# Radial Basis Function Network for Multi-task Learning 

Xuejun Liao<br>Department of ECE<br>Duke University<br>Durham, NC 27708-0291, USA<br>xjliao@ee.duke.edu

Lawrence Carin<br>Department of ECE<br>Duke University<br>Durham, NC 27708-0291, USA<br>lcarin@ee.duke.edu


#### Abstract

We extend radial basis function (RBF) networks to the scenario in which multiple correlated tasks are learned simultaneously, and present the corresponding learning algorithms. We develop the algorithms for learning the network structure, in either a supervised or unsupervised manner. Training data may also be actively selected to improve the network's generalization to test data. Experimental results based on real data demonstrate the advantage of the proposed algorithms and support our conclusions.


## 1 Introduction

In practical applications, one is frequently confronted with situations in which multiple tasks must be solved. Often these tasks are not independent, implying what is learned from one task is transferable to another correlated task. By making use of this transferability, each task is made easier to solve. In machine learning, the concept of explicitly exploiting the transferability of expertise between tasks, by learning the tasks simultaneously under a unified representation, is formally referred to as "multi-task learning" [1].

In this paper we extend radial basis function (RBF) networks [4,5] to the scenario of multitask learning and present the corresponding learning algorithms. Our primary interest is to learn the regression model of several data sets, where any given data set may be correlated with some other sets but not necessarily with all of them. The advantage of multi-task learning is usually manifested when the training set of each individual task is weak, i.e., it does not generalize well to the test data. Our algorithms intend to enhance, in a mutually beneficial way, the weak training sets of multiple tasks, by learning them simultaneously. Multi-task learning becomes superfluous when the data sets all come from the same generating distribution, since in that case we can simply take the union of them and treat the union as a single task. In the other extreme, when all the tasks are independent, there is no correlation to utilize and we learn each task separately.

The paper is organized as follows. We define the structure of multi-task RBF network in Section 2 and present the supervised learning algorithm in Section 3. In Section 4 we show how to learn the network structure in an unsupervised manner, and based on this we demonstrate how to actively select the training data, with the goal of improving the
generalization to test data. We perform experimental studies in Section 5 and conclude the paper in Section 6.

## 2 Multi-Task Radial Basis Function Network

Figure 1 schematizes the radial basis function (RBF) network structure customized to multitask learning. The network consists of an input layer, a hidden layer, and an output layer. The input layer receives a data point $\mathbf{x}=\left[x_{1}, \cdots, x_{d}\right]^{T} \in \mathbb{R}^{d}$ and submits it to the hidden layer. Each node at the hidden layer has a localized activation $\phi^{n}(\mathbf{x})=\phi\left(\left\|\mathbf{x}-\mathbf{c}_{n}\right\|, \sigma_{n}\right)$, $n=1, \cdots, N$, where $\|\cdot\|$ denotes the vector norm and $\phi^{n}(\cdot)$ is a radial basis function (RBF) localized around $\mathbf{c}_{n}$ with the degree of localization parameterized by $\sigma_{n}$. Choosing $\phi(z, \sigma)=\exp \left(-\frac{z^{2}}{2 \sigma^{2}}\right)$ gives the Gaussian RBF. The activations of all hidden nodes are weighted and sent to the output layer. Each output node represents a unique task and has its own hidden-to-output weights. The weighted activations of the hidden nodes are summed at each output node to produce the output for the associated task. Denoting $\mathbf{w}_{k}=\left[w_{0 k}, w_{1 k}, \cdots, w_{N k}\right]^{T}$ as the weights connecting hidden nodes to the $k$-th output node, then the output for the $k$-th task, in response to input $\mathbf{x}$, takes the form

$$
\begin{equation*}
f_{k}(\mathbf{x})=\mathbf{w}_{k}^{T} \phi(\mathbf{x}) \tag{1}
\end{equation*}
$$

where $\phi(\mathbf{x})=\left[\phi^{0}(\mathbf{x}), \phi^{1}(\mathbf{x}), \ldots, \phi^{N}(\mathbf{x})\right]^{T}$ is a column containing $N+1$ basis functions with $\phi^{0}(\mathbf{x}) \equiv 1$ a dummy basis accounting for the bias in Figure 1.


Figure 1: A multi-task structure of RBF Network. Each of the output nodes represents a unique task. Each task has its own hidden-to-output weights but all the tasks share the same hidden nodes. The activation of hidden node $n$ is characterized by a basis function $\phi_{n}(\mathbf{x})=\phi\left(\left\|\mathbf{x}-\mathbf{c}_{n}\right\|, \sigma_{n}\right)$. A typical choice of $\phi$ is $\phi(z, \sigma)=\exp \left(-\frac{z^{2}}{2 \sigma^{2}}\right)$, which gives the Gaussian RBF.

## 3 Supervised Learning

Suppose we have $K$ tasks and the data set of the $k$-th task is $\mathcal{D}_{k}=$ $\left\{\left(\mathbf{x}_{1 k}, y_{1 k}\right), \cdots,\left(\mathbf{x}_{J_{k} k}, y_{J_{k} k}\right)\right\}$, where $y_{i k}$ is the target (desired output) of $\mathbf{x}_{i k}$. By definition, a given data point $\mathbf{x}_{i k}$ is said to be supervised if the associated target $y_{i k}$ is provided and unsupervised if $y_{i k}$ is not provided. The definition extends similarly to a set of data

Table 1: Learning Algorithm of Multi-Task RBF Network
Input: $\left\{\left(\mathbf{x}_{1 k}, y_{2 k}\right), \cdots,\left(\mathbf{x}_{J_{k}, k}, y_{J_{k}, k}\right)\right\}_{k=1: K}, \phi(\cdot, \sigma), \sigma$, and $\rho$; Output: $\phi(\cdot)$ and $\left\{\mathbf{w}_{k}\right\}_{k=1}^{K}$.

1. For $m=1: K$, For $n=1: J_{m}$, For $k=1: K$, For $i=1: J_{k}$

Compute $\widehat{\phi}_{i k}^{n m}=\phi\left(\left\|\mathbf{x}_{n m}-\mathbf{x}_{i k}\right\|, \sigma\right)$;
2. Let $N=0, \phi(\cdot)=1, e_{0}=\sum_{k=1}^{K}\left[\sum_{i=1}^{J_{k}} y_{i k}^{2}-\left(J_{k}+\rho\right)^{-1}\left(\sum_{i=1}^{J_{k}} y_{i k}\right)^{2}\right]$;

For $k=1: K$, compute $\mathbf{A}_{k}=J_{k}+\rho, \mathbf{w}_{k}=\left(J_{k}+\rho\right)^{-1} \sum_{i=1}^{J_{k}} y_{i k}$;
3. For $m=1: K$, For $n=1: J_{m}$ If $\widehat{\phi}^{n m}$ is not marked as "deleted"

For $k=1: K$, compute

$$
\mathbf{c}_{k}=\sum_{i=1}^{J_{k}} \boldsymbol{\phi}_{i k} \widehat{\phi}_{i k}^{n m}, \quad q_{k}=\sum_{i=1}^{J_{k}}\left(\widehat{\phi}_{i k}^{n m}\right)^{2}+\rho-\mathbf{c}_{k}^{T} \mathbf{A}_{k}^{-1} \mathbf{c}_{k} ;
$$

If there exists $k$ such that $q_{k}=0$, mark $\widehat{\phi}^{n m}$ as "deleted";
else, compute $\delta e\left(\phi, \widehat{\phi}^{n m}\right)$ using (5).
4. If $\left\{\widehat{\phi}^{i k}\right\}_{i=1: J_{k}, k=1: K}$ are all marked as "deleted", go to 10 .
5. Let $\left(n^{*}, m^{*}\right)=\arg \max _{\widehat{\phi}^{n m}}$ not marked as "deleted" $\delta e\left(\boldsymbol{\phi}, \widehat{\phi}^{n m}\right)$; Mark $\widehat{\phi}^{n^{*}} m^{*}$ as "deleted".
6. Tune RBF parameter $\sigma_{N+1}=\arg \max _{\sigma} \delta e\left(\phi, \phi\left(\left\|\cdot-\mathbf{x}_{n^{*} m^{*}}\right\|, \sigma\right)\right)$
7. Let $\phi^{N+1}(\cdot)=\phi\left(\left\|\cdot-\mathbf{x}_{n^{*} m^{*}}\right\|, \sigma_{N+1}\right)$; Update $\phi(\cdot) \leftarrow\left[\phi^{T}(\cdot), \phi^{N+1}(\cdot)\right]^{T}$;
8. For $k=1$ : $K$

Compute $\mathbf{A}_{k}^{n e w}$ and $\mathbf{w}_{k}^{\text {new }}$ respectively by (A-1) and (A-3) in the appendix; Update $\mathbf{A}_{k} \leftarrow \mathbf{A}_{k}^{n e w}, \mathbf{w}_{k} \leftarrow \mathbf{w}_{k}^{n e w}$
9. Let $e_{N+1}=e_{N}-\delta e\left(\phi, \phi^{N+1}\right)$; If the sequence $\left\{e_{n}\right\}_{n=0:(N+1)}$ is converged, go to 10 , else update $N \leftarrow N+1$ and go back to 3 .
10. Exit and output $\phi(\cdot)$ and $\left\{\mathbf{w}_{k}\right\}_{k=1}^{K}$.
points. We are interested in learning the functions $f_{k}(\mathbf{x})$ for the $K$ tasks, based on $\cup_{k=1}^{K} \mathcal{D}_{k}$. The learning is based on minimizing the squared error

$$
\begin{equation*}
e(\boldsymbol{\phi}, \mathbf{w})=\sum_{k=1}^{K}\left\{\sum_{i=1}^{J_{k}}\left(\mathbf{w}_{k}^{T} \boldsymbol{\phi}_{i k}-y_{i k}\right)^{2}+\rho\left\|\mathbf{w}_{k}\right\|^{2}\right\} \tag{2}
\end{equation*}
$$

where $\phi_{i k}=\phi\left(\mathbf{x}_{i k}\right)$ for notational simplicity. The regularization terms $\rho\left\|\mathbf{w}_{k}\right\|^{2}, k=$ $1, \cdots, K$, are used to prevent singularity of the A matrices defined in (3), and $\rho$ is typically set to a small positive number. For fixed $\boldsymbol{\phi}$ 's, the w's are solved by minimizing $e(\boldsymbol{\phi}, \mathbf{w})$ with respect to $\mathbf{w}$, yielding

$$
\begin{equation*}
\mathbf{w}_{k}=\mathbf{A}_{k}^{-1} \sum_{i=1}^{J_{k}} y_{i k} \boldsymbol{\phi}_{i k} \quad \text { and } \quad \mathbf{A}_{k}=\sum_{i=1}^{J_{k}} \boldsymbol{\phi}_{i k} \boldsymbol{\phi}_{i k}^{T}+\rho \mathbf{I}, \quad k=1, \cdots, K \tag{3}
\end{equation*}
$$

In a multi-task RBF network, the input layer and output layer are respectively specified by the data dimensionality and the number of tasks. We now discuss how to determine the hidden layer (basis functions $\phi$ ). Substituting the solutions of the w's in (3) into (2) gives

$$
\begin{equation*}
e(\boldsymbol{\phi})=\sum_{k=1}^{K} \sum_{i=1}^{J_{k}}\left(y_{i k}^{2}-y_{i k} \mathbf{w}_{k}^{T} \boldsymbol{\phi}_{i k}\right) \tag{4}
\end{equation*}
$$

where $e(\phi)$ is a function of $\phi$ only because w's are now functions of $\phi$ as given by (3). By minimizing $e(\phi)$, we can determine $\phi$. Recalling that $\phi_{i k}$ is an abbreviation of $\phi\left(\mathbf{x}_{i k}\right)=\left[1, \phi^{1}\left(\mathbf{x}_{i k}\right), \ldots, \phi^{N}\left(\mathbf{x}_{i k}\right)\right]^{T}$, this amounts to determining $N$, the number of basis functions, and the functional form of each basis function $\phi^{n}(\cdot), n=1, \ldots, N$. Consider the candidate functions $\left\{\phi^{n m}(\mathbf{x})=\phi\left(\left\|\mathbf{x}-\mathbf{x}_{n m}\right\|, \sigma\right): n=1, \cdots, J_{m}, m=1, \cdots, K\right\}$. We learn the RBF network structure by selecting $\phi(\cdot)$ from these candidate functions such that $e(\phi)$ in (4) is minimized. The following theorem tells us how to perform the selection in a sequential way; the proof is given in the Appendix.

Theorem 1 Let $\phi(\mathbf{x})=\left[1, \phi^{1}(\mathbf{x}), \ldots, \phi^{N}(\mathbf{x})\right]^{T}$ and $\phi^{N+1}(\mathbf{x})$ be a single basis function. Assume the A matrices corresponding to $\phi$ and $\left[\phi, \phi^{N+1}\right]^{T}$ are all non-degenerate. Then

$$
\begin{equation*}
\delta e\left(\boldsymbol{\phi}, \phi^{N+1}\right)=e(\boldsymbol{\phi})-e\left(\left[\boldsymbol{\phi}, \phi^{N+1}\right]^{T}\right)=\sum_{k=1}^{K}\left(\mathbf{c}_{k}^{T} \mathbf{w}_{k}-\sum_{i=1}^{J_{k}} y_{i k} \phi_{i k}^{N+1}\right)^{2} q_{k}^{-1} \tag{5}
\end{equation*}
$$

where $\phi_{i k}^{N+1}=\phi^{N+1}\left(\phi_{i k}\right), \mathbf{w}_{k}$ and $\mathbf{A}$ are the same as in (3), and

$$
\begin{equation*}
\mathbf{c}_{k}=\sum_{i=1}^{J_{k}} \phi_{i k} \phi_{i k}^{N+1}, \quad d_{k}=\sum_{i=1}^{J_{k}}\left(\phi_{i k}^{N+1}\right)^{2}+\rho, \quad q_{k}=d_{k}-\mathbf{c}_{k}^{T} \mathbf{A}_{k}^{-1} \mathbf{c}_{k} \tag{6}
\end{equation*}
$$

By the conditions of the theorem $\mathbf{A}_{k}^{\text {new }}$ is full rank and hence it is positive definite by construction. By (A-2) in the Appendix, $q_{k}^{-1}$ is a diagonal element of $\left(\mathbf{A}_{k}^{n e w}\right)^{-1}$, therefore $q_{k}^{-1}$ is positive and by (5) $\delta e\left(\phi, \phi^{N+1}\right)>0$, which means adding $\phi^{N+1}$ to $\phi$ generally makes the squared error decrease. The decrease $\delta e\left(\phi, \phi^{N+1}\right)$ depends on $\phi^{N+1}$. By sequentially selecting basis functions that bring the maximum error reduction, we achieve the goal of maximizing $e(\boldsymbol{\phi})$. The details of the learning algorithm are summarized in Table 1.

## 4 Active Learning

In the previous section, the data in $D_{k}$ are supervised (provided with the targets). In this section, we assume the data in $D_{k}$ are initially unsupervised (only $\mathbf{x}$ is available without access to the associated $y$ ) and we select a subset from $\mathcal{D}_{k}$ to be supervised (targets acquired) such that the resulting network generalizes well to the remaining data in $\mathcal{D}_{k}$. The approach is generally known as active learning [6]. We first learn the basis functions $\phi$ from the unsupervised data, and based on $\phi$ select data to be supervised. Both of these steps are based on the following theorem, the proof of which is given in the Appendix.
Theorem 2 Let there be $K$ tasks and the data set of the $k$-th task is $\mathcal{D}_{k} \cup \widetilde{\mathcal{D}}_{k}$ where $\mathcal{D}_{k}=\left\{\left(\mathbf{x}_{i k}, y_{i k}\right)\right\}_{i=1}^{J_{k}}$ and $\widetilde{\mathcal{D}}_{k}=\left\{\left(\mathbf{x}_{i k}, y_{i k}\right)\right\}_{i=J_{k}+1}^{J_{k}+\widetilde{J}_{k}}$. Let there be two multi-task RBF networks, whose output nodes are characterized by $f_{k}(\cdot)$ and $f_{k}^{\sim}(\cdot)$, respectively, for task $k=1, \ldots, K$. The two networks have the same given basis functions (hidden nodes) $\phi(\cdot)=\left[1, \phi^{1}(\cdot), \cdots, \phi^{N}(\cdot)\right]^{T}$, but different hidden-to-output weights. The weights of $f_{k}(\cdot)$ are trained with $\mathcal{D}_{k} \cup \widetilde{\mathcal{D}}_{k}$, while the weights of $f_{k}^{\sim}(\cdot)$ are trained using $\widetilde{\mathcal{D}}_{k}$. Then for $k=1, \cdots, K$, the square errors committed on $\mathcal{D}_{k}$ by $f_{k}(\cdot)$ and $f_{k}^{\sim}(\cdot)$ are related by
$0 \leq\left[\operatorname{det} \boldsymbol{\Gamma}_{k}\right]^{-1} \leq \lambda_{\max , k}^{-1} \leq\left[\sum_{i=1}^{J_{k}}\left(y_{i k}-f_{k}^{\sim}\left(\mathbf{x}_{i k}\right)\right)^{2}\right]^{-1} \sum_{i=1}^{J_{k}}\left(y_{i k}-f_{k}\left(\mathbf{x}_{i k}\right)\right)^{2} \leq \lambda_{\min , k}^{-1} \leq 1$ (7) where $\boldsymbol{\Gamma}_{k}=\left[\mathbf{I}+\boldsymbol{\Phi}_{k}^{T}\left(\rho \mathbf{I}+\widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\mathbf{\Phi}}_{k}^{T}\right)^{-1} \mathbf{\Phi}_{k}\right]^{2}$ with $\boldsymbol{\Phi}=\left[\boldsymbol{\phi}\left(\mathbf{x}_{1 k}\right), \ldots, \boldsymbol{\phi}\left(\mathbf{x}_{J_{k} k}\right)\right]$ and $\widetilde{\boldsymbol{\Phi}}=$ $\left[\phi\left(\mathbf{x}_{J_{k}+1, k}\right), \ldots, \phi\left(\mathbf{x}_{J_{k}+\widetilde{J}_{k}, k}\right)\right]$, and $\lambda_{\max , k}$ and $\lambda_{\min , k}$ are respectively the largest and smallest eigenvalues of $\boldsymbol{\Gamma}_{k}$.
Specializing Theorem 2 to the case $\widetilde{J}_{k}=0$, we have
Corollary 1 Let there be $K$ tasks and the data set of the $k$-th task is $\mathcal{D}_{k}=\left\{\left(\mathbf{x}_{i k}, y_{i k}\right)\right\}_{i=1}^{J_{k}}$. Let the RBF network, whose output nodes are characterized by $f_{k}(\cdot)$ for task $k=$ $1, \ldots, K$, have given basis functions (hidden nodes) $\phi(\cdot)=\left[1, \phi^{1}(\cdot), \cdots, \phi^{N}(\cdot)\right]^{T}$ and the hidden-to-output weights of task $k$ be trained with $\mathcal{D}_{k}$. Then for $k=1, \cdots, K$, the squared error committed on $\mathcal{D}_{k}$ by $f_{k}(\cdot)$ is bounded as $0 \leq\left[\operatorname{det} \boldsymbol{\Gamma}_{k}\right]^{-1} \leq \lambda_{\max , k}^{-1} \leq$ $\left[\sum_{i=1}^{J_{k}} y_{i k}^{2}\right]^{-1} \sum_{i=1}^{J_{k}}\left(y_{i k}-f_{k}\left(\mathbf{x}_{i k}\right)\right)^{2} \leq \lambda_{\min , k}^{-1} \leq 1$, where $\boldsymbol{\Gamma}_{k}=\left(\mathbf{I}+\rho^{-1} \boldsymbol{\Phi}_{k}^{T} \boldsymbol{\Phi}_{k}\right)^{2}$ with $\boldsymbol{\Phi}=\left[\phi\left(\mathrm{x}_{1, k}\right), \ldots, \phi\left(\mathrm{x}_{J_{k}, k}\right)\right]$, and $\lambda_{\max , k}$ and $\lambda_{\min , k}$ are respectively the largest and smallest eigenvalues of $\boldsymbol{\Gamma}_{k}$.
It is evident from the properties of matrix determinant [7] and the definition of $\Phi$ that $\operatorname{det} \boldsymbol{\Gamma}_{k}=\left[\operatorname{det}\left(\rho \mathbf{I}+\boldsymbol{\Phi}_{k} \boldsymbol{\Phi}_{k}^{T}\right)\right]^{2}[\operatorname{det}(\rho \mathbf{I})]^{-2}=\left[\operatorname{det}\left(\rho \mathbf{I}+\sum_{i=1}^{J_{k}} \boldsymbol{\phi}_{i k} \boldsymbol{\phi}_{i k}^{T}\right)\right]^{2}[\operatorname{det}(\rho \mathbf{I})]^{-2}$.

Using (3) we write succinctly $\operatorname{det} \boldsymbol{\Gamma}_{k}=\left[\operatorname{det} \mathbf{A}_{k}^{2}\right][\operatorname{det}(\rho \mathbf{I})]^{-2}$. We are interested in selecting the basis functions $\phi$ that minimize the error, before seeing $y$ 's. By Corollary 1 and the equation $\operatorname{det} \boldsymbol{\Gamma}_{k}=\left[\operatorname{det} \mathbf{A}_{k}^{2}\right][\operatorname{det}(\rho \mathbf{I})]^{-2}$, the squared error is lower bounded by $\sum_{i=1}^{J_{k}} y_{i k}^{2}[\operatorname{det}(\rho \mathbf{I})]^{2}\left[\operatorname{det} \mathbf{A}_{k}\right]^{-2}$. Instead of minimizing the error directly, we minimize its lower bound. As $[\operatorname{det}(\rho \mathbf{I})]^{2} \sum_{i=1}^{L_{k}} y_{i k}^{2}$ does not depend on $\phi$, this amounts to selecting $\phi$ to minimize $\left(\operatorname{det} \mathbf{A}_{k}\right)^{-2}$. To minimize the errors for all tasks $k=1 \cdots, K$, we select $\phi$ to $\operatorname{minimize} \prod_{k=1}^{K}\left(\operatorname{det} \mathbf{A}_{k}\right)^{-2}$.
The selection proceeds in a sequential manner. Suppose we have selected basis functions $\boldsymbol{\phi}=\left[1, \phi^{1}, \cdots, \phi^{N}\right]^{T}$. The associated $\mathbf{A}$ matrices are $\mathbf{A}_{k}=\sum_{i=1}^{J_{k}} \phi_{i k} \boldsymbol{\phi}_{i k}^{T}+$ $\rho \mathbf{I}_{(N+1) \times(N+1)}, k=1, \cdots, K$. Augmenting basis functions to $\left[\phi^{T}, \phi^{N+1}\right]^{T}$, the $\mathbf{A}$ matrices change to $\mathbf{A}_{k}^{\text {new }}=\sum_{i=1}^{J_{k}}\left[\boldsymbol{\phi}_{i k}^{T}, \phi_{i k}^{N+1}\right]^{T}\left[\boldsymbol{\phi}_{i k}^{T}, \phi_{i k}^{N+1}\right]+\rho \mathbf{I}_{(N+2) \times(N+2)}$. Using the determinant formula of block matrices [7], we get $\prod_{k=1}^{K}\left(\operatorname{det} \mathbf{A}_{k}^{n e w}\right)^{-2}=$ $\prod_{k=1}^{K}\left(q_{k} \operatorname{det} \mathbf{A}_{k}\right)^{-2}$, where $q_{k}$ is the same as in (6). As $\mathbf{A}_{k}$ does not depend on $\phi^{N+1}$, the left-hand side is minimized by maximizing $\prod_{k=1}^{K} q_{k}^{2}$. The selection is easily implemented by making the following two minor modifications in Table 1: (a) in step 2, compute $e_{0}=\sum_{k=1}^{K} \ln \left(J_{k}+\rho\right)^{-2}$; in step 3, compute $\delta e\left(\phi, \widehat{\phi}^{n m}\right)=\sum_{k=1}^{K} \ln q_{k}^{2}$. Employing the logarithm is for gaining additivity and it does not affect the maximization.

Based on the basis functions $\phi$ determined above, we proceed to selecting data to be supervised and determining the hidden-to-output weights $\mathbf{w}$ from the supervised data using the equations in (3). The selection of data is based on an iterative use of the following corollary, which is a specialization of Theorem 2 and was originally given in [8].

Corollary 2 Let there be $K$ tasks and the data set of the $k$-th task is $\mathcal{D}_{k}=\left\{\left(\mathbf{x}_{i k}, y_{i k}\right)\right\}_{i=1}^{J_{k}}$. Let there be two RBF networks, whose output nodes are characterized by $f_{k}(\cdot)$ and $f_{k}^{+}(\cdot)$, respectively, for task $k=1, \ldots, K$. The two networks have the same given basis functions $\phi(\cdot)=\left[1, \phi^{1}(\cdot), \cdots, \phi^{N}(\cdot)\right]^{T}$, but different hidden-to-output weights. The weights of $f_{k}(\cdot)$ are trained with $\mathcal{D}_{k}$, while the weights of $f_{k}^{+}(\cdot)$ are trained using $\mathcal{D}_{k}^{+}=\mathcal{D}_{k} \cup\left\{\left(\mathbf{x}_{J_{k}+1, k}, y_{J_{k}+1, k}\right)\right\}$. Then for $k=1, \cdots, K$, the squared errors committed on $\left(\mathbf{x}_{J_{k}+1, k}, y_{J_{k}+1, k}\right)$ by $f_{k}(\cdot)$ and $f_{k}^{+}(\cdot)$ are related by $\left[f_{k}^{+}\left(\mathbf{x}_{J_{k}+1, k}\right)-\right.$ $\left.y_{J_{k}+1, k}\right]^{2}=\left[\gamma\left(\mathbf{x}_{J_{k}+1, k}\right)\right]^{-1}\left[f_{k}\left(\mathbf{x}_{J_{k}+1, k}\right)-y_{J_{k}+1, k}\right]^{2}$, where $\gamma\left(\mathbf{x}_{J_{k}+1, k}\right)=[1+$ $\left.\phi^{T}\left(\mathbf{x}_{J_{k}+1, k}\right) \mathbf{A}_{k}^{-1} \phi\left(\mathbf{x}_{J_{k}+1, k}\right)\right]^{2} \geq 1$ and $\mathbf{A}_{k}=\sum_{i=1}^{J_{k}}\left[\rho \mathbf{I}+\boldsymbol{\phi}\left(\mathbf{x}_{i k}\right) \boldsymbol{\phi}^{T}\left(\mathbf{x}_{i k}\right)\right]$ is the same as in (3).

Two observations are made from Corollary 2. First, if $\gamma\left(\mathbf{x}_{J_{k}+1, k}\right) \approx 1$, seeing $y_{J_{k}+1, k}$ does not effect the error on $\mathbf{x}_{J_{k}+1, k}$, indicating $\mathcal{D}_{k}$ already contain sufficient information about $\left(\mathbf{x}_{J_{k}+1, k}, y_{J_{k}+1, k}\right)$. Second, if $\gamma\left(\mathbf{x}_{i}\right) \gg 1$, seeing $y_{J_{k}+1, k}$ greatly decrease the error on $\mathbf{x}_{J_{k}+1, k}$, indicating $\mathbf{x}_{J_{k}+1, k}$ is significantly dissimilar (novel) to $\mathcal{D}_{k}$ and $\mathbf{x}_{J_{k}+1, k}$ must be supervised to reduce the error. Based on Corollary 2, the selection proceeds sequentially. Suppose we have selected data $\mathcal{D}_{k}=\left\{\left(\mathbf{x}_{i k}, y_{i k}\right)\right\}_{i=1}^{J_{k}}$, from which we compute $\mathbf{A}_{k}$. We select the next data point as $\mathbf{x}_{J_{k}+1, k}=\arg \max _{i>J_{k}, k=1, \cdots, K} \gamma\left(\mathbf{x}_{i k}\right)=$ $\arg \max _{i>J_{k}} k=1, \cdots, K\left[1+\phi^{T}\left(\mathbf{x}_{i k}\right) \mathbf{A}_{k}^{-1} \boldsymbol{\phi}\left(\mathbf{x}_{i k}\right)\right]^{2}$. After $\mathbf{x}_{J_{k}+1, k}$ is selected, the $\mathbf{A}_{k}$ is updated and the next selection begins. As the iteration advances $\gamma$ will decrease until it reaches convergence. We use (3) to compute $\mathbf{w}$ from the selected $\mathbf{x}$ and their associated targets $y$, completing learning of the RBF network.

## 5 Experimental Results

In this section we compare the multi-task RBF network against single-task RBF networks via experimental studies. We consider three types of RBF networks to learn $K$ tasks, each
with its data set $\mathcal{D}_{k}$. In the first, which we call "one RBF network", we let the $K$ tasks share both basis functions $\phi$ (hidden nodes) and hidden-to output weights $\mathbf{w}$, thus we do not distinguish the $K$ tasks and design a single RBF network to learn a union of them. The second is the multi-task RBF network, where the $K$ tasks share the same $\phi$ but each has its own $\mathbf{w}$. In the third, we have $K$ independent networks, each designed for a single task.

We use a school data set from the Inner London Education Authority, consisting of examination records of 15362 students from 139 secondary schools. The data are available at http://multilevel.ioe.ac.uk/intro/datasets.html. This data set was originally used to study the effectiveness of schools and has recently been used to evaluate multi-task algorithms $[2,3]$. The goal is to predict the exam scores of the students based on 9 variables: year of exam (1985, 1986, or 1987), school code (1-139), FSM (percentage of students eligible for free school meals), VR1 band (percentage of students in school in VR band one), gender, VR band of student (3 categories), ethnic group of student (11 categories), school gender (male, female, or mixed), school denomination (3 categories). We consider each school a task, leading to 139 tasks in total. The remaining 8 variables are used as inputs to the RBF network. Following [2,3], we converted each categorical variable to a number of binary variables, resulting in a total number of 27 input variables, i.e., $x \in \mathbb{R}^{27}$. The exam score is the target to be predicted.
The three types of RBF networks as defined above are designed as follows. The multi-task RBF network is implemented as the structure as shown in Figure 1 and trained with the learning algorithm in Table 1. The "one RBF network" is implemented as a special case of Figure 1, with a single output node and trained using the union of supervised data from all 139 schools. We design 139 independent RBF networks, each of which is implemented with a single output node and trained using the supervised data from a single school. We use the Gaussian $\operatorname{RBF} \phi^{n}(\mathbf{x})=\exp \left(-\frac{\left\|\mathbf{x}-\mathbf{c}_{n}\right\|^{2}}{2 \sigma^{2}}\right)$, where the $\mathbf{c}_{n}$ 's are selected from training data points and $\sigma_{n}$ 's are initialized as 20 and optimized as described in Table 1. The main role of the regularization parameter $\rho$ is to prevent the $\mathbf{A}$ matrices from being singular and it does not affect the results seriously. In the results reported here, $\rho$ is set to $10^{-6}$.
Following [2-3], we randomly take $75 \%$ of the 15362 data points as training (supervised) data and the remaining $25 \%$ as test data. The generalization performance is measured by the squared error $\left(f_{k}\left(\mathbf{x}_{i k}\right)-y_{i k}\right)^{2}$ averaged over all test data $\mathbf{x}_{i k}$ of tasks $k=1, \cdots, K$. We made 10 independent trials to randomly split the data into training and test sets and the squared error averaged over the test data of all the 139 schools and the trials are shown in Table 2, for the three types of RBF networks.

Table 2: Squared error averaged over the test data of all 139 schools and the 10 independent trials for randomly splitting the school data into training ( $75 \%$ ) and testing ( $25 \%$ ) sets.

| Multi-task RBF network | Independent RBF networks | One RBF network |
| :--- | :--- | :--- |
| $109.89 \pm 1.8167$ | $136.41 \pm 7.0081$ | $149.48 \pm 2.8093$ |

Table 2 clearly shows the multi-task RBF network outperforms the other two types of RBF networks by a considerable margin. The "one RBF network" ignores the difference between the tasks and the independent RBF networks ignore the tasks' correlations, therefore they both perform inferiorly. The multi-task RBF network uses the shared hidden nodes (basis functions) to capture the common internal representation of the tasks and meanwhile uses the independent hidden-to-output weights to learn the statistics specific to each task.
We now demonstrate the results of active learning. We use the method in Section 4 to actively split the data into training and test sets using a two-step procedure. First we learn the basis functions $\phi$ of multi-task RBF network using all 15362 data (unsupervised). Based on the $\phi$, we then select the data to be supervised and use them as training data to learn
the hidden-to-output weights $\mathbf{w}$. To make the results comparable, we use the same training data to learn the other two types of RBF networks (including learning their own $\phi$ and $\mathbf{w}$ ). The networks are then tested on the remaining data.

Figure 2 shows the results of active learning. Each curve is the squared error averaged over the test data of all 139 schools, as a function of number of training data. It is clear that the multi-task RBF network maintains its superior performance all the way down to 5000 training data points, whereas the independent RBF networks have their performances degraded seriously as the training data diminish. This demonstrates the increasing advantage of multi-task learning as the number of training data decreases. The "one RBF network" seems also insensitive to the number of training data, but it ignores the inherent dissimilarity between the tasks, which makes its performance inferior.


Figure 2: Squared error averaged over the test data of all 139 schools, as a function of the number of training (supervised) data. The data are split into training and test sets via active learning.

## 6 Conclusions

We have presented the structure and learning algorithms for multi-task learning with the radial basis function (RBF) network. By letting multiple tasks share the basis functions (hidden nodes) we impose a common internal representation for correlated tasks. Exploiting the inter-task correlation yields a more compact network structure that has enhanced generalization ability. Unsupervised learning of the network structure enables us to actively split the data into training and test sets. As the data novel to the previously selected ones are selected next, what finally remain unselected and to be tested are all similar to the selected data which constitutes the training set. This improves the generalization of the resulting network to the test data. These conclusions are substantiated via results on real multi-task data.

## References

[1] R. Caruana. (1997) Multitask learning. Machine Learning, 28, p. 41-75, 1997.
[2] B. Bakker and T. Heskes (2003). Task clustering and gating for Bayesian multitask learning. Journal of Machine Learning Research, 4: 83-99, 2003
[3] T. Evgeniou, C. A. Micchelli, and M. Pontil (2005). Learning Multiple Tasks with Kernel Methods. Journal of Machine Learning Research, 6: 615637, 2005
[4] Powell M. (1987), Radial basis functions for multivariable interpolation : A review, J.C. Mason and M.G. Cox, eds, Algorithms for Approximation, pp.143-167.
[5] Chen, F. Cowan, and P. Grant (1991), Orthogonal least squares learning algorithm for radial basis function networks, IEEE Transactions on Neural Networks, Vol. 2, No. 2, 302-309, 1991
[6] Cohn, D. A., Ghahramani, Z., and Jordan, M. I. (1995). Active learning with statistical models. Advances in Neural Information Processing Systems, 7, 705-712.
[7] V. Fedorov (1972), Theory of Optimal Experiments, Academic Press, 1972
[8] M. Stone (1974), Cross-validatory choice and assessment of statistical predictions, Journal of the Royal Statistical Society, Series B, 36, pp. 111-147, 1974.

## Appendix

Proof of Theorem 1:. Let $\phi^{\text {new }}=\left[\phi, \phi^{N+1}\right]^{T}$. By (3), the A matrices corresponding to $\phi^{\text {new }}$ are

$$
\mathbf{A}_{k}^{n e w}=\sum_{i=1}^{J_{k}}\left[\begin{array}{c}
\phi_{i k}  \tag{A-1}\\
\phi_{i k}^{N+1}
\end{array}\right]\left[\begin{array}{ll}
\boldsymbol{\phi}_{i k}^{T} & \phi_{i k}^{N+1}
\end{array}\right]+\rho \mathbf{I}_{(N+2) \times(N+2)}=\left[\begin{array}{cc}
\mathbf{A}_{k} & \mathbf{c}_{k} \\
\mathbf{c}_{k}^{T} & d_{k}
\end{array}\right]
$$

where $\mathbf{c}_{k}$ and $d_{k}$ are as in (6). By the conditions of the theorem, the matrices $\mathbf{A}_{k}$ and $\mathbf{A}_{k}^{n e w}$ are all non-degenerate. Using the block matrix inversion formula [7] we get

$$
\left(\mathbf{A}_{k}^{\text {new }}\right)^{-1}=\left[\begin{array}{cc}
\mathbf{A}_{k}^{-1}+\mathbf{A}_{k}^{-1} \mathbf{c}_{k} q_{k}^{-1} \mathbf{c}_{k}^{T} \mathbf{A}_{k}^{-1} & -\mathbf{A}_{k}^{-1} \mathbf{c}_{k} q_{k}^{-1}  \tag{A-2}\\
-q_{k}^{-1} \mathbf{c}_{k}^{T} \mathbf{A}_{k}^{-1} & q_{k}^{-1}
\end{array}\right]
$$

where $q_{k}$ is as in (6). By (3), the weights $\mathbf{w}_{k}^{\text {new }}$ corresponding to $\left[\phi^{T}, \phi^{N+1}\right]^{T}$ are

$$
\mathbf{w}_{k}^{\text {new }}=\left(\mathbf{A}_{k}^{n e w}\right)^{-1}\left[\begin{array}{c}
\sum_{i=1}^{J_{k}} y_{i k} \boldsymbol{\phi}_{i k}  \tag{A-3}\\
\sum_{i=1}^{J_{k}} y_{i k} \phi_{i k}^{N+1}
\end{array}\right]=\left[\begin{array}{c}
\mathbf{w}_{k}+\mathbf{A}_{k}^{-1} \mathbf{c}_{k} q_{k}^{-1} g_{k} \\
-q_{k}^{-1} g_{k}
\end{array}\right]
$$

with $g_{k}=\mathbf{c}_{k}^{T} \mathbf{w}_{k}-\sum_{i=1}^{J_{k}} y_{i k} \phi_{i k}^{N+1}$. Hence, $\left(\phi_{i k}^{\text {new }}\right)^{T} \mathbf{w}_{k}^{\text {new }}=\phi_{i k}^{T} \mathbf{w}_{k}+\left(\boldsymbol{\phi}_{i k}^{T} \mathbf{A}_{k}^{-1} \mathbf{c}_{k}-\right.$ $\left.\phi_{i k}^{N+1}\right) g_{k} q_{k}^{-1}$, which is put into (4) to get $e\left(\boldsymbol{\phi}^{\text {new }}=\sum_{k=1}^{K} \sum_{i=1}^{J_{k}}\left[y_{i k}^{2}-y_{i k}\left(\boldsymbol{\phi}_{i k}^{\text {new }}\right)^{T} \mathbf{w}_{k}^{\text {new }}\right]=\right.$ $\sum_{k=1}^{K} \sum_{i=1}^{J_{k}}\left[y_{i k}^{2}-y_{i k} \boldsymbol{\phi}_{i k}^{T} \mathbf{w}_{k}-y_{i k}\left(\boldsymbol{\phi}_{i k}^{T} \mathbf{A}_{k}^{-1} \mathbf{c}_{k}-\phi_{i k}^{N+1}\right) g_{k} q_{k}^{-1}\right]=e(\boldsymbol{\phi})-\sum_{k=1}^{K}\left(\mathbf{c}_{k}^{T} \mathbf{w}_{k}-\right.$ $\left.\sum_{i=1}^{J_{k}} y_{i k} \phi_{i k}^{N+1}\right)^{2} q_{k}^{-1}$, where in arriving the last equality we have used (3) and (4) and $g_{k}=$ $\mathbf{c}_{k}^{T} \mathbf{w}_{k}-\sum_{i=1}^{J_{k}} y_{i k} \phi_{i k}^{N+1}$. The theorem is proved.
Proof of Theorem 2: The proof applies to $k=1, \cdots, K$. For any given $k$, define $\boldsymbol{\Phi}=$ $\left[\phi\left(\mathbf{x}_{1 k}\right), \ldots, \phi\left(\mathbf{x}_{J_{k} k}\right)\right], \widetilde{\boldsymbol{\Phi}}=\left[\phi\left(\mathbf{x}_{J_{k}+1, k}\right), \ldots, \phi\left(\mathbf{x}_{J_{k}+\tilde{J}_{k}, k}\right)\right], \mathbf{y}_{k}=\left[y_{1 k}, \ldots, y_{J_{k} k}\right]^{T}, \widetilde{\mathbf{y}}_{k}=$ $\left[y_{J_{k}+1, k}, \ldots, y_{J_{k}+\tilde{J}_{k}, k}\right]^{T}, \mathbf{f}_{k}=\left[f\left(\mathbf{x}_{1 k}\right), \ldots, f\left(\mathbf{x}_{J_{k} k}\right)\right]^{T}, \mathbf{f}_{k}^{\sim}=\left[f_{k}^{\sim}\left(\mathbf{x}_{1 k}\right), \ldots, f_{k}^{\sim}\left(\mathbf{x}_{J_{k} k}\right)\right]^{T}$, and $\widetilde{\mathbf{A}}_{k}=\rho \mathbf{I}+\widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\boldsymbol{\Phi}}_{k}^{T}$. By (1), (3), and the conditions of the theorem, $\mathbf{f}_{k}=\boldsymbol{\Phi}_{k}^{T}\left(\widetilde{\mathbf{A}}_{k}+\right.$ $\left.\boldsymbol{\Phi}_{k} \boldsymbol{\Phi}_{k}^{T}\right)^{-1}\left(\boldsymbol{\Phi}_{k} \mathbf{y}_{k}+\widetilde{\boldsymbol{\Phi}}_{\mathbf{y}}^{k}\right) \stackrel{(a)}{=}\left[\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1}-\left(\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}+\mathbf{I}-\mathbf{I}\right)\left(\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1} \boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1}\right]\left[\boldsymbol{\Phi}_{k} \mathbf{y}_{k}+\right.$ $\left.\widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\mathbf{y}}_{k}\right]=\left[\left(\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1} \boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1}\right]\left[\widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\mathbf{y}}_{k}+\boldsymbol{\Phi}_{k} \mathbf{y}_{k}\right] \stackrel{(b)}{=}\left(\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1} \mathbf{f}_{k}^{\sim}+(\mathbf{I}+$ $\left.\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1}\left(\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}+\mathbf{I}-\mathbf{I}\right) \mathbf{y}_{k}=\mathbf{y}_{k}+\left(\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1}\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)$, where equation (a) is due to the Sherman-Morrison-Woodbury formula and equation (b) results because $\mathbf{f}_{k}^{\sim}=\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\mathbf{y}}_{k}$. Hence, $\mathbf{f}_{k}-\mathbf{y}_{k}=\left(\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right)^{-1}\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)$, which gives

$$
\begin{equation*}
\sum_{i=1}^{J_{k}}\left(y_{i k}-f_{k}\left(\mathbf{x}_{i k}\right)\right)^{2}=\left(\mathbf{f}_{k}-\mathbf{y}_{k}\right)^{T}\left(\mathbf{f}_{k}-\mathbf{y}_{k}\right)=\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)^{T} \boldsymbol{\Gamma}_{k}^{-1}\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right) \tag{A-4}
\end{equation*}
$$

where $\boldsymbol{\Gamma}_{k}=\left[\mathbf{I}+\boldsymbol{\Phi}_{k}^{T} \widetilde{\mathbf{A}}_{k}^{-1} \boldsymbol{\Phi}_{k}\right]^{2}=\left[\mathbf{I}+\boldsymbol{\Phi}_{k}^{T}\left(\rho \mathbf{I}+\widetilde{\boldsymbol{\Phi}}_{k} \widetilde{\boldsymbol{\Phi}}_{k}^{T}\right)^{-1} \boldsymbol{\Phi}_{k}\right]^{2}$.
By construction, $\boldsymbol{\Gamma}_{k}$ has all its eigenvalues no less than 1, i.e., $\boldsymbol{\Gamma}_{k}=\mathbf{E}_{k}^{T} \operatorname{diag}\left[\lambda_{1 k}, \cdots, \lambda_{J_{k} k}\right] \mathbf{E}_{k}$ with $\mathbf{E}_{k}^{T} \mathbf{E}_{k}=\mathbf{I}$ and $\lambda_{1 k}, \cdots, \lambda_{J_{k} k} \geq 1$, which makes the first, second, and last inequality in (7) hold. Using this expansion of $\boldsymbol{\Gamma}_{k}$ in (A-4) we get

$$
\begin{gather*}
\sum_{i=1}^{J_{k}}\left(f_{k}\left(\mathbf{x}_{i k}\right)-y_{i k}\right)^{2}=\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)^{T} \mathbf{E}_{k}^{T} \operatorname{diag}\left[\sigma_{1 k}^{-1}, \ldots, \sigma_{J_{k} k}^{-1}\right]\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right) \\
\leq\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)^{T} \mathbf{E}_{k}^{T}\left[\lambda_{\min , k}^{-1} \mathbf{I}\right] \mathbf{E}_{k}\left(\mathbf{f}_{k}^{\sim}-\mathbf{y}_{k}\right)=\lambda_{\min , k}^{-1} \sum_{i=1}^{J_{k}}\left(f_{k}^{\sim}\left(\mathbf{x}_{i k}\right)-y_{i k}\right)^{2} \tag{A-5}
\end{gather*}
$$

where the inequality results because $\lambda_{\min , k}=\min \left(\lambda_{1, k}, \cdots, \lambda_{J_{k}, k}\right)$. From (A-5) follows the fourth inequality in (7). The third inequality in (7) can be proven in in a similar way.

