of all the components. The augmented statistical models can be related performing a Taylor series expansion on the statistical model [7, 6].

---

[1]For simplicity, in this work the *natural basis* and higher order derivatives are assumed to yield a set of orthogonal basis. Given this assumption it is not necessary to distinguish between covariant and contravariant basis and components [5].

2.3. Augmented Gaussian Mixture Models

One of the standard forms of model used in statistical pattern processing are mixture models, in particular GMMs. The base statistical model, a GMM, has the form

$$
\check{p}(\mathbf{o}; \boldsymbol{\lambda}) = \sum_{m=1}^{M} c_m \mathcal{N}(\mathbf{o}; \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \quad (5)
$$

Considering just the first order derivatives of the mean as an element of the augmented model

$$
\boldsymbol{\nabla}_{\mu_m} \log(\check{p}(\mathbf{o}; \boldsymbol{\lambda})) = P(m|\mathbf{o}; \boldsymbol{\lambda})\boldsymbol{\Sigma}_m^{-1}(\mathbf{o} - \boldsymbol{\mu}_m) \quad (6)
$$

where $P(m|\mathbf{o}; \boldsymbol{\lambda}) = c_m \mathcal{N}(\mathbf{o}; \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m)/\check{p}(\mathbf{o}; \boldsymbol{\lambda})$, and c_m, $\boldsymbol{\mu}_m$ and $\boldsymbol{\Sigma}_m$ are the prior, mean and covariance matrix for component m. The associated A-GMM is then given by

$$
\begin{aligned}
p(\mathbf{o}; \boldsymbol{\lambda}, \boldsymbol{\alpha}) = \frac{1}{\tau} \sum_{m=1}^{M} c_m \Bigg\{ \mathcal{N}(\mathbf{o}; \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \times \\
\exp\left(\sum_{n=1}^{M} P(n|\mathbf{o}; \boldsymbol{\lambda})\boldsymbol{\alpha}_n' \boldsymbol{\Sigma}_n^{-1}(\mathbf{o} - \boldsymbol{\mu}_n)\right) \Bigg\}
\end{aligned} \quad (7)
$$

The parameters $\boldsymbol{\alpha} = \{\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_M\}$ are additional model parameters for the A-GMM.
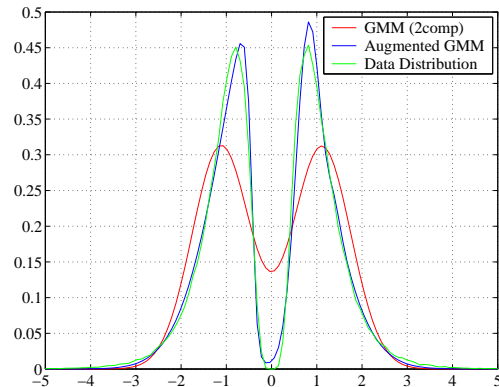

Fig. 1. Modelling of "symmetric" log-normal distribution

Figure 1 shows two distributions used to model data generated using a symmetric log-normal distribution. ML training was used to obtain a two-component GMM for the training data. In addition an A-GMM using first-order derivatives of the mean only was also trained. From the diagram it is clear that the additional power of the A-GMM is able to model the distribution better than the GMM, though using an additional 2 model parameters. This is reflected in the average log-likelihoods, -1.59 for the GMM and -1.45 for the A-GMM. Interestingly even using a 4-component GMM the log-likelihood was only -1.46, still less than the two component A-GMM.

The most commonly used form of acoustic model for speech recognition is the HMM. Augmented forms of HMMs may also be generated. The first order derivatives of the log-likelihood have a similar form to the GMM and may be written as [7]

$$
\boldsymbol{\nabla}_{\mu_{jm}} \log \check{p}(\mathbf{O}; \boldsymbol{\lambda}) = \sum_{t=1}^{T} \gamma_{jm}(t)\boldsymbol{\Sigma}_{jm}^{-1}\left(\mathbf{o}_t - \boldsymbol{\mu}_{jm}\right) \quad (8)
$$

where $\gamma_{jm}(t) = P(\theta_t = \{s_j, m\}|\mathbf{O}; \boldsymbol{\lambda})$, θ_t is the state/component pairing at time t. In contrast to GMMs, HMMs do not assume that the observations are independent of one another. Thus temporal dependencies within the data can be modelled using the higher order derivatives.

3. SUPPORT VECTOR MACHINES

Support Vector Machines (SVMs) [8] are an approximate implementation of structural risk minimisation. They have been found to yield good performance on a wide of range tasks. This section briefly reviews SVMs and the use of generative kernels which are one approach to handling the dynamic nature of speech.

3.1. Maximum Margin Training

SVMs are based upon the intuitive concept of maximising the margin between the decision hyperplane and the closest training examples. This has been shown to be related to minimising an upper bound on the generalisation error [8]. However, in general it is not possible to construct a separating hyperplane between classes with no classification errors. In these situations, an optimal hyperplane is found to minimise the probability of error, averaged over the training set. This is accomplished by allowing *soft margins*. The margin between classes is said to be soft if there exist training examples, $\mathbf{o}_i$ (with labels $y_i \in \{-1, 1\}$), that violate the constraint $y_i(\langle \mathbf{w}, \mathbf{o}_i \rangle + w_0) \geq 1$, where $\mathbf{w}$ is the weight vector, w_0 is the bias of the optimal hyperplane and $\langle \cdot, \cdot \rangle$ indicates the inner product between the two vectors using an appropriate metric[2]. Slack variables, $\epsilon_i \geq 0$, are introduced to measure the deviation of these examples from the ideal condition of pattern separability. The objective function and constraint become,

$$\{\hat{\mathbf{w}}, \hat{w}_0\} = \arg \min_{\mathbf{w}, w_0} \left\{ \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle + C \sum_{i=1}^{n} \epsilon_i \right\} \tag{9}$$

subject to $y_i(\langle \mathbf{w}, \mathbf{o}_i \rangle + w_0) \geq 1 - \epsilon_i$. C acts as a regularisation parameter and controls the trade-off between the margin and the number of misclassified points. For non-linearly separable data, Cover's theorem states that examples may be made linearly separable with a high probability given a non-linear transformation, $\phi(\mathbf{o}; \boldsymbol{\lambda})$, from input-space, $\mathbf{o}$, to a *feature-space* of sufficient dimensionality. Using this mapping, the kernelised form of the dual objective function is defined for the Lagrange multipliers $\boldsymbol{\alpha}^{\mathrm{svm}}$

$$\hat{\boldsymbol{\alpha}}^{\mathrm{svm}} = \tag{10}$$
$$\arg \max_{\boldsymbol{\alpha}^{\mathrm{svm}}} \left\{ \sum_{i=1}^{n} \alpha_i^{\mathrm{svm}} - \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i^{\mathrm{svm}} \alpha_j^{\mathrm{svm}} y_i y_j K(\boldsymbol{\lambda}, \mathbf{o}_i, \mathbf{o}_j) \right\}$$

subject to $\sum_{i=1}^{n} \alpha_i^{\mathrm{svm}} y_i = 0$ and $0 \leq \alpha_i^{\mathrm{svm}} \leq C$. Here

$$K(\boldsymbol{\lambda}, \mathbf{o}_i, \mathbf{o}_j) = \langle \phi(\mathbf{o}_i; \boldsymbol{\lambda}), \phi(\mathbf{o}_j; \boldsymbol{\lambda}) \rangle \tag{11}$$

The upper limit on the Lagrange multipliers, $\boldsymbol{\alpha}^{\mathrm{svm}}$, limits the influence of individual examples (which may be outliers). At optimality, the Karush-Kuhn-Tucker (KKT) conditions ensure that only examples that lie on the margin ($\langle \mathbf{w}, \mathbf{o}_i \rangle + w_0 = 1 - \epsilon_i$) have $\alpha_i^{\mathrm{svm}} > 0$. These examples are known as the *support vectors* [8].

[2]Where a Euclidean metric is used this simply becomes the scalar product of the two vectors

3.2. Generative Kernels

One of the issues with applying SVMs to time varying data, such as speech data, is that the SVM is inherently static in nature. To handle this problem Fisher Kernels [9] and generative kernels [7] have been proposed. Here a generative model that can handle dynamic data is used. An example first-order form of a generative kernel for the set of T observations, $\mathbf{O} = \{\mathbf{o}_1, \ldots, \mathbf{o}_T\}$, may be written as

$$\phi(\mathbf{O}; \boldsymbol{\lambda}) = \frac{1}{T} \begin{bmatrix} \log\left(\breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(1)})\right) - \log\left(\breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(2)})\right) \\ \boldsymbol{\nabla}_{\boldsymbol{\lambda}^{(1)}} \log \breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(1)}) \\ \boldsymbol{\nabla}_{\boldsymbol{\lambda}^{(2)}} \log \breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(2)}) \end{bmatrix} \tag{12}$$

where $\breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(1)})$ and $\breve{p}(\mathbf{O}; \boldsymbol{\lambda}^{(2)})$ are the generative models associated with classes ω_1 and ω_2 respectively.

As SVM training is a distance based learning scheme it is necessary to define an appropriate metric for the distance between two points. The simplest approach is to use a *Euclidean* metric. However, in the same fashion as using the *Mahalanobis*, rather than Euclidean, distances for nearest-neighbour training, an appropriately weighted distance measure may be better. One such metric which is maximally non-committal is given by

$$K(\boldsymbol{\lambda}, \mathbf{O}_i, \mathbf{O}_j) = \phi(\mathbf{O}_i; \boldsymbol{\lambda})' \mathbf{G}^{-1} \phi(\mathbf{O}_j; \boldsymbol{\lambda}) \tag{13}$$

where $\mathbf{O}_i$ and $\mathbf{O}_j$ are sequences of length T_i and T_j respectively, and

$$\mathbf{G} = \mathcal{E}\left\{ \left(\phi(\mathbf{O}; \boldsymbol{\lambda}) - \boldsymbol{\mu}_\phi\right) \left(\phi(\mathbf{O}; \boldsymbol{\lambda}) - \boldsymbol{\mu}_\phi\right)' \right\} \tag{14}$$

where $\boldsymbol{\mu}_\phi = \mathcal{E}\{\phi(\mathbf{O}; \boldsymbol{\lambda})\}$. This will be the form of generative kernel used in this work.

In contrast to many forms of kernel there may be many parameters associated with the generative model. It is therefore sensible to investigate maximum margin training of the generative kernels [10]. Here

$$\{\hat{\boldsymbol{\alpha}}^{\mathrm{svm}}, \hat{\boldsymbol{\lambda}}\} = \arg \max_{\boldsymbol{\alpha}^{\mathrm{svm}}} \min_{\boldsymbol{\lambda}} \left\{ \sum_{i=1}^{n} \alpha_i^{\mathrm{svm}} - \right. \tag{15}$$
$$\left. \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i^{\mathrm{svm}} \alpha_j^{\mathrm{svm}} y_i y_j K(\boldsymbol{\lambda}, \mathbf{O}_i, \mathbf{O}_j) \right\}$$

Unfortunately there is no simple method for optimising the values. A simple iterative process can be used where the support vectors are estimated and then the generative kernel parameters are updated using gradient descent.

4. MAXIMUM MARGIN STATISTICAL MODELS

From the previous section, maximising the margin is a good training criterion when dealing with large dimensional feature-spaces, or little training data. Furthermore as maximum margin training attempts to correctly classify all the training examples, it is inherently discriminatory in nature and is thus an obvious alternative criterion to discriminative criteria such as maximum mutual information (MMI) [11]. The disadvantage of the approach is that it is inherently a binary classification approach. For this work, the binary case will be considered, although schemes have been proposed to handle the multiclass case. For the binary case, there are two sets of augmented model parameters to train, $\{\boldsymbol{\lambda}^{(1)}, \boldsymbol{\alpha}^{(1)}\}$

and $\{\boldsymbol{\lambda}^{(2)}, \boldsymbol{\alpha}^{(2)}\}$. For a binary classification problem, the Bayes' decision is based on

$$\frac{P(\omega_1)\tau^{(2)}\overline{p}(\mathbf{o};\boldsymbol{\lambda}^{(1)},\boldsymbol{\alpha}^{(1)})}{P(\omega_2)\tau^{(1)}\overline{p}(\mathbf{o};\boldsymbol{\lambda}^{(2)},\boldsymbol{\alpha}^{(2)})} \underset{\omega_2}{\overset{\omega_1}{\underset{\lessgtr}{\gtrless}}} 1 \qquad (16)$$

where $P(\omega_1)$ and $P(\omega_2)$ are priors for the two classes. Taking logs of both sides, this may be expressed as

$$\log\left(\frac{\overline{p}(\mathbf{o};\boldsymbol{\lambda}^{(1)},\boldsymbol{\alpha}^{(1)})}{\overline{p}(\mathbf{o};\boldsymbol{\lambda}^{(2)},\boldsymbol{\alpha}^{(2)})}\right) + b \underset{\omega_2}{\overset{\omega_1}{\underset{\lessgtr}{\gtrless}}} 0 \qquad (17)$$

where the class priors and normalisation terms are combined as

$$b = \log\left(\frac{P(\omega_1)\tau^{(2)}}{P(\omega_2)\tau^{(1)}}\right) \qquad (18)$$

Note here the priors for the classes are not trained using ML, but rather using maximum margin. Using equation 2, equation 17 can be rewritten as the scalar product, $(\cdot, \cdot)$,

$$\left(\begin{bmatrix} \mathbf{w} \\ w_0 \end{bmatrix}, \begin{bmatrix} \boldsymbol{\phi}(\mathbf{o};\boldsymbol{\lambda}) \\ 1 \end{bmatrix}\right) \underset{\omega_2}{\overset{\omega_1}{\underset{\lessgtr}{\gtrless}}} 0 \qquad (19)$$

This now has the form of a *score-space*, $\boldsymbol{\phi}(\mathbf{o};\boldsymbol{\lambda})$, which is a function of the base-statistical model parameters

$$\boldsymbol{\phi}(\mathbf{o};\boldsymbol{\lambda}) = \begin{bmatrix} \log\left(\breve{p}(\mathbf{o};\boldsymbol{\lambda}^{(1)})\right) - \log\left(\breve{p}(\mathbf{o};\boldsymbol{\lambda}^{(2)})\right) \\ \boldsymbol{\nabla}_{\boldsymbol{\lambda}}^{(\rho)} \log(\breve{p}(\mathbf{o};\boldsymbol{\lambda})) \end{bmatrix} \qquad (20)$$

and a linear decision boundary which is determined by the augmented model parameters $\boldsymbol{\alpha}$[3]

$$\begin{bmatrix} \mathbf{w} \\ w_0 \end{bmatrix} = \begin{bmatrix} 1 \\ \boldsymbol{\alpha}^{(1)} \\ -\boldsymbol{\alpha}^{(2)} \\ b \end{bmatrix} \qquad (21)$$

One candidate to train the decision boundary is the Support Vector Machine (SVM). This is suitable for these forms of models as the decision boundaries that are estimated achieve good generalisation even when the dimensionality of the feature-space (in this case the score-space) is very large. This may be viewed as maximum margin training of the statistical models. If the SVM is trained, the parameters of the augmented model are given by[4]

$$\alpha_0 \begin{bmatrix} 1 \\ \boldsymbol{\alpha}^{(1)} \\ -\boldsymbol{\alpha}^{(2)} \end{bmatrix} = \alpha_0 \mathbf{w} = \sum_{i=1}^{n} \alpha_i^{\text{svm}} y_i \mathbf{G}^{-1} \boldsymbol{\phi}(\mathbf{o}_i;\boldsymbol{\lambda}) \qquad (22)$$

where $\mathbf{G}$ is given by equation 14. The additional scaling term α_0 has no affect on the decision boundary, but allows standard SVM training to be used[5].

One objection to the use of SVMs is that the distances from the decision boundaries do not have a statistical interpretation [12].

---

[3]Due to the definition of the bias b, there is some interaction between the base statistical model parameters and the $\boldsymbol{\alpha}$

[4]The additional use of the metric $\mathbf{G}$ below is due to training the SVM using an inner product, whereas the augmented model is described in terms of a scalar product.

[5]The values of the bias must also be scaled, hence $b = w_0/\alpha_0$.

This has led to techniques that transform the output so that it is probabilistic in nature [13] and the use of the relevance vector machine [12]. However if generative kernels are used the distance from the decision boundary is directly interpretable as the log-posterior ratio of the two classes. This comes directly from equation 19. It is interesting to contrast this to MMI training [11]. In MMI training the average posterior of the correct label is optimised. In maximum margin training of this form, all correctly labelled points beyond the margin are ignored.
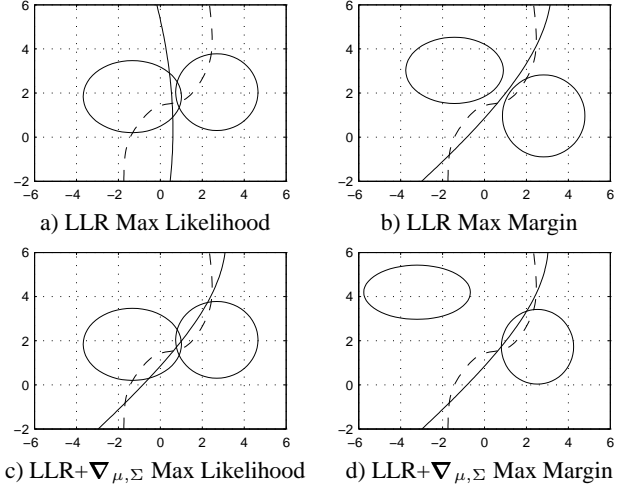

a) LLR Max Likelihood   b) LLR Max Margin

c) LLR+$\boldsymbol{\nabla}_{\mu,\Sigma}$ Max Likelihood   d) LLR+$\boldsymbol{\nabla}_{\mu,\Sigma}$ Max Margin

Fig. 2. (a) Maximum Likelihood (ML) and (b) Maximum Margin (MM) distributions in a Log-Likelihood Ratio (LLR) score-space; and (c) ML base distribution and LLR+$\boldsymbol{\nabla}_{\mu,\Sigma}$ score-space; (d) MM base distribution and a LLR+$\boldsymbol{\nabla}_{\mu,\Sigma}$ score-space.

A simple example of maximum margin training of a statistical model is shown in figure 2 on artificial data generated using three-component GMMs per class. Here single Gaussian component class-conditional models are used as the base distribution. The ML estimate for this model is shown in (a) along with the SVM trained decision boundary from the one-dimension log-likelihood ratio (LLR) score-space

$$\boldsymbol{\phi}(\mathbf{o};\boldsymbol{\lambda}) = \left[\log\left(\breve{p}(\mathbf{o};\boldsymbol{\lambda}^{(1)})\right) - \log\left(\breve{p}(\mathbf{o};\boldsymbol{\lambda}^{(2)})\right)\right] \qquad (23)$$

and the Bayes' decision boundary (the dotted line). The base acoustic model was then trained using maximum margin estimation using equation 15 and the score-space in 23. The decision boundary and positions of the Gaussians are shown in figure (b). The direction of the decision boundary is closer to that of the Bayes' decision boundary and is also reflected in the classification rate in this data. Figures (c) and (d) show the decision boundaries that result from using the LLR with derivatives of the mean and covariances. The decision boundaries that result from (c) and (d) are very similar to that in (b). This is because the class-conditional base distribution was a member of the exponential family (a Gaussian). Since the derivatives with respect to the means and the variances yield first and second order statistics [14], the final distribution should be the same as directly training class-conditional Gaussian distributions using maximum margin.

One of the issues with using maximum margin training with generative kernels is that the final distribution is not guaranteed to

be a valid distribution (though a distance from the decision boundary will always be available). This is best illustrated by examining maximum margin training of a univariate Gaussian class-conditional distribution. Consider a mean and variance first derivative score-space. The resultant maximum margin variance is given by $\sigma^4/(\sigma^2 - \alpha)$ where α is associated with the variance derivative. Thus if $\alpha \geq \sigma^2$ the effective variance will not be positive definite.

For the training to directly correspond to estimating augmented model parameters a linear kernel in the score-space is required. However, non-linear kernels such as polynomial and Gaussian kernels are commonly used. Provided that the normalisation integral (the value of τ) is bounded, the distances from the decision may still be interpreted as a log-posterior ratio. However, this will not have the form of the augmented models described in section 2.

5. LVCSR DECODING

SVMs are inherently a binary classifier. For Large Vocabulary Continuous Speech Recognition (LVCSR) there are a vast number of possible classes making one-v-one binary classification impractical. In order to apply the maximum margin trained statistical models of the previous section to LVCSR it is necessary to map this highly complex classification problem into a set of binary classification problems.

To solve this problem an approach related to that described in [15] is used. A standard, in this case HMM-based, LVCSR system is used to generate a set of word-lattices. These consist of nodes and arcs. The arcs are labelled with words, language and acoustic model likelihoods. The nodes are labelled with time stamps. The word-lattices are then converted to a *confusion network*. This consists of a series of nodes, with a linear graph. Each of the arcs is labelled with a word, a start and end time and a log-posterior, $\mathcal{F}(\omega_i)$. For details of this process see [16]. The confusion network are then pruned so that at each node a maximum of two possible words occur. The pruning is achieved by simply selecting the words with the greatest posteriors.

Once a set of confusion pairs have been generated, it is possible to train a set of statistical models for each pair of data. One issue is the form of score-space to be used. From equation 16 only the unigram prior for the word is available (though the probabilities are trained in a maximum margin fashion). For LVCSR, trigram and higher order language models (LMs) are commonly used. This additional information can be incorporated into the binary classifier using the log-posteriors from the confusion networks. The log-posterior may be treated as an additional information source. Now the decision rule becomes

$$\frac{1}{T} \log \left(\frac{\overline{p}(\mathbf{O}; \boldsymbol{\lambda}^{(1)}, \boldsymbol{\alpha}^{(1)})}{\overline{p}(\mathbf{O}; \boldsymbol{\lambda}^{(2)}, \boldsymbol{\alpha}^{(2)})} \right) + b + \beta \left(\mathcal{F}(\omega_1) - \mathcal{F}(\omega_2) \right) \underset{\omega_2}{\overset{\omega_1}{\gtrless}} 0$$

where β is an empirically set constant. This is a simple process as it only requires combining the log-posterior ratio with the distance from the SVM decision boundary. Alternatively the posterior may be combined into the score-space to give

$$\boldsymbol{\phi}(\mathbf{O}; \boldsymbol{\lambda}) = \begin{bmatrix} \mathcal{F}(\omega_1) - \mathcal{F}(\omega_2) \\ \frac{1}{T} \log \left(\check{p}(\mathbf{O}; \boldsymbol{\lambda}^{(1)}) \right) - \log \left(\check{p}(\mathbf{O}; \boldsymbol{\lambda}^{(2)}) \right) \\ \frac{1}{T} \nabla_{\boldsymbol{\lambda}}^{(\rho)} \log(\check{p}(\mathbf{O}; \boldsymbol{\lambda})) \end{bmatrix} \quad (24)$$

This allows β to be set using maximum margin training. However empirically setting β has a number of advantages. The lattices

used to generate the confusion networks are usually generated using HMMs that have been trained on all the acoustic data. This means that the posteriors obtained from the confusion networks are liable to be biased. Thus the β value will tend to be larger than expected when using held out data.

6. RESULTS

The database used for the LVCSR experiments was a 400 hour subset of the Fisher LDC data. This is the `fsh2004sub` data set used for initial system development [17]. The model set used was based on the standard frontend and models described in [17]. However for this work only ML, rather than discriminatively, trained acoustic models were used. The confusion networks were generated using a bigram language on the same 400hours as used to train the acoustic models. The held-out dataset for these experiments was the eval03 test set, which consists of 6 hours of data.

Word Pair (examples)	Training	CN post.	# Components		
			1	2	4
A/THE (8533)	ML	79.8	58.3	58.4	56.2
	SVM $\boldsymbol{\phi}^{11}()$		61.1	63.0	64.7
	$+\beta$CN		79.8	80.0	80.3
	SVM $\boldsymbol{\phi}^{cn}()$		80.4	80.1	80.6
CAN/CAN'T (3761)	ML	78.5	81.7	86.0	88.2
	SVM $\boldsymbol{\phi}^{11}()$		84.8	89.4	90.5
	$+\beta$CN		88.5	91.2	91.9
	SVM $\boldsymbol{\phi}^{cn}()$		89.0	91.4	91.6
KNOW/NO (4475)	ML	83.1	68.4	69.4	70.8
	SVM $\boldsymbol{\phi}^{11}()$		72.1	73.6	76.6
	$+\beta$CN		84.3	84.5	85.2
	SVM $\boldsymbol{\phi}^{cn}()$		85.7	86.2	86.2

Table 1. 8-Fold cross-validation results (% correct) using variable number of components, ML training of the base model (ML) and SVM training with LLR+$\nabla_{\mu,\Sigma}$ ($\boldsymbol{\phi}^{11}()$) and LLR+$\nabla_{\mu,\Sigma}$+CN posterior ($\boldsymbol{\phi}^{cn}()$) score-spaces with the ML model.

For initial assessment, 8-fold cross-validation experiments were carried out on the training data. For all experiments diagonal covariance matrix GMMs were trained using the longest time-stamps from the confusion networks[6] for the two confusable words. The number of examples for each word pairing were sampled so that the number of positive examples for each word is the same. This means that random selection will yield 50% correct. The GMMs were trained using ML. SVMs were then trained using a first order mean and covariance matrix score-space[7]. A range of SVMs were trained and a few examples are shown in detail in table 1. For all cases using SVMs trained in the likelihood ratio plus derivative score-space gave performance gain over the ML trained base model. For the "CAN/CAN'T" pairing the GMM and SVM sys-

---

[6]By the longest time-stamps the earliest time of the two words and the latest time of the two words is used. This is required as the confusion network times are generated by taking the earliest and latest times that contribute to an arc.

[7]In preliminary experiments the use of maximum margin training of all the model parameters was not found to help performance. Note, for the single component case the SVM trained augmented model is the equivalent of maximum margin training model.

tem were better than the confusion network score used as a baseline. However in general this was not the case. For the "A/THE" pairing the performance was less than 60% for the ML GMM. Using the two forms of combining the confusion network posterior scores from section 5 gains in performance were obtained for most cases. Though schemes where the GMM performance was poor, such as "A/THE", the gains were negligible. Using the score-space including the confusion network posterior gave consistent gains over simply interpolating the information sources.

To examine performance on held-out data the eval03 test set was used. This has a total of 76157 words in the reference transcription. The baseline results using a bigram language model were 34.4% and 33.9% using Viterbi and Confusion Network (CN) decoding respectively, and the baseline numbers using a trigram were 30.8% and 30.1% respectively. As expected the use of CN decoding consistently decreased the error rate. These CN decoding results were used as the baseline for SVM rescoring. Table 2

| SVM Rescoring | #corrected/#pairs (% corrected) | |
	bigram LM	trigram LM
9 SVMs	44/1401 (3.1%)	41/1310 (3.1%)
15 SVMs	55/2116 (2.6%)	43/1954 (2.2%)

Table 2. SVM rescoring giving change in number of errors compared to the CN decoding and total pairs rescored using $\phi^{11}() + \beta$CN on eval03.

shows the results of rescoring with 9 and 15 SVMs trained on confusion pairs from the 400hour training set. All SVM rescoring was based on $\phi^{11}() + \beta$CN, with β roughly tuned to the task[8]. As expected from table 1 there was a range of performances depending on the word pair. The best performance with the bigram LM was using the "CAN/CAN'T" pairing. This reduced the number of errors by 17 in a total of 165 pairs (10.0% reduction). The use of the $\phi^{cn}()$ score-space to find β was worse than the standard CN decoding. This illustrates the dependence of the posterior scores on the exact acoustic/language models used. Overall though the number of errors reduced was small, the percentage of pairs corrected, 3.1%, for 9 pairs indicates that the general approach may be useful. Even using SVMs based on 15 commonly confused pairs less than 3% of the hypothesised words were rescored.

7. CONCLUSION

This paper has described the general form of augmented statistical models. These models are specified by a base distribution and a local exponential family approximation to that distribution at a particular point. The statistics used for the exponential model are based on first, and higher order, derivatives of the base distribution. These models are difficult to train because of the potentially large numbers of model parameters and issues in determining the normalisation term. This paper shows that these augmented models can be trained for a two-class problem using maximum margin training. This is directly related to the use of generative kernels within an SVM framework. Initial experiments using SVM training on a large vocabulary speech recognition task indicate that this form of modelling and training may be useful for speech recognition.

---

[8]The performance was relatively insensitive to β.

8. REFERENCES

[1] M Ostendorff, V V Digalakis, and O A Kimball, "From HMMs to segment models: A unified view of stochastic modelling for speech recognition," *IEEE Transactions Speech and Audio Processing*, vol. 4, pp. 360–378, 1996.

[2] A-V I Rosti and M J F Gales, "Switching linear dynamical systems for speech recogntion," Tech. Rep. CUED/F-INFENG/TR461, Cambridge University, 2003, Available from: svr-www.eng.cam.ac.uk/~mjfg.

[3] J.A. Bilmes, "Buried Markov models: A graphical-modelling approach to automatic speech recognition," *Computer Speech and Language*, vol. 2-3, 2003.

[4] H Nock, *Techniques for modelling Phonological Processes in Automatic Speech Recognition*, Ph.D. thesis, Cambridge University, 2001.

[5] S. Amari and H. Nagaoka, *Methods of Information Geometry*, Oxford University Press, 2000, (Translations of Mathematical Monographs, Volume 191, American Mathematical Society).

[6] N.D. Smith, *Using Augmented Statistical Models and Score Spaces for Classification*, Ph.D. thesis, University of Cambridge, September 2003.

[7] N.D. Smith and M.J.F. Gales, "Speech recognition using SVMs," in *Advances in Neural Information Processing Systems*, 2001.

[8] V.N. Vapnik, *Statistical learning theory*, John Wiley & Sons, 1998.

[9] T. Jaakkola and D. Hausser, "Exploiting generative models in disciminative classifiers," in *Advances in Neural Information Processing Systems 11*, S.A. Solla and D.A. Cohn, Eds. 1999, pp. 487–493, MIT Press.

[10] M. Layton and M.J.F. Gales, "Maximum margin training of generative kernels," Tech. Rep. CUED/F-INFENG/TR.484, Department of Engineering, University of Cambridge, June 2004.

[11] P Brown, *The Acoustic-Modelling Problem in Automatic Speech Recognition*, Ph.D. thesis, IBM T.J. Watson Research Center, 1987.

[12] M.E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, pp. 211–244, 2001.

[13] J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," *Advances in Largin Margin Classifiers*, 2000.

[14] N.D. Smith, M.J.F. Gales, and M. Niranjan, "Data dependent kernels in SVM classification of speech patterns," Tech. Rep. CUED/F-INFENG/TR.387, Department of Engineering, University of Cambridge, April 2001.

[15] V. Venkataramani, S. Chakrabartty, and W Byrne, "Support vector machines for segmental minimum Bayes risk decoding of continuous speech," in *ASRU 2003*, 2003.

[16] L. Mangu, E. Brill, and A. Stolcke, "Finding consensus among words: Lattice-based word error minimization," in *Proc. Eurospeech*, 1999.

[17] G. Evermann, H.Y. Chan, M.J.F. Gales, B. Jia, D. Mrva, P.C. Woodland, and K. Yu, "Training LVCSR systems on thousands of hours of data," Submitted to ICASSP, 2005.