

Multi-Objective Model Type Selection

Alejandro Rosales-Pérez^{a,*}, Jesus A. Gonzalez^a, Carlos A. Coello Coello^b,
Hugo Jair Escalante^a, Carlos A. Reyes-Garcia^a

^a*Instituto Nacional de Astrofísica, Óptica y Electrónica (INAOE), Computer Science
Department, Luis Enrique Erro No.1, Santa María Tonantzintla, Puebla, 72840, Mexico*

^b*Centro de Investigación y de Estudios Avanzados del IPN (CINVESTAV-IPN),
Computer Science Department, Evolutionary Computation Group (EVOCINV), Av. IPN
No. 2508, San Pedro Zacatenco, Mexico City, 07360, Mexico*

Abstract

Classification is a mainstream within the machine learning community. As a result, a large number of learning algorithms have been proposed. The performance of many of these could highly depend on the chosen values of their hyper-parameters. This paper introduces a novel method for addressing the model selection problem for a given classification task. In our model selection formulation, both the learning algorithm and its hyper-parameters are considered. In our proposed approach, model selection is tackled as a multi-objective optimization problem. The empirical error, or training error, and the model complexity are defined as the objectives. We adopt a multi-objective evolutionary algorithm as the search engine, due to its high performance and its advantages for solving multi-objective problems. The model complexity is estimated experimentally, in a general fashion, as normally done for any learning algorithm, through the VC dimension. Strategies for choosing a single model or for constructing an ensemble of models from

*Corresponding author. Tel.: +52 (222) 2663100 x 3413
Email address: arosales@inaoep.mx (Alejandro Rosales-Pérez)

the resulting non-dominated set are also proposed. Experimental results on benchmark data sets indicate the effectiveness of the proposed approach. Furthermore, a comparative study shows that the obtained models are highly competitive, in terms of generalization performance, with other methods in the state of the art that focus on a single-learning algorithm, or a single-objective approach.

Keywords: Model type selection, VC dimension, Multi-objective optimization, Ensemble methods

1. Introduction

Classification is a common task in supervised learning. Its popularity is due to its use in a wide range of applications, such as medical diagnosis, text categorization, etc. In the machine learning community, several learning algorithms to fit a model have been proposed, including decisions trees, artificial neural networks, those based on statistical learning, etc. However, to date there is not a universal “best” model; this is referred to as the **No Free Lunch Theorem** [?]. Moreover, many of these learning algorithms have a set of adjustable parameters, called hyper-parameters, whose fine-tuning can affect their generalization ability. Taking that into consideration, one might ask the questions: what learning algorithm should be used for a specific problem? Also, given a learning algorithm, what hyper-parameters values should be chosen? These questions are related to the issue of model selection.

In the literature, there are several studies that address the model selection problem. Among these, some of them have approached it as an optimization

17 problem, differing in the search technique adopted, including gradient-based
18 methods [? ? ?], grid-search [?], or bio-inspired meta-heuristics such as
19 evolutionary algorithms [? ? ? ? ?], artificial immune systems [?] or
20 particle swarm optimizers [? ? ?], etc. Grid-search is the simplest one, but it
21 could be time-consuming. Although gradient-based methods tend to be more
22 (computationally) efficient, they are very susceptible to the initial search
23 point and they can easily get trapped in a local optimum. Evolutionary
24 algorithms have gained popularity because of their ease of use and their
25 ability to overcome these shortcomings. Indeed, evolutionary algorithms can
26 be less computationally expensive than grid-search, and are less susceptible
27 to their initial search points than gradient-based methods. Furthermore,
28 evolutionary algorithms do not require gradient information and can be easily
29 parallelized.

30 Another major issue in model selection is the criterion used for this pur-
31 pose. In this direction, we can differentiate the works that consider a single-
32 objective criterion and those that consider multiple criteria. The single-
33 objective criterion approaches are generally based on an estimation of the
34 generalization error through the well-known k fold cross validation [? ? ? ?
35]. Attention has also been paid to considering multiple criteria. These works
36 typically consider the model performance and some criterion for penalizing
37 the model complexity [? ?]. Others have considered either to minimize the
38 sensitivity and specificity [? ?], or different estimates of the model perfor-
39 mance [? ? ?]. Alternatively, multiple criteria have also been approached
40 by simplifying the objectives in a weighted linear combination of these [?]
41 instead of simultaneously optimizing the objectives.

42 Despite these efforts, most of the existing studies consider a single model
43 type (i.e., the learning algorithm is fixed *a priori* and the model selection
44 task consists of choosing its hyper-parameters), which could not be the most
45 suitable for a particular problem. To the best of the authors' knowledge,
46 nowadays the works that address both the learning algorithm and the hyper-
47 parameters selection are scarce (e.g. [? ? ?]), and most of them tackle
48 the problem as a single-objective one. Notwithstanding, the disadvantages
49 of using a single-objective approach for hyper-parameters optimization with
50 respect to the generalization performance have been pointed out by several
51 authors [? ? ?].

52 Inspired from previous ideas, we address both the problem of choosing
53 a learning algorithm and its hyper-parameters during the model selection,
54 which is faced as a multi-objective optimization problem. The error on
55 training samples and the model complexity are considered as the objectives
56 in our formulation. Unlike previous works in which the model complexity
57 estimation depends on the learning algorithm (e.g., the number of support
58 vectors in support vector machines), we propose to estimate it through the
59 VC-dimension (for Vapnik-Chervonensky dimension) [?].

60 The main contribution of this paper is a general model selection frame-
61 work, whose formulation makes it applicable to any learning algorithm. Ad-
62 ditional contributions of the paper are as follows: (i) a multi-objective ap-
63 proach for tackling the model type selection problem (i.e., model type plus its
64 hyper-parameters), (ii) the use of the VC-dimension in the model type selec-
65 tion formulation for estimating the model complexity to any model type, and
66 (iii) since the outcome of the multi-objective optimization process is a set of

67 solutions (models), that satisfy an optimal trade-off between the objectives
68 from which a model should be chosen, the strategies proposed for construct-
69 ing a final classification model from the non-dominated solutions set are an
70 additional contribution. The performance of our proposed approach is as-
71 sessed on several binary classification benchmark data sets widely used in the
72 literature. The experimental results and comparisons show that our proposal
73 is able to select highly effective classification models.

74 The remainder of this paper is organized as follows. In Section 2, we
75 describe the VC-dimension theory and the way in which it can be estimated
76 in an experimental fashion. Section 3 presents our proposal, describing in
77 detail how the model selection problem is formulated as a multi-objective one.
78 It also describes the proposal for constructing a final model from solutions
79 in the resulting non-dominated front. Section 4 presents the experiments
80 performed to test the validity of our proposal using benchmark data sets,
81 and the results obtained from these. Finally, the main conclusions and future
82 work direction paths are presented in Section 5.

83 **2. VC Dimension Estimation**

84 Vapnik and Chervonenkis defined the VC dimension [?] as a measure of
85 the capacity of a learning algorithm. The VC Dimension is defined through
86 the notion of “shattering”, which is described as follows: if we have a set of
87 n samples that can be separated by a set of indicator functions F (functions
88 that map a sample to its corresponding binary label) in all 2^n possible ways,
89 we say that the set of samples is shattered by the set of functions F . The
90 VC dimension can be formally defined as [?]:

91 A set of functions F has a VC dimension h if there are h
 92 samples that can be shattered by the set of functions F , but
 93 there are not $h + 1$ samples that can be shattered by the set of
 94 functions F .

95 Notwithstanding that the VC dimension can be seen as a measure of
 96 the model complexity [?], exact analytic estimates of this are only known
 97 for a few classes of functions (linear models), whereas for many others it is
 98 unknown. To overcome this, Vapnik *et al.* [?] proposed a method to exper-
 99 imentally estimate the effective VC dimension of a model. This approach is
 100 based on the best fitting between an analytic formula and measurements of
 101 the maximum deviation between the error rates on two independent data sets
 102 of varying sizes. Conceptually, this approach can be applied to any learning
 103 algorithm [?].

104 The maximum deviation, $\xi(n)$, of the error rates between two indepen-
 105 dent labeled data sets is defined as:

$$\xi(n) = \max_{\omega} (| \text{err}(\mathbf{Z}_n^1) - \text{err}(\mathbf{Z}_n^2) |) \quad (1)$$

106 where \mathbf{Z}_n^1 and \mathbf{Z}_n^2 are two independent labeled data sets of size n , $\text{err}(\mathbf{Z}_n)$ is
 107 the error rate on the data set \mathbf{Z}_n , and ω is the set of parameters of a binary
 108 classifier.

109 As it is stated in [?], $\xi(n)$ is bounded as follows:

$$\xi(n) \leq \Phi(n/h) \quad (2)$$

110 where

$$\Phi(\tau) = \begin{cases} 1 & \text{if } \tau < 0.5 \\ a \frac{\log(2\tau) + 1}{\tau - k} \left(1 + \sqrt{1 + \frac{b(\tau - k)}{\log(2\tau) + 1}} \right) & \text{if } \tau \geq 0.5 \end{cases} \quad (3)$$

111 where $\tau = n/h$, and the values of the parameters $a = 0.16$ and $b = 1.2$ were
 112 empirically determined. The value of $k = 0.14928$ is determined such that
 113 $\Phi(0.5) = 1$.

114 Since the bound in Equation (2) is tight, it can be assumed that

$$\xi(n) \approx \Phi(n/h) \quad (4)$$

115 The VC dimension h can be estimated from Equations (3) and (4). The
 116 maximum deviation $\xi(n)$ can be estimated by simultaneously minimizing the
 117 error rate on one labeled set and maximizing the error rate in the other one.
 118 This can be accomplished through the following procedure [? ?]:

- 119 1. Generate a random labeled set \mathbf{Z}_{2n} of size $2n$.
- 120 2. Split the set \mathbf{Z}_{2n} into two sets of size n : \mathbf{Z}_n^1 and \mathbf{Z}_n^2 .
- 121 3. Flip the labels of the set \mathbf{Z}_n^1 , to form $\bar{\mathbf{Z}}_n^1$.
- 122 4. Merge the two sets: $\bar{\mathbf{Z}} = \bar{\mathbf{Z}}_n^1 \cup \mathbf{Z}_n^2$, and train the binary classifier with
 123 the set $\bar{\mathbf{Z}}$.
- 124 5. Evaluate \mathbf{Z}_n^1 and \mathbf{Z}_n^2 with the trained classifier. Measure the difference
 125 of the error rates between the two sets: $\xi(n) = |\text{err}(\mathbf{Z}_n^1) - \text{err}(\mathbf{Z}_n^2)|$.

126 This procedure gives an estimate of $\xi(n)$ from which an estimate of h can
 127 be obtained. In order to reduce the variability in the estimation, this proce-

128 dure is repeated for different data sets varying the samples sizes n_1, \dots, n_k .
 129 Moreover, to reduce the variability due to the random samples, the procedure
 130 is repeated several times (m_j) for each sample set of size n_i . The average
 131 value for each experiment is taken for each n_i : $\bar{\xi}(n_1), \dots, \bar{\xi}(n_k)$. The effec-
 132 tive VC dimension can be estimated by finding the parameter h^* that best
 133 fits $\xi(n)$ with the theoretical formula $\Phi(n/h)$, as follows:

$$h^* = \operatorname{argmin}_h \sum_{i=1}^k [\bar{\xi}(n_i) - \Phi(n_i/h)]^2 \quad (5)$$

134 3. Multi-Objective Approach for Model Selection

135 The proposed approach formulates the model selection problem as a
 136 multi-objective optimization one, where the training error and the model
 137 complexity are considered as the objectives to be minimized. The general
 138 diagram of our proposal is shown in Figure 1.

139 The process starts by creating an initial population. In our problem,
 140 each individual in the initial population represents a potential model for a
 141 classification task. After that, we compute the components to be optimized:
 142 the training error and the model complexity, which is estimated in a general
 143 fashion via the VC-dimension, as it is explained in Section 2. Next, the
 144 models are evolved through by applying the evolutionary operators to create
 145 an offspring population, which represents new potential models for the given
 146 classification task. Thereafter, the models that satisfy the best trade-off
 147 between the two objectives to be optimized are stored in an external archive.
 148 This process is repeated until a stopping criterion is reached. At the end
 149 of the search, a final classification model is constructed, which is used for

150 predicting the class labels of unknown samples.

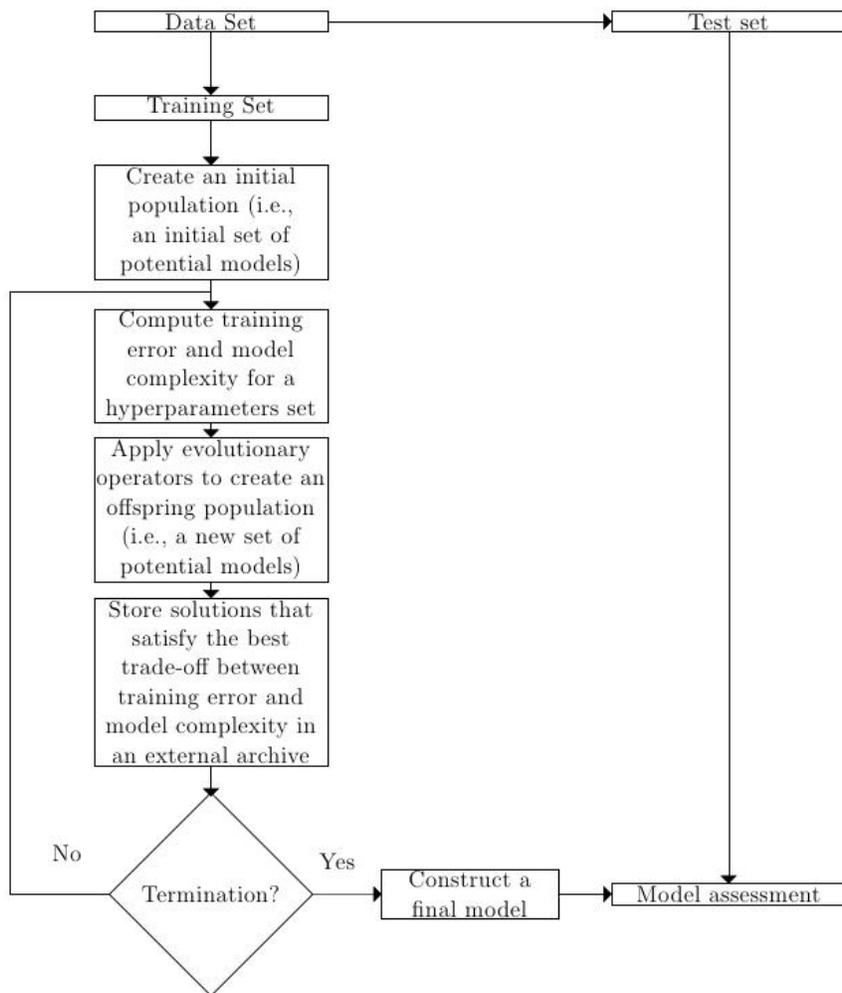


Figure 1: General diagram for the multi-objective model selection process.

151 In the proposed approach, we consider five different model types: support
152 vector machines (SVMs), neural networks (NNs), random forest (RF), j48
153 and random trees (RTs). All of these are available in the WEKA [?] toolbox,
154 and LibSVM [?] for the SVM. Table 1 shows the learning algorithms

155 considered in our study. It also shows for each method the corresponding
 156 hyper-parameters. In the rest of the section we explain our proposal in
 157 detail.

Table 1: Description of the learning methods considered in our study.

Learn. Alg.	Hyper-parameters	Description
J48	Confidence: A confidence threshold for pruning. K : A minimum number of instances per leaf.	It constructs a pruned or unpruned C4.5 decision tree.
NNs	Neurons: Number of neurons in the hidden layer lr : Learning Rate for the backpropagation algorithm. Momentum: Momentum Rate for the backpropagation algorithm. Epochs: Number of epochs to train through. Seed: The value used to seed the random number generator.	It constructs a multi-layer perceptron using the backpropagation algorithm.
RF	Trees: Number of trees to build. K : Number of features to consider. Depth: The maximum depth of the trees. Seed: The value used to seed the random number generator.	It constructs a forest of random trees.
RT	K : Number of features to randomly investigate. Depth: The maximum depth of the tree. Seed: Seed used for the random number generator.	It constructs a tree that considers K randomly chosen attributes at each node. It does not perform a pruning step.
SVMs	Kernel: The kernel type to be used. d : The degree of a polynomial kernel. γ : Gamma value of an RBF kernel. B : A bias value in polynomial kernel. C : The complexity constant. Seed: Seed for the random number generator.	It constructs a support vector classifier.

158 *3.1. Multi-Objective Evolutionary Algorithms*

159 A multi-objective optimization problem (MOOP) can be stated as follows:

$$\begin{aligned}
 & \text{minimize } \mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), \dots, f_l(\mathbf{x})]^T \\
 & \text{subject to } \mathbf{x} \in \mathcal{X}
 \end{aligned} \tag{6}$$

160 where $\mathbf{x} = [x_1, \dots, x_n]^T \in \mathbb{R}^n$ is a decision variables vector, $f_i(\mathbf{x})$, $i =$
 161 $1, \dots, l$, are the objective functions, and \mathcal{X} is the set of feasible solutions.

162 When the objectives in an MOOP are in conflict, there is not a single
 163 solution that would be the best for all of them. Pareto optimality provides a
 164 framework for dealing with such cases. We say that a solution \mathbf{x}^1 dominates
 165 a solution \mathbf{x}^2 (denoted by $\mathbf{x}^1 \preceq \mathbf{x}^2$) if and only if \mathbf{x}^1 is better than \mathbf{x}^2 at least
 166 in one objective and it is not worse in the rest, i.e.,

$$\forall i : f_i(\mathbf{x}^1) \leq f_i(\mathbf{x}^2) \wedge \exists i : f_i(\mathbf{x}^1) < f_i(\mathbf{x}^2) \quad (7)$$

167 A solution \mathbf{x}^* is Pareto optimal if there is not another solution $\mathbf{x}' \in \mathcal{X}$
 168 such that $\mathbf{x}' \preceq \mathbf{x}^*$. The set of all Pareto optimal solutions is called Pareto
 169 optimal set, and the image of this set in objective function space is referred
 170 to as the Pareto Front.

171 Evolutionary algorithms are stochastic search techniques inspired in Dar-
 172 win's evolutionary theory. These algorithms have been successfully used for
 173 solving MOOPs, mainly because they can obtain several elements of the
 174 Pareto optimal set in a single run, and because they are less susceptible than
 175 mathematical programming techniques to the shape and continuity of the
 176 Pareto front [? ?].

177 Since the seminal work of Schaffer [?], a considerable number of multi-
 178 objective evolutionary algorithms (MOEAs) have been proposed, such as:
 179 Multi-Objective Genetic Algorithm (MOGA) [?], Niche Pareto Genetic
 180 Algorithm (NPGA) [?], Strength Pareto Evolutionary Algorithm (SPEA) [?
 181], and its improved version SPEA2 [?], Pareto Archived Evolutionary Strat-
 182 egy (PAES) [?], Non-dominated Sorting Genetic Algorithm (NSGA) [?]

183 and NSGA-II [?], and Multi-Objective Evolutionary Algorithm based on
184 Decomposition (MOEA/D) [?], among others. A comprehensive review of
185 MOEAs can be found in [? ? ?].

186 In the context of multi-objective model selection, evaluating the objec-
187 tives is computationally expensive, inasmuch as each candidate model has
188 to be trained and tested, possibly several times. Furthermore, in the model
189 selection problem, the optimal model is unknown *a priori*. The latter makes
190 necessary that the generated solutions are diverse to each other in order to
191 use *a posteriori* processing for constructing a final model. Taking this infor-
192 mation into account, in this study we used MOEA/D, due to its high perfor-
193 mance over different difficult problems [?]. Additionally, MOEA/D has a
194 lower computational complexity than other MOEAs (such as the NSGA-II),
195 and is able to provide well-distributed solutions along the Pareto front [?].

196 3.1.1. MOEA/D

197 MOEA/D [?] is one of the most recent MOEAs reported in the state
198 of the art. It is based on the idea of decomposing an MOOP into a number
199 of scalar objective optimization problems, also called subproblems, through
200 a weighted aggregation of the objectives. MOEA/D minimizes all these sub-
201 problems iteratively in a single run. A neighborhood relation based on the
202 distance of the aggregation weights vectors is defined among the subprob-
203 lems. The optimal solutions to two neighboring subproblems should be very
204 similar. Each subproblem has its best solution found so far in the population
205 and is optimized in MOEA/D by using information from its neighbors.

206 A description of MOEA/D is presented in Algorithm 1. MOEA/D starts
207 by creating an empty external population (*EP*) (step 1), which is used

208 to store the non-dominated solutions found so far during the search. In
209 MOEA/D, the T closest weight vectors in $\{\lambda^1, \dots, \lambda^N\}$ to a weight vector λ^i
210 constitute the neighborhood of λ^i . Thus, for each vector λ^i , it is computed
211 the Euclidean distance between it and the others, and their T closest weight
212 vectors are determined, where T defines the neighborhood size. The indexes
213 of such T closest weight vectors are assigned to $B(i)$ (step 2). Next, the
214 initial population of N individuals is randomly created (step 3). The indi-
215 viduals of the initial population are evaluated by using the fitness functions.
216 For each objective, the lowest value attained by the individuals in the initial
217 population is used to initialize a reference vector \mathbf{z} (step 4).

218 The process to generate a new solution \mathbf{y} comes here. To do this, the
219 parents are randomly selected from the neighborhood, to which evolutionary
220 operators (such as crossover and mutation) are applied to create \mathbf{y} (step 7).
221 In case \mathbf{y} violates any constraint, the next step consists of applying some
222 repair heuristic in order to make \mathbf{y} a feasible solution \mathbf{y}' (step 8). Next,
223 reference vector \mathbf{z} is updated in case an objective with a lower value is found
224 (step 9). After that, the neighboring solutions are updated by considering
225 all the neighbors of the i^{th} subproblem and replacing \mathbf{x}^j by \mathbf{y}' if \mathbf{y}' performs
226 better than \mathbf{x}^j (step 10). The external population EP that was initialized in
227 step 1 is updated by the new generated solution if and only if this solution is
228 non-dominated with respect to those that are in EP . Moreover, if the new
229 solution dominates any of those stored in EP , such solutions are removed
230 from EP (step 11). Steps 7 to 11 are repeated while a stopping criterion is
231 not reached. A detailed description of MOEA/D is beyond the scope of this
232 paper, but interested readers are referred to [?] for more information about

233 this approach.

Algorithm 1 MOEA/D [?]

Require: A stopping criterion,

N : number of subproblems considered in MOEA/D,

A uniform spread of N weight vectors: $\lambda^1, \dots, \lambda^N$,

T : the number of weight vectors in the neighborhood of each weight vector

Ensure: EP : an external population

1: Initialize $EP \rightarrow \emptyset$

2: Compute the Euclidean distance between any two weight vectors and then work out the T closest weight vectors to each weight vector. For each $i = 1, \dots, N$, set $B(i) = \{i_1, \dots, i_T\}$, where $\lambda^{i_1}, \dots, \lambda^{i_T}$ are the closest weight vectors to λ^i .

3: Generate an initial population $\mathbf{x}^1, \dots, \mathbf{x}^N$

4: Initialize $\mathbf{z} = [z_1, \dots, z_m]$ by setting $z_j = \min_{1 \leq i \leq N} f_j(\mathbf{x}^i)$

5: **while** stopping criterion is not satisfied **do**

6: **for** $i = 1$ **to** N **do**

7: Randomly select two indexes k, l from $B(i)$, and then generate a new solution \mathbf{y} from \mathbf{x}^k and \mathbf{x}^l by using genetic operators.

8: Apply a repair/improvement heuristic on \mathbf{y} to produce \mathbf{y}' .

9: Update \mathbf{z} , for each $j = 1, \dots, m$ if $z_j > f_j(\mathbf{y})$, then set $z_j = f_j(\mathbf{y})$

10: Update of neighboring solutions: For each index $j \in B(i)$, if $g^{te}(\mathbf{y}'\lambda^j, \mathbf{z}) \leq g(\mathbf{x}^j\lambda^j, \mathbf{z})$, then set $\mathbf{x}^j = \mathbf{y}'$, $FV^j = F(\mathbf{y}')$

11: Update of EP: Add $F(\mathbf{y}')$ to EP if it is non-dominated with respect to the vectors stored in EP, and remove from EP the vectors dominated by $F(\mathbf{y}')$.

12: **end for**

13: **end while**

234 As evolutionary operators we used a differential evolution crossover-mechanism [?

235], and polynomial-based mutation [?]. In the differential evolution operator

236 adopted ,each element \bar{y}_j of a new solution $\bar{\mathbf{y}} = [\bar{y}_1, \dots, \bar{y}_n]$ is generated as

237 follows:

$$\bar{y}_j = \begin{cases} x_j^i + F \times (x_j^k - x_j^l) & \text{with probability } CR, \\ x_j^i & \text{with probability } 1 - CR \end{cases} \quad (8)$$

238 where CR and F are two control parameters.

239 Polynomial-based mutation generates the new solution, $\mathbf{y} = [y_1, \dots, y_n]$
 240 as follows:

$$y_j = \begin{cases} \bar{y}_j + \Delta_j \times (U_b - L_b) & \text{with probability } pm \\ \bar{y}_j & \text{with probability } 1 - pm, \end{cases} \quad (9)$$

241 where pm is the probability of mutation, U_b and L_b are the upper and lower
 242 bounds, respectively, and Δ_j is a polynomial distribution for random numbers
 243 generation in the following way:

$$\Delta_j = \begin{cases} (2 \times \text{rand})^{\frac{1}{\eta+1}} - 1 & \text{if } \text{rand} < 0.5 \\ 1 - [2 \times (1 - \text{rand})]^{\frac{1}{\eta+1}} & \text{otherwise} \end{cases} \quad (10)$$

244 where “rand” is a uniform random number in $[0, 1]$, and η is the distribution
 245 index for the mutation operator.

246 One of the key issues in MOEA/D is the method used for decomposing
 247 the MOOP into a number of scalar objective problems. A simple method
 248 in this regard is the weighted sum approach, but it has the disadvantage of
 249 not being able to generate concave portions of a Pareto front [?]. We used
 250 instead, the Tchebycheff approach [?], due to the fact that it is more robust
 251 to a concave shape of the Pareto front than the weighted sum approach.
 252 However, any other decomposition approach could be used in MOEA/D.

253 Using the Tchebycheff approach, an MOOP is decomposed into a N scalar
 254 optimization subproblem as follows:

$$\text{minimize } g(\mathbf{x} \mid \lambda, \mathbf{z}^*) = \max_{1 \leq i \leq m} \{\lambda_i \mid f_i(\mathbf{x}) - \mathbf{z}_i^* \mid\} \quad (11)$$

255 where $\lambda = [\lambda_1, \dots, \lambda_m]$ is a weight vector, $\mathbf{z}^* = [z_1, \dots, z_m]$ is a reference
 256 point, and m is the number of objectives in the problem.

257 In the literature, several repair heuristics have been proposed [?]. Nev-
 258 ertheless, we formulate the multi-objective model selection problem as an
 259 unconstrained one. For this reason, a repair heuristic is not used in our
 260 study; therefore, step 8 is not performed. The following sections explain the
 261 proposed approach for multi-objective model selection using MOEA/D.

262 3.2. Representation

263 Evolutionary Algorithms work with a population of solutions. In our pro-
 264 posed approach, each solution, also called individual, represents a potential
 265 model for the classification task. As previously stated, the task approached
 266 by our model selection proposal is to choose among a pool of learning al-
 267 gorithms and their corresponding hyper-parameters. To achieve this task,
 268 each model (the learning algorithm plus its hyper-parameters) should be en-
 269 coded in a D -dimensional vector. In this study, each solution is encoded in
 270 a 7-dimensional vector as follows:

$$\mathbf{x}^i = [x_m^i, x_{hp_1}^i, \dots, x_{hp_{D-1}}^i] \quad (12)$$

271 where x_m^i controls the learning algorithm, and $[x_{hp_1}^i, \dots, x_{hp_{D-1}}^i]$ represents
 272 the hyper-parameters for the learning algorithm.

273 Since the hyper-parameters are numerical values, and in order to have
274 them as accurate as possible, we used a real encoding for the individuals.
275 By applying the evolutionary operators, such as crossover (Eq. (8)) and
276 mutation, the individuals are evolved in an iterative process. One should
277 note that there are some discrete variables, such as x_m^i , which represent a
278 learning algorithm, or $x_{hp_x}^i$, which could represent a kernel type in the SVM
279 case. For the evolutionary operators, this type of variable is internally treated
280 as a real number, but we round it off to its nearest allowable discrete value.

281 From Table 1, we can observe that different learning algorithms require
282 different hyper-parameters. For example, in J48 two hyper-parameters are
283 considered, whilst in SVMs there are six hyper-parameters. Thus, x_m^i and
284 the six hyper-parameters are the seven variables in our representation. The
285 configuration given by an individual and the training set are used to fit a
286 model.

287 The seven variables constitute the search space for our problem. An initial
288 population is created using the Latin Hypercube sampling technique [?]]
289 with the aim of having a representative distribution of solutions in the search
290 space. Once the initial population is created, it is used for producing an
291 offspring population by applying the evolutionary operators until a stopping
292 criterion is satisfied, and a set of non-dominated solutions is obtained.

293 *3.3. Fitness Functions*

294 In the proposed approach, the model selection problem is tackled as a
295 multi-objective optimization problem, and an MOEA is used for solving it.
296 Since the search is based on a population of solutions, it is required to have
297 a way to measure how well a model performs in order to choose the best one.

298 The fitness function is in charge of this, and its definition is a crucial issue
 299 in model selection. One could try to estimate the effectiveness of the model
 300 based on the error on the training samples, also known as empirical error,
 301 and the optimization problem would try to minimize that error. Nonetheless,
 302 this would result in optimistic estimations of model performance, and could
 303 lead to highly complex models, causing the problem known as over-fitting.
 304 In other words, the model has a good performance on the training samples,
 305 but not on unseen samples (see [? ? ?] for more information about this
 306 problem). To overcome this handicap, the model complexity should also be
 307 controlled. Taking this into account, in this paper we propose not only to
 308 minimize the error on the training data, but also to minimize the model
 309 complexity.

310 The VC-dimension is a measure of the capacity of the model, which is
 311 related to its complexity, and it is used in the present study. The fitness
 312 functions defined for our problem are stated as follows:

$$\begin{aligned}
 err &= \frac{1}{N} \sum_{i=1}^N \mathcal{L}(y_i, y_i^*) \\
 complexity &= \operatorname{argmin}_h \sum_{i=1}^k [\bar{\xi}(n_i) - \Phi(n_i/h)]^2
 \end{aligned} \tag{13}$$

313 where N is the number of samples in the training set, y_i is the class label, y_i^*
 314 is the class predicted by the model, $\mathcal{L}(y_i, y_i^*)$ is a loss function, $\xi(n_i)$ is the
 315 experimental maximum deviation error rate of two observed independent
 316 labeled data sets, and $\Phi(n_i/h)$ is the expectation of the largest deviation
 317 error between two sets (see Section 2 for details about complexity estimation).
 318 We used the 0/1 loss function because it is well suited for classification tasks.

319 The 0/1 loss function is defined as:

$$\mathcal{L}(y_i, y_i^*) = \begin{cases} 1 & \text{if } y_i^* \neq y_i \\ 0 & \text{if } y_i^* = y_i \end{cases} \quad (14)$$

320 In consequence, the goal of performing this optimization is to simulta-
321 neously minimize the training error and model complexity. The outcome
322 of this optimization step is a set of potential models that satisfies the best
323 trade-off between the objectives, from which a model should be chosen. The
324 next section explains how we approach this issue.

325 3.4. *Constructing a Final Model*

326 Once the evolutionary search is completed, a set of non-dominated so-
327 lutions is obtained. Mathematically, all of them are equally acceptable so-
328 lutions of the multi-objective optimization problem and, in our case, each
329 of them represents a potential model for a given classification task. There-
330 fore, it is desirable to select one model to be used to predict new samples
331 from such set. In model selection for classification tasks, we have to choose
332 the model with the highest possible generalization capability. Nevertheless,
333 it is not clear what classification model from the non-dominated set is the
334 “best” one. In this paper, we studied three strategies for constructing a final
335 classification model, which are explained in the rest of this section.

336 3.4.1. *Choosing a Single Model*

337 As we previously stated, for our problem we seek the solution with the
338 best possible generalization ability. In order to identify such solution, we
339 studied the performance of the non-dominated solutions on unseen samples.

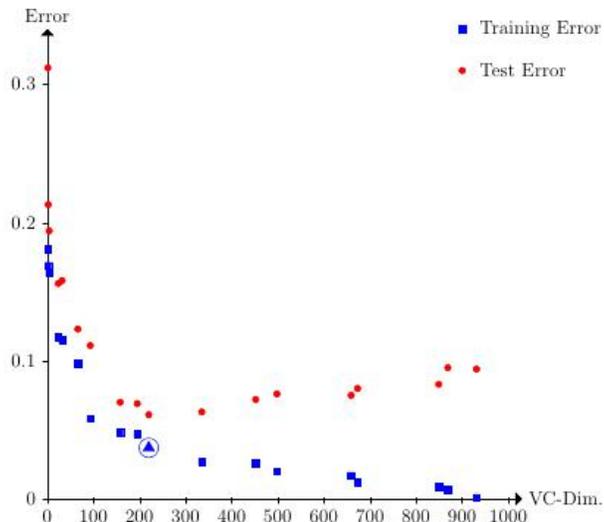


Figure 2: Behavior of non-dominated solutions on training samples and test samples

340 We noticed that the best solutions are those located at the knee of the curve,
 341 while solutions with low complexity and high complexity lead to models
 342 with a poor generalization performance. Both problems are well-known in
 343 machine learning as under-fitting and over-fitting, respectively. Figure 2
 344 depicts an example of this behavior for a particular case. It also shows the
 345 trade-off between the training error and the model complexity, such that by
 346 increasing the model complexity, the training error is reduced.

347 We empirically found that in most cases, the solution with a good gen-
 348 eralization performance is the one nearest to a reference point z^* , which is
 349 defined as:

$$z_j^* = \min_{1 \leq i \leq L} f_j(x^i) \text{ for } j = \{1, 2\} \quad (15)$$

350 where L is the cardinality of the non-dominated set.

351 As it is shown in Figure 2, the objectives are measured in different scales.

352 In order to avoid that one objective has a higher impact than the other one in
353 the distance computation, both objectives are firstly normalized in the range
354 0 to 1. Subsequently, the Euclidean distance is computed on the normalized
355 objective vector. In the end, the closest solution is chosen¹, and is used
356 to predict future samples of the problem. One should note that since the
357 objectives are normalized, the reference vector \mathbf{z}^* corresponds to the (0,0)
358 point. Figure 2 shows with a triangle the solution that would be chosen with
359 this strategy.

360 3.4.2. *Ensemble of the Whole Non-Dominated Front*

361 Ensembles of classifiers are based on the idea of combining the predicted
362 outputs from different individual classification models. They have been suc-
363 cessfully used for improving the performance of individual models [? ?]. One
364 should remember that the output of the MOEA is a set of non-dominated
365 solutions. Based on this, one might ask why not to construct an ensemble
366 with the potential models (solutions) in the non-dominated front instead of
367 choosing a single model.

368 Now the problem is to determine which models should be used in the en-
369 semble. In the absence of knowledge about the preferences, all non-dominated
370 solutions are equally good. With these ideas in mind, we used all of them for
371 constructing an ensemble. One should recall that the non-dominated front
372 could contain models with a very low complexity or a very high complexity,
373 which could lead to over-fitted and under-fitted models, respectively. In a

¹One should note that this is a suggestion. If the model, however, does not satisfy the performance requirements, any other solution can be chosen from the non-dominated set without performing a new search.

374 majority vote scheme all models are equally important, even those whose
 375 generalization performance could not be good enough. For that reason, we
 376 argue in favor of a weighted linear combination approach over a majority vote
 377 one². This is because models with a less satisfactory performance are kept in
 378 the ensemble, but with a lower weight. Therefore, the final prediction given
 379 by the classification model would consist of a weighted linear combination of
 380 the individual predictions, as follows:

$$y^* = \sum_{i=1}^L \omega_i y_i^* \quad (16)$$

381 where L is the number of single classification models, and is equal to the
 382 cardinality of the non-dominated set, y_i^* is the prediction given by the i^{th}
 383 single model, and ω_i is the weight associated to that model. The weight
 384 vector $\omega = [\omega_1, \dots, \omega_L]$ is a normalized vector, whose values depend on the
 385 distances between the reference point (defined by Equation (15)) and the
 386 potential solution. The normalized weight vector is computed as follows:

$$\omega_j = \frac{\frac{1}{d_j}}{\sum_{i=1}^L \frac{1}{d_i}} \quad (17)$$

387 where d_j is the Euclidean distance between the j^{th} solution and the reference
 388 point \mathbf{z}^* .

389 Alternatively, one could think in exploring other well-known ensemble
 390 techniques, such as bagging or boosting, or even trying to optimize the en-
 391 semble performance. Nevertheless, the main focus in this stage of our study

²Please note that the majority vote ensemble is a special case of the weighted linear combination approach.

392 is not the ensemble itself, but the optimization of the hyper-parameters for
393 the classifiers. Our aim is to find the hyper-parameters that satisfy the best
394 trade-off between the objectives, as well as to find ways to construct a classi-
395 fication model from the resulting non-dominated front, which could then be
396 used for the prediction of unknown samples. These issues could be explored
397 as part of our future work.

398 *3.4.3. Ensemble of Some Solutions in the Non-Dominated Front*

399 It is well-known from machine learning that for constructing an ensemble,
400 two conditions have to be satisfied: the individual models should be accurate
401 (i.e., the performance should be better than a random guessing), and they
402 have to be diverse (i.e., single models should incur in different errors on new
403 samples) among them [?]. This issue is explored in the third strategy for
404 the final model construction. Therefore, for constructing an ensemble in this
405 third strategy, we need to choose a subset of potential models in the non-
406 dominated front, such that these models are accurate and diverse among
407 them.

408 We would like to remark that the models were optimized during the op-
409 timization step, and the ones that satisfy the best trade-off are obtained as
410 a result of this. Thus, we can assume that the models in the resulting non-
411 dominated set are accurate (i.e., their performance is better than a random
412 one). By making this, the problem is reduced to choosing a subset of these
413 models that are as diverse as possible among them, which are used in the
414 ensemble. In order to determine such subset, a forward aggregation approach
415 is used. In the forward aggregation approach, we start by adding the solution
416 closest to the reference point, z^* (as it was defined in Equation (15)). After

417 that, a second model that maximizes the diversity is added, followed by a
 418 third model and so on. This process is repeated while the diversity among
 419 the models is not deteriorated.

420 Under the adopted approach, a diversity measure is required. There are
 421 a number of diversity measures reported in the literature, and a review of
 422 these can be found in [?]. In this study, we used one based on entropy, but
 423 any other can also be used. This measure is defined as follows:

$$E = \frac{1}{N} \sum_{i=1}^N \frac{1}{L - \lceil \frac{L}{2} \rceil} \min \{l(\mathbf{s}_i), L - l(\mathbf{s}_i)\} \quad (18)$$

424 where N is the number of samples, L is the number of individual models,
 425 and $l(\mathbf{s}_i)$ is the number of models that correctly predict the sample \mathbf{s}_i . This
 426 measure ranges between 0 and 1, where 0 indicates no difference and 1 is the
 427 highest possible diversity.

428 Finally, the prediction given by the ensemble is based on the weighted
 429 linear combination of the predictions of the individual models, as it is shown
 430 in Equation (16).

431 3.5. Final Remarks

432 One should note that under the proposed approach the expert’s knowl-
 433 edge is not exploited. This could be a key issue in order to improve the
 434 performance of the models. In the agnostic learning vs. prior knowledge
 435 challenge [?] it was shown that, even when prior knowledge outperforms
 436 agnostic learning for most of the problems, there were some problems for
 437 which agnostic learning performs better than prior knowledge. In conse-
 438 quence, it is difficult to know when it is going to be better to use this kind of

439 knowledge. Based on the results of this challenge, the organizers concluded
440 that the agnostic learning approach is very powerful. Furthermore, there are
441 cases in which this knowledge could not be available. For these reasons, we
442 bet in favor of not using expert’s knowledge.

443 Notwithstanding, if prior domain-knowledge is available, this could be
444 integrated in several manners in the proposed approach. For example, based
445 on the characteristics of the problem at hand, an expert could suggest that a
446 particular learning algorithm would be more suitable than the others. This
447 suggestion could be used for fixing *a priori* the learning algorithm. Thus,
448 the search would be performed under its hyper-parameters set, reducing the
449 search space. The expert’s knowledge could also be used for choosing a single
450 solution from the non-dominated front. Another way in which prior knowl-
451 edge could be used would be during the ensemble construction, through the
452 assignment of weights to each classifier. For our experiments, we assumed
453 that expert’s knowledge is not available. Next section describes the experi-
454 ments and results obtained by our proposal.

455 **4. Experiments and Results**

456 In this section, we describe the experiments performed as well as the re-
457 sults obtained by our proposal using a benchmark test suite. We present
458 a comparative study between the three proposed strategies for construct-
459 ing a final classification model from the resulting non-dominated front. We
460 also present statistical tests to validate our results when compared to other
461 approaches reported in the specialized literature.

462 *4.1. Experimental Settings*

463 In order to evaluate the feasibility of our proposal in the model selection
464 problem, we used the IDA benchmark³ data sets introduced by [?]. This
465 benchmark is well-suited for this purpose and it has been widely used in
466 several related studies (e.g. [? ? ? ? ? ?]). Table 2 describes the suite
467 of thirteen benchmark data sets, which are diverse in the number of samples
468 and features. These data sets correspond to binary classification problems⁴,
469 and have been previously pre-processed in [?], in which the samples with
470 missing values have been removed and all features have been standardized,
471 i.e., all features have mean zero and a standard deviation of one.

472 The typical experimental protocol used with this benchmark was intro-
473 duced by [?], and is sometimes called the median protocol. The median
474 protocol consists on performing the model selection on the first five parti-
475 tions. After that, the median values of the hyper-parameters resulting from
476 those partitions are taken, which are used to estimate the error rate for each
477 partition. However, this protocol can introduce an optimistic bias into the
478 performance estimation [?]. In order to overcome this bias in the perfor-
479 mance evaluation, the model selection process is performed independently
480 for each partition of each data set; this protocol is known as the internal
481 protocol. The use of the internal protocol leads to a total of 1140 model
482 selection experiments.

483 The parameters configuration used in our experiments is the following.

³Available at <http://www.raetschlab.org/Members/raetsch/benchmark>

⁴Without loss of generality, the experiments are performed on binary classification problems. Multi-class classification problems can be approached with multiple binary classifiers.

Table 2: Details of the data sets used in our experiments. The table shows the number of features for each data set and the number of instances for training and testing for each replication of each data set.

ID	Data set	Feat.	Training Samples	Testing Samples	Replications
1	Banana	2	400	4900	100
2	Breast Cancer	9	200	77	100
3	Diabetes	8	468	300	100
4	Flare Solar	9	666	400	100
5	German	20	700	300	100
6	Heart	13	170	100	100
7	Image	20	1300	1010	20
8	Ringnorm	20	400	7000	100
9	Splice	60	1000	2175	20
10	Thyroid	5	140	75	100
11	Titanic	3	150	2051	100
12	Twonorm	20	400	7000	100
13	Waveform	21	400	4600	100

484 For the differential evolution crossover, we fixed the value of $F = 0.5$,
485 $CR = 0.7$. With respect to the mutation operator, the mutation proba-
486 bility pm was fixed to 0.1 and index distribution to 20. These parameters
487 were experimentally tuned by evaluating the performance under each con-
488 figuration of $pm = \{0.1, 0.2, 0.3\}$, $CR = \{0.5, 0.6, 0.7, 0.8, 0.9\}$, and $F =$
489 $\{0.3, 0.4, 0.5, 0.6, 0.7\}$ on the first five partitions of splice data set, one of the
490 largest both in number of training samples and features. The stopping crite-
491 rion is defined as performing 1,000 fitness functions evaluations. To achieve
492 this, the population size is set to 20, and the number of generations to 50.
493 Moreover, the VC-dimension for each model is estimated experimentally.
494 Thus, it is required to train and to test a number of times each model. In
495 our experiments, we fixed this number to 10. Next, we present the results

496 reached by our proposal, comparing the proposed strategies for a final model
497 construction and with other evolutionary and non-evolutionary approaches
498 for model selection.

499 *4.2. Experimental Results and Discussion*

500 In this section, we present the results obtained by our proposal (MOMTS,
501 Multi-Objective Model Type Selection) so as to demonstrate its feasibility
502 for the model selection problem. Table 3 shows the average error rates and
503 standard deviations on the test sets attained for the three proposed strategies
504 for constructing a final model, i.e., choosing a single model (MOMTS-S1),
505 ensemble of the whole non-dominated front (MOMTS-S2), and the ensemble
506 of some solutions in the non-dominated front (MOMTS-S3). As a baseline,
507 we report the results obtained by using random forest (RF) with its default
508 hyper-parameters, which is a standard learning algorithm based on an en-
509 semble of decision trees.

510 We compare our results with those reported by Cawley and Talbot [?
511], who used Bayesian regularization at the second level of inference, adding
512 a regularization term in the model selection criterion. Furthermore, in or-
513 der to make a fair comparison, we also performed experiments consider-
514 ing approaches that consider different learning algorithms and their hyper-
515 parameters during the model selection process. For that sake, we used
516 PSMS [?] and SUMO [?], which are two evolutionary approaches that were
517 proposed for model selection. PSMS is a single-objective approach based on
518 a particle swarm optimizer that minimizes the error rate estimated through
519 k fold cross validation. SUMO adopts a genetic algorithm as a search engine
520 and the fitness function can be defined as minimizing some measure obtained

521 via some evaluation strategy; in our case the measure was fixed to be the er-
522 ror rate and the evaluation strategy to be the k fold cross validation. In both
523 cases, the number of particles/individuals was set to 20, and the number of
524 iterations/generations to 50, resulting in 1,000 fitness function evaluations.
525 This is the same number of fitness function evaluations set for our proposed
526 approach. The reference results used the same 100 partitions (20 in case of
527 the image and splice data sets) for training and testing, and also used the
528 same experimental protocol (i.e., the internal protocol), making the results
529 directly comparable.

530 Figure 3 shows the non-dominated fronts generated by our proposal for
531 some data sets in a particular trial. It is expected that these non-dominated
532 fronts are an approximation to the true Pareto front. We can observe that
533 different solutions are distributed along the non-dominated front. We can
534 also note that the non-dominated front is formed by solutions that represent
535 different learning algorithms. Each one of these solutions corresponds to
536 models with different levels of complexity. Although, in some cases, a learning
537 algorithm is represented by more than one solution, these correspond to
538 different configurations of its hyper-parameters, which could lead to diverse
539 models. Thus, in the resulting non-dominated front there are models, which
540 are learned by different learning algorithms with a different hyper-parameters
541 configuration⁵.

⁵The full list of the models generated by our proposed method for each partition of each data set is available at <http://ccc.inaoep.mx/~arosalesp/Resources/Models.zip>

Table 3: Results obtained by the proposed approach, and those obtained by random forest (RF), LS-SVM using Bayesian regularization, PSMS, and SUMO. The reported results are the average and standard error on test sets over the 100 or 20 replications of each data set. The best result for each data set is shown in **boldface**.

Data Set	RF	LS-SVM-BR [?]	PSMS [?]	SUMO [?]	MOMTS-S1	MOMTS-S2	MOMTS-S3
Banana	13.14 ± 0.069	10.59 ± 0.050	11.08 ± 0.083	10.88 ± 0.074	14.34 ± 0.105	10.48 ± 0.046	12.91 ± 0.160
Breast Cancer	27.94 ± 0.412	27.08 ± 0.494	33.01 ± 0.658	26.27 ± 0.448	29.89 ± 0.736	25.61 ± 0.593	27.82 ± 0.676
Diabetes	25.83 ± 0.212	23.14 ± 0.166	27.06 ± 0.259	23.49 ± 0.177	28.34 ± 0.318	23.08 ± 0.174	25.66 ± 0.214
Flare Solar	36.44 ± 0.173	34.07 ± 0.171	34.81 ± 0.173	38.47 ± 0.573	34.90 ± 0.224	34.59 ± 0.189	34.52 ± 0.214
German	25.16 ± 0.240	23.59 ± 0.216	30.10 ± 0.720	23.83 ± 0.213	28.30 ± 0.274	23.67 ± 0.224	25.89 ± 0.218
Heart	20.26 ± 0.382	16.19 ± 0.348	20.69 ± 0.634	17.67 ± 0.355	23.14 ± 0.542	16.48 ± 0.241	18.75 ± 0.351
Image	2.09 ± 0.095	2.90 ± 0.154	2.90 ± 0.112	2.45 ± 0.126	3.79 ± 0.226	2.24 ± 0.123	3.03 ± 0.246
Ringnorm	9.57 ± 0.095	1.61 ± 0.015	7.98 ± 0.660	1.72 ± 0.071	2.66 ± 0.079	2.49 ± 0.074	3.02 ± 0.164
Splice	8.50 ± 0.210	10.91 ± 0.154	14.63 ± 0.324	10.94 ± 0.146	7.43 ± 0.373	4.84 ± 0.156	6.71 ± 0.269
Thyroid	5.89 ± 0.273	4.63 ± 0.218	4.32 ± 0.235	4.85 ± 0.224	6.48 ± 0.350	4.00 ± 0.194	6.11 ± 0.347
Titanic	22.51 ± 0.124	22.59 ± 0.120	24.18 ± 0.193	34.99 ± 0.523	26.53 ± 0.127	22.08 ± 0.085	22.22 ± 0.100
Twonorm	9.02 ± 0.074	2.84 ± 0.021	3.09 ± 0.127	2.55 ± 0.022	5.21 ± 0.555	3.73 ± 0.179	5.70 ± 0.679
Waveform	13.75 ± 0.071	9.78 ± 0.044	12.80 ± 0.325	9.78 ± 0.060	11.34 ± 0.180	9.93 ± 0.043	10.95 ± 0.256

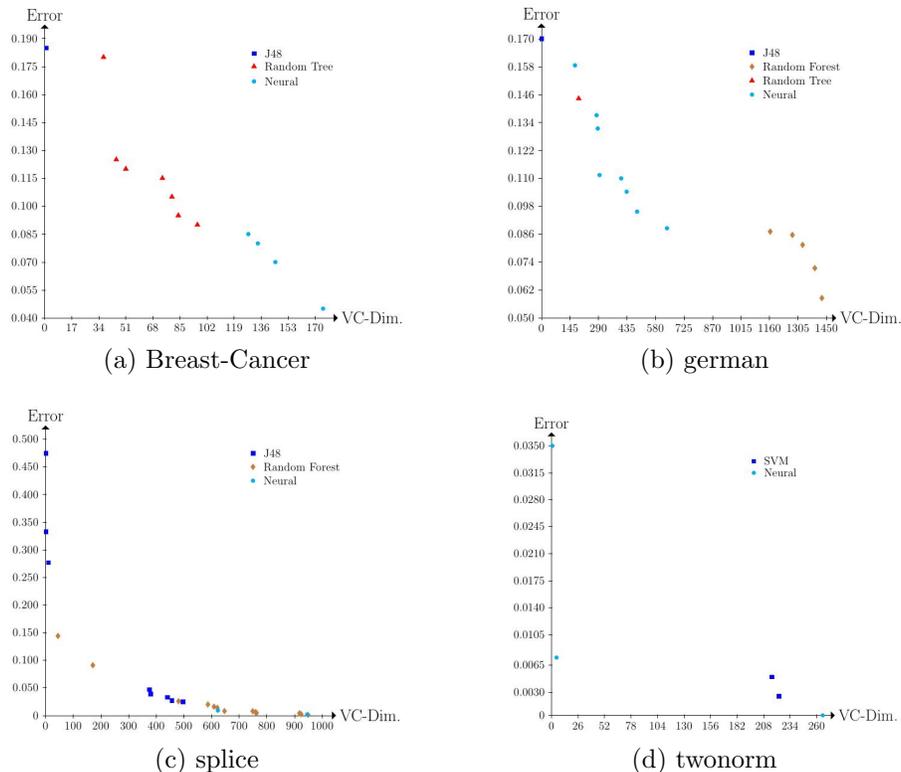


Figure 3: Non-dominated fronts generated from a particular trial of the proposed method. The solutions in the non-dominated front represent different learning algorithms with different hyper-parameter configurations.

542 *4.2.1. Comparison of Strategies for Constructing the Final Model*

543 The results of the proposed strategies for constructing a final model are
 544 shown in the last three columns of Table 3. The results of the ensemble ap-
 545 proaches outperformed those obtained when a single model is chosen in most
 546 cases. The single model was better than the ensemble of some solutions in
 547 the non-dominated front in 2 out of 13 data sets (ringnorm and twonorm
 548 data sets). This seems reasonable insomuch as it is well-known that using an
 549 ensemble of models helps to improve the predictions. Between the two en-

550 sembles approaches, the one based on the whole non-dominated front showed
551 better results in 12 out of 13 data sets.

552 An ANOVA statistical test with a 95% of confidence is applied so as
553 to determine if the difference between the proposed strategies is significant,
554 and Tukey's test is used as a post-hoc test. The results obtained by this
555 test are shown in Table 4. In this table, we can note that the analysis of
556 variance showed a statistical significance difference for most of the data sets,
557 except for the flare solar one, to which the post-hoc test is not applied.
558 According to the pairwise comparisons, we can also note that the ensemble
559 of some solutions of the front (MOMTS-S3) performs significantly better
560 than the single model approach in 6 out of 13 data sets (banana, diabetes,
561 german, heart, image, and titanic). On the other hand, the ensemble of the
562 whole front (MOMTS-S2) showed to be significantly better than the single
563 model approach in 10 out of 13 benchmark data sets (banana, breast cancer,
564 diabetes, german, heart, image, splice, thyroid, titanic, and waveform), and
565 also significantly outperformed the ensemble of some solutions approach in
566 10 out of 13 data sets (banana, diabetes, german, heart, image, ringnorm,
567 splice, thyroid, twonorm, and waveform). It seems clear that the ensemble
568 of the whole front approach is the best of the three approaches. However,
569 for assessing the statistical difference among them over the different data
570 sets, Demšar [?] recommends the Friedman's test for comparing multiple
571 classifiers over multiple data sets. This test is performed with a 95% of
572 confidence, and the Nemenyi test as the post hoc test. According to these
573 tests, the ensemble of the whole front approach is statistically superior to the
574 others.

Table 4: F -statistic obtained from the ANOVA test and q -values from the Tukey HSD test for performing all possible pairwise comparisons among the proposed strategies for a final model construction. The critical values at the 0.05 level for ANOVA test are 3.16 ($F(2, 57)$) for the image and splice data sets and 3.03 ($F(2, 297)$) for the rest. The critical values at 0.05 level for the Tukey HSD test are 3.41 for the image and splice data sets (57 degrees of freedom) and 3.34 for the rest of the data sets (297 degrees of freedom). Cases that exceed the critical value are considered as a difference that is statistically significant at the fixed level and are marked with an asterisk (*).

Data Set	ANOVA	q Tukey HSD		
	F	MOMTS-S1 vs. MOMTS-S2	MOMTS-S1 vs. MOMTS-S3	MOMTS-S2 vs. MOMTS-S3
Banana	294.899*	33.967*	12.584*	21.384*
Breast Cancer	10.178*	6.380*	3.085	3.294
Diabetes	117.12*	21.643*	11.027*	10.616*
Flare Solar	0.932	--	--	--
German	93.107*	19.293*	10.042*	9.251*
Heart	72.394*	16.734*	11.032*	5.705*
Image	14.221*	7.542*	3.698*	3.844*
Ringnorm	5.690*	1.499	3.173	4.672*
Splice	22.736*	9.238*	2.568	6.670*
Thyroid	19.140*	8.110*	1.210	6.900*
Titanic	575.616*	42.203*	40.876*	1.328
Twonorm	3.939*	2.864	0.948	3.812*
Waveform	15.937*	7.731*	2.138	5.593*

575 *4.2.2. Comparison with Other Model Selection Approaches*

576 Table 3 also shows the performance of random forest (RF), LS-SVM with
577 Bayesian Regularization (LS-SVM-BR), which uses a radial basis function
578 kernel reported by [?], as well as the results obtained with the application
579 of PSMS and SUMO in the benchmark data sets. Due to the fact that the
580 best results of our proposal were reached with the ensemble of the whole non-
581 dominated front (MOMTS-S2), this approach is used for the comparison.

582 First, we compare with random forest (RF), which is used as a baseline
583 to evaluate the benefits of performing model selection. From the reported
584 results in Table 3, we can note that our proposal outperformed RF in 12 out
585 of 13 data sets, being the image data set the only one in which RF performed
586 better than our proposal.

587 Comparing the ensemble of the whole non-dominated approach (MOMTS-
588 S2) with LS-SVM-BR, we can note that our proposal obtained better results
589 in 7 out of 13 data sets (banana, breast cancer, diabetes, image, splice, thy-
590 roid, and titanic), but it was outperformed in the rest of the data sets. In
591 addition, it is worth noting that an improvement of more than a 6% was
592 reached in the splice data set.

593 With respect to PSMS, a single-objective approach that considers dif-
594 ferent learning algorithms and hyper-parameters selection, we note that our
595 approach performed better than PSMS in 12 out of 13 benchmark data sets.

596 Comparing MOMTS-S2 with SUMO, another evolutionary approach for
597 model selection, we note that MOMTS-S2 got better generalization perfor-
598 mance on 10 out of 13 data sets.

599 Regarding statistical assessment, we applied the ANOVA test with a 95%

600 of confidence to compare the performance of the model selection approaches:
601 LS-SVM-BR, PSMS, SUMO, and MOMTS-S2. Inasmuch as our goal is to
602 compare the performance of the proposed approach with the reference results,
603 the Dunnett’s test is used as the post-hoc test. These statistical tests were
604 conducted independently for each data set. The results of these are shown in
605 Table 5, from which we can note that, for all cases, the analysis of variance
606 revealed that there is a statistically significance difference at the 0.05 level,
607 i.e., $p < 0.05$. Thus, the post-hoc test is applied.

608 According to the results shown in Table 5, statistical tests indicated that
609 the proposed approach significantly outperformed LS-SVM-BR in 2 data sets
610 (image and splice), and it was significantly outperformed in one data set
611 (twonorm). Regarding SUMO, our method performed significantly better in
612 5 out of 13 data sets (banana, flare solar, splice, thyroid, and titanic), and it
613 was significantly outperformed in the twonorm data set. On the other hand,
614 our approach significantly outperformed PSMS in 10 of the benchmarks data
615 sets (banana, breast cancer, diabetes, german, heart, image, ringnorm, splice,
616 titanic, and waveform data sets), but it was significantly worse than PSMS
617 in the twonorm data set.

618 Overall, our ensemble approach was able to get lower error rates than the
619 other model selection methods in 7 out of 13 data sets, while the Bayesian
620 regularization approach did the same in 5 out of 13 data sets, and SUMO
621 in 2 out of 13 data sets. There is not a clear advantage of LS-SVM-BR and
622 MOMTS-S2 when multiple data sets are considered. In order to statistically
623 assess the four model selection approaches over the suite of 13 benchmark
624 data sets, the Friedman test with a 95% of confidence was used. As a post-

Table 5: F -statistic obtained from the ANOVA test and t_d -values from Dunnett’s test when MOMTS-S2 is compared with LS-SVM-BR, PSMS, and SUMO. The critical values at the 0.05 level for the ANOVA test are 2.72 ($F(3, 76)$) for the image and splice data sets and 2.63 ($F(3, 396)$) for the rest. For Dunnett’s test, the critical values at the 0.05 level are 2.40 for the image and splice data sets (76 degrees of freedom) and 2.36 for the rest of the data sets (396 degrees of freedom). Cases that exceed the critical value are considered as a difference that is statistically significant at the fixed level and are marked with an asterisk (*).

Data Set	ANOVA	MOMTS-S2 versus		
	(F-value)	$(t_d \text{ Dunnett})$		
	F	LS-SVM-BR	PSMS	SUMO
Banana	17.594*	1.191	6.513*	4.346*
Breast Cancer	37.582*	1.875	9.438*	0.841
Diabetes	94.396*	0.216	14.239*	1.466
Flare Solar	38.338*	1.130	0.478	8.435*
German	62.129*	0.138	11.188*	0.280
Heart	23.938*	0.488	7.080*	2.001
Image	6.543*	3.605*	3.605*	1.156
Ringnorm	83.108*	1.862	11.621*	1.629
Splice	377.328*	20.560*	33.152*	20.660*
Thyroid	2.881*	2.042	1.036	2.752*
Titanic	445.726*	1.250	5.150*	31.664*
Twonorm	20.588*	5.683*	4.086*	7.534*
Waveform	78.231*	0.631	12.074*	0.631

625 hoc test, we used the Bonferroni-Dunn test, to compare the performance of
626 our proposal (MOMTS-S2) with the references. According to these tests,
627 MOMTS-S2 is statistically better than PSMS, but there is not a statistical
628 significance difference between MOMTS-S2 and LS-SVM-BR and MOMTS-
629 S2 and SUMO.

630 Another aspect to take into consideration is the computational cost of
631 the methods. In this regard, we compare the execution time required by
632 our proposal against PSMS and SUMO. The average execution time of our
633 proposal (MOTMS) was 54.29 minutes, whilst PSMS and SUMO required,

634 respectively, 30.36 and 31.90 minutes on average. As one could note, our
635 proposal is more time-consuming than the others. This is due to the fact that
636 under the proposed approach two objectives have to be evaluated, while in
637 PSMS and SUMO only a single objective is evaluated. In our case, estimating
638 the model complexity through the VC-dimension implies to train and to test
639 a model a number of times (10 times, according to the parameter that we
640 used). Measuring the training error also implies to train and test such model.
641 Hence, evaluating both objectives involves training and testing the model.
642 This could represent a disadvantage with respect to the others, in terms of
643 computational cost. Notwithstanding, we can argue that the task of model
644 selection can be performed off-line. Moreover, since the models are in the
645 non-dominated set, several strategies for constructing a final classification
646 model can be performed without significantly increasing the computational
647 cost. In addition to this, our proposal (MOTMS) has the advantage of getting
648 highly competitive models, outperforming SUMO and PSMS in most of the
649 data sets.

650 *4.2.3. Discussion*

651 From the experimental results shown in Table 3, we can note how over-
652 fitting can be present in model selection. Among the three strategies for
653 constructing a final model, those based on ensembles proved being benefi-
654 cial, reducing the over-fitting effect. In spite of this, we cannot say that
655 ensemble approaches completely solve the problem. We can also note that in
656 most cases, the use of the solutions in the whole non-dominated front in an
657 ensemble achieved a better generalization performance than when a subset
658 of these are considered for the ensemble. This is a surprising result, since it

659 was expected that by taking into account the diversity as a criterion for the
660 ensemble construction, a better performance would be attained than when
661 not doing so. Observing the diversity between both approaches, we noted
662 that the whole non-dominated approach has better diversity, whereas the
663 ensemble of a subset of solutions approach gets trapped in a local optimal
664 solution.

665 A comparison with random forest (RF) showed the benefits of performing
666 model selection against not doing so. This is specially remarkable in the
667 ringnorm, splice, and twonorm data sets, in which an improvement above a
668 4% is reached. Even though a simple RF outperformed our proposal in the
669 image data set, a pairwise comparison did not show a statistical significant
670 difference between both. Therefore, it is worth performing the computational
671 effort in order to construct a reliable classification model.

672 The ensemble of the whole front of the proposed approach (MOMTS-S2)
673 significantly outperformed LS-SVM-BR on three benchmark data sets, but
674 it was significantly worse in one data set. The greatest improvement was
675 obtained in the splice data set, reducing the error rate in 6.07%. The great-
676 est degradation was on the twonorm data set, with a difference of 0.89%.
677 In spite of this, the overall performance of both approaches was similar.
678 Neither the reference nor the proposed approach were significantly better
679 than each other. It is interesting to note that MOMTS-S2 does not out-
680 perform LS-SVM-BR, which is a model selection method of the state of the
681 art. This is due to the fact that MOMTS-S2 deals with different model
682 types and their corresponding hyper-parameters. Nevertheless, we can argue
683 that in LS-SVM-BR there are only two parameters to be optimized, while in

684 MOMTS-S2, seven parameters are taken into consideration, which consider-
685 ably increase the search space and makes it harder to reach the “optimal”
686 solutions with a lower number of iterations. Moreover, we gain generality
687 without significantly over-fitting the models.

688 The experimental evaluations showed that MOMTS-S2 significantly out-
689 performed PSMS. Although there was not a statistical significant difference
690 between MOMTS-S2 and SUMO, when different data sets were considered,
691 MOMTS-S2 significantly outperformed SUMO on several data sets. This
692 gives evidence about the suitability of using a multi-objective approach in
693 contrast to a single-objective approach, in spite of the computational cost of
694 doing so. The experimental results showed that only minimizing the error
695 rate estimated through k -fold cross validation could lead to choose a model
696 with a small degree of over-fitting. In spite of this, the k -fold cross validation
697 approach has the advantage of being free from the model assumptions, which
698 makes it applicable to any learning algorithm and feasible to the full model
699 selection formulation⁶. On the other hand, the use of the VC-dimension for
700 controlling the model complexity, and avoiding over-fitting, as much as pos-
701 sible, also shows its potential for being applicable to different model types.
702 Therefore, we believe that this approach can also be applicable to the full
703 model selection formulation.

⁶The full model selection formulation consists of the task of finding the best combi-
nation of pre-processing, feature selection, and learning algorithms together with their
parameters [? ?].

704 5. Conclusions and Future Work

705 In this paper, we have proposed a multi-objective approach for dealing
706 with the problem of model selection. Our model selection approach takes
707 into account both the learning algorithm and the hyper-parameters during
708 the search process. We defined the training error, or empirical error, and
709 the model complexity, which is estimated through the VC dimension, as the
710 objectives to be optimized. The adopted formulation showed the following
711 advantages: (i) the experimental way for measuring the VC dimension allows
712 us to consider different learning algorithms in a general framework, and makes
713 the method applicable to the full model selection problem; (ii) our proposal
714 had a competitive performance over different benchmark data sets, making
715 it applicable to problems from diverse domains; and (iii) the multiple non-
716 dominated solutions obtained through the multi-objective formulation makes
717 it easy to extend it to ensembles of models.

718 The experimental results showed that constructing an ensemble of mod-
719 els performs better than choosing a single model. Furthermore, the ensemble
720 approach showed to be effective for reducing the effect of over-fitting. The ad-
721 vantages of the multi-objective approach over a single-objective formulation
722 such as PSMS were also supported by the experiments. The experimental re-
723 sults also show that highly competitive classification models were generated
724 by our proposal, without significantly degrading the performance in most
725 cases. Hence, we can conclude that our proposed approach can be an useful
726 framework for model selection in real world problems.

727 In the proposed approach, the VC dimension is experimentally estimated,
728 making it computationally expensive. Alternatives such as approximating

729 this value through surrogate-assisted evolutionary computation, or comput-
730 ing it using parallel computing would be interesting paths of future research.
731 Other future research directions also include the extension to the full model
732 selection problem, i.e., considering feature selection and data pre-processing
733 into the model selection process. Studying more effective ways for construct-
734 ing an ensemble (possibly) by using a second level of optimization would be
735 another interesting direction for future research.

736 **Acknowledgments**

737 The authors would like to show gratitude to the anonymous reviewers,
738 whose valuable comments helped to improve this paper. The first author is
739 grateful for the support from CONACyT scholarship no. 329013.