



Comparing error minimized extreme learning machines and support vector sequential feed-forward neural networks[☆]

Enrique Romero^{a,*}, René Alquézar^b

^a *Departament de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya, Barcelona, Spain*

^b *Institut de Robòtica i Informàtica Industrial, CSIC Universitat Politècnica de Catalunya, Barcelona, Spain*

ARTICLE INFO

Article history:

Received 23 June 2010

Received in revised form 16 May 2011

Accepted 16 August 2011

Keywords:

Error minimized extreme learning machines

Support vector sequential feed-forward neural networks

Sequential approximations

ABSTRACT

Recently, error minimized extreme learning machines (EM-ELMs) have been proposed as a simple and efficient approach to build single-hidden-layer feed-forward networks (SLFNs) sequentially. They add random hidden nodes one by one (or group by group) and update the output weights incrementally to minimize the sum-of-squares error in the training set. Other very similar methods that also construct SLFNs sequentially had been reported earlier with the main difference that their hidden-layer weights are a subset of the data instead of being random. These approaches are referred to as support vector sequential feed-forward neural networks (SV-SFNNs), and they are a particular case of the sequential approximation with optimal coefficients and interacting frequencies (SAOCIF) method. In this paper, it is firstly shown that EM-ELMs can also be cast as a particular case of SAOCIF. In particular, EM-ELMs can easily be extended to test some number of random candidates at each step and select the best of them, as SAOCIF does. Moreover, it is demonstrated that the cost of the computation of the optimal output-layer weights in the originally proposed EM-ELMs can be improved if it is replaced by the one included in SAOCIF. Secondly, we present the results of an experimental study on 10 benchmark classification and 10 benchmark regression data sets, comparing EM-ELMs and SV-SFNNs, that was carried out under the same conditions for the two models. Although both models have the same (efficient) computational cost, a statistically significant improvement in generalization performance of SV-SFNNs vs. EM-ELMs was found in 12 out of the 20 benchmark problems.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Feed-forward neural networks (FNNs) are a popular machine learning approach for classification and regression problems with very interesting properties (see, for example, Bishop, 1995). As a specific type of FNNs, the single-hidden-layer feed-forward networks (SLFNs) play an important role in practical applications. Since the optimal number of hidden nodes is problem dependent and unknown in advance, users often choose the number of hidden nodes by trial-and-error. Once the architecture is fixed, an iterative learning algorithm such as back-propagation gradient descent is normally applied to adjust the weights in the output and hidden layers simultaneously.

There exist, however, FNN models that construct the network sequentially, so that the number of hidden units is a result

of the learning process rather than being fixed a priori. For a review of constructive FNNs see, for example, (Kwok & Yeung, 1997). Recently, error minimized extreme learning machines (EM-ELMs) have been proposed as a simple and efficient approach to build SLFNs sequentially (Feng, Huang, Lin, & Gay, 2009). EM-ELMs are an incremental extension of the previously presented extreme learning machines (ELMs) (Huang, Zhu, & Siew, 2006). Both methods use random hidden nodes and find the output weights to minimize the sum-of-squares error in the training set by solving a linear system of equations. The specific features of EM-ELMs with respect to ELMs are that they add random hidden nodes one by one (or group by group) and update the output weights incrementally in an efficient way by taking advantage of the incremental construction of the hidden-layer output matrix involved in the linear system. Other recent extensions of ELMs can be found in Liang, Huang, Saratchandran, and Sundararajan (2006) and Miche et al. (2010).

Very similar methods that also construct SLFNs sequentially had been reported earlier (Chen, Cowan, & Grant, 1991; Romero & Alquézar, 2006; Vincent & Bengio, 2002). They all find the optimal linear weights of the output layer by solving the same linear system. In fact, the idea of adding random hidden units

[☆] This work was supported in part by the Ministerio de Ciencia e Innovación (MICINN), under project TIN2009-13895-C02-01.

* Corresponding author.

E-mail addresses: eromero@lsi.upc.edu (E. Romero), ralquezar@iri.upc.edu (R. Alquézar).

was already stated in the sequential approximation with optimal coefficients and interacting frequencies (SAOCIF) algorithm (Romero & Alquézar, 2002, 2006) as a possible strategy to be used and, as shown in Romero and Alquézar (2006), the solution of the linear system can be computed efficiently thanks to the incremental construction of the hidden-layer output matrix. The EM-ELMs and SAOCIF with *random* selection strategy can easily be shown to be essentially equivalent (see Section 2), although it will be demonstrated in Section 3.1 that the cost of the output-layer weight computation described in Feng et al. (2009) is greater than the cost of the corresponding one in SAOCIF described in Romero and Alquézar (2006).

Another strategy proposed in Romero and Alquézar (2002, 2006) to be used within SAOCIF was to take hidden-layer weights always as a subset of the data (*input* strategy). In this case, the resulting method is equivalent to the orthogonal least squares learning (OLSL) algorithm (Chen et al., 1991) and to kernel matching pursuit with *pre-fitting* (KMP-*prefit*) (Vincent & Bengio, 2002). All of them select the hidden-layer weights among the input vectors, a property shared with support vector machines (SVMs) (Vapnik, 1995). By this reason, these approaches are referred to as support vector sequential feed-forward neural networks (SV-SFNNs) (Romero & Toppo, 2007).

SV-SFNNs and SVMs were compared experimentally in Romero and Toppo (2007). Very similar accuracies were found, although computational times were lower for SVMs. Regarding the number of support vectors, SV-SFNNs constructed models with less hidden units than standard SVMs and in the same range as “sparse” SVMs (Keerthi, Chapelle, & DeCoste, 2006). On the other hand, EM-ELMs were compared in Feng et al. (2009) with other sequential algorithms, namely resource allocation networks (RAN) (Platt, 1991) and minimum resource allocation networks (MRAN) (Yingwei, Sundararajan, & Saratchandran, 1997), as well as with the original ELMs (Huang, Zhu et al., 2006). EM-ELMs obtained better performance and less training time than RAN and MRAN, and a similar performance but less training time than ELMs.

This work shows a comparison between EM-ELMs (i.e. SAOCIF with *random* strategy) and SV-SFNNs (i.e. SAOCIF with *input* strategy). An experimental study on 10 benchmark data sets for classification problems and 10 benchmark data sets for regression problems is presented in which the two methods are compared in the same conditions and using the same software. Since both approaches can be adjusted to have the same computational cost (each candidate weight vector for a hidden unit is either generated randomly or selected randomly among the input vectors), the goal is finding out whether there is any difference in generalization performance between EM-ELMs and SV-SFNNs. In other words, does the use of inputs (support vectors) as hidden unit weights provide any advantage over pure random selection?

Intuitively, the *input* strategy should perform better than the *random* one, since the *input* strategy samples the weights from the underlying distribution of the data. This information is lost in the *random* strategy, that ignores the statistical properties of the data. By contrast, the *random* strategy could take profit of its high flexibility to explore areas of the space of weights that the *input* strategy cannot. The growing interest in models that randomly select the weights of the hidden units (see Feng et al., 2009; Liang et al., 2006; Liu & Wang, 2010; Miche et al., 2010, and references therein) makes the comparison of the *input* and *random* strategies relevant. Moreover, it should be noted that, on one hand, EM-ELMs have been proposed to build SLFNs as a superior alternative (due to their simplicity, efficiency and effectiveness) to common neural net approaches like back-propagation gradient descent (Feng et al., 2009) and, on the other hand, input vectors are rarely used as hidden-layer weights in SLFNs, especially in the case of additive units (i.e. two-layer perceptrons).

2. Background

The output function of an SLFN (i.e. a fully connected FNN with a single hidden layer of N_h units and m linear output units) can be expressed as a linear combination of simple (basis) functions:

$$f_{N_h}(\mathbf{x}) = \lambda_0 + \sum_{i=1}^{N_h} \lambda_i \varphi(\omega_i, b_i, \mathbf{x}) \quad (1)$$

where $\omega_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ are the learning parameters of the hidden units, $\lambda_i \in \mathbb{R}^m$ are the output-layer weights connecting the i -th hidden unit to the m output units, φ is the activation function of the hidden units, $\varphi(\omega_i, b_i, \mathbf{x})$ is the output of the i -th hidden unit with respect to the input \mathbf{x} , and $\lambda_0 \in \mathbb{R}^m$ denotes the bias terms (if any) of the linear output units.

Although a lot of activation functions φ (even not neuron alike) can be used that allow universal approximation capability, the more usual choices are the Gaussian RBF applied to a distance between an input vector and a center

$$\varphi(\omega_i, b_i, \mathbf{x}) = \exp(-b_i \|\mathbf{x} - \omega_i\|^2 / 2) \quad (2)$$

and the sigmoid (e.g. hyperbolic tangent) applied to a scalar product of the input and weight vectors (this will be referred to as a sigmoid additive unit (Feng et al., 2009))

$$\varphi(\omega_i, b_i, \mathbf{x}) = \tanh(\omega_i \cdot \mathbf{x} + b_i). \quad (3)$$

For Gaussian RBF units, $\omega_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}^+$ are the center and the impact factor of the i -th RBF unit. For sigmoid additive units, $\omega_i \in \mathbb{R}^n$ is the weight vector connecting the input layer to the i -th hidden unit and $b_i \in \mathbb{R}$ is the bias of the i -th hidden unit. In our experiments presented in Section 3, a third activation function has also been tested, the sine applied to the scalar product (i.e. a sine additive unit)

$$\varphi(\omega_i, b_i, \mathbf{x}) = \sin(\omega_i \cdot \mathbf{x} + b_i). \quad (4)$$

For a given set of training examples $\{(\mathbf{x}_j, \mathbf{t}_j)\}_{j=1}^L \subset \mathbb{R}^n \times \mathbb{R}^m$, if the outputs of the network are equal to the targets, we have

$$f_{N_h}(\mathbf{x}_j) = \sum_{i=1}^{N_h} \lambda_i \varphi(\omega_i, b_i, \mathbf{x}_j) = \mathbf{t}_j, \quad j = 1, \dots, L. \quad (5)$$

Eq. (5) can be written compactly as

$$\mathbf{H}\boldsymbol{\lambda} = \mathbf{T} \quad (6)$$

where \mathbf{H} is an $L \times N_h$ matrix called the hidden-layer output matrix of the network ($\mathbf{H}_{ji} = \varphi(\omega_i, b_i, \mathbf{x}_j)$), $\boldsymbol{\lambda}$ is an $N_h \times m$ matrix containing the output-layer weights, and \mathbf{T} is an $L \times m$ matrix containing the target values in the training set. The output layer biases can be added by including in \mathbf{H} a first column with a fixed value of 1 (and increasing N_h by 1).

Normally, the number of training examples L will be much greater than the number of hidden units N_h and an exact solution of (6) cannot be expected. Then, the usual cost function in SLFNs (and in general in FNNs) is the sum-of-squares error

$$E = \frac{1}{2} \sum_{j=1}^L \|f_{N_h}(\mathbf{x}_j) - \mathbf{t}_j\|^2. \quad (7)$$

It is well known (e.g. Bishop, 1995) that, to minimize E , the optimal output-layer weights can be computed as

$$\hat{\boldsymbol{\lambda}} = \mathbf{H}^\dagger \mathbf{T} \quad \text{where } \mathbf{H}^\dagger = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \quad (8)$$

is the pseudo-inverse (or Moore–Penrose generalized inverse) of the hidden-layer output matrix \mathbf{H} . The sum-of-squares error can be expressed as

$$E(\mathbf{H}) = \frac{1}{2} \|\mathbf{H}\boldsymbol{\lambda} - \mathbf{T}\|^2 = \frac{1}{2} \|\mathbf{H}\mathbf{H}^\dagger \mathbf{T} - \mathbf{T}\|^2. \quad (9)$$

Given a set of training examples $\{(\mathbf{x}_j, \mathbf{t}_j)\}_{j=1}^L \subset \mathbb{R}^n \times \mathbb{R}^m$, the maximum number of hidden units N_{\max} , and the expected learning accuracy ϵ :

1) *Initialization phase*:

- Initialize the SLFN with N_0 random hidden units (ω_i, b_i) , $i = 1, \dots, N_0$
- Calculate the hidden-layer output matrix \mathbf{H}_0
- Calculate the corresponding output error

2) *Recursively growing phase*:

- Let $k := 0$
- **While** $N_k < N_{\max}$ and $E(\mathbf{H}_k) > \epsilon$ **do**
 - Let $k := k + 1$
 - Randomly add δN_k hidden units to the existing SLFN; the total number of hidden units becomes $N_k = N_{k-1} + \delta N_k$ and the corresponding hidden-layer output matrix $\mathbf{H}_k = [\mathbf{H}_{k-1}, \delta \mathbf{H}_k]$, where $\delta \mathbf{H}_k$ contains the new δN_k columns
 - The output-layer weights λ are updated as

$$\lambda_k = \mathbf{H}_k^\dagger \mathbf{T} = \begin{bmatrix} \mathbf{U}_k \\ \mathbf{D}_k \end{bmatrix} \mathbf{T}, \text{ where}$$

$$\mathbf{D}_k = \left((\mathbf{I} - \mathbf{H}_{k-1} \mathbf{H}_{k-1}^\dagger) \delta \mathbf{H}_k \right)^\dagger \text{ and } \mathbf{U}_k = \mathbf{H}_{k-1}^\dagger (\mathbf{I} - \delta \mathbf{H}_k \mathbf{D}_k)$$

- **End while**

Fig. 1. Algorithm for EM-ELMs.

2.1. Error minimized extreme learning machines (EM-ELMs)

Huang, Chen, and Siew (2006) have shown that SLFNs with random weights in the hidden layer have universal approximation capability for many different choices of the activation function, including the ones stated in Eqs. (2)–(4). Based on this result, they propose the ELMs learning algorithm (Huang, Zhu et al., 2006), which can be summarized as follows:

Algorithm for ELMs: Given a set of training examples $\{(\mathbf{x}_j, \mathbf{t}_j)\}_{j=1}^L \subset \mathbb{R}^n \times \mathbb{R}^m$, the hidden-layer activation function $\varphi(\omega, b, \mathbf{x})$, and an *a-priori* fixed number N_h of hidden units:

- (1) Randomly assign hidden-unit parameters (ω_i, b_i) , $i = 1, \dots, N_h$.
- (2) Calculate the hidden-layer output matrix \mathbf{H} .
- (3) Calculate the output-layer weight matrix λ using (8).

In order to avoid the need of setting in advance the number N_h of hidden units and to reduce the training computational time, a fast sequential extension of the ELMs algorithm called EM-ELMs has been recently reported by Feng et al. (2009). The EM-ELM is described in Fig. 1.

In the reported experiments, both ELMs and EM-ELMs sample the random values for the weights from the uniform distribution.

2.2. Sequential approximation with optimal coefficients and interacting frequencies (SAOCIF)

A rather general constructive method for SLFNs, called SAOCIF, was proposed in Romero and Alquézar (2002, 2006). The specific features of SAOCIF are: (i) the optimal (in a least squares sense) output-layer weights are recalculated each time a hidden unit is added by solving a linear equation system, and (ii) the added hidden unit is selected among a set of candidates taking into account its interaction with the previously added hidden units (i.e. to minimize together the training error). The SAOCIF algorithm is described in Fig. 2.

Note that if we set $C_{\max} = 1$ and the candidates are generated with a *random* strategy, then the resulting particularization of SAOCIF is essentially equivalent to EM-ELMs with $\delta N_k = 1$ (for all k). In this case, N in SAOCIF is equivalent to both k and N_k in EM-ELMs.

It might be argued that EM-ELMs add hidden nodes one by one ($\delta N_k = 1$, for all k) or group by group ($\delta N_k > 1$), while SAOCIF does

not mention group based addition. Although this last statement is certainly true, it is obvious that in the original EM-ELMs formulation adding a group of $\delta N_k > 1$ hidden units is completely equivalent to δN_k additions of one hidden unit. This is because EM-ELMs only select one candidate ($C_{\max} = 1$) for every hidden unit. However, if several candidates per hidden unit are allowed ($C_{\max} > 1$, something that EM-ELMs do not foresee) then it is clear that one-by-one and group-by-group additions are not equivalent. Anyway, it is also quite obvious that SAOCIF may work for group addition (see Section 3.1), in the same way as EM-ELMs.

2.3. Support vector sequential feed-forward neural networks (SV-SFNNs)

Apart from the *random* strategy, other possibilities are allowed in the SAOCIF approach to generate candidates. In particular, let us define the *input* strategy as the one in which the candidates are only selected among the input examples in the training set; more precisely, $\omega = \mathbf{x}_j$, for some j not already used, and b is a constant depending on the activation function (e.g. $b = 1$ for RBF units and $b = 0$ for additive units). Then, if we set $C_{\max} = L - N + 1$ and the candidates are generated using the *input* strategy, the resulting method selects the best of the input examples as the hidden-layer weights of the new hidden unit. This method, which has been called SV-SFNNs (Romero & Toppo, 2007), is equivalent to the OLSL algorithm (Chen et al., 1991) and to KMP-*prefit* (Vincent & Bengio, 2002). Actually, OLSL was only proposed for RBF units and KMP-*prefit* for kernel-based activation functions, while SAOCIF with *input* strategy permits as well any other activation function with universal approximation capabilities (e.g. sinusoidal additive units).

3. Comparing EM-ELMs and SV-SFNNs

This section compares EM-ELMs with SV-SFNNs and explains the methodology followed in the experiments.

3.1. Analysis of the computational cost

The computational complexity of the algorithm for EM-ELMs is dominated by the computation of \mathbf{D}_k and \mathbf{U}_k , which involves, as intermediate steps, the computation of $L \times L$ matrices, that is

Given a set of training examples $\{(\mathbf{x}_j, \mathbf{t}_j)\}_{j=1}^L \subset \mathbb{R}^n \times \mathbb{R}^m$, the maximum number of hidden units N_{\max} , a strategy to generate the candidates, the maximum number of candidates for any hidden unit C_{\max} , and the expected learning accuracy ϵ :

```

- Let  $N := 0$ 
- Let  $\mathbf{H}_0 = []$ 
- Repeat
  - Let  $N := N + 1$ 
  - Let  $c := 0 // c$  is the number of valid candidates tested for the  $N^{\text{th}}$  hidden unit
  - While  $c < C_{\max}$  do
    - Generate a candidate  $(\omega, b)$  for the  $N^{\text{th}}$  hidden unit with the given strategy,
      and store in a temporary matrix  $\mathbf{H}$  the corresponding hidden-layer output
      matrix  $\mathbf{H} = [\mathbf{H}_{N-1}, \delta\mathbf{H}]$ , where  $\delta\mathbf{H}$  contains the new column computed
    - If  $\mathbf{H}^T\mathbf{H}$  is well conditioned then
      - Let  $c := c + 1$ 
      - Find the optimal output-layer weights  $\lambda = \mathbf{H}^\dagger\mathbf{T}$  for the current candidate
         $(\omega, b)$  using the incremental method in [Romero & Alquézar 2006]
        (equivalent in essence to the incremental method described in the algorithm
        for EM-ELMs, see Section 3.1)
      - Calculate the corresponding output error  $E(\mathbf{H}) = \frac{1}{2} \|\mathbf{H}\lambda - \mathbf{T}\|^2$ 
      - If  $E(\mathbf{H})$  is minimum in the current loop then
        - Let  $(\omega_N, b_N) = (\omega, b)$ ;  $\lambda_N = \lambda$ ;  $\mathbf{H}_N = \mathbf{H}$ 
      - End if
    - End if
  - End while
- Until  $N = N_{\max}$  or  $E(\mathbf{H}_N) \leq \epsilon$ 

```

Fig. 2. Algorithm for SAOCIF.

expensive in terms of time and memory. Since L will normally be much larger than N_{\max} , and assuming that $N_{k-1} > \delta N_k$, the computational cost to obtain \mathbf{D}_k and \mathbf{U}_k is $O(L^2 \cdot N_{k-1})$. However, this computation can be done with lower computational cost, as explained next.

Defining $\mathbf{A} = \mathbf{H}_{k-1}^T \mathbf{H}_{k-1}$, $\mathbf{v} = \mathbf{H}_{k-1}^T \delta\mathbf{H}_k$, $\mathbf{u} = \delta\mathbf{H}_k^T \delta\mathbf{H}_k$, $\alpha = \mathbf{H}_{k-1}^T \mathbf{T}$ and $\beta = \delta\mathbf{H}_k^T \mathbf{T}$, it is very easy to verify that

$$\mathbf{D}_k \mathbf{T} = (\mathbf{u} - \mathbf{v}^T \mathbf{A}^{-1} \mathbf{v})^{-1} (\beta - \mathbf{v}^T \mathbf{A}^{-1} \alpha) = \eta \quad (10)$$

$$\mathbf{U}_k \mathbf{T} = \mathbf{A}^{-1} (\alpha - \mathbf{v} \eta) \quad (11)$$

that is a generalization (for m outputs and adding δN_k hidden units in the same step) of the incremental method described in Romero and Alquézar (2006) for the addition of one hidden unit. To check this, simply replace in (10) \mathbf{u} by γ , \mathbf{v} by \mathbf{a} and α by \mathbf{b} , and you will obtain Eq. (6) in Romero and Alquézar (2006). The computation of (10) and (11) is faster than those of the original algorithm for EM-ELMs because its computational cost is $O(L \cdot N_{k-1}^2)$.

Therefore, although the computational cost of the original algorithms EM-ELMs and SV-SFNNs is not the same, they can be easily made equivalent.

For $C_{\max} > 1$, \mathbf{A} can be computed for the first candidate, kept in memory and recovered for the rest of candidates. Therefore, the computational cost for $C_{\max} > 1$ is lower than C_{\max} times the computational cost for $C_{\max} = 1$. More precisely, with this optimization the computational cost of the first iteration in the inner loop of SAOCIF is $O(L \cdot N_{k-1}^2)$, and $O(L \cdot N_{k-1} \cdot \max(\delta N_k, m))$ for the rest of iterations.

3.2. Compared methods and settings

At first, the original algorithms for EM-ELMs and SV-SFNNs can only be directly compared if we set $C_{\max} = 1$, just because in EM-ELMs the inner loop of SV-SFNNs (While $c < C_{\max}$) is not carried out. However, that choice would limit the flexibility of SV-SFNNs. Since the aim of this paper is to make a general comparison

between the *random* and *input* strategies, two settings have been defined.

In the former, $C_{\max} = 1$, so the original EM-ELMs are confronted with a very limited version of SV-SFNNs in which only one randomly selected input (not the best) yields the single candidate. In the latter, $C_{\max} = 59$, so an extended version of EM-ELMs (with the upgrade of selecting the best random candidate among C_{\max} at each step) is confronted with a not so limited version of SV-SFNNs in which not the best of the remaining candidates but the best of a randomly selected subset (of size C_{\max}) is added. The choice of $C_{\max} = 59$ is justified because, in order to obtain a candidate that is with probability 0.95 among the best 5% of all candidates, a random subset of size $\lceil \log 0.05 / \log 0.95 \rceil = 59$ suffices (Smola & Schölkopf, 2000).

For every setting, the only difference resides on whether the candidates are randomly generated or taken randomly from the input patterns. These settings allow to make a fair comparison of EM-ELMs and SV-SFNNs, since they work in the same conditions and take the same computation time.

3.3. Software

We have used our own implementation in C setting the algorithm parameters as explained in the preceding paragraph. The optimal output-layer weights were computed using (10) and (11).

3.4. Data sets

The comparison was performed using 20 benchmark data sets, 10 for classification and 10 for regression problems. The classification data sets were *Australian Credit*, *Splice-junction Gene Sequences*, *German Credit*, *Ionosphere*, *Iris*, *Landsat Satellite (Satimage)*, *Image Segmentation*, *Sonar*, *Vehicle Silhouettes* and *Wine*, and can be found in the UCI repository (Asuncion & Newman, 2007). The features of these data sets are summarized in Table 1. The regression data sets were *Abalone*, *Auto Price*, *Boston Housing*, *California Housing*, *Census House*, *Delta Ailerons*,

Table 1
Features of the classification benchmark data sets.

Data set	#Inputs	#Exa.	#Classes
<i>Australian</i>	43	690	2
<i>Gene</i>	120	3175	3
<i>German</i>	56	1000	2
<i>Ionosphere</i>	34	351	2
<i>Iris</i>	4	150	3
<i>Satimage</i>	36	6435	6
<i>Segmentation</i>	16	2310	7
<i>Sonar</i>	60	208	2
<i>Vehicle</i>	18	846	4
<i>Wine</i>	13	178	3

Table 2
Features of the regression benchmark data sets.

Data set	#Inputs	#Exa.	Target mean	StdDev
<i>Abalone</i>	8	4 177	9.93	3.22
<i>Auto Price</i>	15	159	11 445.73	5877.85
<i>Boston Housing</i>	13	506	22.53	9.20
<i>Calif. Housing</i>	8	20640	206 854.97	115 395.58
<i>Census House</i>	8	22 784	50 073.10	52 846.16
<i>Delta Ailerons</i>	5	7 129	−0.000007	0.0003
<i>Delta Elevators</i>	6	9517	−0.000133	0.0023
<i>Machine CPU</i>	6	209	105.62	160.83
<i>Servo</i>	4	167	1.39	1.56
<i>Stock</i>	9	950	46.99	6.54

Delta Elevators, *Machine CPU*, *Servo* and *Stock*, that can be found at <http://www.liaad.up.pt/~ltorgo/Regression/DataSets.html>. The features of these data sets are summarized in Table 2.

3.5. Methodology

- **Preprocessing.** Categorical attributes were converted to dummy variables (the p different categories were represented with p input variables, so that only the input variable associated to its category is one, and all others are zero). The rest of the attributes (including the target variable for regression data sets) were scaled to mean zero and variance one.
- **Random weights.** In the *random* strategy, hidden-layer weights were uniformly chosen within the same range of values that the input values (after scaling). In this way, the ranges of the hidden-layer weights were the same for both strategies.
- **Activation functions.** Three types were used: Gaussian RBF (2), sigmoid additive (3) and sine additive (4) units, but with a further multiplicative positive parameter γ introduced for a wider search. Specifically, γ multiplies the distance $\|\mathbf{x} - \omega_i\|$ in the RBF units and the scalar product $\omega_i \cdot \mathbf{x}$ in the additive units.
- **Parameters and model selection.** A hidden-unit candidate weight vector was not considered valid if the associated linear equations system could not be solved or if the 1-norm of the solution (the output-layer weights) was greater than a certain value M . This can be seen as a form of regularization. M was set to 1024. We fixed $N_{\max} = 99$ and $\epsilon = 0$, so that N_{\max} hidden units were always added. These values were selected according to our previous experience with these models (recall that the aim of the experiments is to compare EM-ELMs and SV-SFNNs under the same conditions for the two models). In order to get an adequate value for the γ parameter, much more problem-dependent, a search was performed ranging γ from 2^{-10} to 2^5 . The same search was performed for all the models, and repeated for every activation function.
- **Model training and testing.** The methods were trained and tested over 30 training-validation-test different random partitions (80% training, 10% validation, 10% test) of the whole data set. For every configuration (defined by a given strategy, C_{\max} , activation function and γ), the networks with the lowest errors

Table 3
Comparison of average test accuracy for classification data sets—one candidate.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Australian</i>	83.86	83.00	83.67	83.14	84.25	83.48
<i>Gene</i>	84.71	83.38	84.75	83.13	84.82	82.95
<i>German</i>	77.37	78.23	77.73	77.20	77.83	77.10
<i>Ionosphere</i>	93.87	90.19	90.10	88.67	90.19	88.67
<i>Iris</i>	100	99.11	100	99.78	100	99.56
<i>Satimage</i>	82.98	80.57	79.17	77.73	79.40	77.62
<i>Segmentation</i>	86.62	81.65	86.42	86.08	86.70	86.39
<i>Sonar</i>	89.84	80.83	77.17	76.17	77.50	75.83
<i>Vehicle</i>	85.48	85.52	86.11	84.80	85.87	85.20
<i>Wine</i>	99.61	100	99.80	100	100	100

Table 4
Comparison of average test accuracy for classification data sets—best of 59 candidates.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Australian</i>	84.49	83.48	83.77	83.43	85.02	83.82
<i>Gene</i>	86.60	86.05	86.34	85.86	86.36	85.92
<i>German</i>	77.13	78.33	77.60	77.30	78.03	77.37
<i>Ionosphere</i>	93.90	90.67	89.52	88.67	90.00	88.83
<i>Iris</i>	100	100	100	100	100	100
<i>Satimage</i>	86.35	83.31	82.83	77.56	81.73	77.50
<i>Segmentation</i>	88.56	83.39	86.61	86.64	87.30	86.83
<i>Sonar</i>	96.50	81.83	87.67	76.67	75.00	74.17
<i>Vehicle</i>	86.67	85.56	86.87	86.11	86.75	86.63
<i>Wine</i>	100	100	100	100	100	100

in the validation subsets were selected as the final models. For classification data sets, the accuracies of the final models were given by the average accuracies measured in the test subsets. For regression data sets, the performance of the final models was measured by the following normalized squared error (NSE) (Bishop, 1995):

$$\text{NSE} = \frac{\sum_{j=1}^K \|f(\mathbf{x}_j) - \mathbf{t}_j\|^2}{\sum_{j=1}^K \|\bar{\mathbf{t}} - \mathbf{t}_j\|^2},$$

where K is the number of examples in the test subset, $f(\mathbf{x})$ is the final model and $\bar{\mathbf{t}}$ is the mean target value in the test subset. The sizes of the final models are defined by their average number of hidden units.

3.6. Experimental results

Tables 3 and 4 show the average accuracies of the best final models (among all γ) for the two strategies (*input* and *random*) and the three activation functions tried (Gaussian RBF, sine additive, sigmoid additive) using the methodology previously described for the 10 classification data sets studied. Table 3 displays the results of the methods for $C_{\max} = 1$, where the *input* strategy is fully comparable to EM-ELMs, and Table 4 displays the results for $C_{\max} = 59$. For the 10 regression data sets studied, Tables 5 and 6 show the average NSE of the best final models for $C_{\max} = 1$ and $C_{\max} = 59$ respectively. Tables 7–10 show the average number of hidden units in the final models selected for each combination of strategy and activation function for $C_{\max} = 1$ and $C_{\max} = 59$.

It can be observed that *Iris* and *Wine* data sets correspond to easy problems that have been learned perfectly using both strategies. For the other data sets, test results look similar between the two strategies in some cases and a superior performance of the *input* strategy can be appreciated in the rest, except for the *Stock* data set. Not surprisingly, the best values for each strategy and data

Table 5
Comparison of average test NSE for regression data sets—one candidate.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Abalone</i>	0.517	0.526	0.521	0.553	0.514	0.550
<i>Auto Price</i>	0.341	1.570	0.270	0.870	0.280	0.917
<i>Boston Housing</i>	0.725	0.726	0.941	1.166	0.871	1.162
<i>Calif. Housing</i>	0.344	0.369	0.367	0.367	0.367	0.366
<i>Census House</i>	0.404	1.070	0.484	0.580	0.457	0.502
<i>Delta Ailerons</i>	0.292	0.297	0.305	0.301	0.304	0.302
<i>Delta Elevators</i>	0.376	0.377	0.375	0.377	0.376	0.376
<i>Machine CPU</i>	0.364	0.404	0.358	0.394	0.330	0.378
<i>Servo</i>	0.214	0.248	0.228	0.329	0.246	0.337
<i>Stock</i>	1.689	1.689	2.023	9.883	8.897	7.796

Table 6
Comparison of average test NSE for regression data sets—best of 59 candidates.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Abalone</i>	0.512	0.524	0.512	0.537	0.511	0.537
<i>Auto Price</i>	0.177	0.840	0.275	0.456	0.299	0.517
<i>Boston Housing</i>	0.657	0.620	0.762	0.898	0.745	0.921
<i>Calif. Housing</i>	0.332	0.345	0.368	0.366	0.369	0.365
<i>Census House</i>	0.383	0.508	0.425	0.528	0.403	0.470
<i>Delta Ailerons</i>	0.293	0.295	0.303	0.303	0.299	0.301
<i>Delta Elevators</i>	0.375	0.376	0.374	0.375	0.375	0.375
<i>Machine CPU</i>	0.299	0.348	0.370	0.385	0.345	0.380
<i>Servo</i>	0.158	0.206	0.185	0.289	0.171	0.300
<i>Stock</i>	1.714	1.608	2.896	2.147	3.318	1.807

Table 7
Average number of hidden units for classification data sets—one candidate.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Australian</i>	43.53	30.33	25.17	37.40	61.33	44.90
<i>Gene</i>	97.07	98.87	97.40	98.60	97.20	98.00
<i>German</i>	40.47	43.23	44.53	33.87	36.77	36.80
<i>Ionosphere</i>	73.10	40.63	19.40	24.30	14.97	24.60
<i>Iris</i>	55.97	31.63	3.87	3.60	3.87	3.67
<i>Satimage</i>	95.30	96.80	97.93	97.13	95.47	97.73
<i>Segmentation</i>	41.53	25.90	63.10	74.30	53.40	82.07
<i>Sonar</i>	66.53	64.00	47.67	43.97	38.13	42.10
<i>Vehicle</i>	74.77	38.00	67.30	61.73	64.03	63.97
<i>Wine</i>	11.57	11.77	7.73	10.00	7.90	10.90

Table 8
Average number of hidden units for classification data sets—best of 59 candidates.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Australian</i>	50.63	23.33	15.53	32.67	24.93	21.77
<i>Gene</i>	95.23	89.87	90.33	87.07	88.33	84.80
<i>German</i>	12.20	12.20	12.97	14.23	13.83	16.90
<i>Ionosphere</i>	65.73	29.03	10.40	12.10	12.53	12.67
<i>Iris</i>	19.13	4.50	3.87	2.97	3.97	2.77
<i>Satimage</i>	97.67	89.73	96.23	96.83	95.40	97.43
<i>Segmentation</i>	95.83	32.00	75.33	63.57	91.87	68.27
<i>Sonar</i>	88.27	55.43	31.17	10.57	25.13	21.13
<i>Vehicle</i>	81.73	37.13	44.80	50.77	58.37	54.27
<i>Wine</i>	4.90	7.13	5.00	6.13	5.13	6.23

set are included in Tables 4 and 6 (i.e. they have been obtained using $C_{\max} = 59$) except, again, for the *Stock* data set. Note that the low number of hidden units of the best final models for the *Stock* data set (see Tables 9 and 10) is a clear indication of a strong tendency to overfitting in this data set.

Although we do not claim that our results are optimal for the tested data sets, they are competitive with the results of other state-of-the-art methods (see, for example Van Gestel et al., 2004 for classification data sets and Yu, Zhou, & Ting, 2007 for regression data sets).

Table 9
Average number of hidden units for regression data sets—one candidate.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Abalone</i>	55.60	22.27	61.13	10.17	78.23	78.70
<i>Auto Price</i>	8.40	5.07	7.43	10.70	7.53	10.80
<i>Boston Housing</i>	18.57	23.23	4.03	5.00	3.80	4.50
<i>Calif. Housing</i>	36.40	34.13	8.10	30.17	8.67	31.43
<i>Census House</i>	87.67	6.00	94.70	82.80	93.03	92.20
<i>Delta Ailerons</i>	81.40	51.63	32.03	39.70	33.67	36.33
<i>Delta Elevators</i>	66.10	80.33	52.10	50.37	58.00	64.13
<i>Machine CPU</i>	4.33	5.73	7.63	4.97	7.50	5.17
<i>Servo</i>	60.43	42.27	65.67	14.40	65.50	14.57
<i>Stock</i>	1.07	1.07	1.80	4.53	3.20	2.97

Table 10
Average number of hidden units for regression data sets—best of 59 candidates.

Data set	Gaussian RBF		Sine MLP		Sigmoid MLP	
	Input	Rand.	Input	Rand.	Input	Rand.
<i>Abalone</i>	34.60	48.17	54.67	40.13	31.70	56.23
<i>Auto Price</i>	7.37	4.20	3.47	5.43	3.00	4.87
<i>Boston Housing</i>	11.93	14.90	1.53	2.33	1.30	8.80
<i>Calif. Housing</i>	88.87	54.00	7.57	19.07	7.10	14.47
<i>Census House</i>	95.13	37.00	96.43	36.97	91.23	78.27
<i>Delta Ailerons</i>	54.73	34.43	21.83	30.07	72.20	72.20
<i>Delta Elevators</i>	40.17	57.90	31.50	28.97	36.20	29.20
<i>Machine CPU</i>	5.37	3.97	5.97	3.07	5.27	3.53
<i>Servo</i>	61.57	21.17	54.00	32.47	57.37	30.63
<i>Stock</i>	4.93	1.00	1.10	2.07	1.13	1.33

In order to obtain an objective statistical measure, a Student's t -test was applied to each data set to check if the difference between the best mean results of the two strategies was statistically significant (p -value = 0.05, i.e. confidence of 95%). The test was applied to the best models obtained with $C_{\max} = 59$, which are marked in bold (all of them are the best results for each strategy and data set, except for the *Stock* data set). Results of the test are shown in Table 11. In six of the classification data sets (*Australian Credit*, *Gene*, *Ionosphere*, *Satimage*, *Segmentation*, *Sonar*) and six of the regression data sets (*Auto Price*, *California Housing*, *Census House*, *Delta Elevators*, *Machine CPU*, *Servo*) the t -test gave a significant difference with a superior mean accuracy of the *input* strategy, whereas no significant difference was found in the other ones (*German Credit* and *Vehicle* for classification and *Abalone*, *Boston Housing*, *Delta Ailerons* and *Stock* for regression). Small differences (see *Gene*, *California Housing*, *Delta Elevators*, for example) are sometimes more significant than larger ones (*Stock*, for example) because the former have very small variances.

Although no clear trend is observed about the number of hidden units selected by both strategies (it depends quite a lot on the specific activation function), the *input* strategy seems to need more units than the *random* strategy in the case of Gaussian RBF hidden units (this can be seen easily in Table 8). Regarding the number of candidates, final models obtained with $C_{\max} = 1$ normally have more hidden units than those obtained with $C_{\max} = 59$.

As a reference, the mean execution times to obtain every final model were 58.8 s for *Satimage* and 99 s for *Census House* (carried out in a node of a computation cluster with Intel® Xeon® CPUs at 2.66 GHz).

3.7. Discussion

There are several reasons that can explain why the *input* strategy performs better than the *random* one, as explained next. On the one hand, the structural risk minimization principle defines a trade-off between the quality of the approximation to the data and the complexity of the model (Vapnik, 1995). According to this principle, given two models with the same empirical

Table 11
Student's *t*-test for all data sets—best of 59 candidates and best activation function.

Data set	Input		Random		t-test
	Mean	Std	Mean	Std	<i>p</i> -value
<i>Australian</i>	85.02	1.6304	83.82	1.3237	0.002616
<i>Gene</i>	86.60	0.9905	86.05	0.8409	0.022246
<i>German</i>	78.03	1.9205	78.33	2.0398	0.559828
<i>Ionosphere</i>	93.90	1.9466	90.67	2.7003	0.000002
<i>Iris</i>	100	0.0	100	0.0	–
<i>Satimage</i>	86.35	0.6592	83.31	1.2996	0.000000
<i>Segmentation</i>	88.56	0.9289	86.83	0.8342	0.000000
<i>Sonar</i>	96.50	8.6253	81.83	9.6921	0.000000
<i>Vehicle</i>	86.87	2.3103	86.63	3.2014	0.744651
<i>Wine</i>	100	0.0	100	0.0	–
<i>Abalone</i>	0.511	0.0118	0.524	0.0348	0.053008
<i>Auto Price</i>	0.177	0.0619	0.456	0.3090	0.000032
<i>Boston Housing</i>	0.657	0.1101	0.620	0.0954	0.172238
<i>Calif. Housing</i>	0.332	0.0092	0.345	0.0182	0.000894
<i>Census House</i>	0.383	0.0045	0.470	0.0323	0.000000
<i>Delta Ailerons</i>	0.293	0.0057	0.295	0.0031	0.057887
<i>Delta Elevators</i>	0.374	0.0020	0.375	0.0020	0.023594
<i>Machine CPU</i>	0.299	0.0431	0.348	0.0981	0.015547
<i>Servo</i>	0.158	0.0308	0.206	0.0538	0.000116
<i>Stock</i>	1.714	0.0004	1.608	0.5498	0.298702

risk, the model with lower complexity (approximation capability) should be preferred. In our experiments we did not observe differences between the empirical risks of the *input* and *random* strategies, although it is clear that the *random* strategy has greater approximation capability than the *input* one. On the other hand, choosing the hidden-layer weights among the input vectors is more likely to obtain better candidates (in the sense of informative candidates, such as prototypes, border points, etc.) than sampling it randomly. In some sense, choosing the hidden-layer weights among the input vectors is related to sampling the underlying distribution.

4. Conclusions and future work

First, it has been shown that EM-ELMs can be cast as a particular case of the SAOCIF method (with *Random* strategy) for constructing SLFNs. The only two real differences between the original EM-ELMs proposed in Feng et al. (2009) and the SAOCIF method with *Random* strategy proposed in Romero and Alquézar (2006) are: (i) the incremental method in which the same optimal (in a least squares sense) output-layer weights are recalculated each time a random hidden unit is added (or tested), and (ii) the number of random candidates tested for each hidden unit. Regarding the first difference, it has been demonstrated in Section 3.1 that the cost of the computation described in Feng et al. (2009) is greater than the corresponding one in the (generalized) method described in Romero and Alquézar (2006). Since both methods compute the same optimal output-layer weights, there is no problem in using the more efficient one also for EM-ELMs, thus removing the difference. The second difference can also be eliminated either by restricting SAOCIF to test a single candidate at each step (as the original EM-ELMs do) or extending EM-ELMs to test some number of random candidates at each step and select the best of them (as SAOCIF does). Both possibilities have been explored in the experimental study carried out in this work.

Seemingly, a third difference between EM-ELM algorithm and SAOCIF with *Random* strategy is that the former allows group by group addition of hidden units whereas the latter does not. However, this is not a real difference, because when a single random candidate is used per hidden unit, adding a group of *N* hidden units is completely equivalent to *N* additions of one hidden unit, and SAOCIF may also perform group addition in the same way that EM-ELMs. Note, however, that if several candidates per hidden

unit are allowed (something that EM-ELMs do not foresee) then one-by-one and group-by-group additions are not equivalent. We affirm that it is much more significant to allow several candidates per unit in one-by-one addition than adding single-candidate units group by group.

Second, we have claimed that an alternative sequential method to construct SLFNs can be based on selecting the hidden-layer weights among the input vectors in the training set. This method, which has been referred to as SV-SFNNs (Romero & Toppo, 2007) or SAOCIF with *Input* strategy (Romero & Alquézar, 2006), is essentially equivalent to the OLS algorithm and to KMP-*prefit*. In order to assess the relative performance of both approaches (EM-ELMs vs. SV-SFNNs) in a fair manner, an empirical study has been realized on twenty benchmark data sets, 10 for classification and 10 for regression, under the same conditions and using the same software.

The experimental comparison between EM-ELMs and SV-SFNNs presented in the paper draws two interesting conclusions that can be further investigated in future research. The first one is that selecting the hidden-layer weights as a subset of the input data, even if this selection is done randomly, yields better generalization results than selecting the hidden-layer weights in a purely random manner from scratch (like EM-ELMs do). As discussed at the end of Section 1, this is not an obvious result. Indeed, no statistically significant difference between the average performances obtained by the two strategies was found in eight of the benchmark problems, but SV-SFNNs showed a statistically significant improvement in generalization performance in the other twelve. One might ask whether there is any noticeable difference between these two groups of problems. For classification problems, data sets with a higher number of variables (see Table 1 and imagine for instance an arbitrary threshold of 20 inputs) were the ones in which SV-SFNNs outperformed EM-ELMs (with the exceptions of the *German credit* and the *Segmentation* problems). Although this can be considered as a reasonable result, which may be justified by the difficulty in finding adequate decision boundaries in high-dimensional input spaces from randomly distributed hidden-layer weights, the underlying hypothesis needs further validation in future studies. For regression problems this trend is not so clear.

The second conclusion of the experimental study is that, independently of the strategy used (*input* or *random*), the number of candidates for the hidden-layer weights is a parameter that controls the trade-off between the generalization performance, the computational cost and the number of hidden units of the final models. In general terms, by increasing the number of candidates at each step of the sequential algorithm (recall that in the originally proposed EM-ELMs this number is 1), the generalization is improved and the final number of hidden units is reduced, at the expense of a higher training time. However, as pointed in Section 3.1, the computational cost of trying *C* candidates is lower than *C* times the cost of trying a single candidate, due to the incremental way in which the optimal output-layer weights are calculated.

Both strategies can be further improved. The *random* strategy could be modified so as to obtain the weights taking into account the underlying distribution of the data. The *input* one could take profit of sample selection methods (see Plutowski, Cottrell, & White, 1996, for example) to select the input vectors in a better way than randomly. In addition, regularization techniques can also be used for both strategies.

References

- Asuncion, A., & Newman, D. J. (2007). UCI machine learning repository. University of California, Irvine, School of Information and Computer Science. <http://www.ics.uci.edu/~mlearn/MLRepository.html>.
- Bishop, C. M. (1995). *Neural networks for pattern recognition*. New York: Oxford University Press Inc.

- Chen, S., Cowan, C. F. N., & Grant, P. M. (1991). Orthogonal least squares learning algorithm for radial basis function networks. *IEEE Transactions on Neural Networks*, 2, 302–309.
- Feng, G., Huang, G. B., Lin, Q., & Gay, R. (2009). Error minimized extreme learning machine with growth of hidden nodes and incremental learning. *IEEE Transactions on Neural Networks*, 20, 1352–1357.
- Huang, G. B., Chen, L., & Siew, C. K. (2006). Universal approximation using incremental constructive feedforward networks with random hidden nodes. *IEEE Transactions on Neural Networks*, 17, 879–892.
- Huang, G. B., Zhu, Q. Y., & Siew, C. K. (2006). Extreme learning machine: theory and applications. *Neurocomputing*, 70, 489–501.
- Keerthi, S. S., Chapelle, O., & DeCoste, D. (2006). Building support vector machines with reduced classifier complexity. *Journal of Machine Learning Research*, 7, 1493–1515.
- Kwok, T. Y., & Yeung, D. Y. (1997). Constructive algorithms for structure learning in feedforward neural networks for regression problems. *IEEE Transactions on Neural Networks*, 8, 630–645.
- Liang, N. Y., Huang, G. B., Saratchandran, P., & Sundararajan, N. (2006). A fast and accurate online sequential learning algorithm for feedforward networks. *IEEE Transactions on Neural Networks*, 17, 1411–1423.
- Liu, N. L., & Wang, H. (2010). Ensemble based extreme learning machine. *IEEE Signal Processing Letters*, 17, 754–757.
- Miche, Y., Sorjamaa, A., Bas, P., Simula, O., Jutten, C., & Lendasse, A. (2010). OP-ELM: optimally pruned extreme learning machine. *IEEE Transactions on Neural Networks*, 21, 158–162.
- Platt, J. (1991). A resource-allocating network for function interpolation. *Neural Computation*, 3, 213–225.
- Plutowski, M., Cottrell, G., & White, H. (1996). Experience with selecting exemplars from clean data. *Neural Networks*, 9, 273–294.
- Romero, E., & Alquézar, R. (2002). A new incremental method for function approximation using feed-forward neural networks. In *International joint conference on neural networks: Vol. 2* (pp. 1968–1973).
- Romero, E., & Alquézar, R. (2006). A sequential algorithm for feed-forward neural networks with optimal coefficients and interacting frequencies. *Neurocomputing*, 69, 1540–1552.
- Romero, E., & Toppo, D. (2007). Comparing support vector machines and feed-forward neural networks with similar hidden-layer weights. *IEEE Transactions on Neural Networks*, 18, 959–963.
- Smola, A. J., & Schölkopf, B. (2000). Sparse greedy matrix approximation for machine learning. In *International conference on machine learning* (pp. 911–918).
- Van Gestel, T., Suykens, B., Baesens, J. A. K., Viane, S., Vanthienen, J., Dedene, G., et al. (2004). Benchmarking least squares support vector machine classifiers. *Machine Learning*, 54, 5–32.
- Vapnik, V. N. (1995). *The nature of statistical learning theory*. NY: Springer-Verlag.
- Vincent, P., & Bengio, Y. (2002). Kernel matching pursuit. In *New methods for model combination and model selection Machine Learning*, 48, 165–187 (special issue).
- Yingwei, L., Sundararajan, N., & Saratchandran, P. (1997). A sequential learning scheme for function approximation using minimal radial basis function neural networks. *Neural Computation*, 9, 461–478.
- Yu, Y., Zhou, Z. H., & Ting, K. M. (2007). Cocktail ensemble for regression. In *International conference on data mining* (pp. 721–726).