

Kernel Approximation Methods for Speech Recognition

Avner May, Alireza Bagheri Garakani, Zhiyun Lu, Dong Guo, Kuan Liu,
Aurélien Bellet, Linxi Fan, Michael Collins, Daniel Hsu, Brian Kingsbury, et
al.

► **To cite this version:**

Avner May, Alireza Bagheri Garakani, Zhiyun Lu, Dong Guo, Kuan Liu, et al.. Kernel Approximation Methods for Speech Recognition. [Research Report] INRIA Lille. 2017, pp.1-31. <hal-01665417>

HAL Id: hal-01665417

<https://hal.inria.fr/hal-01665417>

Submitted on 15 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Kernel Approximation Methods for Speech Recognition

Avner May^{1†}, Alireza Bagheri Garakani^{2‡}, Zhiyun Lu^{2‡}, Dong Guo^{2‡}, Kuan Liu^{2‡},
Aurélien Bellet³, Linxi Fan⁴, Michael Collins^{1*}, Daniel Hsu¹, Brian Kingsbury⁵,
Michael Picheny⁵, Fei Sha²

¹Dept. of Computer Science, Columbia University, New York, NY 10027, USA
{avnermay, mcollins, djhsu}@cs.columbia.edu, lf2422@columbia.edu

²Dept. of Computer Science, University of Southern California, Los Angeles, CA 90089, USA
{bagherig, zhiyunlu, dongguo, kuanl, feisha}@usc.edu

³INRIA, 40 Avenue Halley, 59650 Villeneuve d’Ascq, France
aurelien.bellet@inria.fr

⁴Dept. of Computer Science, Stanford University, Stanford, CA 94305, USA
jimfan@cs.stanford.edu

⁵IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA
{bedk, picheny}@us.ibm.com

^{† ‡}: *Contributed equally as the first and second co-authors, respectively*

Abstract

We study large-scale kernel methods for acoustic modeling in speech recognition and compare their performance to deep neural networks (DNNs). We perform experiments on four speech recognition datasets, including the TIMIT and Broadcast News benchmark tasks, and compare these two types of models on frame-level performance metrics (accuracy, cross-entropy), as well as on recognition metrics (word/character error rate). In order to scale kernel methods to these large datasets, we use the random Fourier feature method of Rahimi and Recht [2007]. We propose two novel techniques for improving the performance of kernel acoustic models. First, in order to reduce the number of random features required by kernel models, we propose a simple but effective method for feature selection. The method is able to explore a large number of non-linear features while maintaining a compact model more efficiently than existing approaches. Second, we present a number of frame-level metrics which correlate very strongly with recognition performance when computed on the heldout set; we take advantage of these correlations by monitoring these metrics during training in order to decide when to stop learning. This technique can noticeably improve the recognition performance of both DNN and kernel models, while narrowing the gap between them. Additionally, we show that the linear bottleneck method of Sainath et al. [2013a] improves the performance of our kernel models significantly, in addition to speeding up training and making the models more compact. Together, these three methods dramatically improve the performance of kernel acoustic models, making their performance comparable to DNNs on the tasks we explored.

Keywords: Kernel Methods, Deep Neural Networks, Acoustic Modeling, Automatic Speech Recognition, Feature Selection, Logistic Regression.

*On leave at Google Inc. New York.

1 Introduction

In recent years, deep learning techniques have significantly advanced state-of-the-art performance in automatic speech recognition (ASR), achieving large drops in word error rates [Seide et al., 2011a, Hinton et al., 2012, Mohamed et al., 2012, Xiong et al., 2016]. Deep neural networks (DNNs) are able to gracefully scale to very large datasets, and can successfully leverage this additional data to achieve strong empirical performance. In stark contrast, kernel methods, which are attractive due to their powerful modeling of highly nonlinear data, as well as for their theoretical learning guarantees and tractability [Schölkopf and Smola, 2002], do not scale well. In particular, with data sets of size N , the $\Theta(N^2)$ size of the kernel matrix makes training prohibitively slow, while the typical $\Theta(N)$ size of the resulting models [Steinwart, 2004] makes their deployment impractical.

Much recent effort has been devoted to the development of approximations to kernel methods, primarily via the Nyström approximation [Williams and Seeger, 2001] or via random feature expansion [e.g., Rahimi and Recht, 2007, Kar and Karnick, 2012]. These methods yield explicit feature representations on which linear learning methods can provide good approximations to the original non-linear kernel method. However, there have been very few successful applications of these methods to ASR, let alone any “head-on” comparisons to DNNs, except for a few efforts which were limited in scope [Deng et al., 2012, Cheng and Kingsbury, 2011, Huang et al., 2014].

In this paper, we investigate empirically how kernel methods can be scaled to tackle typical ASR tasks. We focus on four datasets: the IARPA Babel Program Cantonese (IARPA-babel101-v0.4c) and Bengali (IARPA-babel103b-v0.4b) limited language packs, a 50-hour subset of Broadcast News (BN-50) [Kingsbury, 2009, Sainath et al., 2011], and TIMIT [Garofolo et al., 1993]. We present several results: First, we show that kernel methods can be efficiently scaled to large-scale ASR tasks, using the above-mentioned random Fourier feature technique [Rahimi and Recht, 2007]. Our contribution is to demonstrate the practical utility of this method in constructing large-scale classifiers for acoustic modeling. Second, we have found that when leveraging the novel techniques discussed in this paper, our kernel-based acoustic models are generally competitive with layer-wise discriminatively pre-trained DNN-based models [Seide et al., 2011b].

In order to attain strong performance for the kernel acoustic models, we have developed a few new methods. First, we propose a simple feature selection algorithm, which effectively reduces the number of random features required. We iteratively select features from large pools of random features, using learned weights in the selection criterion. This has two clear benefits: (i) the subsequent training on the selected features is considerably faster than training on the entire pool of random features, and (ii) the resulting model is also much smaller. For certain kernels, this feature selection approach—which is applied at the level of the random features—can be regarded as a non-linear method for feature selection at the level of the input features, and we use this observation to motivate the design of a new kernel function.

Second, we present several novel frame-level metrics which correlate very strongly with the token error rate (TER),¹ and which can thus be monitored on the heldout set during training in order to determine when to stop learning. Using this method, we achieve notable gains in TER for both kernels and DNNs. This method partially mitigates a well-known problem in acoustic modeling; namely, that the training criterion (cross-entropy) often does not align well with the true objective (TER). In our case, we noticed that although our kernel and DNN models would often attain very similar cross-entropy values on the heldout set, the DNNs would generally perform better, sometimes by a wide margin, in terms of TER. Although sequence training techniques can also be used to address this issue [e.g., Kingsbury, 2009, Veselý et al., 2013], they are

¹For our Cantonese dataset, ‘token error rate’ corresponds to ‘character error rate.’ For our Bengali and Broadcast News datasets, it corresponds to ‘word error rate.’ For TIMIT, it corresponds to ‘phone error rate.’

very computationally expensive, and they generally depend on frame-level training for initialization; thus, our proposed method can be used in conjunction with existing sequence training techniques, by providing them with a better initial model.

Lastly, we demonstrate the importance of using a linear bottleneck [Sainath et al., 2013a] in the parameter matrix of our kernel models. Not only does this method improve the performance of our kernel models significantly, it also makes training faster, and reduces the size of the models learned.

This paper builds on the previous works of Lu et al. [2016] and May et al. [2016]. Lu et al. [2016] provide comparisons between DNN and kernel acoustic models on 3 datasets (Cantonese, Bengali, and Broadcast News); they additionally present the “entropy regularized log loss” (ERLL) metric, and show how using it as a model selection criterion can yield TER improvements.² May et al. [2016] present the feature selection algorithm described in this paper, along with ASR experiments on two datasets (Cantonese and Bengali); comparisons to DNNs are also performed. The work in the current paper builds on this existing work in several ways. First, we provide a more extensive set of experiments, including results on the ASR benchmark TIMIT dataset, in addition to updated results on the other datasets (Cantonese, Bengali, and Broadcast News). Second, we have extended the work on ERLL by presenting a larger set of metrics which correlate strongly with TER; these additional metrics help explain the unusual correlation between ERLL and TER, as well as the poor correlation between cross-entropy and TER. In this paper, we show how these metrics can be evaluated on the heldout set during training in order to decide when to decay the learning rate and stop training. Lastly, we provide an extensive set of experiments showing the importance of using a linear bottleneck for attaining strong TER performance for our kernel methods.

The rest of the paper is organized as follows. We review related work in §2. We provide some background for kernel approximation methods, as well as for acoustic modeling, in §3. We present our feature selection algorithm in §4. In §5, we present several novel metrics which correlate strongly with TER, and show how they can be used during training to improve TER performance. In §6, we report extensive experiments comparing DNNs and kernel methods, including results using the methods discussed above. We conclude in §7.

2 Related Work

Scaling up kernel methods has been a long-standing and actively studied problem [Bottou et al., 2007, Smola, 2014, DeCoste and Schölkopf, 2002, Platt, 1998, Tsang et al., 2005, Clarkson, 2010]. For kernels with sparse feature expansions, Sonnenburg and Franc [2010] show how to efficiently scale kernel SVMs to datasets with up to 50 million training samples by using sparse vector operations for parameter updates. Approximating kernels by constructing explicit finite-dimensional feature representations, where dot products between these representations approximate the kernel function, has emerged as a powerful technique [e.g., Williams and Seeger, 2001, Rahimi and Recht, 2007]. The Nyström method constructs these feature maps, for arbitrary kernels, via a low-rank decomposition of the kernel matrix [Williams and Seeger, 2001]. For shift-invariant kernels, the random Fourier feature technique of Rahimi and Recht [2007] uses random projections in order to generate the features. Random projections can also be used to approximate a wider range of kernels [Kar and Karnick, 2012, Vedaldi and Zisserman, 2012, Hamid et al., 2014, Pennington et al., 2015]. Many recent works have been developed to speed-up the random Fourier feature approach to kernel approximation. One line of work attempts to reduce the time (and memory) needed to compute the random feature expansions by imposing structure on the random projection matrix [Le et al., 2013, Yu

²In Lu et al. [2016] this metric was called “entropy regularized perplexity” (ERP).

et al., 2015]. It is also possible to use doubly-stochastic methods to speed-up stochastic gradient training of models based on the random features [Dai et al., 2014].

Despite much progress in kernel approximation, there have been only a few reported empirical studies of these techniques on speech recognition tasks [Deng et al., 2012, Cheng and Kingsbury, 2011, Huang et al., 2014]. However, those tasks were relatively small-scale (for instance, on the TIMIT dataset). For the most part, a detailed evaluation of these methods on large-scale ASR tasks, together with a thorough comparison with DNNs, is lacking. Our work fills this gap, tackling challenging large-scale acoustic modeling problems, where deep neural networks achieve strong performance [Hinton et al., 2012, Dahl et al., 2012]. Additionally, we provide a number of important improvements to the kernel methods, which boost their performance significantly.

One contribution of our work is to introduce a feature selection method that works well in conjunction with random Fourier features in the context of large-scale multi-class classification problems. Recent work on feature selection methods with random Fourier features, for binary classification and regression problems, includes the Sparse Random Features algorithm of Yen et al. [2014]. This algorithm is a coordinate descent method for smooth convex optimization problems in the (infinite) space of non-linear features: each step involves solving a batch ℓ_1 -regularized convex optimization problem over randomly generated non-linear features (note that a natural extension of this method to multi-class problems is to use mixed norms such as ℓ_1/ℓ_2). Here, the ℓ_1 -regularization may cause the learned solution to only depend on a subset of the generated features. A drawback of this approach is the computational burden of fully solving many batch optimization problems, which is prohibitive for large data sets. In our attempts to implement an online variant of this method, using FOBOS [Duchi and Singer, 2009] and ℓ_1/ℓ_2 -regularization for the multi-class setting, we observed that very strong regularization was required to obtain any intermediate sparsity, which in turn severely hurt prediction performance. Effectively, the regularization was so strong that it made the learning meaningless, and the selected features were basically random. Our approach for selecting random features is more efficient, and more directly ensures sparsity, than regularization.

Another improvement we propose alters the frame-level training of the acoustic model in order to improve the recognition performance (TER) of the final model. A set of methods, typically referred to as *sequence training* techniques, share our goal of tuning the acoustic model for the purpose of improving its recognition performance. There are a number of different sequence training criteria which have been proposed, including maximum mutual information (MMI) [Bahl et al., 1986, Valtchev et al., 1997], boosted MMI (BMMI) [Povey et al., 2008], minimum phone error (MPE) [Povey and Woodland, 2002], or minimum Bayes risk (MBR) [Kaiser et al., 2000, Gibson and Hain, 2006, Povey and Kingsbury, 2007]. These methods, though originally proposed for training Gaussian mixture model (GMM) acoustic models, can also be used for neural network acoustic models [Kingsbury, 2009, Veselý et al., 2013]. Nonetheless, all of these methods are quite computationally expensive and are typically initialized with an acoustic model trained via the frame-level cross-entropy criterion. Our method, by contrast, is very simple, only making a small change to the frame-level training process. Furthermore, it can be used in conjunction with the above-mentioned sequence training techniques, by providing a better initial model. Recently, Povey et al. [2016] showed that it is possible to train an acoustic model using *only* sequence-training methods, with the lattice-free version of the MMI criterion. For future work, we would like to see how much our kernel models can benefit from the various sequence training methods mentioned above, relative to DNNs.

This work also contributes to the debate on the relative strengths of deep and shallow neural networks. As explained in Section 3.4, many types of kernels (including popular kernels like the Gaussian kernel and the Laplacian kernel) can be understood as shallow neural networks. As such, comparing kernel methods to DNNs is also in a sense comparing shallow and deep neural networks. There is much literature on this topic.

Classic results show that both deep and shallow neural networks are “universal approximators,” meaning that they can approximate any real-valued continuous function with bounded support to an arbitrary degree of precision [Cybenko, 1989, Hornik et al., 1989]. However, a number of papers have argued that there exist functions which deep neural networks can express with exponentially fewer parameters than shallow neural networks [Montúfar et al., 2014, Bianchini and Scarselli, 2014]. In Ba and Caruana [2014], the authors show that the performance of shallow neural networks can be increased considerably by training them to match the outputs of deep neural networks. In showing that kernel methods can compete with DNNs on large-scale speech recognition tasks, this paper adds credence to the argument that shallow networks can perform on par with deep networks.

3 Background

3.1 Kernel Methods and Random Features

Kernel methods, broadly speaking, are a set of machine learning techniques which either explicitly or implicitly map data from the input space \mathcal{X} to some feature space \mathcal{H} , in which a linear model is learned. A “kernel function” $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is then defined³ as the function which takes as input $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, and returns the dot-product of the corresponding points in \mathcal{H} . If we let $\phi: \mathcal{X} \rightarrow \mathcal{H}$ denote the map into the feature space, then $K(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle$. Standard kernel methods avoid inference in \mathcal{H} , because it is generally a very high-dimensional, or even infinite-dimensional, space. Instead, they solve the dual problem by using the N -by- N kernel matrix, containing the values of the kernel function applied to all pairs of N training points. When $\dim(\mathcal{H})$ is far greater than N , this “kernel trick” provides a nice computational advantage. However, when N is exceedingly large, the $\Theta(N^2)$ size of the kernel matrix makes training impractical.

Rahimi and Recht [2007] address this problem by leveraging Bochner’s Theorem, a classical result in harmonic analysis, in order to provide a fast way to approximate any positive-definite *shift-invariant* kernel K with *finite*-dimensional features. A kernel $K(\mathbf{x}, \mathbf{y})$ is *shift-invariant* if and only if $K(\mathbf{x}, \mathbf{y}) = \hat{K}(\mathbf{x} - \mathbf{y})$ for some function $\hat{K}: \mathbb{R}^d \rightarrow \mathbb{R}$. We now present Bochner’s Theorem:

Theorem 1 (*Bochner’s theorem, adapted from Rahimi and Recht [2007]*): *A continuous shift-invariant kernel $K(\mathbf{x}, \mathbf{y}) = \hat{K}(\mathbf{x} - \mathbf{y})$ on \mathbb{R}^d is positive-definite if and only if \hat{K} is the Fourier transform of a non-negative measure $\mu(\boldsymbol{\omega})$.*

Thus, for any positive-definite shift-invariant kernel $\hat{K}(\boldsymbol{\delta})$, we have that

$$\hat{K}(\boldsymbol{\delta}) = \int_{\mathbb{R}^d} \mu(\boldsymbol{\omega}) e^{-j\boldsymbol{\omega}^T \boldsymbol{\delta}} d\boldsymbol{\omega}, \tag{1}$$

where

$$\mu(\boldsymbol{\omega}) = (2\pi)^{-d} \int_{\mathbb{R}^d} \hat{K}(\boldsymbol{\delta}) e^{j\boldsymbol{\omega}^T \boldsymbol{\delta}} d\boldsymbol{\delta} \tag{2}$$

is the inverse Fourier transform⁴ of $\hat{K}(\boldsymbol{\delta})$, and where $j = \sqrt{-1}$. By Bochner’s theorem, $\mu(\boldsymbol{\omega})$ is a non-negative measure. As a result, if we let $Z = \int_{\mathbb{R}^d} \mu(\boldsymbol{\omega}) d\boldsymbol{\omega}$, then $p(\boldsymbol{\omega}) = \frac{1}{Z} \mu(\boldsymbol{\omega})$ is a proper probability

³It is also possible to define the kernel function prior to defining the feature map; then, for positive-definite kernel functions, Mercer’s theorem guarantees that a corresponding feature map ϕ exists such that $K(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle$.

⁴There are various ways of defining the Fourier transform and its inverse. We use the convention specified in Equations (1) and (2), which is consistent with Rahimi and Recht [2007].

Kernel name	$K(\mathbf{x}, \mathbf{y})$	$p(\boldsymbol{\omega})$	Density name
Gaussian	$e^{-\ \mathbf{x}-\mathbf{y}\ _2^2/2\sigma^2}$	$(2\pi(1/\sigma^2))^{-d/2} e^{-\frac{\ \boldsymbol{\omega}\ _2^2}{2(1/\sigma^2)}}$	Normal($\mathbf{0}_d, \frac{1}{\sigma^2} \mathbb{1}_d$)
Laplacian	$e^{-\lambda\ \mathbf{x}-\mathbf{y}\ _1}$	$\prod_{i=1}^d \frac{1}{\lambda\pi(1+(\boldsymbol{\omega}_i/\lambda)^2)}$	Cauchy($\mathbf{0}_d, \lambda$)

Table 1: Gaussian and Laplacian Kernels, together with their sampling distributions $p(\boldsymbol{\omega})$

distribution, and we get that

$$\frac{1}{Z} \hat{K}(\boldsymbol{\delta}) = \int_{\mathbb{R}^d} p(\boldsymbol{\omega}) e^{-j\boldsymbol{\omega}^T \boldsymbol{\delta}} d\boldsymbol{\omega}.$$

For simplicity, we will assume going forward that \hat{K} is properly-scaled, meaning that $Z = 1$. Now, the above equation allows us to rewrite this integral as an expectation:

$$\hat{K}(\boldsymbol{\delta}) = \hat{K}(\mathbf{x} - \mathbf{y}) = \int_{\mathbb{R}^d} p(\boldsymbol{\omega}) e^{j\boldsymbol{\omega}^T (\mathbf{x} - \mathbf{y})} d\boldsymbol{\omega} = \mathbb{E}_{\boldsymbol{\omega}} \left[e^{j\boldsymbol{\omega}^T \mathbf{x}} e^{-j\boldsymbol{\omega}^T \mathbf{y}} \right]. \quad (3)$$

This can be further simplified as

$$\hat{K}(\mathbf{x} - \mathbf{y}) = \mathbb{E}_{\boldsymbol{\omega}, b} \left[\sqrt{2} \cos(\boldsymbol{\omega}^T \mathbf{x} + b) \cdot \sqrt{2} \cos(\boldsymbol{\omega}^T \mathbf{y} + b) \right],$$

where $\boldsymbol{\omega}$ is drawn from $p(\boldsymbol{\omega})$, and b is drawn uniformly from $[0, 2\pi]$. See Appendix A for details on why this specific functional form is correct.

This motivates a sampling-based approach for approximating the kernel function. Concretely, we draw $\{\boldsymbol{\omega}_1, \boldsymbol{\omega}_2, \dots, \boldsymbol{\omega}_D\}$ independently from the distribution $p(\boldsymbol{\omega})$, and $\{b_1, b_2, \dots, b_D\}$ independently from the uniform distribution on $[0, 2\pi]$, and then use these parameters to approximate the kernel, as follows:

$$K(\mathbf{x}, \mathbf{y}) \approx \frac{1}{D} \sum_{i=1}^D \sqrt{2} \cos(\boldsymbol{\omega}_i^T \mathbf{x} + b_i) \cdot \sqrt{2} \cos(\boldsymbol{\omega}_i^T \mathbf{y} + b_i) = z(\mathbf{x})^T z(\mathbf{y}),$$

where $z_i(\mathbf{x}) = \sqrt{\frac{2}{D}} \cos(\boldsymbol{\omega}_i^T \mathbf{x} + b_i)$ is the i^{th} element of the D -dimensional random vector $z(\mathbf{x})$. In Table 1, we list two popular (properly-scaled) positive-definite kernels with their respective inverse Fourier transforms.

Using these random feature maps, in conjunction with linear learning algorithms, can yield huge gains in efficiency relative to standard kernel methods on large datasets. Learning with a representation $z(\cdot) \in \mathbb{R}^D$ is relatively efficient provided that D is far less than the number of training samples N . For example, in our experiments (see Section 6), we have 2 million to 16 million training samples, while $D \approx 25,000$ often leads to good performance.

Rahimi and Recht [2007, 2008] prove a number of important theoretical results about these random feature approximations. First, they show that if $D = \tilde{\Omega}(\frac{d}{\epsilon^2})$, then with high probability $z(\mathbf{x})^T z(\mathbf{y})$ will be within ϵ of $K(\mathbf{x}, \mathbf{y})$ for all \mathbf{x}, \mathbf{y} in some compact subset $\mathcal{M} \in \mathbb{R}^d$ of bounded diameter.⁵ See claim 1 of Rahimi and Recht [2007] for the more precise statement and proof of this result.

In their follow-up work [Rahimi and Recht, 2008], the authors prove a generalization bound for models learned using these random features. They show that with high-probability, the excess risk⁶ assumed from using this approximation, relative to using the ‘‘oracle’’ kernel model (the exact kernel model with the lowest

⁵We are using the $\tilde{\Omega}$ notation to hide logarithmic factors.

⁶The ‘‘risk’’ of a model is defined as its expected loss on unseen data.

risk), is bounded by $O(\frac{1}{\sqrt{N}} + \frac{1}{\sqrt{D}})$ (see the main result of Rahimi and Recht [2008] for more details). Given that the generalization error of a model trained using exact kernel methods is known to be within $O(\frac{1}{\sqrt{N}})$ of the oracle model [Bartlett et al., 2002], this implies that in the worst case, $D = \Theta(N)$ random features may be required in order for the approximated model to achieve generalization performance comparable to the exact kernel model. Empirically, however, fewer than $\Theta(N)$ features are often needed in order to attain strong performance [Yu et al., 2015].

3.2 Using Neural Networks for Acoustic Modeling

Neural network acoustic models provide a conditional probability distribution $p(y|\mathbf{x})$ over C possible acoustic states, conditioned on an acoustic frame \mathbf{x} encoded in some feature representation. The acoustic states correspond to context-dependent phoneme states [Dahl et al., 2012], and in modern speech recognition systems, the number of such states is of the order 10^3 to 10^4 . The acoustic model is used within probabilistic systems for decoding speech signals into word sequences. Typically, the probability model used is a hidden Markov model (HMM), where the model’s emission and transition probabilities are provided by an acoustic model together with a language model. We use Bayes’ rule in order to compute the probability $p(\mathbf{x}|y)$ of emitting a certain acoustic feature vector \mathbf{x} from state y , given the output $p(y|\mathbf{x})$ of the neural network:

$$p(\mathbf{x}|y) = \frac{p(y|\mathbf{x})p(\mathbf{x})}{p(y)}.$$

Note that $p(\mathbf{x})$ can be ignored at inference time because it doesn’t affect the relative scores assigned to different word sequences, and $p(y)$ is simply the prior probability of HMM state y . The Viterbi algorithm can then be used to determine the most likely word sequence (see Gales and Young [2007] for an overview of using HMMs for speech recognition).

3.3 Using Random Fourier Features for Acoustic Modeling

In order to train an acoustic model using random Fourier features, we can simply plug the random feature vector $z(\mathbf{x})$ (for an acoustic frame \mathbf{x}) into a multinomial logistic regression model:

$$p(y|\mathbf{x}) = \frac{\exp(\langle \boldsymbol{\theta}_y, z(\mathbf{x}) \rangle)}{\sum_{y'} \exp(\langle \boldsymbol{\theta}_{y'}, z(\mathbf{x}) \rangle)}. \quad (4)$$

The label y can take any value in $\{1, 2, \dots, C\}$, each corresponding to a context-dependent phonetic state label, and the parameter matrix $\Theta = [\boldsymbol{\theta}_1 | \dots | \boldsymbol{\theta}_C]$ is learned. Note that we also include a bias term, by appending a 1 to $z(\mathbf{x})$ in the equation above.

3.4 Viewing Kernel Acoustic Models as Shallow Neural Networks

The model in Equation (4) can be seen as a shallow neural network, with the following properties: (1) the parameters from the inputs (i.e., acoustic feature vectors) to the hidden units are set randomly, and are not learned; (2) the hidden units use $\cos(\cdot)$ as their activation function; (3) the parameters from the hidden units to the output units are learned (can be optimized with convex optimization); and (4) the softmax function is used to normalize the outputs of the network. See Figure 1 for a visual representation of this model architecture.




Figure 1: Kernel-acoustic model seen as a shallow neural network

3.5 Linear Bottlenecks

The number of phonetic state labels can be very large. This will significantly increase the number of parameters in Θ . We can reduce this number with a *linear bottleneck* layer between the hidden layer and the output layer; the linear bottleneck corresponds to a low-rank factorization $\Theta = UV$ of the parameter matrix [Sainath et al., 2013a]. This is particularly important for our kernel models, where the number of trainable parameters is $D \times C$, where D is the number of random features, and C is the number of output classes. Using a linear bottleneck of size r , this can be reduced to $(D \times r + r \times C)$, which is significantly less than $D \times C$ when $r \ll \min(D, C)$. This strictly decreases the capacity of the resulting model, while unfortunately rendering the optimization problem non-convex.

4 Random Feature Selection

In this section, we first motivate and describe our proposed feature selection algorithm. We then introduce a new “Sparse Gaussian” kernel, which performs well in conjunction with the feature selection algorithm.

4.1 Proposed Feature Selection Algorithm

Our proposed random feature selection method, shown in Algorithm 1, is based on a general iterative framework. In each iteration, random features are generated and added to the current set of features; a subset of these features are selected, while the rest are discarded. The selection process works as follows: first, a model is trained on the current set of features using a single pass of stochastic gradient descent (SGD) on a subset of the training data. Then, the features whose corresponding rows in Θ have the largest ℓ_2 norms are kept. Note that the row of weights corresponding to the feature $z_i(\mathbf{x})$ in the model $f(\mathbf{x}) = h(\Theta^T z(\mathbf{x}))$ are those in the i^{th} row of Θ . In the case where we are using a linear bottleneck to decompose Θ into UV , we perform the SGD training using this decomposition. After we complete the training in a given iteration, we compute $\Theta = UV$, and then select features based on the ℓ_2 -norms of the rows of Θ .

This feature selection method has the following advantages: The overall computational cost is mild, as it requires just T passes through subsets of the data of size R (equivalent to $\sim TR/N$ full SGD epochs). In fact, in our experiments, we find it sufficient to use $R = O(D)$. Moreover, the method is able to explore a large number of non-linear features, while maintaining a compact model. If $s_t = Dt/T$, then the learning algorithm is exposed to $D(T + 1)/2$ random features throughout the feature selection process; this is the

Algorithm 1 Random feature selection

input Target number of random features D , data subset size R ,
Integers (s_1, \dots, s_{T-1}) such that $0 < s_1 < \dots < s_{T-1} < D$, specifying selection schedule.

- 1: **initialize** set of selected indices $S := \emptyset$.
- 2: **for** $t = 1, 2, \dots, T$ **do**
- 3: **for** $i \in \{1, \dots, D\} \setminus S$ **do**
- 4: $\omega_i \sim p(\omega)$.
- 5: $b_i \sim \mathcal{U}(0, 2\pi)$.
- 6: **end for**
- 7: **if** $t \neq T$ **then**
- 8: Initialize parameter matrix Θ .⁷
- 9: Learn weights $\Theta \in \mathbb{R}^{D \times C}$ using a single pass of SGD over R randomly selected training examples, using the projection vectors $(\omega_1, \dots, \omega_D)$, and the biases (b_1, \dots, b_D) , to generate the random Fourier features.
- 10: $S := \{i \mid \Theta_i \text{ is amongst the } s_t \text{ rows of } \Theta \text{ with highest } \ell_2 \text{ norm}\}$.
- 11: **end if**
- 12: **end for**
- 13: **return** The selected projection vectors $(\omega_1, \dots, \omega_D)$, and the selected biases (b_1, \dots, b_D) .

selection schedule we used in all our experiments. We show in Section 6 that this empirically increases the predictive quality of the selected non-linear features.

It is important to note the similarities between this method, and the FOBOS method with ℓ_1/ℓ_2 -regularization [Duchi and Singer, 2009]. In the latter method, one solves the ℓ_1/ℓ_2 -regularized problem in a stochastic fashion by alternating between taking unregularized stochastic gradient descent (SGD) steps, and then “shrinking” the rows of the parameter matrix; each time the parameters are shrunk, the rows whose ℓ_2 -norms are below a threshold are set to 0. After training completes, the solution will likely have some rows which are all zero, at which point the features corresponding to those rows can be discarded. In our method, on the other hand, we take many consecutive unregularized SGD steps, and only thereafter do we choose to discard the rows whose ℓ_2 -norm is below a threshold. As mentioned in the Related Work section, our attempts at using FOBOS for feature selection failed, because the magnitude of the regularization parameter needed in order to produce a sparse model was so large that it dominated the learning process; as a result, the models learned performed badly, and the selected features were essentially random.

One disadvantage of our method is that the index used for selection may misrepresent the features’ actual predictive utilities. For instance, the presence of some random feature may increase or decrease the weights for other random features relative to what they would be if that feature were not present. An alternative would be to consider features in isolation, and add features one at a time (as in stagewise regression methods and boosting), but this would be significantly more computationally expensive. For example, it would require $O(D)$ passes through the data, relative to $O(T)$ passes, which would be prohibitive for large D values. We find empirically that the influence of the additional random features in the selection criterion is tolerable, and it is still possible to select useful features with this method.

⁷See Section 6.2 for details on how Θ is initialized.

4.2 A Sparse Gaussian Kernel

Recall from Table 1 that for the Laplacian kernel, the sampling distribution used for the random Fourier features is the multivariate Cauchy density $p(\boldsymbol{\omega}) \propto \prod_{i=1}^d (1 + \omega_i^2)^{-1}$ (we let $\lambda = 1$ here for simplicity). If we draw $\boldsymbol{\omega} = (\omega_1, \dots, \omega_d)$ from p , then each ω_i has a two-sided fat tail distribution, and hence $\boldsymbol{\omega}$ will typically contain some entries much larger than the rest.

This property of the sampling distribution implies that many of the random features generated in this way will effectively concentrate on only a few of the input features. We can thus regard such random features as being non-linear combinations of a small number of the original input features. Thus, the proposed feature selection method effectively picks out useful non-linear interactions between small sets of input features.

We can also directly construct sparse non-linear combinations of the input features. Instead of relying on the properties of the Cauchy distribution, we can actually choose a small number k of coordinates $F \subseteq \{1, 2, \dots, d\}$, say, uniformly at random, and then choose the random vector $\boldsymbol{\omega}$ so that it is always zero in positions outside of F ; the same non-linearity (e.g., $\mathbf{x} \mapsto \cos(\boldsymbol{\omega}^T \mathbf{x} + b)$) can be applied once the sparse random vector is chosen. Compared to the random Fourier feature approximation to the Laplacian kernel, the vectors $\boldsymbol{\omega}$ chosen in this way are truly sparse, which can make the random feature expansion more computationally efficient to apply (if efficient sparse matrix operations are used).

Note that random Fourier features with such sparse sampling distributions in fact correspond to shift-invariant kernels that are rather different from the Laplacian kernel. For instance, if the non-zero entries of $\boldsymbol{\omega}$ are drawn i.i.d. from $\mathcal{N}(0, \sigma^{-2})$, then the corresponding kernel is

$$K(\mathbf{x}, \mathbf{y}) = \binom{d}{k}^{-1} \sum_{\substack{F \subseteq \{1, \dots, d\} \\ |F|=k}} \exp\left(-\frac{\|\mathbf{x}_F - \mathbf{y}_F\|_2^2}{2\sigma^2}\right), \quad (5)$$

where \mathbf{x}_F is a vector composed of the elements x_i for $i \in F$. The kernel in Equation (5) puts equal emphasis on all input feature subsets F of size k . However, the feature selection process may effectively bias the distribution of the feature subsets to concentrate on some small family \mathcal{F} of input feature subsets.

5 New Early Stopping Criteria

As discussed in the introduction, there is a well-known problem in the training of acoustic models; namely, that the training criterion (cross-entropy) does not perfectly correlate with the true objective (TER). Consequently, lowering the cross-entropy performance on a heldout set does not necessarily result in better TER performance. For example, we noticed that our DNNs were often attaining stronger TER performance than our kernel models, even though they had comparable cross-entropy performance. In order to partially address this problem, in this section we present several new metrics whose empirical correlation with TER, amongst the fully trained DNN and kernel models we trained, was high. We then leverage these metrics during training by evaluating them on the heldout set after each epoch in order to decide when to decay the learning rate and stop training (see Section 6.2 for details on how we decay the learning rate). Note that the reason we use these metrics as proxies for the TER, instead of directly using the TER, is that it is very expensive to compute the TER on the development set.

The common thread which unites all the metrics we will present is that they do not penalize very incorrect examples (meaning, examples for which the model assigned a probability very close to 0 to the correct label) as strongly as cross-entropy does. Notice, for instance, that there is no limit to how much the

cross-entropy loss (e.g., log loss) can penalize a single incorrect example. Our metrics are more lenient. We present them now:

1. **“Entropy Regularized Log Loss (ERLL):”** This loss rewards models for being confident (e.g., low entropy), by considering a weighted sum of the cross entropy loss (CE) and the average entropy (ENT) of the model on the heldout data. Specifically, for any $\beta \in \mathbb{R}$ (typically we take $\beta = 1$), we define the loss as follows:

$$CE + \beta \cdot ENT = -\frac{1}{N} \sum_{i=1}^N \sum_{y=1}^C [\mathbb{I}(y = y_i) + \beta \cdot p(y|x_i)] \log p(y|x_i)$$

This metric encourages models to be more confident, even if it means having a worse cross-entropy loss as a result.

2. **“Capped Log Loss:”** For any value of $\lambda \geq 0$, we can define:

$$-\frac{1}{N} \sum_{i=1}^N \log(p(y_i|x_i) + \lambda).$$

Effectively, this loss ensures that no single example contributes more than $-\frac{1}{N} \log(\lambda)$ to the loss. If λ is a small positive number, this loss is very similar to the normal log loss for values of $p(y_i|x_i)$ close to 1, while affecting the loss dramatically for values close to 0 (for example, when $p(y_i|x_i) < \lambda$).

3. **“Top-k Log Loss:”** For this loss, assume that the heldout examples (\mathbf{x}_i, y_i) are sorted in descending order of their $p(y_i|x_i)$ values. Now, for any positive integer $k \leq N$, we can define the “Top-k Log Loss” as follows:

$$-\frac{1}{k} \sum_{i=1}^k \log p(y_i|x_i).$$

This metric judges a model based on how well it does on the k heldout examples to which it assigns highest probabilities.

Notice that for $\beta = 0$, $\lambda = 0$, and $k = N$, these metrics all simplify to the standard log loss. In Figure 2, we show plots of the empirical correlations of these metrics with TER values, as a function of each metric’s “hyperparameters,” based on models we have trained. More specifically, we *fully* train a large number of kernel and DNN models, and then evaluate the TER performance of these models on the development set, as well as compute the heldout performance of these models in terms of the 3 metrics described above (for various settings of β , λ , and k). The precise set of models we used are those in Tables 3 and 4 (see Section 6.3 for details). We then compute the empirical correlations between these values, and plot them as a function of each metric’s hyperparameters. Note that for the Top-k Log Loss, we plot the correlation with TER as a function of the fraction $1 - \frac{k}{N}$ of the heldout dataset which is *ignored*.

As can be seen from these plots, for certain ranges of values of the metric hyperparameters, the correlation of these metrics with TER is quite high. For example, for $\beta = 1$, the correlations between entropy regularized log loss and TER are 0.91, 0.93, 0.90, and 0.95 for Bengali, BN-50, Cantonese, and TIMIT




Figure 2: Empirical correlations of (left to right) Entropy Regularized Log Loss, Capped Log Loss, and Top-k Log Loss with TER, as a function of β , λ , and $1 - \frac{k}{N}$, respectively.

respectively. This is compared to correlations of 0.31, -0.02 , -0.23 , and -0.21 for the cross-entropy objective.

Based on this analysis, one reasonable thing to do would be to use the heldout Entropy Regularized Log Loss as a stopping criteria, instead of the standard cross-entropy loss. In practice, this results in the training continuing past the point of lowest heldout cross-entropy, and producing models with lower heldout entropy (and lower ERL), which hopefully have a lower TER as well. Similarly, one could use the heldout Entropy Regularized Log Loss, instead of the heldout cross-entropy, in order to decide when to decay the learning rate. This is what we do in our experiments. The results are reported in Section 6. We use $\beta = 1$ in all our experiments. We note that we could have also used the other metrics (Capped, Top-k) for this purpose, but chose to use ERL with $\beta = 1$, since we observed that it attained high correlation values with TER on models we trained, across all 4 datasets (see Figure 2).

6 Experiments

In this section, we first provide a description of the datasets we use, and our evaluation criteria. We then give an overview of our training procedure, and provide details regarding hyperparameter choices. We then present our experimental results comparing the performance of kernel approximation methods to DNNs, demonstrating the effectiveness of using linear bottlenecks, performing feature selection, and using the new early stopping criteria in bringing down the TER. Lastly, we take a deeper look at the dynamics of the feature selection process.

6.1 Tasks, Datasets, and Evaluation Metrics

We train both DNNs and kernel-based multinomial logistic regression models, as described in §3, to predict context-dependent HMM state labels from acoustic feature vectors. We test these methods on four datasets.

Each dataset is partitioned in four: a training set, a heldout set, a development set, and a test set. We use the heldout set to tune the hyperparameters of our training procedure (e.g., the learning rate). We then run decoding on the development set, using IBM’s Attila speech recognition toolkit [Soltau et al., 2010], to select a small subset of models which perform best in terms of TER (e.g., the best kernel model, and the best DNN model, per dataset). We tune the acoustic model weight in order to optimize the relative contributions of the language model and the acoustic model to the final score our system assigns to a given word sequence. Finally, we decode the test set using this select group of models (using, for each model, the best acoustic

model weight on the development set), in order to get a fair comparison between the methods we are using. Having a separate development set helps us avoid the risk of over-fitting to the test set.

The first two datasets we use are the IARPA Babel Program Cantonese (IARPA-babel101-v0.4c) and Bengali (IARPA-babel103b-v0.4b) limited language packs. Each pack contains training and development sets of approximately 20 hours, and an approximately 5 hour test set. We designate about 10% of the training data as a heldout set. The training, heldout, development, and test sets all contain different speakers. Babel data is challenging because it is two-person conversations between people who know each other well (family and friends) recorded over telephone channels (in most cases with mobile telephones) from speakers in a wide variety of acoustic environments, including moving vehicles and public places. As a result, it contains many natural phenomena such as mispronunciations, disfluencies, laughter, rapid speech, background noise, and channel variability. An additional challenge in Babel is that the only data available for training language models is the acoustic transcripts, which are comparatively small.

The third dataset is a 50-hour subset of Broadcast News (BN-50), which is a well-studied benchmark task in the ASR community [Kingsbury, 2009, Sainath et al., 2011]. 45 hours of audio are used for training, and 5 hours are used as a heldout set. For the development set, we use the “Dev04F” dataset provided by LDC, which consists of 2 hours of broadcast news from various news shows. We use the DARPA EARS RT-03 English Broadcast News Evaluation Set [Fiscus et al., 2003] as our test set, consisting of 72 5-minute conversations.

The last dataset we use is TIMIT [Garofolo et al., 1993], which contains recordings of 630 speakers, of various English dialects, each reciting ten sentences, for a total of 5.4 hours of speech. The training set (from which the heldout set is then taken) consists of data from 462 speakers each reciting 8 sentences (SI and SX sentences). The development set consists of speech from 50 speakers. For evaluation, we use the “core test set”, which consists of 192 utterances total from 24 speakers (SA sentences are excluded). For reference, we use the exact same features, labels, and divisions of the dataset as Huang et al. [2014], which allows direct comparison of our results with theirs.

The language models we use are all n -gram language models estimated using modified Kneser-Ney smoothing, with n values of 2,4,3, and 3 for Bengali, Broadcast News, Cantonese, and TIMIT, respectively. The TIMIT language model is a phone-level model. The Bengali and Cantonese language models are particularly small (approximately 60,000 bigrams and 136,000 trigrams, respectively), trained using only the provided audio transcripts. The Broadcast News model is small as well, containing only 3.3 million 4-grams.

The acoustic features, representing 25 ms acoustic frames with context, are real-valued dense vectors. A 10 ms shift is used between adjacent frames (except on TIMIT, where a 5 ms shift is used). For the Cantonese, Bengali, and Broadcast News datasets we use a standard 360-dimensional speaker-adapted representation used by IBM [Kingsbury et al., 2013]. The state labels are obtained via forced alignment using a GMM/HMM system. For the TIMIT experiments, we use 40 dimensional feature space maximum likelihood linear regression (fMLLR) features [Gales, 1998], and concatenate the 5 neighboring frames in either direction, for a total of 11 frames and 440 features.

The Cantonese and Bengali datasets each have 1000 labels, corresponding to quinphone context-dependent HMM states clustered using decision trees. For Broadcast News, there are 5000 such states. The TIMIT dataset has 147 context-independent labels, corresponding to the beginning, middle, and end of 49 phonemes.

For all datasets, the number of training points significantly exceeds typical machine learning tasks tackled by kernel methods. In particular, our training sets all contain between 2 and 16 millions frames. Additionally, the large number of output classes for our datasets presents a scalability challenge, given that the size of the kernel models scales linearly with the number of output classes (if no bottleneck is used). Table 2

Dataset	Train	Heldout	Dev	Test	# Features	# Classes
Beng.	21 hr (7.7M)	2.8 hr (1.0M)	20 hr (7.1M)	5 hr (1.7M)	360	1000
BN-50	45 hr (16M)	5 hr (1.8M)	2 hr (0.7M)	2.5 hr (0.9M)	360	5000
Cant.	21 hr (7.5M)	2.5 hr (0.9M)	20 hr (7.2M)	5 hr (1.8M)	360	1000
TIMIT	3.2 hr (2.3M)	0.3 hr (0.2M)	0.15 hr (0.1M)	0.15 hr (0.1M)	440	147

Table 2: Dataset details. We report the size of each dataset partition in terms of the number of hours of speech, and in terms of the number of acoustic frames (in parentheses).

provides details on the sizes of all the datasets, as well as on their number of features and classes.

We use five metrics to evaluate the acoustic models:

1. **Cross-entropy:** Given examples, $\{(\mathbf{x}_i, y_i), i = 1 \dots N\}$, the cross-entropy is defined as

$$-\frac{1}{N} \sum_{i=1}^N \log p(y_i | \mathbf{x}_i).$$

2. **Average Entropy:** The average entropy of a model is defined as

$$-\frac{1}{N} \sum_{i=1}^N \sum_{y=1}^C p(y | x_i) \log p(y | x_i)$$

If a model has low average entropy, it is generally confident in its predictions.

3. **Entropy Regularized Log Loss (ERLL):** Defined in Section 5. We use $\beta = 1$ unless specified otherwise.
4. **Classification Error:** The classification error is defined as

$$\frac{1}{N} \sum_{i=1}^N \mathbb{1} \left[y_i \neq \arg \max_{y \in \{1, 2, \dots, C\}} p(y | \mathbf{x}_i) \right].$$

5. **Token Error Rate (TER):** We feed the predictions of the acoustic models, which correspond to probability distributions over the phonetic states, to the rest of the ASR pipeline and calculate the misalignment between the decoder’s outputs and the ground-truth transcriptions. For Bengali and BN-50, we measure the error in terms of the word error rate (WER), for Cantonese we use the character error rate (CER), and for TIMIT we use the phone error rate (PER). We use the term “token error rate” (TER) to refer, for each dataset, to its corresponding metric.

6.2 Details of Acoustic Model Training

All our kernel models were trained with either the Laplacian, the Gaussian, or the Sparse Gaussian (§4.2) kernel. These kernel models typically have 3 hyperparameters: the kernel bandwidth (σ for the Gaussian kernels, λ for the Laplacian kernel; see Table 1), the number of random projections, and the initial learning rate of the optimization procedure. As a rule of thumb, good values for the kernel bandwidths (specifically, $2\sigma^2$ for the Gaussian kernels, and $1/\lambda$ for the Laplacian kernel) range from 0.3-5 times the median of

the pairwise distances in the data.⁸ We try various numbers of random features, ranging from 5000 to 200,000. Using more random features leads to a better approximation of the kernel function, as well as to more powerful models, though there are diminishing returns as the number of features increases. The Sparse Gaussian kernel additionally has the hyperparameter k which specifies the sparsity of each random projection vector ω_i . For all experiments, we use $k = 5$.

For all DNNs, we tune hyperparameters related to both the architecture and the optimization. This includes the number of layers, the number of hidden units in each layer, and the learning rate. We perform 1 epoch of layer-wise discriminative pre-training [Seide et al., 2011b, Kingsbury et al., 2013], and then train the entire network jointly using SGD. We find that 4 hidden layers is generally the best setting for our DNNs, so all the DNN results we present in this paper use this setting. Additionally, all our DNNs use the tanh activation function. We vary the number of hidden units per layer (1000, 2000, or 4000).

For both DNN and kernel models, we use stochastic gradient descent (SGD) as our optimization algorithm, with a mini-batch size of 250 or 256 samples. We use the heldout set to tune the other hyperparameters (e.g., learning rate). We use the learning rate decay scheme described in [Morgan and Bourlard, 1990, Sainath et al., 2013a,b], which monitors performance on the heldout set in order to decide when to decay the learning rate. This method divides the learning rate in half at the end of an SGD epoch if the heldout cross-entropy doesn't improve by at least 1%; additionally, if the heldout cross-entropy gets *worse*, it reverts the model back to its state at the beginning of the epoch. Instead of using the heldout cross-entropy, in some of our experiments we use the heldout ERLI in order to decide when to decay the learning rate.

As mentioned in Section 3.5, one effective way of reducing the number of parameters in our models is to impose a low-rank constraint on the output parameter matrix; we refer to this as a "linear bottleneck" [Sainath et al., 2013a]. We use bottlenecks of size 1000, 250, 250, and 100 for BN-50, Bengali, Cantonese, and TIMIT, respectively. We train models both with and without this technique; the only exception is that we are unable to train BN-50 kernel models *without* the bottleneck of size 1000, due to memory constraints on our GPUs.

We initialize our DNN parameters uniformly at random in the range $[-\frac{\sqrt{6}}{\sqrt{d_{in}+d_{out}}}, \frac{\sqrt{6}}{\sqrt{d_{in}+d_{out}}}]$, as suggested by Glorot and Bengio [2010]; here, d_{in} and d_{out} refer to the dimensionality of the input and output of a DNN layer, respectively. For our kernel models, we initialize the random projection matrix as discussed in Section 3, and we initialize the parameter matrix Θ as the zero matrix. When using a linear bottleneck to decompose the parameter matrix, we initialize the resulting two matrices randomly, like we do for our DNNs.

For each iteration of random feature selection, we draw a random subsample of the training data of size $R = 10^6$ (except when $D \geq 10^5$, in which case we use $R = 2 \times 10^6$, to ensure a safe n to D ratio), but ultimately we use all N training examples once the random features are selected. Thus, each iteration of feature selection has equivalent computational cost to a R/N fraction of an SGD epoch (roughly 1/7 or 2/7 for $D < 10^5$ and $D \geq 10^5$ respectively, on the Babel speech data sets, for example). We use $T = 50$ iterations of feature selection, and in iteration t , we select $s_t = t \cdot (D/T) = 0.02Dt$ random features. Thus, the total computational cost we incur for feature selection is equivalent to approximately seven (or 14) epochs of training on the Babel data sets. For the Broadcast News dataset, it corresponds to the cost of approximately 6 full epochs of training (when using $R = 2 \times 10^6$).

All our training code is written in MATLAB, leveraging its GPU features. We execute our code on

⁸For the Gaussian kernel, we take the median of the ℓ_2 squared distances between a large number of random pairs of training examples. For the Laplacian kernel, we use ℓ_1 distances instead. For the Sparse Gaussian kernel, we use the median ℓ_2 squared distances between randomly chosen sub-vectors of size k of random pairs of training points.

	Laplacian				Gaussian				Sparse Gaussian			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	74.5	72.1	74.5	71.4	72.6	72.0	72.6	71.8	73.0	71.5	73.0	70.9
+FS	72.9	71.1	72.8	70.4	74.1	71.4	74.2	70.3	72.9	71.2	72.8	70.7
BN-50	N/A	17.9	N/A	17.7	N/A	17.3	N/A	17.1	N/A	17.3	N/A	17.0
+FS	N/A	17.1	N/A	16.7	N/A	17.5	N/A	17.0	N/A	17.1	N/A	16.7
Cant.	69.9	68.2	69.2	67.4	70.2	67.6	70.0	67.1	68.6	67.5	68.1	67.1
+FS	68.4	67.5	68.5	66.7	69.9	67.7	69.8	66.9	68.6	67.4	68.5	66.8
TIMIT	20.6	19.2	20.4	18.9	19.8	18.9	19.6	18.6	19.9	18.8	19.6	18.4
+FS	19.5	18.6	19.3	18.4	19.5	18.6	19.4	18.4	19.3	18.4	19.1	18.2

Table 3: Kernel TER Results (development set): This table shows TER results for our kernel experiments using either the Laplacian, Gaussian, or Sparse Gaussian kernels. ‘NT’ specifies that no “tricks” were used during training (no bottleneck, no feature selection, no special learning rate decay). A ‘B’ specifies that a linear bottleneck was used for the parameter matrix; an ‘R’ specifies that entropy regularized log loss was used for learning rate decay (so ‘BR’ means both were used). ‘+FS’ specifies that feature selection was used for the experiments in that row. The best result for each row is in bold.

	1000				2000				4000			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	72.3	71.6	71.7	70.9	71.5	71.1	70.7	70.3	71.1	70.6	70.5	70.2
BN-50	18.0	17.3	17.8	17.1	17.4	16.7	17.1	16.4	16.8	16.7	16.7	16.5
Cant.	68.4	68.1	67.9	67.5	67.7	67.7	67.2	67.1	67.7	67.1	67.2	67.2
TIMIT	19.5	19.3	19.4	19.2	19.0	18.9	19.2	19.2	18.6	18.6	18.7	18.9

Table 4: DNN TER Results (development set): This table shows TER results for DNNs with 1000, 2000, or 4000 hidden units per layer. ‘NT’ specifies that no “tricks” were used while training the DNN (no bottleneck, no special learning rate decay). A ‘B’ specifies that a linear bottleneck was used for the output parameter matrix; an ‘R’ specifies that entropy regularized log loss was used for learning rate decay (so ‘BR’ means both were used). The best result for each language is in bold.

Amazon EC2 machines, with instances of type g2.2xlarge. We use StarCluster⁹ to more easily manage our clusters of EC2 machines.

6.3 Results

In this section, we report results from experiments comparing kernel methods to deep neural networks (DNNs) on ASR tasks. We report results on all 4 datasets, using various combinations of the methods discussed previously. For both DNN and kernel methods, we train models with and without linear bottlenecks, and with and without using ERLI to determine the learning rate decay. For our kernel methods, we additionally train models with and without using feature selection. We run experiments with all three kernels (Laplacian, Gaussian, Sparse Gaussian) and we use 100,000 random features on all datasets except for TIMIT, where we are able to use 200,000 random features (because the output dimensionality is lower).

⁹<http://star.mit.edu/cluster>

	Beng. (D/K)	BN-50 (D/K)	Cant. (D/K)	TIMIT (D/K)
CE	1.243 / 1.256	2.001 / 2.004	1.916 / 1.883	1.056 / 0.9182
ENT	0.9079 / 1.082	1.274 / 1.457	1.375 / 1.516	0.447 / 0.5756
ERLL	2.302 / 2.406	3.548 / 3.625	3.459 / 3.493	1.671 / 1.607
ERR	0.2887 / 0.2936	0.4887 / 0.4931	0.4353 / 0.4287	0.324 / 0.3085
TER (dev)	70.2 / 70.3	16.4 / 16.7	67.1 / 66.7	18.6 / 18.2
TER (test)	69.1 / 69.2	11.7 / 11.9	63.7 / 63.2	20.5 / 20.4

Table 5: Table comparing the Best DNN (‘D’) and kernel (‘K’) results, across 4 datasets and 6 metrics. The first 4 metrics are on the heldout set, the fifth is on the development set, and the last metric is reported on the test set.

	Test TER (DNN)	Test TER (Kernel)
Huang et al. [2014]	20.5	21.3
This work	20.5	20.4

Table 6: Table comparing the Best DNN and kernel results from this work to those from Huang et al. [2014], on the TIMIT test set.

As mentioned in the previous section, for our DNN experiments, we train models with 4 hidden layers,¹⁰ using the tanh activation function, and using either 1000, 2000, or 4000 hidden units per layer. We focus on comparing the performance of these methods in terms of TER, but we also report results for other metrics. Unless specified otherwise, all TER results are on the development set, and all cross-entropy, entropy, classification error, and ERLL results are on the heldout set.

In Tables 3 and 4, we show our TER results for our kernel and DNN models, respectively, across all datasets. There are many things to notice about these results. Within the kernel models, we see that incorporating a linear bottleneck brings large drops in TER across the board.¹¹ Performing feature selection generally improves TER as well; we see that it improves TER considerably for the Laplacian kernel, and modestly for the Sparse Gaussian kernel. For the Gaussian kernel, it typically helps, though there are several instances in which feature selection hurts TER (see Section 6.5 for discussion). Second, we see that using heldout ERLL to determine when to decay the learning rate helps all our kernel models attain lower TER values. Next, we see that without using feature selection, the Sparse Gaussian kernel has the best performance across the board. After we include feature selection, it performs very comparably to the Laplacian kernel with feature selection. It is interesting to note that without using feature selection, the Gaussian kernel is generally better than the Laplacian kernel; however, with feature selection, the Laplacian kernel surpasses the Gaussian kernel (see Section 6.5). In general, the kernel function which performed best, across the majority of settings, was the Sparse Gaussian kernel.

For our DNN models, linear bottlenecks almost always lower TER values, though in a few cases they have no effect on TER. Using ERLL to determine when to decay the learning rate generally helps lower TER values for our DNNs, but in a few cases it actually hurts (Cantonese with 4000 hidden units, and TIMIT with 2000 and 4000). The DNNs with 4000 hidden units typically attain the best results, though on a couple of datasets they are matched or narrowly beaten by the 2000 models.

In Table 5, we compare for each dataset the performance of the best DNN model with the best kernel

¹⁰As mentioned in Section 6.2, we find that this is generally the best setting.

¹¹Recall that we are unable to train BN-50 kernel models *without* using a bottleneck because the resulting models would not fit on our GPUs.

model, across 6 metrics. In terms of heldout cross-entropy and classification error, kernels and DNNs performed similarly, with kernels outperforming DNNs on Cantonese and TIMIT, while the DNNs beat the kernels on Bengali and BN-50. In terms of the average heldout entropy of the models, the DNNs were consistently more confident in their predictions (lower entropy) than the kernels. Significantly, we observe that the best development set TER results for our DNN and kernel models are quite comparable; on Cantonese and TIMIT, the kernel models outperform the DNNs by 0.4% absolute, whereas on Bengali and BN-50, the DNN does better by 0.1% and 0.3%, respectively.

We will now discuss the results on the test sets. First of all, in order to avoid overfitting to the test set, for each dataset we only performed test set evaluations for the DNN and kernel models which performed best in terms of the development set TER. The final row of Table 5 thus contains all the test results we collected. As one can see, the relative performance of the DNN and kernel models is very similar to the development set results, with the DNNs performing better on Bengali and BN-50, and the kernels performing better on Cantonese and TIMIT. For direct comparison, we include in Table 6 the test results for the best DNN and kernel models from Huang et al. [2014]. As mentioned in Section 6.1, we use the same features, labels, data set partitions (train/heldout/dev/test), and decoding script as Huang et al., and thus our results are directly comparable. We achieve a 0.9% absolute improvement in TER with our kernel model relative to Huang et al. [2014]; our DNN performs the same as theirs. Furthermore, while their DNN beat their kernel by 0.8% TER, our kernel beats our DNN by 0.1% TER.

In Appendix B, we include more detailed tables comparing the various models we trained across all the abovementioned metrics. Some important things to take note of in those tables are as follows:

- The linear bottleneck typically causes large drops in the average entropy of kernel models, while not having as strong or consistent an effect on cross-entropy. For DNNs, the bottleneck typically causes increases in cross-entropy, and relatively modest decreases in entropy.
- Using ERLI to determine learning rate decay typically causes increases in cross-entropy, and decreases in entropy, with the decrease in entropy typically being larger than the increase in cross-entropy. As a result, the ERLI is typically lower for models that use this method (with the exception of TIMIT DNN models).
- Feature selection typically results in large drops in cross-entropy, especially for Laplacian and Sparse Gaussian kernels, while its effect on entropy is quite small. It thus helps lower heldout ERLI across the board, as well as TER in the vast majority of cases.

6.4 Importance of the Number of Random Features

We will now illustrate the importance of the number of random features D on the final performance of the model. For this purpose, we trained a number of different models on the BN-50 dataset, using $D \in \{5000, 10000, 25000, 50000, 100000\}$. We trained models using the 3 different kernels, with and without feature selection. We used a linear bottleneck of size 1000 for all these models, and used heldout cross-entropy to determine the learning rate decay. In Figure 3, we show how increasing the number of features dramatically improves the performance of the learned model, both in terms of cross-entropy and TER; there are diminishing returns, however, with very small improvements in TER when increasing D from 50,000 to 100,000. Furthermore, the size of the gap between the dashed and solid lines (representing experiments with and without feature selection, respectively), indicates the importance of feature selection in attaining strong performance. This gap is very large for the Laplacian kernel, modest for the Sparse Gaussian kernel, and relatively insignificant for the Gaussian kernel.




Figure 3: Performance of kernel acoustic models on BN-50 dataset, as a function of the number of random features D used. Results are reported in terms of heldout cross-entropy as well as development set TER. Dashed lines signify that feature selection was performed, while solid lines mean it was not. The color and shape of the markers indicate the kernel used.




Figure 4: Fraction of the s_t features selected in iteration t that are in the final model (survival rate) for Cantonese dataset.

6.5 Effects of Random Feature Selection

We now explore the dynamics of the feature selection process. In our method, there is no guarantee that a feature selected in one iteration will be selected in the next. In Figure 4, we plot the fraction of the s_t features selected in iteration t that actually remain in the model after all T iterations. We only show the results for Cantonese (models without linear bottleneck, and without using entropy regularized log loss for LR decay), as the plots for other datasets are qualitatively similar. In nearly all iterations and for all kernels, over half of the selected features survive to the final model. For instance, over 90% of the Laplacian kernel features selected at iteration 10 survive the remaining 40 rounds of selection. For comparison, we also plot the expected fraction of the s_t features selected in iteration t that would survive until the end if the selected features in each iteration were chosen uniformly at random from the pool. Since we use $s_t = Dt/T$, the expected fraction at iteration t is $T!/(t! \cdot T^{T-t})$, which is exponentially small in T when $t \leq \beta T$ for any fixed $\beta < 1$.¹² For example, at $t = 10$ the expected survival rate is approximately 9×10^{-11} with $T = 50$.

Finally, we consider how the random feature selection process can be regarded as selecting non-linear

¹²This can be shown using Stirling’s formula. See Jameson [2015] for a useful review.




Figure 5: The relative weight of each input feature in the random matrix Θ , for Cantonese dataset.

combinations of input features. Consider the final matrix of random vectors $\Theta := [\theta^{(1)}|\theta^{(1)}|\dots|\theta^{(D)}] \in \mathbb{R}^{d \times D}$ after random feature selection. A coarse measure of how much influence an input feature $i \in \{1, 2, \dots, d\}$ has in the final feature map is the relative “weight” of the i -th row of Θ . In Figure 5, we plot $\sum_{j=1}^D |\Theta_{i,j}|/Z$ for each input feature $i \in \{1, 2, \dots, d\}$. Here, $Z = \frac{1}{d} \sum_{i,j} |\Theta_{i,j}|$ is a normalization term.¹³ There is a strong periodic effect as a function of the input feature number. The reason for this stems from the way the acoustic features are generated. Recall that the features are the concatenation of nine 40-dimensional acoustic feature vectors for nine audio frames. An examination of the feature pipeline from Kingsbury et al. [2013] reveals that these 40 features are ordered by a measure of discriminative quality (via linear discriminant analysis). Thus, it is expected that the features with low $(i - 1) \bmod 40$ value may be more useful than the others; indeed, this is evident in the plot. Note that this effect exists, but is extremely weak, with the Gaussian kernel. We believe this is because Gaussian random vectors in \mathbb{R}^d are likely to have *all* their entries be bounded in magnitude by $O(\sqrt{\log(d)})$.

6.6 Other Possible Improvements to DNNs and Kernels

It is important to mention a few things regarding other ways the performance of our DNN and kernel models could be improved, and why they are not investigated at length in this work. For the kernel methods, given that the optimization is convex when no bottleneck is used, it would be possible to get stronger convergence guarantees using the Stochastic Average Gradient (SAG) algorithm instead of SGD for training [Le Roux et al., 2012]. In fact, in Lu et al. [2016] we did this on Cantonese and Bengali, and attained strong recognition performance.¹⁴ Unfortunately, it is challenging to scale this algorithm to larger tasks, since it requires storing, for every training example, the previous gradient of the loss function at that example. Because $\frac{\partial L(x_i, y_i)}{\partial \theta_y} = z(x_i)[\mathbb{I}(y = y_i) - p(y|x_i)]$, and because $z(x_i)$ is fixed, the gradient information can be stored by simply storing, for each training example, the vector $\mathbf{p}_i = [p(1|x_i), \dots, p(C|x_i)]$. However, this still takes NC storage, which is quite expensive when there are millions of training examples N and thousands of output classes C (320 GB for the Broadcast News dataset, for example). Unfortunately, once a bottleneck is introduced, not only is the optimization problem non-convex, but we must also store the full gradients, thus making the memory requirement too large. As a result, for scalability reasons, as well as for

¹³For the Laplacian kernel, we discard the largest element in each of the d rows of Θ , because there are sometimes outliers which dominate the entire sum for their row.

¹⁴A few more details regarding the experiments in Lu et al. [2016]: we did not use feature selection in that work, and we only used ERL as a model selection criterion (not for learning rate decay). Additionally, instead of training the large kernel models jointly, we trained them in blocks of 25,000 random features, and then combined the models via logit averaging (final models had 200,000 random features).

consistency across all our experiments, we have used SGD for all our kernel experiments. Additionally, we did not investigate the use of sequence training techniques for our kernel methods, leaving this for future work.

For our DNN models, we have observed that restricted Boltzmann machine (RBM) pre-training [Hinton et al., 2006] often improves recognition performance [Lu et al., 2016]. Additionally, there are various other deep architectures (e.g., Convolutional Neural Networks [Abdel-Hamid et al., 2014], Long Short Term Memory Networks [Sak et al., 2014]), as well as numerous training techniques (e.g., momentum [Sutskever et al., 2013], dropout [Srivastava et al., 2014], batch normalization [Ioffe and Szegedy, 2015]), which can further improve the performance of neural networks. Our intention for this paper was to provide a comparison between kernel methods and a strong DNN baseline (DNN with tanh activation, and discriminative pre-training), not to provide an exhaustive comparison against all possible deep learning architectures and optimization methods.

7 Conclusion

In this paper, we explore the performance of kernel methods on large-scale ASR tasks, leveraging the kernel approximation technique of Rahimi and Recht [2007]. We propose two new methods (feature selection, new early stopping criteria) which lead to large improvements in the performance of kernel acoustic models. We further show that using a linear bottleneck [Sainath et al., 2013a] to decompose the parameter matrix of these kernel models leads to significant improvements in TER as well. We replicate these findings on four different datasets, including the Broadcast News (50 hour) and TIMIT benchmark tasks. The linear bottleneck, as well as the learning rate decay method, also typically improve the performance of our DNN acoustic models. Using all these methods in conjunction, the kernel methods attain comparable TER values to DNNs across our four test sets; on Cantonese and TIMIT, the kernel models outperform the DNNs by 0.5% and 0.1% absolute, respectively, whereas on Bengali and BN-50, the DNN does better by 0.1% and 0.2%.

For future work, we are interested in a number of questions: (1) Can we develop techniques other than feature selection for “learning” the kernel function more effectively? (2) Are kernel methods as robust as DNNs to different types of input (e.g., log-mel filterbank features)? (3) How much does the performance of the kernel models improve using sequence training methods? (4) Can sequence kernels be used to improve the recognition performance of kernel acoustic models, in a manner analogous to how LSTMs can give improvements over DNNs? (5) Can kernel methods compete well with DNNs in domains outside of speech recognition? (6) Broadly speaking, what are the biggest limitations of kernel methods, and how can they be overcome?

Acknowledgments

This research is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Defense U.S. Army Research Laboratory (DoD / ARL) contract number W911NF-12-C-0012. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoD/ARL, or the U.S. Government.

F. S. is grateful to Lawrence K. Saul (UCSD), Léon Bottou (Facebook), Alex Smola (Amazon), and

Chris J. C. Burges (Microsoft Research) for many fruitful discussions and pointers to relevant work.

Computation for the work described in this paper was partially supported by the University of Southern California’s Center for High-Performance Computing (<http://hpc.usc.edu>).

Additionally, A. B. G. is partially supported by a USC Provost Graduate Fellowship. F. S. is partially supported by a NSF IIS-1065243, 1451412, 1139148 a Google Research Award, an Alfred. P. Sloan Research Fellowship, an ARO YIP Award (W911NF-12-1-0241) and ARO Award W911NF-15-1-0484. A.B. is partially supported by a grant from CPER Nord-Pas de Calais/FEDER DATA Advanced data science and technologies 2015-2020.

Appendices

A Derivation of Functional Form for Random Fourier Features

In this appendix, we will prove that for a properly-scaled (i.e., $Z = 1$) positive-definite shift-invariant kernel k ,

$$k(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{\omega, b} \left[\sqrt{2} \cos(\omega^T \mathbf{x} + b) \cdot \sqrt{2} \cos(\omega^T \mathbf{y} + b) \right], \quad (6)$$

where ω is drawn from $p(\omega)$, the inverse Fourier transform of k , and b is drawn uniformly from $[0, 2\pi]$. We begin this proof using Equation (3) from Section 3.1:

$$\begin{aligned} k(\mathbf{x}, \mathbf{y}) &= \int_{R^d} p(\omega) e^{j\omega^T(\mathbf{x}-\mathbf{y})} d\omega \\ &= \mathbb{E}_{\omega} \left[e^{j\omega^T \mathbf{x}} e^{-j\omega^T \mathbf{y}} \right] \\ &= \mathbb{E}_{\omega} \left[(\cos(\omega^T \mathbf{x}) + j \sin(\omega^T \mathbf{x})) (\cos(\omega^T \mathbf{y}) - j \sin(\omega^T \mathbf{y})) \right] \\ &= \mathbb{E}_{\omega} \left[\cos(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) + \sin(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \right] \\ &\quad + j \cdot \mathbb{E}_{\omega} \left[\sin(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) - \sin(\omega^T \mathbf{y}) \cos(\omega^T \mathbf{x}) \right] \\ &= \mathbb{E}_{\omega} \left[\cos(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) + \sin(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \right] \end{aligned} \quad (7)$$

Note the Equation (7) is true because we know that $k(\mathbf{x}, \mathbf{y})$ is a real-valued function, and thus the imaginary part of the expectation must disappear. We now show that the right-hand side of Equation (6) is equal to this same expression:

$$\begin{aligned}
& \mathbb{E}_{\omega, b} \left[\sqrt{2} \cos(\omega^T \mathbf{x} + b) \cdot \sqrt{2} \cos(\omega^T \mathbf{y} + b) \right] \\
&= 2 \cdot \mathbb{E}_{\omega, b} \left[\left(\cos(\omega^T \mathbf{x}) \cos(b) - \sin(\omega^T \mathbf{x}) \sin(b) \right) \cdot \right. \\
&\quad \left. \left(\cos(\omega^T \mathbf{y}) \cos(b) - \sin(\omega^T \mathbf{y}) \sin(b) \right) \right] \tag{8} \\
&= 2 \cdot \mathbb{E}_{\omega, b} \left[\cos(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) \cos^2(b) \right. \\
&\quad - \cos(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \cos(b) \sin(b) \\
&\quad - \sin(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) \cos(b) \sin(b) \\
&\quad \left. + \sin(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \sin^2(b) \right] \\
&= 2 \cdot \mathbb{E}_{\omega} \left[\frac{1}{2} \cos(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) + \frac{1}{2} \sin(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \right] \tag{9} \\
&= \mathbb{E}_{\omega} \left[\cos(\omega^T \mathbf{x}) \cos(\omega^T \mathbf{y}) + \sin(\omega^T \mathbf{x}) \sin(\omega^T \mathbf{y}) \right] \\
&= k(\mathbf{x}, \mathbf{y})
\end{aligned}$$

Equation (8) is true by the cosine sum of angles formula, and Equation (9) is true because $\mathbb{E}_b [\cos^2(b)] = \mathbb{E}_b [\sin^2(b)] = \int_0^{2\pi} \frac{1}{2\pi} \sin^2(b) = \frac{1}{2}$, and because $\mathbb{E}_b [\sin(b) \cos(b)] = 0$. This concludes the proof. \square

B Detailed Results

In this appendix, we include tables comparing the models we trained in terms of 4 different metrics (CE, ENT, ERR, and ERLI). The notation is the same as in Tables 4 and 3. For both DNN and kernel models, ‘NT’ specifies that no ‘tricks’ were used during training (no bottleneck, no feature selection, no special learning rate decay). A ‘B’ specifies that a linear bottleneck was used for the output parameter matrix, while an ‘R’ specifies that entropy regularized log loss was used for learning rate decay (so ‘BR’ means both were used). For kernel models, ‘+FS’ specifies that feature selection was performed for the corresponding row. The best result for each metric and language is in bold.

	1000				2000				4000			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	1.25	1.26	1.24	1.27	1.24	1.26	1.26	1.32	1.24	1.25	1.30	1.39
BN-50	2.05	2.05	2.04	2.08	2.01	2.04	2.05	2.22	2.00	2.03	2.09	2.27
Cant.	1.92	1.96	1.92	1.98	1.93	1.94	1.97	2.06	1.92	1.97	2.03	2.10
TIMIT	1.06	1.08	1.20	1.28	1.08	1.09	1.25	1.31	1.10	1.11	1.25	1.33

Table 7: DNN: Metric CE

	Laplacian				Gaussian				Sparse Gaussian			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	1.34	1.32	1.35	1.39	1.35	1.33	1.36	1.34	1.31	1.29	1.34	1.33
+FS	1.28	1.26	1.29	1.27	1.35	1.31	1.36	1.35	1.28	1.26	1.31	1.27
BN-50	N/A	2.15	N/A	2.43	N/A	2.05	N/A	2.16	N/A	2.05	N/A	2.19
+FS	N/A	2.01	N/A	2.07	N/A	2.04	N/A	2.13	N/A	2.00	N/A	2.06
Cant.	1.93	1.95	1.95	2.04	1.99	1.98	2.00	2.04	1.93	1.94	1.95	2.00
+FS	1.88	1.90	1.89	1.95	1.97	1.97	1.98	2.03	1.90	1.91	1.91	1.96
TIMIT	0.97	0.99	0.97	1.07	0.94	0.96	0.94	1.02	0.94	0.95	0.94	1.03
+FS	0.92	0.95	0.92	1.03	0.93	0.96	0.93	1.02	0.92	0.96	0.92	1.03

Table 8: Kernel: Metric CE

	1000				2000				4000			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	1.23	1.17	1.18	1.09	1.18	1.13	1.09	0.99	1.14	1.11	1.02	0.91
BN-50	1.95	1.77	1.90	1.68	1.76	1.68	1.65	1.40	1.65	1.60	1.48	1.27
Cant.	1.71	1.67	1.67	1.57	1.66	1.64	1.55	1.42	1.63	1.55	1.43	1.38
TIMIT	0.72	0.70	0.58	0.53	0.63	0.63	0.50	0.48	0.57	0.57	0.48	0.45

Table 9: DNN: Metric ENT

	Laplacian				Gaussian				Sparse Gaussian			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	1.43	1.23	1.41	1.08	1.36	1.31	1.35	1.28	1.35	1.23	1.30	1.10
+FS	1.32	1.21	1.28	1.14	1.44	1.27	1.45	1.13	1.32	1.22	1.26	1.14
BN-50	N/A	1.89	N/A	1.46	N/A	1.83	N/A	1.53	N/A	1.81	N/A	1.48
+FS	N/A	1.81	N/A	1.56	N/A	1.84	N/A	1.55	N/A	1.80	N/A	1.57
Cant.	1.84	1.67	1.76	1.52	1.94	1.73	1.91	1.58	1.77	1.69	1.70	1.55
+FS	1.75	1.66	1.73	1.54	1.91	1.72	1.87	1.57	1.75	1.68	1.72	1.54
TIMIT	0.95	0.72	0.91	0.61	0.88	0.73	0.86	0.62	0.89	0.76	0.85	0.61
+FS	0.86	0.70	0.82	0.58	0.86	0.70	0.83	0.61	0.84	0.69	0.82	0.58

Table 10: Kernel: Metric ENT

	1000				2000				4000			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	2.48	2.43	2.43	2.37	2.42	2.39	2.35	2.31	2.39	2.37	2.32	2.30
BN-50	3.99	3.82	3.94	3.76	3.77	3.72	3.70	3.63	3.65	3.63	3.58	3.55
Cant.	3.63	3.63	3.58	3.56	3.59	3.58	3.51	3.48	3.55	3.52	3.46	3.47
TIMIT	1.77	1.77	1.77	1.81	1.71	1.72	1.76	1.79	1.67	1.68	1.73	1.78

Table 11: DNN: Metric ERL

	Laplacian				Gaussian				Sparse Gaussian			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	2.77	2.55	2.76	2.47	2.71	2.65	2.71	2.62	2.67	2.52	2.64	2.44
+FS	2.60	2.47	2.57	2.41	2.79	2.58	2.80	2.48	2.60	2.48	2.57	2.41
BN-50	N/A	4.04	N/A	3.88	N/A	3.88	N/A	3.69	N/A	3.86	N/A	3.67
+FS	N/A	3.82	N/A	3.63	N/A	3.88	N/A	3.67	N/A	3.80	N/A	3.62
Cant.	3.77	3.62	3.71	3.56	3.94	3.71	3.91	3.62	3.71	3.63	3.65	3.54
+FS	3.63	3.56	3.63	3.49	3.88	3.69	3.86	3.60	3.64	3.58	3.63	3.50
TIMIT	1.92	1.71	1.87	1.68	1.82	1.70	1.80	1.65	1.83	1.71	1.79	1.64
+FS	1.78	1.65	1.74	1.61	1.79	1.67	1.76	1.64	1.76	1.64	1.74	1.61

Table 12: Kernel: Metric ERLI

	1000				2000				4000			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.30	0.29	0.29	0.29	0.30
BN-50	0.50	0.50	0.50	0.50	0.49	0.50	0.50	0.51	0.49	0.49	0.50	0.51
Cant.	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44
TIMIT	0.33	0.33	0.34	0.34	0.33	0.33	0.34	0.34	0.33	0.32	0.33	0.33

Table 13: DNN: Metric ERR

	Laplacian				Gaussian				Sparse Gaussian			
	NT	B	R	BR	NT	B	R	BR	NT	B	R	BR
Beng.	0.30	0.30	0.31	0.31	0.31	0.31	0.31	0.31	0.30	0.30	0.30	0.30
+FS	0.29	0.29	0.30	0.29	0.31	0.31	0.31	0.31	0.30	0.29	0.30	0.30
BN-50	N/A	0.52	N/A	0.54	N/A	0.50	N/A	0.51	N/A	0.50	N/A	0.51
+FS	N/A	0.49	N/A	0.50	N/A	0.50	N/A	0.50	N/A	0.49	N/A	0.49
Cant.	0.43	0.44	0.44	0.44	0.45	0.45	0.45	0.45	0.43	0.44	0.44	0.44
+FS	0.43	0.43	0.43	0.44	0.44	0.44	0.44	0.45	0.43	0.44	0.43	0.44
TIMIT	0.32	0.32	0.32	0.33	0.31	0.32	0.31	0.32	0.31	0.31	0.31	0.32
+FS	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.32	0.31	0.31	0.31	0.31

Table 14: Kernel: Metric ERR

References

- Ossama Abdel-Hamid, Abdel-rahman Mohamed, Hui Jiang, Li Deng, Gerald Penn, and Dong Yu. Convolutional neural networks for speech recognition. *IEEE/ACM Trans. Audio, Speech & Language Processing*, 22(10):1533–1545, 2014.
- Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2654–2662, 2014.
- L. Bahl, P. Brown, P. de Souza, and R. Mercer. Maximum mutual information estimation of hidden Markov model parameters for speech recognition. In *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '86.*, volume 11, pages 49–52, Apr 1986.
- Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Localized Rademacher complexities. In *Proceedings of the 15th Annual Conference on Computational Learning Theory, COLT '02*, pages 44–58, London, UK, UK, 2002. Springer-Verlag. ISBN 3-540-43836-X.
- Monica Bianchini and Franco Scarselli. On the complexity of neural network classifiers: A comparison between shallow and deep architectures. *IEEE Trans. Neural Netw. Learning Syst.*, 25(8):1553–1565, 2014.
- Léon Bottou, Olivier Chapelle, Dennis DeCoste, and Jason Weston, editors. *Large Scale Kernel Machines*. MIT Press, Cambridge, MA., 2007.
- C.-C. Cheng and B. Kingsbury. Arccosine Kernels: Acoustic Modeling with Infinite Neural Networks. In *Proc. ICASSP*, pages 5200–5203, 2011.
- Kenneth L. Clarkson. Coresets, Sparse Greedy Approximation, and the Frank-Wolfe Algorithm. *ACM Trans. Algorithms*, 6(4):63:1–63:30, 2010.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *MCSSS*, 2(4):303–314, 1989.
- G. E. Dahl, D. Yu, L. Deng, and A. Acero. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *Audio, Speech, and Language Processing, IEEE Transactions on*, 20(1): 30–42, 2012.
- Bo Dai, Bo Xie, Niao He, Yingyu Liang, Anant Raj, Maria-Florina Balcan, and Le Song. Scalable kernel methods via doubly stochastic gradients. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3041–3049, 2014.
- Dennis DeCoste and Bernhard Schölkopf. Training Invariant Support Vector Machines. *Mach. Learn.*, 46: 161–190, 2002.
- Li Deng, Gökhan Tür, Xiaodong He, and Dilek Z. Hakkani-Tür. Use of Kernel Deep Convex Networks and End-to-end Learning for Spoken Language Understanding. In *2012 IEEE Spoken Language Technology Workshop (SLT), Miami, FL, USA, December 2-5, 2012*, pages 210–215, 2012.

- John C. Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *Journal of Machine Learning Research*, 10:2899–2934, 2009.
- Jonathan Fiscus, George Doddington, Audrey Le, Greg Sanders, Mark Przybocki, and David Pallett. 2003 NIST Rich Transcription evaluation data. <https://catalog.ldc.upenn.edu/LDC2007S10>, 2003. Linguistic Data Consortium Catalog No. LDC2007S10.
- Mark Gales and Steve Young. The application of hidden Markov models in speech recognition. *Found. Trends Signal Process.*, 1(3):195–304, January 2007. ISSN 1932-8346.
- M.J.F. Gales. Maximum likelihood linear transformations for HMM-based speech recognition. *Computer Speech & Language*, 12(2):75 – 98, 1998. ISSN 0885-2308.
- J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, and N. L. Dahlgren. DARPA TIMIT acoustic phonetic continuous speech corpus CDROM, 1993. URL <http://www.ldc.upenn.edu/Catalog/LDC93S1.html>.
- Matthew Gibson and Thomas Hain. Hypothesis spaces for minimum Bayes risk training in large vocabulary speech recognition. In *INTERSPEECH 2006 - ICSLP, Ninth International Conference on Spoken Language Processing, Pittsburgh, PA, USA, September 17-21, 2006*, 2006.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, pages 249–256, 2010.
- Raffay Hamid, Ying Xiao, Alex Gittens, and Dennis DeCoste. Compact random feature maps. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 19–27, 2014.
- Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, and Brian Kingsbury. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006.
- Kurt Hornik, Maxwell B. Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- Po-Sen Huang, Haim Avron, Tara N. Sainath, Vikas Sindhwani, and Bhuvana Ramabhadran. Kernel methods match deep neural networks on TIMIT. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2014, Florence, Italy, May 4-9, 2014*, pages 205–209. IEEE, 2014.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 448–456, 2015.
- G. J. O. Jameson. A simple proof of Stirling’s formula for the gamma function. 99(544):68–74, March 2015. ISSN 0025-5572 (print), 2056-6328 (electronic). doi: <http://dx.doi.org/10.1017/mag.2014.9>.

- Janez Kaiser, Bogomir Horvat, and Zdravko Kacic. A novel loss function for the overall risk criterion based discriminative training of HMM models. In *Sixth International Conference on Spoken Language Processing, ICSLP 2000 / INTERSPEECH 2000, Beijing, China, October 16-20, 2000*, pages 887–890, 2000.
- Purushottam Kar and Harish Karnick. Random feature maps for dot product kernels. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2012, La Palma, Canary Islands, April 21-23, 2012*, pages 583–591, 2012.
- B. Kingsbury, J. Cui, X. Cui, M. J. F. Gales, K. Knill, J. Mamou, L. Mangu, D. Nolden, M. Picheny, B. Ramabhadran, R. Schlüter, A. Sethy, and P. C. Woodland. A High-performance Cantonese Keyword Search System. pages 8277–8281, 2013.
- Brian Kingsbury. Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2009, 19-24 April 2009, Taipei, Taiwan*, pages 3761–3764, 2009.
- Quoc V. Le, Tamás Sarló, and Alexander J. Smola. Fastfood – approximating kernel expansions in loglinear time. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 244–252, 2013.
- Nicolas Le Roux, Mark W. Schmidt, and Francis R. Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States.*, pages 2672–2680, 2012.
- Zhiyun Lu, Dong Quo, Alireza Bagheri Garakani, Kuan Liu, Avner May, Aurélien Bellet, Linxi Fan, Michael Collins, Brian Kingsbury, Michael Picheny, and Fei Sha. A comparison between deep neural nets and kernel acoustic models for speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016, Shanghai, China, March 20-25, 2016*, pages 5070–5074. IEEE, 2016.
- Avner May, Michael Collins, Daniel J. Hsu, and Brian Kingsbury. Compact kernel models for acoustic modeling via random feature selection. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016, Shanghai, China, March 20-25, 2016*, pages 2424–2428. IEEE, 2016.
- Abdel-rahman Mohamed, George Dahl, and Geoffrey Hinton. Acoustic Modeling Using Deep Belief Networks. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):14–22, 2012.
- Guido F. Montúfar, Razvan Pascanu, KyungHyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2924–2932, 2014.
- N. Morgan and H. Bourlard. Generalization and parameter estimation in feedforward nets: Some experiments. In *Advances in Neural Information Processing Systems 2*, 1990.
- Jeffrey Pennington, Felix X. Yu, and Sanjiv Kumar. Spherical random features for polynomial kernels. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors,

- Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1846–1854, 2015.
- John C. Platt. Fast Training of Support Vector Machines using Sequential Minimal Optimization. In *Advances in Kernel Methods - Support Vector Learning*. MIT Press, 1998.
- D. Povey and P. C. Woodland. Minimum phone error and i-smoothing for improved discriminative training. In *Acoustics, Speech, and Signal Processing (ICASSP), 2002 IEEE International Conference on*, volume 1, pages I–105–I–108, May 2002.
- D. Povey, D. Kanevsky, B. Kingsbury, B. Ramabhadran, G. Saon, and K. Visweswariah. Boosted MMI for model and feature-space discriminative training. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4057–4060, March 2008.
- Daniel Povey and Brian Kingsbury. Evaluation of proposed modifications to MPE for large scale discriminative training. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2007, Honolulu, Hawaii, USA, April 15-20, 2007*, pages 321–324, 2007.
- Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahramani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur. Purely sequence-trained neural networks for asr based on lattice-free mmi. In *Interspeech*, 2016.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, pages 1177–1184, 2007.
- Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, pages 1313–1320, 2008.
- Tara N Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arısoy, and Bhuvana Ramabhadran. Low-rank Matrix Factorization for Deep Neural Network Training with High-dimensional Output Targets. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6655–6659. IEEE, 2013a.
- T.N. Sainath, B. Kingsbury, B. Ramabhadran, P. Fousek, P. Novak, and A. Mohamed. Making deep belief networks effective for large vocabulary continuous speech recognition. In *Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on*, pages 30–35. IEEE, 2011.
- T.N. Sainath, B. Kingsbury, H. Soltau, and B. Ramabhadran. Optimization techniques to improve training speed of deep neural networks for large speech tasks. *Audio, Speech, and Language Processing, IEEE Transactions on*, 21(11):2267–2276, Nov 2013b. ISSN 1558-7916.
- Hasim Sak, Andrew W. Senior, and Françoise Beaufays. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *INTERSPEECH 2014, 15th Annual Conference of the International Speech Communication Association, Singapore, September 14-18, 2014*, pages 338–342, 2014.

- B. Schölkopf and A. Smola. *Learning with kernels*. MIT Press, 2002.
- Frank Seide, Gang Li, Xie Chen, and Dong Yu. Feature engineering in context-dependent deep neural networks for conversational speech transcription. In *2011 IEEE Workshop on Automatic Speech Recognition & Understanding, ASRU 2011, Waikoloa, HI, USA, December 11-15, 2011*, pages 24–29, 2011a.
- Frank Seide, Gang Li, and Dong Yu. Conversational speech transcription using context-dependent deep neural networks. In *INTERSPEECH 2011, 12th Annual Conference of the International Speech Communication Association, Florence, Italy, August 27-31, 2011*, pages 437–440, 2011b.
- Alex Smola. Personal communication, 2014.
- Hagen Soltau, George Saon, and Brian Kingsbury. The IBM attila speech recognition toolkit. In *2010 IEEE Spoken Language Technology Workshop, SLT 2010, Berkeley, California, USA, December 12-15, 2010*, pages 97–102, 2010.
- Sören Sonnenburg and Vojtech Franc. COFFIN: A computational framework for linear svms. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 999–1006. Omnipress, 2010.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014. URL <http://dl.acm.org/citation.cfm?id=2670313>.
- I. Steinwart. Sparseness of support vector machines—some asymptotically sharp bounds. In *Advances in Neural Information Processing Systems 16*, 2004.
- Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 1139–1147, 2013.
- Ivor W. Tsang, James T. Kwok, and Pak-Ming Cheung. Core Vector Machines: Fast SVM Training on Very Large Data Sets. *Journal of Machine Learning Research*, 6:363–392, 2005.
- V Valtchev, J.J Odell, P.C Woodland, and S.J Young. MMIE training of large vocabulary recognition systems. *Speech Communication*, 22(4):303 – 314, 1997. ISSN 0167-6393.
- A. Vedaldi and A. Zisserman. Efficient Additive Kernels via Explicit Feature Maps. *IEEE Trans. on Pattern Anal. & Mach. Intell.*, 34(3):480–492, 2012.
- Karel Veselý, Arnab Ghoshal, Lukás Burget, and Daniel Povey. Sequence-discriminative training of deep neural networks. In *INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France, August 25-29, 2013*, pages 2345–2349, 2013.
- C.K.I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In T.K. Leen, T.G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 682–688. MIT Press, 2001.
- W. Xiong, Jasha Droppo, Xuedong Huang, Frank Seide, Mike Seltzer, Andreas Stolcke, Dong Yu, and Geoffrey Zweig. Achieving human parity in conversational speech recognition. *CoRR*, abs/1610.05256, 2016. URL <http://arxiv.org/abs/1610.05256>.

E.-H. Yen, T.-W. Lin, S.-D. Lin, P.K. Ravikumar, and I.S. Dhillon. Sparse random feature algorithm as coordinate descent in Hilbert space. In *Advances in Neural Information Processing Systems 27*, 2014.

Felix X. Yu, Sanjiv Kumar, Henry A. Rowley, and Shih-Fu Chang. Compact nonlinear maps and circulant extensions. *CoRR*, abs/1503.03893, 2015. URL <http://arxiv.org/abs/1503.03893>.