

SPMS-ALS: A Single-Point Memetic Structure with Accelerated Local Search for Instance Reduction

Hoang Lam Le^{a,*}, Ferrante Neri^a, Isaac Triguero^a

^a*The Optimisation and Learning (COL) Lab at the School of Computer Science,
University of Nottingham, Nottingham NG8 1BB, United Kingdom*

Abstract

Real-world optimisation problems pose domain specific challenges that often require an ad-hoc algorithmic design to be efficiently addressed. The present paper investigates the optimisation of a key stage in data mining, known as instance reduction, which aims to shrink the input data prior to applying a learning algorithm. Performing a smart selection or creation of a reduced number of samples that represent the original data may become a complex large-scale optimisation problem, characterised by a computationally expensive objective function, which has been often tackled by sophisticated population-based metaheuristics that suffer from a high runtime.

Instead, by following the Ockham's Razor in Memetic Computing, we propose a Memetic Computing approach that we refer to as fast Single-Point Memetic Structure with Accelerated Local Search (SPMS-ALS). Using the k -nearest neighbours algorithm as base classifier, we first employ a simple local search for large-scale problems that exploits the search logic of Pattern Search, perturbing an n -dimensional vector along the directions identified by its design variables one by one. This point-by-point perturbation mechanism allows us to design a strategy to re-use most of the calculations previously made to compute the objective function of a candidate solution. The proposed Accelerated Local Search is integrated within a single-point memetic framework and coupled with a resampling mechanism and a crossover. A thorough experimental analysis shows that SPMS-ALS, despite its simplicity, displays an excellent performance which is as good as that of the state-of-the-art while reducing up to approximately 85% of the runtime with respect

*Corresponding author

Email address: `hoang.le@nottingham.ac.uk` (Hoang Lam Le)

to any other algorithm that performs the same number of function calls.

Keywords: Memetic Algorithm, Pattern Search, Instance Reduction, Classification, Data Science, Ockham’s Razor in Memetic Computing

1. Introduction

Since their earliest definition [35, 34] Memetic Algorithms (MAs) were introduced to enhance upon the performance of algorithms, such as Genetic Algorithms and Simulated Annealing. Unlike the majority of other algorithms, MAs are not fixed to a specific structure but are flexible and thus versatile optimisation frameworks, see [38]. This flexibility is one of the main features of MAs which likely inspired numerous subsequent studies that shaped, over the past three decades, the field of Memetic Computing (MC).

By following the visionary ideas reported in [21] and the classification in [12], three groups/generations of MC approaches have been identified:

- **Simple Hybrids:** this group includes hybrid algorithms generated by two or more algorithms joined together in a synergistic manner. Usually, the algorithms of this type combine a global search and at least one local search. Some examples of successful hybridisations are reported in e.g. [30, 50, 31]
- **Adaptive Hybrids:** this includes hybrid algorithms where multiple local search algorithms are coordinated by an adaptive mechanism that selects the algorithmic elements at runtime. Popular selection criteria are performance-based like in hyperheuristics [44] and meta-Lamarckian learning [26, 43], diversity-based [6] or self-adaptive [42].
- **(Future) Memetic Automation:** this kind reinterprets MAs as a combination of “agents” without a predefined structure [1, 63] and investigates mechanisms to attain fully self-generated MAs. Although this design approach is still under investigation, some interesting domain-specific frameworks [15, 29] and prototypes [7] have been proposed.

The flexibility of the subject facilitating domain-specific algorithmic design is one of the reasons of the success of MAs in real-world applications, see [3, 14, 61]. In other words, while robust algorithmic design and testing on multiple abstract mathematical functions is fundamental for the development of novel memetic structures (as well as for any optimisation algorithm)

31 [17, 33] real-world problems often pose specific challenges which may be ad-
32 dressed by ad-hoc representations and specific operators [21]. Among the
33 plethora of MC structures the need to design simple algorithm on a limited
34 hardware inspired **Single-Point Memetic Structures** which are the focus
35 of the present study. For example, in [39] memetic structures using virtual
36 populations (statistical models of populations) have been implemented di-
37 rectly in the control cards of robots. In [9], a simplistic single-solution MC
38 approach composed of a global evolutionary operator and a local search has
39 been proven to be competitive with complex metaheuristics and has been
40 successfully implemented in the control card of an helicopter robot.

41 The latter approach is part a family of MAs designed according to the so-
42 called *Ockham's Razor in Memetic Computing* principle formulated in [22]:
43 simple algorithmic structures designed by combining memes in a bottom-up
44 approach while addressing the knowledge of the problem (prior or available
45 at run-time) often have a high performance despite their simplicity. This idea
46 links to other areas of optimisation research such as the pioneering studies
47 in [60, 11] and the work on Fitness Landscape Analysis [32, 23].

48 The present article addresses a real-world problem in the field of data
49 science, known as **instance reduction** [58]. Datasets can be extremely large
50 and usually require the use of pre-processing techniques to enable data mining
51 and machine learning techniques to learn from a cleaner and smaller dataset
52 that is free of noise, redundant or irrelevant samples (the so-called, *Smart*
53 *Data* [53]). Instance reduction is an important pre-processing procedure that
54 pursues to shrink the original dataset and keep it as informative as by either
55 selecting (**instance selection**) [19] or generating (**instance generation**)
56 [51] representative instances from a very large raw dataset. This is not a
57 trivial task, and it is essential to properly select or artificially generate those
58 representative instances.

59 Instance reduction can be conceived as an optimisation problem and be
60 tackled by search algorithms as either a binary search problem in the case
61 of instance selection [5], or as a continuous search problem to artificially
62 generate representative instances. In both cases, MAs have been preeminent
63 in comparison with other approaches in terms of performance [18, 52]. The
64 vast majority of the existing instance reduction approaches were proposed to
65 improve the performance of the well-known Nearest Neighbour (NN) classifier
66 [13]. However, the resulting reduced dataset may be used by any classifier
67 [5]. In this work, we will also focus on the NN classifier.

68 The main issue for current instance reduction solutions is related to the

69 high cost of evaluating candidate solutions. When tackling bigger datasets,
70 their runtime may become excessive and we can find in the specialised lit-
71 erature parallelisation approaches for instance reduction [55], which allow
72 them to be executed, whilst increasing the need for additional computa-
73 tional resources. Reducing the computational cost of the fitness evaluation
74 is an under-explored area in instance reduction, and just a few approximation
75 approaches exist (e.g. windowing [4] or surrogate models [41]).

76 Bearing in mind the elevated computational cost of the fitness function,
77 we propose a simple and yet effective domain-specific MC approach for in-
78 stance reduction. The proposed MC approach is composed of a novel domain-
79 specific implementation of local search hybridised with a global evolutionary
80 operator. The local search exploits the logic of the Generalised Pattern
81 Search that performs an implicit variable decomposition technique and per-
82 turbs the elements of a candidate solution one by one [40]. In contradistinc-
83 tion with existing population-based approaches that create new solutions
84 perturbing multiple variables at once, we exploit the fact that the proposed
85 local search produces candidate solutions that are only “slightly” different
86 w.r.t the previous fitness evaluation. Based on this fact, we devise a mech-
87 anism to drastically reduce the cost of the objective function when using
88 the NN algorithm as base classifier. The global search operator is a simple
89 resampling mechanism followed by crossover while an elite memory slot re-
90 tains the solution with the best performance. The key idea lies in keeping a
91 single-point approach to highly accelerate the objective function evaluation
92 while using a global operator to avoid getting stuck in local optima.

93 The remainder of this article is organised in the following way. Section 2
94 provides the background about instance reduction, formalises it as an optimi-
95 sation problem and provides an explanation why the problem is unavoidably
96 large scale and why calculation of the objective function is computationally
97 expensive. Section 3 describes and justifies the proposed method. Section
98 4 presents the experimental setup while Section 5 shows and discusses the
99 results. Finally, Section 6 provides the conclusion of this study.

100 **2. Problem Formulation and Challenges Associated with it**

101 In a supervised classification problem, the data is usually split into train-
102 ing (**TR**) and test (**TS**) sets. Each instance belongs to a class w , which
103 is known for **TR** and unknown for **TS**. Both datasets can be viewed as a

104 matrix in which instances \mathbf{I}_i are displayed on the rows whilst features \mathbf{f}_i are
 105 shown on the columns:

$$\mathbf{TR} = \begin{bmatrix} & \mathbf{f}_1 & \mathbf{f}_2 & \dots & \mathbf{f}_m \\ \mathbf{I}_1 & a_{11} & a_{12} & \dots & a_{1m} \\ \mathbf{I}_2 & a_{21} & a_{22} & \dots & a_{2m} \\ \dots & \dots & \dots & \dots & \dots \\ \mathbf{I}_l & a_{l1} & a_{l2} & \dots & a_{lm} \end{bmatrix} \quad (1)$$

106 The main purpose of an instance reduction technique is to clean and
 107 compress \mathbf{TR} into a reduced set \mathbf{RS} , by either selecting or generating new
 108 representative instances, so that, it preserves and provides valuable infor-
 109 mation for a machine learning algorithm to learn useful insights about a
 110 classification problem. Thus, the resulting \mathbf{RS} should satisfy several condi-
 111 tions such as well-representing the distributions of the classes, significantly
 112 reducing in size to minimise the required storage, which would be beneficial
 113 to the posterior classification phase.

$$\mathbf{RS} = \begin{bmatrix} & \mathbf{f}_1 & \mathbf{f}_2 & \dots & \mathbf{f}_m \\ \mathbf{I}_1 & b_{11} & b_{12} & \dots & b_{1m} \\ \mathbf{I}_2 & b_{12} & b_{22} & \dots & b_{2m} \\ \dots & \dots & \dots & \dots & \dots \\ \mathbf{I}_p & b_{p1} & b_{p2} & \dots & b_{pm} \end{bmatrix} \quad (2)$$

114 with $p \ll l$. In this study we choose to treat p as a parameter that
 115 signifies the compression of the data with respect to the size of the entire
 116 set of training set (number of rows of \mathbf{TR}). More specifically, we use the
 117 reduction rate $\frac{l}{p}$ as a parameter of our problem. In both matrices \mathbf{TR} and
 118 \mathbf{RS} each row is associated with its class label, that is each instance \mathbf{I}_i is
 119 assigned to its class on the basis of its features.

120 2.1. Evaluation of an \mathbf{RS}

121 The development of many data pre-processing techniques such as instance
 122 reduction was initially motivated by the imprecision and inefficiencies of the
 123 well-known nearest neighbour(s) (NN) algorithm [13]. These weaknesses have
 124 turned into strengths and made the NN rule a core algorithm to preprocess
 125 raw data [53]. Thus, most instance reduction techniques verify how well a
 126 candidate matrix \mathbf{RS} represents the entire training dataset, \mathbf{TR} , by using the
 127 NN algorithm as base classifier. To do so, this approach essentially checks

Distance matrix		1	2	3	---	p
	1	0.93	0.22	0.35		1.03
	2					
	3					

	l	0.21	1.08	1.08	---	1.50

Figure 1: Distance matrix of l instances in **TR** and p instances in **RS**. The instance at the first row is verified by instance at column 2, while the instance at the last row is checked by the one at column 1. Blue entries represent the shortest distance among the neighbours.

128 how well we can classify, the large dataset **TR** using the small dataset **RS**
 129 as training data, and consists of the following steps. The Euclidean distance
 130 between each instance \mathbf{I}_i (row vector) of **RS** and each instance \mathbf{I}_j (row vec-
 131 tor) of **TR** is calculated. This process yields $l \times p$ distance computations.
 132 Typically, the nearest neighbours (smallest distances) are computed “on the
 133 fly”, just by keeping the shortest distance and instance ID/number, and any
 134 intermediate distance computations are disregarded. As part of the strategy
 135 we will devise in Section 3.2, we could store all computed distances on a
 136 *distance matrix*; Figure 1 shows an example of a distance matrix. An entry
 137 $D_{i,j}$ of the distance matrix in position i, j indicates the distance of the i^{th}
 138 instance in **TR** to the j^{th} instance in the **RS**:

$$D_{i,j} = \sqrt{(b_{i,1} - a_{j,1})^2 + (b_{i,2} - a_{j,2})^2 + \dots + (b_{i,m} - a_{j,m})^2}.$$

139 When the distance matrix \mathbf{D} is calculated, for each row (i.e. each instance
 140 of **TR**), the smallest entry is detected, e.g. 0.22 in the first row of Figure
 141 1, and that instance is given the class label w of the closest instance in **RS**.
 142 When all instances in **TR** have been classified, there are different ways to
 143 assign a score (objective function) to the performance of **RS** [2, 59].

As we are dealing with mostly balanced datasets, in this study we use the Accuracy Rate Acc [2, 59], that is

$$Acc = \frac{\text{number of correct classifications by means of RS}}{\text{total number of examined samples}}$$

Thus, for an input **RS** the objective function value is Acc , that is

$$f(\mathbf{RS}) = Acc.$$

144 Algorithm 1 describes step-by-step the calculation of the objective func-
145 tion based on the distance matrix.

Algorithm 1 Objective Function

```
1: INPUT matrices  $\mathbf{TR} = [a_{i,j}]$  and  $\mathbf{RS} = [b_{i,j}]$ 
2: Build the matrix of Euclidean distances  $\mathbf{D} = [D_{i,j}]$ 
3: for each row of the matrix  $\mathbf{D}$  do
4:   Find the smallest number and save its row and column indices
5:   Select, from  $\mathbf{TR}$  and  $\mathbf{RS}$ , the instances corresponding to the calcu-
   lated indices
6:   Check the corresponding labels
7:   if the labels coincide then
8:     Update the number of correct classifications
9:   end if
10: end for
11: Calculate  $Acc$ 
12: OUTPUT the objective function value  $Acc$ 
```

146 Of course, since higher values of Acc correspond to a better classification,
147 the objective function needs to be maximised. The maximisation occurs
148 within a $(p \times m)$ -dimensional space where each variables can continuously
149 vary in a normalised interval. Hence, the search for the optimal solution
150 occurs in the set $[0, 1]^{p \times m}$.

151 2.2. Computational Cost of the Objective Function

152 Regardless of the score used to measure the quality of \mathbf{RS} , the procedure
153 described above requires the calculation of $l \times p$ Euclidean distances. This
154 operation can be computationally expensive especially when large datasets
155 are under examination. The two main problems that arise when tackling
156 larger datasets are runtime (due to the large number of distance computations
157 required) and memory consumption (e.g. when the size of \mathbf{TR} does not allow
158 us to store it in main memory). However, the required runtime to evaluate
159 candidate solutions tends to be the most important factor to enable instance
160 reduction of large datasets.

161 In the literature, two popular types of approach to accelerate the process-
162 ing of instance reduction techniques are:

- 163 • **divide-and-conquer:** the execution of instance reduction approaches
164 is parallelised, splitting the training data into a number of chunks, typ-
165 ically through big data technologies, see [55]. Whilst they are necessary
166 when the **TR** set does not fit in main memory, the main limitation of
167 this approach is that it does not address the computational complexity
168 of the problem, but its processing time, by using additional computa-
169 tional resources. In addition, a trade-off between the number of splits
170 and the accuracy that can be obtained exist, and must be experimen-
171 tally found for the dataset at hand.
- 172 • **approximation:** to reduce the complexity of the objective function,
173 the quality of **RS** may also be estimated by an approximation function.
174 For example, a windowing approach that uses a different partition of
175 **TR** to evaluate an **RS** at each iteration of a search algorithm [4]. Other
176 more sophisticated approximation (also known as surrogate) models to
177 reduce the number of evaluations for instance reduction algorithms have
178 been recently investigated [25, 41]. While this approach reduces the
179 runtime, its main limitation is that an approximated objective function
180 may mislead the search of the optimisation algorithm.

181 In the present paper, we propose a new mechanism that while exploiting
182 the structure of the optimisation algorithm allows a substantial reduction of
183 the computational complexity (i.e. number of distance computations) of the
184 objective function without approximations, see Section 3. Thus, the goal of
185 this work is not to tackle big datasets and the memory limitations associated
186 to it, but to devise a very fast and reliable instance reduction process that
187 could be couple together with the approaches provided in [55] when very big
188 datasets need to be addressed.

189 **3. Single-Point Memetic Structure with Accelerated Local Search** 190 **for Instance Reduction**

191 From the description in Section 2, we may characterise Instance Reduc-
192 tion as an optimisation problem with the following considerations:

- 193 • the problem is large-scale and its number of variables ($p \times m$) can be
194 extremely high depending on the size of the dataset
- 195 • due to the large number of variables, the problem is likely to be hard
196 to solve and the fitness landscape could be highly multimodal

- 197 • even if it were multimodal, an excessive exploitation of the basin of
198 attraction may yield an overfitted solution, that is a solution that per-
199 forms well on the training set but not on the test set
- 200 • each objective function call (or fitness evaluation) is computationally
201 expensive due to calculation of multiple Euclidean distances

202 In order to address the Instance Reduction problem, a domain-specific
203 MC approach that takes into account the considerations above is here pro-
204 posed. The proposed MC approach, namely Single-Point Memetic Structure
205 with Accelerated Local Search (SPMS-ALS) is population-less and designed
206 according to the bottom-up logic reported in [22]. SPMS-ALS perturbs a
207 single solution and makes use of one more memory slot to store the elite
208 solution, that is the best solution ever found. A novel domain-specific accel-
209 erated local search implementation is here proposed. Section 3.1 describes
210 the local search operator employed in SPMS-ALS while Section 3.2 illustrates
211 how the local search logic is exploited to accelerate the calculation of the ob-
212 jective function. The proposed SPMS-ALS makes also use of a simple global
213 search operator illustrated in Section 3.3. Finally, Section 3.4 discusses and
214 justifies the design of SPMS-ALS.

215 *3.1. Local Search Operator*

216 With the purpose of effectively describing SPMS-ALS, let us slightly re-
217 define the notation. As introduced in Equation 2, $\mathbf{RS} = [b_{i,j}]$ is a matrix
218 of size $p \times m$ which can be rewritten as a vector \mathbf{x} of length $n = p \times m$
219 containing all the rows of \mathbf{RS} arranged sequentially:

$$\mathbf{x} = (b_{11}, b_{12}, \dots, b_{1m}, b_{21}, b_{22}, \dots, b_{2m}, \dots, b_{p1}, b_{p2}, \dots, b_{pm}) = (x_1, x_2, \dots, x_n)$$

220 where x_i represents the design variables of the optimisation problem.

Let \mathbf{e}^i be the i^{th} versor (that is a vector of modulus equal to 1) of a
basis in an n -dimensional space, that is a vector whose elements are all zeros
except from the i^{th} element which is one [37]:

$$\mathbf{e}^i = (0, 0, \dots, 1, \dots, 0, 0)$$

221 The local search works on the candidate solution \mathbf{x} to locally improve it.
222 The following greedy implementation of a Generalised Pattern Search has

223 been used, see [40]. The algorithm perturbs each feature value of an instance
 224 at a time in its feasible range and then check if any improvement is found.
 225 Specifically, let \mathbf{x} be the base vector (the best solution found at the time),
 226 for each design variable i from 1 to n the algorithm explores at first

$$\mathbf{x}^t = \mathbf{x} - \rho \cdot \mathbf{e}^i$$

where \mathbf{x}^t is a trial vector and the scalar ρ is the step-size (exploratory radius).
 For each index i , the algorithm attempts to explore the opposite orientation
 of the direction identified by \mathbf{e}^i if the first attempt fails, that is

$$\mathbf{x}^t = \mathbf{x} + \frac{\rho}{2} \cdot \mathbf{e}^i$$

227 As a remark, the asymmetric step-size is designed to avoid to revisit
 228 the same solution (vector), see [40]. As soon as \mathbf{x}^t outperforms \mathbf{x} , that is
 229 $f(\mathbf{x}^t) \geq f(\mathbf{x})$, the trial vector \mathbf{x}^t replaces the base vector \mathbf{x} .

Note that when applying the above perturbations, the resulting values
 in the vector x^t could be outside of the bounds $[x_{low}, x_{high}]$. On the basis
 of preliminary tests we employed a toroidal handling of the bounds, i.e. for
 $x_i \in [x_{low}, x_{high}]$, if $x_i > x_{high}$ it is reinserted by reassignment:

$$x_i = x_{low} + \left((x_i - x_{high}) \lfloor \frac{(x_i - x_{high})}{(x_{high} - x_{low})} \rfloor (x_{high} - x_{low}) \right)$$

230 while if $x_i < x_{low}$ it is reinserted by reassignment

$$x_i = x_{high} - \left((x_{low} - x_i) - \lfloor \frac{(x_{low} - x_i)}{(x_{high} - x_{low})} \rfloor (x_{high} - x_{low}) \right)$$

231 where the parentheses $\lfloor \cdot \rfloor$ indicate the truncation to the lower integer.

232 As an example, if we are in the range $[0,1]$, and the resulting value x_i of
 233 the perturbation is 1.1, the toroidal handling will begin from the beginning of
 234 the range, producing a 0.1. Conversely, if x_i were to be below 0, e.g. -0.1, this
 235 circular handling would provide 0.9. This ensures that the investigated values
 236 are within the range. Also, by forcing the perturbation to go to the other
 237 side of the bound, we increase the exploratory abilities of the method before
 238 reducing the radius ρ . This strategy provided good results in preliminary
 239 tests in comparison with other alternatives. If after the entire exploration
 240 along the n directions no improved solution \mathbf{x}^t is found, then the radius ρ
 241 is reduced by a reduction rate. The local search is interrupted when either a

242 budget condition is met or when the radius ρ is smaller than a pre-arranged
 243 precision. For sake of clarity Algorithm 2 shows the local search operator
 244 used in SPMS-ALS.

Algorithm 2 Local Search of the family of Pattern Search used by SPMS-ALS

```

1: INPUT  $\mathbf{x}$ 
2: while local budget and precision conditions are not met do
3:    $\mathbf{x}^t = \mathbf{x}$ 
4:   for  $i = 1 : n$  do
5:      $\mathbf{x}^t = \mathbf{x} - \rho \cdot \mathbf{e}^i$ 
6:     Apply toroidal handling of the bounds
7:     if  $f(\mathbf{x}^t) \geq f(\mathbf{x})$  then
8:        $\mathbf{x} = \mathbf{x}^t$ 
9:     else
10:       $\mathbf{x}^t = \mathbf{x} + \frac{\rho}{2} \cdot \mathbf{e}^i$ 
11:      Apply toroidal handling of the bounds
12:      if  $f(\mathbf{x}^t) \geq f(\mathbf{x})$  then
13:         $\mathbf{x} = \mathbf{x}^t$ 
14:      end if
15:    end if
16:  end for
17:  if  $\mathbf{x}$  has not been updated then
18:    reduce  $\rho$ 
19:  end if
20: end while
21: RETURN  $\mathbf{x}$ 

```

245 *3.2. Accelerated Local Search*

The proposed local search makes use of the search logic outlined in Algorithm 2 and integrates within it a domain-specific procedure to reduce the computational time of the algorithm. As highlighted in Section 2, when the NN algorithm is used as based classifier, most of the high computational cost of the Instance Reduction problem is due to the calculation of $l \times p$ Euclidean distances. However, the local search moves

$$\mathbf{x}^t = \mathbf{x} - \rho \cdot \mathbf{e}^i$$

and

$$\mathbf{x}^t = \mathbf{x} + \frac{\rho}{2} \cdot \mathbf{e}^i$$

246 affect only one design variable that is only one entry of the **RS** matrix.

247 As a consequence, if we build a distance matrix **D** associated with \mathbf{x}^t ,
248 this differs by only one column from the matrix **D** associated with \mathbf{x} . When
249 the objective function $f(\mathbf{x}^t)$ is calculated according to Algorithm 1, there is
250 no need to recompute $l \times p$ Euclidean distances since $l \times (p - 1)$ elements
251 have already been computed and appropriately stored.

252 Thus, when Algorithm 2 is applied, each objective function call requires
253 the calculation of only l Euclidean distances. This fact can be effectively rep-
254 resented as the modified objective function used by the local search outlined
255 in Algorithm 3.

Algorithm 3 Objective Function $f(\mathbf{x}^t)$ of the Accelerated Local Search

- 1: **INPUT** matrix **TR** = $[a_{i,j}]$, matrix **D** associated with the base vector \mathbf{x} , and trial vector \mathbf{x}^t
 - 2: Build the matrix **RS** = $[b_{i,j}]$ from \mathbf{x}^t
 - 3: Update the matrix of Euclidean distances **D** = $[D_{i,j}]$ by recalculating the l elements of the pertinent column
 - 4: **for** each row of the matrix **D** **do**
 - 5: Find the smallest number and save its row and column indices
 - 6: Select, from **TR** and **RS**, the instances corresponding to the calculated indices
 - 7: Check the corresponding labels
 - 8: **if** the labels coincide **then**
 - 9: Update the number of correct classifications
 - 10: **end if**
 - 11: **end for**
 - 12: Calculate Acc
 - 13: **OUTPUT** the objective function value Acc
-

256 Our proposed local search performs once at the beginning the objective
257 function call as in Algorithm 1 and then integrates Algorithm 3 into each
258 $f(\mathbf{x}^t)$ function call for the rest of its execution.

259 *3.3. Evolutionary Global Search Operator*

260 At the beginning of the optimisation, a matrix \mathbf{RS} (i.e. Equation 2) is
 261 randomly sampled from the matrix \mathbf{TR} (i.e. Equation 1) and from \mathbf{RS} the
 262 corresponding base vector \mathbf{x} constructed and inputted into the local search
 263 operator. The local search is continued until the stopping criteria conditions
 264 on budget and precision are met. The local search returns a (possibly im-
 265 proved) solution \mathbf{x} . The best solution ever found is saved and stored in an
 266 elite slot and called \mathbf{x}^e . Then, a new solution \mathbf{x}^r is generated by randomly
 267 sampling a new \mathbf{RS} matrix from \mathbf{TR} and constructing the corresponding
 268 vector. A uniform crossover is applied to \mathbf{x}^r and \mathbf{x} to generate a new trial
 269 vector \mathbf{x}^t .

In order to explain the functioning of this crossover, let us consider a
 candidate solution \mathbf{x} and let us remind it that it corresponds to a matrix \mathbf{RS}
 whose rows are instances and columns are features. By applying a matrix
 partitioning we may represent \mathbf{RS} as a vector of row vectors

$$\mathbf{RS} = \begin{bmatrix} \mathbf{I}_1 \\ \mathbf{I}_2 \\ \dots \\ \mathbf{I}_p \end{bmatrix}$$

Similarly, we may consider the random solution \mathbf{x}^r and represent the
 corresponding \mathbf{RS}^r matrix

$$\mathbf{RS}^r = \begin{bmatrix} \mathbf{I}_1^r \\ \mathbf{I}_2^r \\ \dots \\ \mathbf{I}_p^r \end{bmatrix}$$

270 whose instances are randomly selected from \mathbf{TR} .

271 The proposed crossover generates a trial vector \mathbf{x}^t by randomly selecting
 272 some rows from \mathbf{RS} and some rows from \mathbf{RS}^r . Each row of the resulting
 273 matrix \mathbf{RS}^t has a gene-resampling probability Gr to be selected from \mathbf{RS}
 274 and $1 - Gr$ probability to be selected from \mathbf{RS}^r . It must be remarked that
 275 a crossover that perturbs single elements of \mathbf{RS} instead of entire rows would
 276 yield a candidate solution which could be noisy (i.e. not have the right class
 277 label), and therefore, not meaningful from a classification point of view.

278 The gene-resampling probability Gr expresses the rate of the instances
 279 in \mathbf{RS} which are replaced by other instances sampled from \mathbf{TR} . Algorithm
 280 4 describes the crossover mechanism.

Algorithm 4 Crossover between \mathbf{x} and \mathbf{x}^r

```
1: INPUT base vector  $\mathbf{x}$  and random vector  $\mathbf{x}^r$ 
2: Build the matrices  $\mathbf{RS} = [\mathbf{I}_i]$  and  $\mathbf{RS}^r = [\mathbf{I}_i^r]$ 
3:  $\mathbf{RS}^t = [\mathbf{I}_i^t] = \mathbf{RS}$ 
4: for  $i = 1 : p$  do
5:   Generate a random number  $rand$ 
6:   if  $rand < Gr$  then
7:      $\mathbf{I}_i^t = \mathbf{I}_i^r$ 
8:   end if
9: end for
10: From  $\mathbf{RS}^t$  calculate  $\mathbf{x}^t$ 
11: OUTPUT the trial vector  $\mathbf{x}^t$ 
```

281 The local and global search operators are repeated until the global budget
282 conditions are met. The framework of the proposed SPMS-ALS is illustrated
283 in Algorithm 5.

284 3.4. Algorithmic Design

285 The proposed SPMS-ALS follows a bottom-up strategy as suggested in
286 [22]: we implemented within the algorithmic operators the necessary coun-
287 termeasures to address each challenge associated with the problem.

288 The structure of the local search has been selected to address the large
289 scale nature of the instance reduction problem, that is for a large dataset, the
290 matrix \mathbf{RS} can easily have hundreds if not thousands of rows. The proposed
291 local search perturbs the variables separately and thus implicitly performs a
292 variable decomposition. Approaches of this type have been proved effective
293 for large scale problems, see [47, 56, 28].

294 This observation was reported in the experimental study in [8]. Large
295 scale problems are by no means easier than low-dimensional problems. How-
296 ever, since in practice the computational budget cannot grow exponentially
297 with the problem dimensionality only a very limited portion of the decision
298 space is explored. Under these experimental conditions, the algorithm “sees”
299 the problem as separable: average Pearson and Spearman coefficients of the
300 variables approach zero independently on the problem when the number of
301 dimensions grows, see [8].

302 The high computational cost of each function call is addressed by the
303 acceleration mechanism outlined above: only the elements of one column of

Algorithm 5 Framework of the SPMS-ALS for Instance Reduction

- 1: Randomly generate a base vector \mathbf{x} in $[0, 1]^n$ and calculate $f(\mathbf{x})$ according to Algorithm 1
 - 2: Assign the elite $\mathbf{x}^e = \mathbf{x}$
 - 3: **while** global budget conditions are met **do**
 - 4: Apply the Accelerated Local Search to the base vector \mathbf{x} according to Algorithm 2 with the objective function $f(\mathbf{x}^t)$ calculated according to Algorithm 3
 - 5: **if** $f(\mathbf{x}) \geq f(\mathbf{x}^e)$ **then**
 - 6: Update the elite $\mathbf{x}^e = \mathbf{x}$
 - 7: **end if**
 - 8: Randomly generate a vector \mathbf{x}^r in $[0, 1]^n$
 - 9: Apply Crossover between \mathbf{x} and \mathbf{x}^r according to Algorithm 4 and generate a new trial vector \mathbf{x}^t
 - 10: Calculate $f(\mathbf{x}^t)$ according to Algorithm 1
 - 11: **if** $f(\mathbf{x}^t) \geq f(\mathbf{x}^e)$ **then**
 - 12: Update the elite $\mathbf{x}^e = \mathbf{x}^t$
 - 13: **end if**
 - 14: Assign $\mathbf{x} = \mathbf{x}^t$
 - 15: **end while**
-

304 the Euclidean matrix \mathbf{D} and not those of the entire matrix are calculated
305 at each function call. The population-less structure of SPMS-ALS has also
306 been chosen taking into consideration the computational cost. The proposed
307 SPMS-ALS naturally devotes most of the computational budget (in terms
308 of function calls) to the local search. On the contrary, the global search
309 operator performs only sporadic function calls. This logic perfectly suits the
310 needs of reducing the computational cost since the global operator requires
311 the expensive objective function as in Algorithm 1 the local search operator
312 uses its computationally cheaper version as in Algorithm 3.

313 In order to address the multimodality of the fitness landscape and pre-
314 vent that the algorithm converges to a suboptimal solution, we combined the
315 Accelerated Local Search with the simplistic global search described above.
316 It must be noted that the global search makes use of part of decision vari-
317 ables (genotype) of previously improved solutions. The elitism guarantees
318 that previously detected promising solutions are available at the end of the
319 run. Furthermore, the gene-resampling mechanism, happening at the in-

320 stance level (considering the rows as building blocks) complements the local
321 search that happens at the level of the elements of **RS**.

322 At last, the restarting local search logic combined with a limited local
323 search budget is an important countermeasure to prevent from overfitting: an
324 excessive local search budget is likely to yield an overly specialised solution
325 that performs poorly when the solution is tested on a new dataset. This
326 characteristic is experimentally analysed in Section 5.3.

327 4. Experimental Framework

328 This section presents the used datasets (Section 4.1) and is followed by
329 introducing several instance reduction techniques that will be used for com-
330 parison with our proposal (Section 4.2). Finally, the parameter configuration
331 is explained (Section 4.3).

332 4.1. Datasets

333 In the experimental study, we have examined 40 small and 17 medium
334 multi-class datasets from the KEEL dataset repository [54]. The properties
335 of these datasets including name (**Dataset**), the number of samples (**Samp**),
336 the number of attributes (**Att**), the number of classes (**%Class**) are sum-
337 marised in Table 1.

338 As defined in [51], small datasets have less than 2000 instances while
339 medium datasets have at least 2000 instances. Each dataset is partitioned
340 using a 10-fold stratified cross-validation (10-fcv) procedure, see [45]. Thus,
341 the performance of each dataset is reported by an average of the 10 folds. All
342 of the experiments with these datasets have been conducted on computers
343 at which each has 2 x 20 core processors (Intel Xeon Gold 6138 20C 2.0GHz
344 CPU) and 192 GB RAM.

345 4.2. Comparison Algorithms

346 In order to understand the benefits of the proposed MC approach, we
347 first define two baselines:

- 348 • Nearest Neighbour (1NN): we use the NN algorithm ($k=1$) employing
349 the entire **TR** for training, without any pre-processing. The perfor-
350 mance of the NN in **TR** is calculated following a *leave-one-out* vali-
351 dation scheme. This serves of a baseline to understand the benefits of
352 instance reduction.

- Local Search Instance Reduction (LSIR): the local search presented and used in [41]. LSIR is essentially the basic pattern search shown in Algorithm 2 without any acceleration, that is the local search by using the basic fitness function as in Algorithm 1.

Table 1: Summary description for small (Sample < 2000) and medium (Sample >= 2000) datasets.

Dataset	Samp	Att	Class	Dataset	Samp	Att	Class
Abalone	4174	8	28	Monks	432	6	2
Appendicitis	106	7	2	Movement_libras	360	90	15
Australian	690	14	2	Newthyroid	215	5	3
Autos	205	25	6	Nursery	12960	8	5
Balance	625	4	3	Page-blocks	5472	10	5
Banana	5300	2	2	Penbased	10992	16	10
Bands	539	19	2	Phoneme	5404	5	2
Breast	286	9	2	Pima	768	8	2
Bupa	345	6	2	Ring	7400	20	2
Car	1728	6	4	Saheart	462	9	2
Chess	3196	36	2	Satimage	6435	36	7
Cleveland	297	13	5	Segment	2310	19	7
Contraceptive	1473	9	3	Sonar	208	60	2
Crx	125	15	2	Spambase	4597	57	2
Dermatology	366	33	6	Spectheart	267	44	2
Ecoli	336	7	8	Splice	3190	60	3
Flare-solar	1066	9	2	Tae	151	5	3
German	1000	20	2	Texture	5500	40	11
Glass	214	9	7	Thyrod	7200	21	3
Haberman	306	3	2	Tic-tac-toe	958	9	2
Hayes-roth	133	4	3	Titanic	2201	3	2
Heart	270	13	2	Twonorm	7400	20	2
Hepatitis	155	19	2	Vehicle	846	18	4
Housevotes	435	16	2	Vowel	990	13	11
Iris	150	4	3	Wine	178	13	3
Led7digit	500	7	10	Wisconsin	683	9	2
Lymphography	148	18	4	Yeast	1484	8	10
Magic	19020	10	2	Zoo	101	16	7
Mammographic	961	5	2				

In addition, we test the performance of the proposed approach against the current state-of-the-art in instance reduction. SPMS-ALS belongs to the family of positioning adjustment methods (see [51]), which are, to date, the best performing instance reductions methods in the literature and follow a similar algorithmic structure to the proposed approach. In [41], we showed the classification performance of the local search against the entire family

363 of positioning adjustment methods. For the sake of simplicity, here we only
364 report the comparison against the most competitive methods. SPMS-ALS
365 can be categorised as a pure instance generation approach, as we perform
366 a continuous search. Thus, we choose the following metaheuristics instance
367 generation methods to compete against SPMS-ALS:

- 368 • Scale Factor Local Search Differential Evolution (SFLSDE): this memetic
369 approach optimises the positioning of prototypes using an implemen-
370 tation of differential evolution [52].
- 371 • Particle Swarm Optimisation (PSO): this algorithm modifies the posi-
372 tion of an initial set using PSO rules, aiming to maximise the classifi-
373 cation performance [36].

374 Additionally, to compare against more recent meta-heuristics, we have
375 adapted a recent metaheuristic, proposed for the continuous domain, to
376 tackle instance reduction.

- 377 • Linear Population Size Reduction of the Success-History based Adap-
378 tive Differential Evolution (LSHADE) [49]: this approach is developed
379 from Success-History based Adaptive Differential Evolution (SHADE)
380 [48] and Adaptive Differential Evolution with Optional External Archive
381 JADE [62]. It makes use of success-history and also applies the popula-
382 tion size reduction to progress the search. Note that this metaheuristic
383 has not been previously used for instance reduction, but due to its simi-
384 larity to JADE, we used the design ideas from [52] to adapt it to solve
385 the instance reduction problem.

386 These approaches evolve a population of solutions, whilst our method only
387 evolves a single solution (or more precisely two solutions the trial solution
388 \mathbf{x}^t and the elite \mathbf{x}^e). However, similar to our method, both approaches
389 start off from a random (stratified) subset of the training set \mathbf{TR} (one for
390 each individual of their population), which keeps the original distribution of
391 instances per class. Thus, the reduction rate is also defined by a parameter
392 that determines how much we want to reduce \mathbf{TR} .

393 As a further remark, while PSO, SFLSDE were existing metaheuristics
394 that have been adapted to solve the instance generation problem, the pro-
395 posed SPMS-ALS has been expressly designed to solve this problem effec-
396 tively in terms of accuracy efficiently in terms of runtime. This design ap-
397 proach follows the bottom-up design logic of MC [22, 38] and can be ob-
398 served in both accelerated local search logic and crossover operator. Another

399 remark is that LSHADE was used in the continuous domain to solve bench-
400 mark functions, this metaheuristic design is first time adapted to solve the
401 instance generation problem in this study.

402 In [52], the authors showed that using a random selection as initialisation
403 mechanism is not usually appropriate, and the hybridisation of an instance
404 selection step followed by instance generation was suggested to replace this
405 random initialisation. More specifically, the use of a Steady-State Memetic
406 Algorithm (SSMA) [18] demonstrated empirically to provide an excellent
407 starting point, which means a good selection of instances per class and a good
408 reduction rate (automatically determined by the instance selection step). To
409 the best of our knowledge, the hybrid instance reduction algorithm, SSMA-
410 SFLSDE [52], remains the best performing method for Instance Reduction in
411 both accuracy and reduction rate. To establish a fair comparison against it,
412 we will also hybridise the proposed MC approach and the local search with
413 SSMA (see Section 5.5).

414 Whilst the hybrid instance selection/generation method, SSMA-SFLSDE
415 has not been outperformed to date, in order to assess the potential of the
416 proposed approach, we add a comparison with recently published algorithms
417 belonging to the family of instance selection. We included an approach based
418 on local sets [27] and a method based on instance ranking [10]. Note that
419 these methods follow a completely different approach to produce a reduced
420 set from the training set. Thus, we cannot set up the same computational
421 budget that we do for the rest of the comparison algorithms, as they do not
422 follow an optimisation-based approach (Section 4.3.1).

423 The study in [27] contains three instance selection methods, namely Lo-
424 cal Set Smoother (LSSm), Local Set Core (LSCo) and Local Set Border
425 (LSBo). LSSm aims at achieving the highest accuracy regardless of the re-
426 duction, while LSCo seeks at obtaining the highest reduction with acceptable
427 accuracy. LSBo addresses both accuracy and reduction rate with the same
428 priority. For this reason, LSBo has been selected for comparison against the
429 proposed memetic approach.

430 The main idea of [10] is to exploit the relationship among members in the
431 training set by computing a rank for each element. A rank of an instance
432 introduces the correlation between itself and others in the training set. In-
433 stances with higher ranks are likely to be selected compared to those holding
434 a low rank. In Section 5.4.2, we report the performance of Ranking-based In-
435 stance Selection (RIS1) as it showed in [10] to display the best performance,
436 among multiple variants, w.r.t both measures of accuracy and reduction rate.

437 *4.3. Parameter Settings*

438 This section presents the parameter configuration for all the methods
439 employed in this study, including the accelerated local search outlined in
440 Algorithms 2 and 3 and the entire memetic framework SPMS-ALS shown in
441 Algorithm 5. Subsection 4.3.1 focuses on the computational budget while
442 Subsection 4.3.2 discusses the other parameters.

443 *4.3.1. Computational Budget*

444 In the previous studies on instance reduction, the computational budget
445 has usually been set empirically to a number of iterations (in search-like algo-
446 rithms), which remained fixed for all datasets [52, 41]. Just like in scalability
447 studies for ordinary optimisation problems, in the instance reduction prob-
448 lem, the complexity of the search space grows exponentially with the problem
449 size, [8]. On the other hand, instance reduction poses a further challenge that
450 is the risk of overfitting and underfitting. An incorrect local search budget
451 allocation is likely to lead to overfitting in small datasets and underfitting
452 appears in larger datasets. In order to overcome this challenge and propose
453 a standard for setting the computational budget, we have here conducted an
454 extensive experimental study. Note that keeping the same number of eval-
455 uations through the different comparison methods will also help establish a
456 fairer comparison (which has not been the case in previous studies).

457 Among the various properties of a dataset, the number of instances in
458 training data, i.e. the number of rows l of the matrix **TR** and the number
459 of features (**Features**) in an instance at each dataset are the two important
460 factors that define the size of the problem and need to be considered when the
461 budget is allocated. We acknowledge that other factors may be also required
462 into consideration such as the number of classes or the ratio of samples among
463 classes. However, this simple yet effective approach of parameter setting has
464 proven to highly reduce the unnecessary allocated number of evaluations and
465 thus can help mitigate overfitting in small datasets and underfitting in larger
466 ones. Since RIS1 and LSBo do not perform any evaluation of their reduced
467 set against the training set, we cannot apply a computational budget.

468 In the original setting based on forty small datasets, SFLSDE [52] and
469 LSIR [41] use approximately 20,000 and 30,000 evaluations, respectively. We
470 took these values as a reference and set three levels in our experimental study:
471 lower, comparable, and greater than the reference ones. Table 2 displays, for
472 all the algorithms considered in this study that employ local search, the
473 three local search budgets scenarios. From the total number of evaluations

474 presented in Table 2, we split the evaluations into two parts when SSMA is
 475 included. SSMA takes $3 \times l$ evaluations in Setting 1 and $5 \times l$ in Setting 2 and
 476 3. The budget allocated to SSMA is indicated as SSMA_Eval. The rest of
 477 the evaluations is used for instance generation methods (LSIR, SPMS-ALS,
 478 PSO, SFLSDE, LSHADE).

Table 2: Changing the number of evaluation considering training size and features for fairer comparison.

Computational Budget			
Algorithm	Setting 1	Setting 2	Setting 3
SSMA	SSMA_Eval = $3 \times l$	SSMA_Eval = $5 \times l$	SSMA_Eval = $5 \times l$
SFLSDE			
LSIR	$2.5 \times \mathbf{Features} \times l$	$5 \times \mathbf{Features} \times l$	$10 \times \mathbf{Features} \times l$
SPMS-ALS	– SSMA_Eval	– SSMA_Eval	– SSMA_Eval
PSO			
LSHADE			

479 4.3.2. Parameters

480 The proposed SPMS-ALS contains some parameters to set to coordinate
 481 global and local search. In particular, the following parameters are funda-
 482 mental to coordinate the interruption of accelerated local search and restart
 483 of global search.

- 484 • N_{\max} : the maximum number of times the local search accepts a new
 485 trial solution \mathbf{x}^t with the same objective function (Acc) as that of the
 486 previous trial solution i.e. maximum number of search moves allowed
 487 on a plateau
- 488 • ρ_{Red} : the reduction rate of the exploratory step ρ after the same fitness
 489 has been calculated N_{\max} times
- 490 • ρ_{Thr} : the threshold after which the local search is interrupted
- 491 • Gr : the gene-resampling probability as in Algorithm 4

492 Since small datasets have only few samples per class, a large Gr value
 493 is required to make a significant refresh of the candidate solution. On the
 494 contrary, medium datasets inherently pose a highly multivariate problems.
 495 Hence smaller Gr values result into a major alternation of the candidate
 496 solution. We may consider this effect analogous to the setting of the crossover

497 rate in Differential Evolution with respect to the number of dimensions of
498 the problem, see [39]. On the other hand, numerous configurations have been
499 examined to find a set of parameters that can guarantee a robust performance
500 of SPMS-ALS on both small and medium datasets. In this study, we report
501 the performance of SPMS-ALS using the following parameters: initial radius
502 $\rho = 0.4$, $N_{\max} = 3$, $\rho_{Red} = 0.25$, $\rho_{Thr} = 0.005$ and $Gr = 0.5$ for small and 0.05
503 for medium datasets, respectively. Apart from the budget condition, which
504 is investigated for all the comparison algorithms as described above, the rest
505 of the parameters for all the algorithms are established as recommended
506 by the authors. All the details are presented in Table 3. Following the
507 experimental setup in [51], the reduction rate parameter is set to 95% for
508 small size datasets, and 98% for medium datasets.

Table 3: Parameters used for comparison algorithms

Algorithms	Parameter setting
SFLSDE	PopulationSize = 40, iterSFGSS =8, iterSFHC = 20, Fl = 0.1, Fu = 0.9
LSIR	initial $\rho = 0.4$
SSMA	Population = 40, Cross = 0.5, Mutation=0.001
PSO	SwarmSize = 40, C1 = 1, C2 = 3, Vmax = 0.25, Wstart = 1.5, Wend = 0.5
NN	k = 1, Euclidean distance
RIS	Thresholds = [0.0, 0.1, 0.2, ..., 0.9, 1.0]
LSBo	–
LSHADE	ArchiveSize = 1.4, PopulationSize = 40, MemorySize = 5

509 5. Analysis of results

510 In this section, we analyse the results obtained from different sets of
511 experiment, divided into multiple subsections, to empirically examine the
512 individual effect of each component we propose in our algorithm. In the
513 analysis, our aims are:

- 514 • To understand how well LSIR works in different settings of evaluations
515 (Subsection 5.1). We discuss multiple aspects of LSIR such as the
516 change in performance measured by accurate rate to see the overfitting
517 or underfitting effects on the learning process. In addition, the number
518 of evaluations that has been used and saved for each dataset is reported.
- 519 • To measure the actual savings in terms of runtime when the proposed
520 acceleration is integrated within LSIR (Subsection 5.2). We report in
521 detail the absolute and percentage figures of the runtime savings.

- 522 • To examine the performance enhancement due to the proposed memetic
523 components (Subsection 5.3), in comparison with the local search. We
524 report the accuracy rate depending on the number of evaluations and
525 we analyse the statistical significance of the improvements.
- 526 • To compare the performance of SPMS-ALS with the state-of-the-art
527 techniques in the family of instance reduction, with a focus on instance
528 generation (Subsection 5.4.1) considering the 1NN rule as a baseline and
529 recent instance selection methods (Subsection 5.4.2). In addition, the
530 average runtime required by each algorithm is contrasted to highlight
531 the substantial computational saving in the proposed method.
- 532 • To establish a fair comparison between the proposed approach and
533 the state-of-the-art algorithm in the family of instance reduction with
534 hybrid instance selection and generation algorithm, SSMA-SFLSDE,
535 using the same memetic instance selection algorithm as initialisation
536 mechanism (Subsection 5.5).
- 537 • To contextualise the results presented in this paper by comparing the
538 performance of SPMS-ALS with a recently proposed classifier (obRaF(H))
539 which represents a robust algorithm in the field of classification [24]
540 (Subsection 5.6).

541 For the sake of space, this section will only present summary results, and
542 all the detailed results can be found in the Supplementary Material and the
543 associated GitHub repository¹.

544 *5.1. LSIR Running with Different Computational Budgets*

545 The Local Search LSIR as Shown in Algorithm 2 has been run with
546 the three budget settings outlined in Table 2 to understand the influence
547 of the budget allowance in the performance of this algorithm. Its average
548 classification performance in the training and test phase is displayed in Table
549 4 on small and medium datasets. Note that the reported performance is
550 obtained from using Algorithm 1 changing **TR** by **TS** to evaluate LSIR in
551 the test phase. Analysing these average results and the detailed results in
552 the supplementary material, we can make the following comments:

¹<https://github.com/lehoanglam20000/SPMS-ALS>

- 553 • In the training phase the computational budget has a major impact on
554 the performance. However, while the performance grows consistently
555 from Setting 1 to Setting 3, this improvement is not major when the
556 performance of Setting 2 and Setting 3 are compared. This may infer
557 that changing each feature value cannot help the search seek a better
558 solution after a certain number of function calls.
- 559 • Regarding the test performance, we observe that the results are dra-
560 matically different to those achieved during the training phase: Setting
561 1 achieves most of the wins in test data overall. In small datasets,
562 Setting 1 has 22 wins out of 40, while Setting 2 and 3 win 14 and 11
563 times, respectively. In medium datasets, Setting 1 has 9 wins out of 17,
564 while Setting 2 and 3 win 6 and 8 times, respectively. We conclude that
565 overfitting is likely to happen for LSIR (possibly due to its exploitative
566 structure) in Setting 2 and 3. This tendency appears evident in small
567 size datasets.

Table 4: Average training and test performance in different settings of LSIR over small and medium datasets.

	Training		Test	
	Small	Medium	Small	Medium
Setting 1	0.8521 ± 0.0128	0.9005 ± 0.0039	0.7411 ± 0.0605	0.8612 ± 0.0139
Setting 2	0.8657 ± 0.0128	0.9049 ± 0.0039	0.7419 ± 0.0614	0.8610 ± 0.0133
Setting 3	0.8693 ± 0.014	0.9052 ± 0.0041	0.7415 ± 0.0607	0.8609 ± 0.0136

568 In summary, we can conclude that this local search does not seem to
569 benefit from using a larger budget and, as possibly expected, may be falling
570 into local optima, which do not generalise well in terms of classification per-
571 formance. This is especially noticeable in medium size datasets in which the
572 average training performance does not seem to increase much in respect to
573 the number of evaluations.

574 To further illustrate the behaviour of the LSIR approach with respect to
575 the number of evaluations, Table 5 shows the effect of its stopping criteria.
576 More specifically, when LSIR does not succeed at enhancing upon the trial
577 solution \mathbf{x}^t (see Algorithm 2), the exploratory step decreases by the factor
578 ρ_{Red} until a threshold value ρ_{Thr} is met. When these conditions are met the
579 run of LSIR is interrupted. Table 5 displays the computational budget saving
580 caused by the interruption of the run. The savings are shown for small and
581 medium datasets and for each of the setting under consideration. For each

582 configuration of dataset and setting the number of function calls used by the
 583 algorithm is also shown.

Table 5: Number of evaluations used and saved by LSIR in different settings and datasets.

	Small datasets			Medium datasets		
	Used	Saved	(%) Saved	Used	Saved	(%) Saved
Setting 1	15398	117	0.76	290777	0	0.00
Setting 2	30795	562	1.83	581554	54094	9.30
Setting 3	61591	2966	4.82	1163108	588637	50.61

584 Table 5 shows that Setting 1 mostly uses up the allocated number of eval-
 585 uations, whilst Setting 2 saves 1.83% and 9.3% in small and medium datasets,
 586 respectively. Setting 3 spends most of the evaluations in small datasets but
 587 only consumes nearly half of the allocated number of evaluations. These fig-
 588 ures may help optimise the number of evaluations used for each dataset based
 589 on their size and features. On the other hand, the allocation of a very large
 590 budget to the local search (like Setting 3) may not be always beneficial, and
 591 as mentioned above, the algorithm seems to get trapped into local optima.

592 5.2. Runtime Reduced in the Accelerated Version of LSIR

593 This subsection reports the runtime used by LSIR and how much it is
 594 reduced from its accelerated version using Algorithm 2 and the fitness in
 595 Algorithm 3, here referred to as Accelerated Local Search for Instance Re-
 596 duction (ALSIR).

597 Of course, the time required by the local search depends directly on the
 598 allocated budget and when the stopping criteria is reached. It is also impor-
 599 tant to remember that ALSIR always provides exactly the same classification
 600 performance as LSIR, this is because ALSIR focuses on accelerating the ex-
 601 ecution of the proposed method but it does not change the behaviour of the
 602 algorithm at all. The objective of the section is therefore to show how much
 603 we can accelerate LSIR with the proposed acceleration strategy.

604 Details of the runtime of both LSIR vs ALSIR in small and medium
 605 datasets, respectively, with respect to each setting of the number of evalua-
 606 tions can be found in the Supplementary material. Table 6 summarises the
 607 average runtime for each setting and the average percentage of time saved
 608 by the proposed acceleration. On average in small datasets, the objective
 609 function of the accelerated local search as in Algorithm 3 enables a time re-
 610 duction from at least 66% to 84.29%. However, the average runtime saved in

611 medium datasets settles around 90% in the three settings. Thus, the larger
 612 the dataset the more we can benefit from the proposed acceleration strategy,
 613 as distance computations become the most dominant part of the execution
 614 of the local search.

Table 6: Average runtime (in seconds) saved in different settings of LSIR and ALSIR, smaller values are in bold.

	Small datasets			Medium datasets		
	LSIR	ALSIR	(%) Time saved	LSIR	ALSIR	(%) Time saved
Setting 1	6.98	2.35	66.25	8676.67	856.05	90.13
Setting 2	19.47	3.60	81.54	15957.94	1508.85	90.54
Setting 3	37.68	5.92	84.29	18061.49	1784.64	90.12

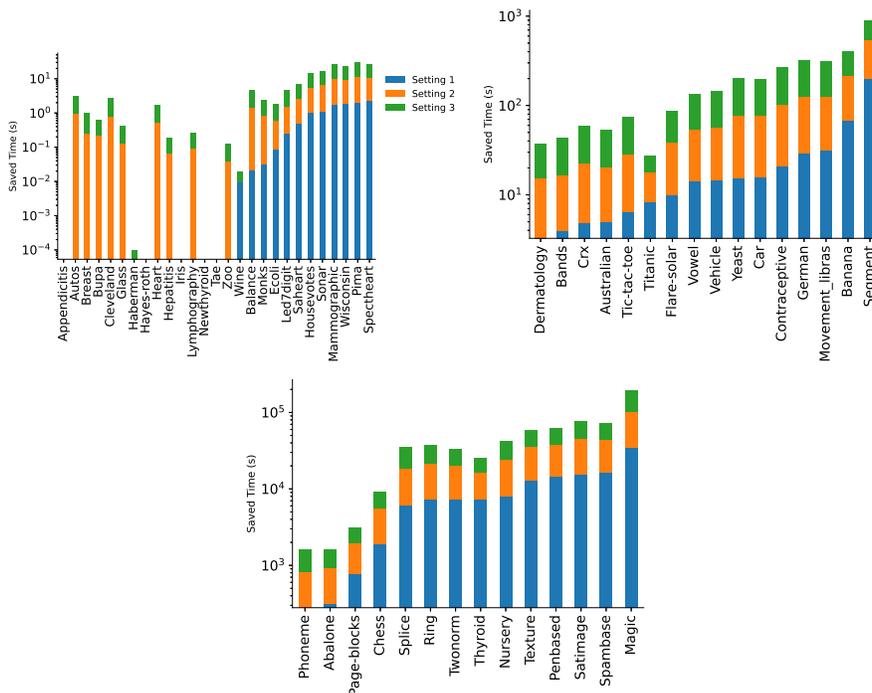


Figure 2: Runtime saved across 57 datasets, sorted by the ascending order of time gaps in Setting 1.

615 To illustrate the runtime reduction depending on the dataset size, Figure
 616 2 depicts the difference in runtime between LSIR and ALSIR for all the
 617 datasets, providing a graphical representation of the average time saving
 618 for each dataset. In order to enhance the readability of the diagram, the

619 logarithmic scale has been used. Those datasets which appear to have no
620 value represent those scenarios where the search can be completed is less
621 than a second. Hence, the acceleration may not be essential in these cases.

622 5.3. Validation of the Memetic Framework of SPMS-ALS

623 In this section, we compare the performance of LSIR and SPMS-ALS to
624 demonstrate the effectiveness of the proposed memetic framework. Table 7
625 provides a full summary of this comparison, presenting the average accuracy
626 values (over all the datasets) and the corresponding standard deviations in
627 the three settings of computational budget in both training and test phases.
628 The best average results in training and test are highlighted in bold face.

629 Furthermore, the Wilcoxon test [57] is also applied to detect the statisti-
630 cal differences between the two methods. The corresponding p -values are
631 also shown in the last column of Table 7. When one algorithm significantly
632 outperforms the other, the p -value is less than the confidence level 0.05. We
633 highlight in italic these p -values.

Table 7: Comparison in average training and test performance between LSIR and SPMS-ALS over small and medium datasets. Wilcoxon p -value is obtained from the comparison between SPMS-ALS and LSIR.

		SPMS-ALS				LSIR				Wilcoxon
		TRAINING		TEST		TRAINING		TEST		p -value
	Evaluations	Acc	Std	Acc	Std	Acc	Std	Acc	Std	
SMALL	Setting 1	0.8598	0.0130	0.7477	0.0633	0.8521	0.0128	0.7411	0.0605	0.1015
	Setting 2	0.8665	0.0122	0.7512	0.0625	0.8657	0.0128	0.7419	0.0614	0.0867
	Setting 3	0.8733	0.0110	0.7549	0.0615	0.8693	0.0140	0.7415	0.0607	<i>0.0132</i>
MEDIUM	Setting 1	0.9126	0.0034	0.8625	0.0127	0.9005	0.0039	0.8612	0.0139	>0.2
	Setting 2	0.9129	0.0033	0.8626	0.0126	0.9049	0.0039	0.8610	0.0133	>0.2
	Setting 3	0.9199	0.0028	0.8668	0.0110	0.9052	0.0041	0.8609	0.0136	<i>0.0577</i>

634 Numerical results in Table 7 show that for both training and test phases,
635 the memetic framework outperforms on a regular basis LSIR. We may observe
636 that in training phase and small datasets, SPMS-ALS slightly outperforms
637 LSIR while the difference in performance is larger for medium datasets. Ac-
638 cording to our interpretation, this shows the effectiveness of the global search
639 component in complex spaces: while the local search exploits the space and is
640 likely to achieve a suboptimal point (we may see that different computational
641 budgets do not yield major changes in LSIR performance), the crossover al-
642 lows the search a further chance to detect a solution closer to the global
643 optimum. The results in the test phase display a consistent better perfor-
644 mance of the memetic framework across the datasets. This finding can be

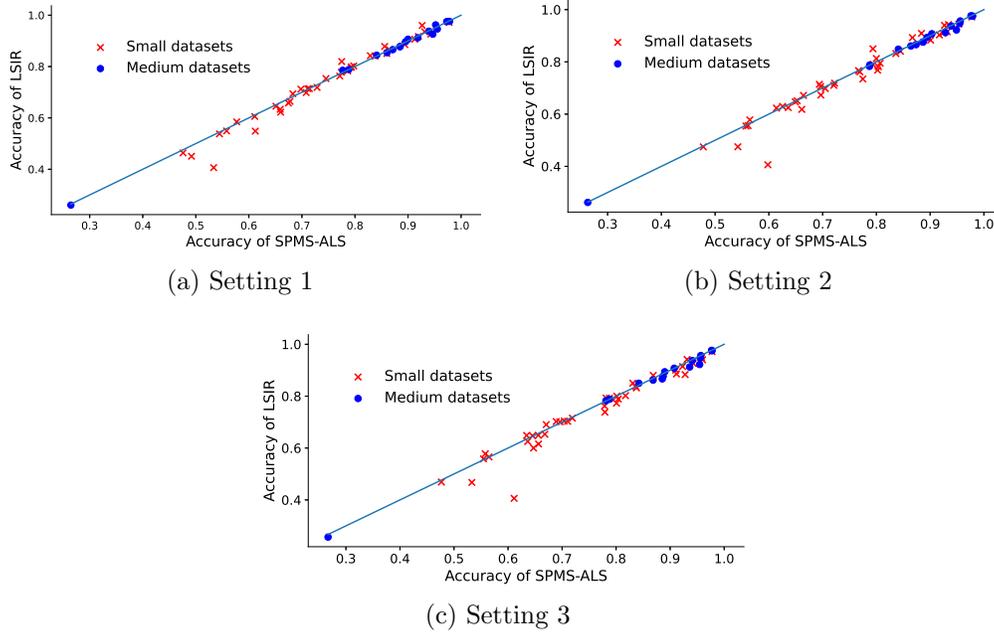


Figure 3: Accuracy scatter plots over 40 small and 17 medium datasets in the test phase.

645 interpreted as a better performance of SPMS-ALS in terms of overfitting: the
 646 deterministic and exploitative nature of LSIR may lead to overfitting while
 647 the degree of randomisation introduced by the crossover-based global search
 648 element reduces the risk of overfitting hence improving upon the performance
 649 of the algorithm in test phase. Finally we may observe that SPMS-ALS sta-
 650 tistically outperforms LSIR in Setting 3 in small datasets and shows improved
 651 progress in medium datasets. This fact is expected since longer runs tend
 652 to be more stable and thus be associated with lower standard deviation val-
 653 ues. On the contrary, with Setting 1 and 2 we are more likely to observe
 654 “lucky” or “unlucky” runs that may jeopardise the statistical significance of
 655 the results.

656 The test results are also graphically presented in Figure 3 which contains
 657 scatter plots of the accuracy of the methods. Each point compares the test
 658 performance of SPMS-ALS and LSIR algorithm on a single dataset. The
 659 accuracy of SPMS-ALS is shown on the x-axis position of the point, while
 660 that of LSIR is on the y-axis position. Thus, points below the $y = x$ line
 661 correspond to datasets for which SPMS-ALS achieves better performance

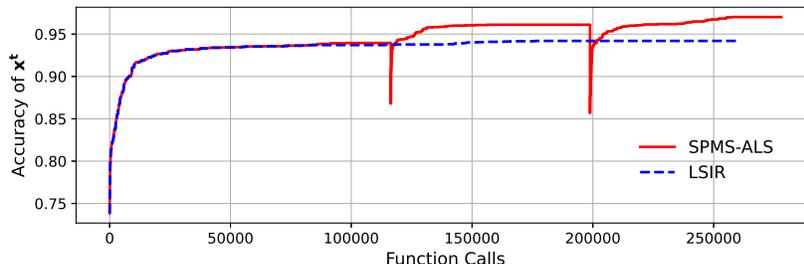


Figure 4: Functioning (Accuracy) of SPMS-ALS and LSIR on the Chess dataset.

662 than the compared algorithm. In most of the cases, the points are plotted
 663 on or below the separating line, inferring greater performance of SPMS-ALS.
 664 In this plot, we can also see that the biggest improvements have been made
 665 in small datasets, but in turn, there are a few datasets in which SPMS-ALS
 666 performs slightly worse. However, in medium size datasets the improvements
 667 are less significant, especially in settings 1 and 2, but consistently better.

668 In order to emphasise the different behaviour of LSIR vs SPMS-ALS we
 669 plot in Figure 4 the accuracy of the trial solution \mathbf{x}^t against function calls
 670 (evaluations) of the two algorithms on the Chess dataset, using Setting 1. To
 671 show the functioning of the crossover the plot of SPMS-ALS refers to the local
 672 solution (and not the elite). We may observe that the crossover functions as
 673 a restart which then quickly reaches a solution with a good performance.

674 In conclusion, whilst the local search seemed to get stuck after a number
 675 of evaluations, the proposed MC approach, despite its simplicity, benefits
 676 from larger computation budgets, outperforming the local search.

677 5.4. Comparison with the State-of-the-art Methods in Instance Reduction

678 This section consists of two sub-sections: Section 5.4.1 covers the compar-
 679 ison of our proposal with related instance generation methods; Section 5.4.2
 680 presents the comparison with recently published instance selection methods.

681 5.4.1. Comparison against similar instance generation techniques

682 In order to compare the performance of SPMS-ALS against that of the
 683 other global optimisers for instance generation (PSO [36], SFLSDE [52], and
 684 the adapted LSHADE [49]), we will focus on the maximum number of eval-
 685 uations (Setting 3) for all the algorithms. As a baseline, we also include the
 686 1NN algorithm as a comparison algorithm.

687 Table 8 summarises the performance of the comparison algorithm in all
688 (57) datasets (small and medium). We have employed the Friedman proce-
689 dure [20] plus a Holm post-hoc test to perform a ranking-based statistical
690 analysis on the performance of the algorithms for small and medium datasets,
691 respectively. The last two columns of Table 8 provide the results of these
692 tests, including the rankings and the resulting p -values. Note that the con-
693 trol method will obtain the lowest ranking, and therefore, the p -value shows
694 if the differences are significant comparing the control algorithm against the
695 rest of the methods.

696 As shown in Table 8, LSHADE and SFLSDE are reported as the con-
697 trol method in small and medium datasets, respectively, since they hold the
698 smallest ranking values. In small datasets, our proposal SPMS-ALS ranks
699 third after SFLSDE, while it ranks second in the medium datasets and its
700 ranking value is not far away from that of the control algorithm.

701 The Holm post-hoc test is used to detect if there is any significant sta-
702 tistical differences between the control algorithm (LSHADE and SFLSDE)
703 with respect to the remaining methods. Considering a level of significance
704 of $\alpha = 0.05$, LSHADE statistically outperforms only 1NN in small datasets,
705 and PSO in medium datasets. The statistical tests have not reported sig-
706 nificant differences between our proposal and the control algorithm in either
707 small or medium datasets.

Table 8: Summary of the performance of SPMS-ALS against SFLSDE, PSO, LSHADE and 1NN for instace reduction over 57 datasets. The best performance in the column is shown in bold.

	Algorithm	TRAINING		TEST		Friedman+Holm	
		Acc	Std	Acc	Std	Ranking	p_{Holm}
SMALL	LSHADE	0.8401	0.0165	0.7541	0.0612	2.425	–
	SFLSDE	0.8480	0.0092	0.7615	0.0634	2.525	0.7773
	SPMS-ALS	0.8733	0.0110	0.7549	0.0615	3.125	0.1017
	PSO	0.8147	0.0156	0.7414	0.0606	3.175	0.1017
	1NN	0.7369	0.0088	0.7369	0.0088	3.750	0.0007
MEDIUM	SFLSDE	0.8887	0.0048	0.8608	0.0122	2.177	–
	SPMS-ALS	0.9199	0.0028	0.8668	0.0110	2.353	0.7448
	LSHADE	0.8859	0.0138	0.8503	0.0147	3.294	0.1180
	1NN	0.8316	0.0045	0.8316	0.0045	3.294	0.1180
	PSO	0.8537	0.0066	0.8319	0.0137	3.882	0.0066

708 According to our interpretation, in training phase an exploitative action
709 guarantees a better performance of the algorithm especially in high dimen-
710 sions [8]. However, the exploitative pressure should be counterbalanced by
711 a certain degree of randomisation to prevent the algorithm from overfitting

712 and pay off with a deteriorated performance in test phase. This feature of
 713 the instance generation problem makes it especially suitable to be tackled by
 714 memetic frameworks. Albeit reasonable, the excessively exploratory nature
 715 of PSO does not appear to effectively address the large dimensional space.

716 In comparison with the baseline, 1NN, which uses all the data to clas-
 717 sify the test set, we have conducted the Wilcoxon test to conduct a pairwise
 718 comparison to our method. Although SPMS-ALS shows better average per-
 719 formance, the Wilcoxon test compute p -value 0.0415 for small datasets and
 720 >0.2 for medium datasets. These numeric p -values indicates that our algo-
 721 rithm statistically outperform 1NN in small dataset, but has no significant
 722 different in medium datasets, considering a level of significance of $\alpha = 0.05$.

723 Finally, Table 9 displays the average runtime of the four global optimisers
 724 for the small and medium datasets, respectively. On average, the runtime
 725 spent for small datasets is 6.25s, saving 92.52%, 83.93% and 28.17% with
 726 respect to LSHADE, SFLSDE and PSO, respectively. For medium datasets,
 727 the percentage of saving time is slightly higher than what it did in small
 728 datasets, but the absolute value is more meaningful. Specifically, SPMS-ALS
 729 only consumes roughly 5000s on average, while SFLSDE and PSO experience
 730 about 35000s, and LSHADE used up to approximate 160.000s. In other
 731 words, for medium datasets, SPMS-ALS achieves similar if not better results
 732 of SFLSDE and LSHADE in one seventh and less than one thirtieth of the
 733 runtime, respectively.

Table 9: Comparison of the runtime (in seconds) consumed in SPMS-ALS and other approaches. Min values are in bold.

	Small datasets	Medium datasets
SFLSDE	38.89 ± 1.31	34738.82 ± 188.55
PSO	19.63 ± 0.80	34941.65 ± 188.86
LSHADE	83.63 ± 4.12	159009.10 ± 1154.75
SPMS-ALS	6.25 ± 0.39	5006.93 ± 405.33

734 5.4.2. Comparison against recent Instance Selection

735 As mentioned in Section 4.2, two recent instance selection algorithms
 736 LSBo [27] and RIS1 [10] have been selected for comparison with our proposal.
 737 This section reports the experimental results of the these two algorithms
 738 against SPMS-ALS over 57 datasets with reference to test performance and
 739 reduction rate.

740 Details of the classification performance of RIS1, LSBo and SPMS-ALS
 741 in the test phase on small and medium datasets can be found in the Supple-

742 mentary Material, while the summary information is displayed at Table 10.
743 Overall, RIS1 obtains a majority of wins in both small and medium datasets,
744 SPMS-ALS ranks second and LSBo lies at the lowest position. Particularly,
745 RIS1 has 25 wins (18 small and 7 medium), SPMS-ALS achieves the best
746 results 18 times (15 small and 3 medium), while LSBo obtains 15 wins (8
747 small and 7 medium). However, the average test performance of LSBo and
748 SPMS-ALS are the highest for small and medium datasets, respectively.

749 The overall goal of instance reduction is to reduce the original dataset
750 as much as possible whilst keeping (or improving) the accuracy. Therefore,
751 to establish a fairer comparison between the two instance selection methods
752 and our method, we will also provide an additional metric to consider both
753 test accuracy and reduction as equally important. Following [52], we simply
754 multiply the accuracy in test **Acc** and the reduction rate **Red** to form a
755 new metric **Acc*Red**. Table 10 presents the overall average performance
756 in accuracy **Acc**, reduction rate **Red** and both metrics **Acc*Red**. Fur-
757 thermore, Table 11 provides the set of rankings and p -values obtained from
758 Friedman+Holm tests for the three contrasted algorithms.

759 In the previous section, all algorithms (PSO, SFLSDE, LSHADE, and
760 SPMS-ALS) yield the same reduction rate. In particular, in this experiment,
761 SPMS-ALS fixes the rate up to 95% and 98% for small and mediums datasets,
762 respectively. However, LSBo and RIS1 do not specify the reduction rate as
763 a parameter, but their reduction depends on a particular dataset. RIS1
764 yields an average reduction rate of 59.06% and 70.52% in small and medium
765 datasets, respectively; whilst LSBo reduces 78.59% and 87.39% in small and
766 medium datasets. Thus, both algorithms achieve a smaller reduction rate
767 than the instance generation approach investigated in this paper.

Table 10: Summary of the performance of SPMS-ALS against RIS1 and LSBo consider-
ing **Acc** in the test phase, **Red** and **Acc*Red** measures for instance reduction over 57
datasets. The best performance in the column is shown in bold.

	Algorithm	Acc (Test)	Std	Red	Acc*Red
SMALL	RIS1	0.7319	0.0583	0.5906	0.4499
	LSBo	0.7605	0.0576	0.7859	0.6064
	SPMS-ALS	0.7549	0.0615	0.9500	0.7172
MEDIUM	RIS1	0.7972	0.0141	0.7052	0.6071
	LSBo	0.8540	0.0099	0.8739	0.7639
	SPMS-ALS	0.8668	0.0110	0.9800	0.8495

768 On average, the test performance of LSBo is the highest on small datasets,

769 while SPMS-ALS reports the highest average on medium datasets. However,
770 apart from having the best reduction rate, SPMS-ALS obtains the best bal-
771 ance between accuracy and reduction. The results from the non-parametric
772 tests (Friedman+Holm) in Table 11 reveal the advantage of SPMS-ALS look-
773 ing at the Ranking and p -value columns. The p -values reported for the com-
774 parison in terms of accuracy column do not reflect any significant differences
775 between the three methods in either small or medium datasets. However,
776 when the reduction rate is taken into consideration the proposed technique
777 stands out significantly.

Table 11: Friedman+Holm statistical test results in both **Acc** and **Acc*Rec** metrics for small and medium datasets. The best performance in the column is shown in bold.

		Acc			Acc * Red		
		Algorithm	Ranking	p -value	Algorithm	Ranking	p -value
SMALL	SPMS-ALS	1.912		–	SPMS-ALS	1.175	–
	RIS1	1.938	0.9110		LSBo	2.000	0
	LSBo	2.150	0.2882		RIS1	2.825	2.25E-04
MEDIUM	LSBo	1.882		–	SPMS-ALS	1.176	–
	RIS1	2.000	0.7316		LSBo	2.059	1.01E-02
	SPMS-ALS	2.117	0.4927		RIS1	2.765	4.00E-06

778 5.5. Hybridisation with Instance Selection

779 As stated in Section 4.2, to perform a fair comparison against hybrid
780 instance reduction method SSMA-SFLSDE [52], the initialisation process
781 of the proposed algorithm must be replaced with a smarter approach. In
782 particular, we use the same instance selection algorithm, SSMA [18], as tested
783 in [52]. This section compares LSIR, SFLSDE, LSHADE and SPMS-ALS
784 after using SSMA as initialisation. The resulting algorithms are indicated as
785 SSMA-LSIR, SSMA-SPMS-ALS, SSMA-LSHADE and the state-of-the-art
786 algorithm SSMA-SFLSDE [52]. The detailed accuracy results for small and
787 medium datasets in both training and test can be found in the Supplementary
788 Material. Table 12 shows the average results obtained from the compared
789 algorithms in conjunction with SSMA and the ranking plus p -values from the
790 Friedman + Holm test.

791 The detailed results show that for small datasets and in training phase
792 SSMA-SPMS-ALS outperforms SSMA-SFLSDE and SSMA-LSHADE, and is
793 outperformed by SSMA-LSIR. However these results are not confirmed in test
794 phase where SSMA-LSHADE achieves the best performance, SSMA-SPMS-
795 ALS the third best performance after SSMA-SFLSDE, and SSMA-LSIR the
796 worst performance over the four algorithms considered in this section. This

797 ranking is statistically significant and confirmed by the Friedman + Holm
798 test. According to our interpretation, the deterministic and exploitative local
799 search logic in LSIR causes overfitting. The restriction of the search space
800 caused by SSMA increases the risk of overfitting. In the proposed memetic
801 framework, the resampling and crossover mechanism seems to mitigate the
802 overfitting.

803 Our interpretation is confirmed by the results for medium datasets. Since
804 the search space is naturally large, the exploitative local search is beneficial
805 [8] and is improved by the memetic framework. Hence, SSMA-SPMS-ALS
806 achieves the best performance in training phase. This ranking is confirmed
807 in test phase where SSMA-SPMS-ALS slightly outperforms SSMA-SFLSDE
808 and is established as the control algorithm in the Friedman test.

809 In summary, and similar to what we saw when comparing against purely
810 instance generation methods, the proposed memetic framework can obtain
811 a very competitive classification performance, especially in larger datasets,
812 whilst reducing drastically the required runtime.

Table 12: Summary performance between four hybrid models over 57 datasets.

		TRAINING		TEST		Friedman + Holm	
Algorithm		Acc	Std	Acc	Std	Ranking	p_{Holm}
SMALL	SSMA-LSHADE	0.8687	0.0101	0.7792	0.0570	2.000	–
	SSMA-SFLSDE	0.8684	0.0108	0.7767	0.0594	2.200	0.4884
	SSMA-SPMS-ALS	0.8727	0.0134	0.7670	0.0574	2.700	0.0306
	SSMA-LSIR	0.8911	0.0148	0.7642	0.0604	3.100	0.0004
MEDIUM	SSMA-SPMS-ALS	0.9264	0.0033	0.8700	0.0107	2.265	–
	SSMA-SFLSDE	0.9059	0.0040	0.8675	0.0125	2.441	1.0000
	SSMA-LSHADE	0.9069	0.0035	0.8706	0.0118	2.647	1.0000
	SSMA-LSIR	0.9245	0.0039	0.8682	0.0127	2.647	1.0000

813 At last we report some considerations about future improvements that
814 can be applied. We will investigate the extension of our approach to big
815 data frameworks [55]. In addition, we plan to expand our approach to new
816 promising classifiers, such as Heterogeneous oblique random forest [24]. An
817 initial comparison with this kind of classifier can be found in the Supple-
818 mentary Material. Further investigation is however required to perform an
819 appropriate instance reduction for those classifiers.

820 5.6. Contextualising the results and limitations of the proposal

821 Experimental results in Sections 5.1 to 5.5, show that the proposed SPMS-
822 ALS and its hybrid form, SSMA-SPMS-ALS, are effective at reducing the size

823 of the training data whilst maintaining, or even improving, the performance
824 of the base classifier; in our case, the 1NN rule. The goal of this section
825 is to contextualise the classification results presented in this paper with the
826 1NN as base classifier, and let the reader know where we are going in our
827 future research. To do so, we compare the results of the proposed SPMS-
828 ALS and 1NN against the popular Random Forest (RaF) algorithm and a
829 state-of-the-art classifier also based on Trees, obRaF(H) [24].

830 It is important to note that the classification performance of a classifier
831 is not only influenced by the pre-processing techniques but also its inherent
832 robustness. For example, an ensemble classifier is likely to outperform a single
833 model [46], or tree-based approaches handle categorical attributes better than
834 the NN classifier. Thus, whilst the reader may expect the results of RaF and
835 obRaF(H) to be superior to the ones presented with the 1NN rule, we believe
836 it is beneficial to still observe the performance gap and understand potential
837 future research lines for preprocessing technique for more robust classifiers.

838 To establish a fair comparison against methods, we have re-run SPMS-
839 ALS over the 121 UCI datasets used in [16, 24], following the exact same
840 experimental framework, including depth of a tree (i.e. 57), number of trees
841 (i.e. 500), and a 4-fold cross validation scheme. Details of the comparison on
842 each single dataset can be found in the Supplementary Material, while the
843 summary information and results of the statistical test are displayed in Table
844 13. As expected, both RaF and obRaF(H) display a higher performance than
845 NN-based results. On the other hand, the proposed data reduction does not
846 only reduce the storage need but also becomes much more efficient in terms
847 of runtime as we only preserve 2 to 5% of the training data. For this reason,
848 it is essential to highlight that the contribution of our proposal lies in the
849 reduction of the training set, as a preprocessing technique, whilst maintaining
850 (or improving) the classification performance of 1NN.

Table 13: Summary the average performance of 4-fold cross validation between the basic models (1NN and RaF) and their improved versions (SPMS-ALS and obRaF(H)) over 121 datasets.

	TRAINING	TEST	Friedman + Holm	
			Ranking	P_{Holm}
obRaF(H)	–	0.8336	146.24	–
RaF(Scikit-learn)	0.9892	0.8286	173.13	0.05
SPMS-ALS	0.8926	0.7597	324.99	0.03
1NN	0.7487	0.7534	325.64	0.02

851 The reader might wonder if the result of the preprocessing performed by
852 the technique proposed in this paper could be used directly by any other
853 classifier like RaF or obRaF(B). Although this pre-processing approach is
854 intended for 1NN, as highlighted in [5] the resulting set could potentially be
855 used by any other classifiers. However, it is not straightforward to directly
856 use the reduced set obtained from SPMS-ALS in another classifier. We have
857 performed some preliminary experiments using the resulting reduced set as
858 training data for RaF in 89 small datasets (from the set of 121). The average
859 results in the test phase sets at 0.5656, whilst it is 1.000 on training. This
860 suggests that RaF is overfitting the training data with the parameters used
861 (e.g. depth of the trees, or number of trees). This could be expected as the
862 reduction on small datasets may end up having as few as 2-15 samples in some
863 extremely small datasets. Whilst NN technique would work well with such
864 amount of data, Tree-like technique will not. Thus, as future work we plan
865 to explore the interaction between the proposed SPMS-ALS and more robust
866 classifiers like obRaF(B), for example, by adding it as base classifier or fine
867 tuning the parameters to use smaller training datasets without overfitting.

868 6. Conclusion

869 This paper proposes a single-point MC approach for instance reduction.
870 The proposed algorithm is composed of a novel accelerated local search and a
871 crossover based global search. The local search is deterministic and exploita-
872 tive belonging to the family of Pattern Search methods whilst the global
873 search is stochastic, based on resampling and crossover. By making some
874 considerations about the functioning of the NN classifier in instance gener-
875 ation and exploiting the search logic of Pattern Search, the local search has
876 been redesigned and implemented in an accelerated version. The accelerated
877 local search uses most of the calculations performed at the previous step and
878 thus lead to a major saving in terms of runtime with respect to the existing
879 algorithms in the literature.

880 Numerical results performed with and without instance selection as ini-
881 tialisation mechanism show that the proposed MC approach tends to be
882 slightly worse than only one instance reduction algorithm in small datasets.
883 On the other hand, on medium datasets, the proposed MC approach achieves
884 the best accuracy performance in both training and test phases. These re-
885 sults are extremely valuable when we consider that the proposed approach
886 is up to seven times faster than the other algorithms.

887 Besides the proposed domain-specific MC approach this article offers an
888 extra contribution about experimentalism in data reduction. More specifi-
889 cally, in this paper we perform a thorough parameter setting of the compu-
890 tational budget and display the results in multiple scenarios. These results
891 aim to offer some guidelines to data scientists to set their experimental con-
892 ditions in a fair and effective manner to detect the desired trade-off between
893 accuracy and runtime.

894 7. Acknowledgement

895 The work of H. Lam Le was funded by a Ph.D. scholarship from the School
896 of Computer Science of the University of Nottingham. All experiments in
897 this study were conducted with the High Performance Computing system
898 at the University of Nottingham. We thank George Darmiton da Cunha
899 Cavalcanti and Rodolfo Jose de Oliveira Soares, the authors of paper [10] for
900 the provided code and their support to run the experiment.

901 References

- 902 [1] G. Acampora, V. Loia, M. Gaeta, Exploring e-Learning Knowledge
903 Through Ontological Memetic Agents, *IEEE Computational Intelligence*
904 *Magazine* 5 (2) (2010) 66–77.
- 905 [2] E. Alpaydin, *Introduction to Machine Learning*, MIT press, 2014.
- 906 [3] J. E. Amaya, C. Cotta, A. J. Fernández, P. García-Sánchez, Deep
907 memetic Models for Combinatorial Optimization Problems: Applica-
908 tion to the Tool Switching Problem, *Memetic Computing* 12 (1) (2020)
909 3–22.
- 910 [4] J. Bacardit, D. E. Goldberg, M. V. Butz, X. Llorà, J. M. Garrell,
911 Speeding-up Pittsburgh Learning Classifier Systems: Modeling Time
912 and Accuracy, in: *International Conference on Parallel Problem Solving*
913 *from Nature*, Springer, 1021–1031, 2004.
- 914 [5] J. R. Cano, F. Herrera, M. Lozano, Using Evolutionary Algorithms as
915 Instance Selection for Data reduction in KDD: an experimental study,
916 *IEEE Transactions on Evolutionary Computation* 7 (6) (2003) 561–575.

- 917 [6] A. Caponio, G. L. Cascella, F. Neri, N. Salvatore, M. Sumner, A Fast
918 Adaptive Memetic Algorithm for On-line and Off-line Control Design of
919 PMSM Drives, *IEEE Transactions on System Man and Cybernetics-part*
920 *B, Special Issue on Memetic Algorithms* 37 (1) (2007) 28–41.
- 921 [7] F. Caraffini, F. Neri, M. G. Epitropakis, HyperSPAM: A Study on
922 Hyper-heuristic Coordination Strategies in the Continuous Domain, *In-*
923 *formation Sciences* 477 (2019) 186–202.
- 924 [8] F. Caraffini, F. Neri, G. Iacca, Large Scale Problems in Practice:
925 The Effect of Dimensionality on the Interaction Among Variables, in:
926 G. Squillero, K. Sim (Eds.), *Applications of Evolutionary Computation*,
927 Springer, 636–652, 2017.
- 928 [9] F. Caraffini, F. Neri, B. Passow, G. Iacca, Re-sampled Inheritance
929 Search: High Performance Despite the Simplicity, *Soft Computing*
930 17 (12) (2014) 2235–2256.
- 931 [10] G. D. Cavalcanti, R. J. Soares, Ranking-based Instance Selection for
932 Pattern Classification, *Expert Systems with Applications* 150 (2020)
933 113269.
- 934 [11] Chang-Yong Lee, Xin Yao, Evolutionary Programming Using Mutations
935 Based on the Levy Probability Distribution, *IEEE Transactions on Evo-*
936 *lutionary Computation* 8 (1) (2004) 1–13.
- 937 [12] X. Chen, Y. Ong, M. Lim, K. C. Tan, A Multi-Facet Survey on Memetic
938 Computation, *IEEE Transactions on Evolutionary Computation* 15 (5)
939 (2011) 591–607.
- 940 [13] T. M. Cover, P. E. Hart, Nearest Neighbor Pattern Classification, *IEEE*
941 *Transactions on Information Theory* 13 (1) (1967) 21–27.
- 942 [14] A. Elola, J. Del Ser, M. N. Bilbao, C. Perfecto, E. Alexandre, S. Salcedo-
943 Sanz, Hybridizing Cartesian Genetic Programming and Harmony Search
944 for Adaptive Feature Construction in Supervised Learning Problems,
945 *Applied Soft Computing* 52 (2017) 760–770.
- 946 [15] L. Feng, Y. Ong, M. Lim, I. W. Tsang, Memetic Search With Interdo-
947 main Learning: A Realization Between CVRP and CARP, *IEEE Trans-*
948 *actions on Evolutionary Computation* 19 (5) (2015) 644–658, ISSN 1941-
949 0026.

- 950 [16] M. Fernández-Delgado, E. Cernadas, S. Barro, D. Amorim, Do We Need
951 Hundreds of Classifiers to Solve Real World Classification Problems?,
952 The Journal of Machine Learning Research 15 (1) (2014) 3133–3181.
- 953 [17] I. Fister, A. Iglesias, A. Gálvez, J. Del Ser, E. Osaba, I. F. Jr., M. Perc,
954 M. Slavinec, Novelty Search for global Optimization, Applied Mathe-
955 matics and Computation 347 (2019) 865–881.
- 956 [18] S. García, J. R. Cano, F. Herrera, A Memetic Algorithm for Evolution-
957 ary Prototype Selection: A Scaling up Approach, Pattern Recognition
958 41 (8) (2008) 2693–2709.
- 959 [19] S. García, J. Derrac, J. Cano, F. Herrera, Prototype Selection for Near-
960 est Neighbor Classification: Taxonomy and Empirical Study, IEEE
961 Transactions on Pattern Analysis and Machine Intelligence 34 (3) (2012)
962 417–435.
- 963 [20] S. García, A. Fernández, J. Luengo, F. Herrera, Advanced Nonpara-
964 metric Tests for Multiple Comparisons in the Design of Experiments in
965 Computational Intelligence and Data Mining: Experimental Analysis of
966 Power, Information Sciences 180 (10) (2010) 2044–2064.
- 967 [21] A. Gupta, Y.-S. Ong, Memetic Computation, Adaptation, Learning, and
968 Optimization, Springer, 2019.
- 969 [22] G. Iacca, F. Neri, E. Mininno, Y. S. Ong, M. H. Lim, Ockham’s Razor
970 in Memetic Computing: Three Stage Optimal Memetic Exploration,
971 Information Sciences 188 (2012) 17–43.
- 972 [23] N. D. Jana, J. Sil, S. Das, Continuous Fitness Landscape Analysis Using
973 a Chaos-Based Random Walk Algorithm, Soft Computing 22 (2018)
974 921–948.
- 975 [24] R. Katuwal, P. N. Suganthan, L. Zhang, Heterogeneous Oblique Ran-
976 dom Forest, Pattern Recognition 99 (2020) 107078.
- 977 [25] H. L. Le, D. Landa-Silva, M. Galar, S. Garcia, I. Triguero, EUSC:
978 A Clustering-based Surrogate Model to Accelerate Evolutionary Un-
979 dersampling in Imbalanced Classification, Applied Soft Computing 101
980 (2021) 107033.

- 981 [26] M. N. Le, Y. S. Ong, Y. Jin, B. Sendhoff, Lamarckian Memetic Algo-
982 rithms: Local Optimum and Connectivity Structure Analysis, *Memetic*
983 *Computing Journal* 1 (3) (2009) 175–190.
- 984 [27] E. Leyva, A. González, R. Pérez, Three New Instance Selection Methods
985 Based on Local Sets: A Comparative Study with Several Approaches
986 from a Bi-objective Perspective, *Pattern Recognition* 48 (4) (2015) 1523–
987 1537.
- 988 [28] X. Li, X. Yao, Cooperatively Coevolving Particle Swarms for Large Scale
989 Optimization, *Evolutionary Computation, IEEE Transactions on* 16 (2)
990 (2012) 210–224.
- 991 [29] P. López-García, E. Onieva, E. Osaba, A. D. Masegosa, A. Perallos,
992 GACE: A Meta-heuristic Based in the Hybridization of Genetic Algo-
993 rithms and Cross Entropy Methods for Continuous Optimization, *Ex-
994 pert Systems with Applications* 55 (2016) 508–519.
- 995 [30] L. Ma, J. Li, Q. Lin, M. Gong, C. A. Coello Coello, Z. Ming, Cost-Aware
996 Robust Control of Signed Networks by Using a Memetic Algorithm,
997 *IEEE Transactions on Cybernetics* (2020) 1–14 To appear.
- 998 [31] X. Ma, X. Li, Q. Zhang, K. Tang, Z. Liang, W. Xie, Z. Zhu, A Sur-
999 vey on Cooperative Co-Evolutionary Algorithms, *IEEE Transactions on*
1000 *Evolutionary Computation* 23 (3) (2019) 421–441.
- 1001 [32] K. M. Malan, A. P. Engelbrecht, A Survey of Techniques for Character-
1002 ising Fitness Landscapes and Some Possible Ways Forward, *Information*
1003 *Sciences* 241 (2013) 148 – 163.
- 1004 [33] A. D. Martinez, E. Osaba, I. Oregi, I. F. Jr., I. Fister, J. Del Ser, Hy-
1005 bridizing differential evolution and novelty Search for multimodal Op-
1006 timization Problems, in: *Proceedings of the Genetic and Evolutionary*
1007 *Computation Conference Companion, GECCO 2019, Prague, Czech Re-
1008 public, July 13-17, 2019, 1980–1989, 2019.*
- 1009 [34] P. Moscato, *On Evolution, Search, Optimization, Genetic Algorithms*
1010 *and Martial Arts: Towards Memetic Algorithms*, Tech. Rep. 826, Cal-
1011 tech, 1989.

- 1012 [35] P. Moscato, M. Norman, A Competitive and Cooperative Approach to
1013 Complex Combinatorial Search, Tech. Rep. 790, Caltech, 1989.
- 1014 [36] L. Nanni, A. Lumini, Particle Swarm Optimization for Prototype Re-
1015 duction, *NeuroComputing* 72 (4-6) (2009) 1092–1097.
- 1016 [37] F. Neri, *Linear Algebra for Computational Sciences and Engineering*,
1017 Springer, second edn., 2019.
- 1018 [38] F. Neri, C. Cotta, Memetic Algorithms and Memetic Computing Opti-
1019 mization: A Literature Review, *Swarm and Evolutionary Computation*
1020 2 (2012) 1–14.
- 1021 [39] F. Neri, G. Iacca, E. Mininno, Disturbed Exploitation Compact Differential
1022 Evolution for Limited Memory Optimization Problems, *Information*
1023 *Sciences* 181 (12) (2011) 2469 – 2487.
- 1024 [40] F. Neri, S. Rostami, Generalised Pattern Search Based on Covariance
1025 Matrix Diagonalisation, *SN Comput. Sci.* 2 (3) (2021) 171.
- 1026 [41] F. Neri, I. Triguero, A Local Search with a Surrogate Assisted Option
1027 for Instance Reduction, in: P. A. Castillo, J. L. J. Laredo, F. F. de Vega
1028 (Eds.), *Applications of Evolutionary Computation*, vol. 12104 of *Lecture*
1029 *Notes in Computer Science*, Springer, 578–594, 2020.
- 1030 [42] Q. H. Nguyen, Y. S. Ong, L. M. Hiot, N. Krasnogor, Adaptive Cellular
1031 Memetic Algorithms, *Evolutionary Computation* 17 (2) (2009) 231–256.
- 1032 [43] R. Nogueras, C. Cotta, Studying Self-balancing Strategies in Island-
1033 based Multimemetic Algorithms, *Journal of Computational and Applied*
1034 *Mathematics* 293 (2016) 180–191.
- 1035 [44] E. Özcan, B. Bilgin, E. E. Korkmaz, A Comprehensive Analysis of
1036 Hyper-heuristics, *Intelligent Data Analysis* 12 (1) (2008) 3–23.
- 1037 [45] P. Refaeilzadeh, L. Tang, H. Liu, Cross-Validation, in: L. LIU, M. T.
1038 ÖZSU (Eds.), *Encyclopedia of Database Systems*, Springer US, Boston,
1039 MA, 532–538, 2009.
- 1040 [46] L. Rokach, Ensemble-based Classifiers, *Artificial Intelligence Review*
1041 33 (1) (2010) 1–39.

- 1042 [47] R. Ros, N. Hansen, A Simple Modification in CMA-ES Achieving Linear
1043 Time and Space Complexity, in: G. Rudolph, T. Jansen, N. Beume,
1044 S. Lucas, C. Poloni (Eds.), *Parallel Problem Solving from Nature –*
1045 *PPSN X*, Springer Berlin Heidelberg, Berlin, Heidelberg, 296–305, 2008.
- 1046 [48] R. Tanabe, A. Fukunaga, Success-history Based Parameter Adaptation
1047 for Differential Evolution, in: *2013 IEEE Congress on Evolutionary*
1048 *Computation (CEC)*, IEEE, 71–78, 2013.
- 1049 [49] R. Tanabe, A. S. Fukunaga, Improving the Search Performance of
1050 SHADE Using Linear Population Size Reduction, in: *2014 IEEE*
1051 *Congress on Evolutionary Computation (CEC)*, IEEE, 1658–1665, 2014.
- 1052 [50] C. Ting, R. Liaw, T. Wang, T. Hong, Mining Fuzzy Association
1053 Rules Using a Memetic Algorithm Based on Structure Representation,
1054 *Memetic Computing* 10 (1) (2018) 15–28.
- 1055 [51] I. Triguero, J. Derrac, S. García, F. Herrera, A Taxonomy and Exper-
1056 imental Study on Prototype Generation for Nearest Neighbor Classifi-
1057 cation, *IEEE Transactions on Systems, Man, and Cybernetics–Part C*
1058 42 (1) (2012) 86–100.
- 1059 [52] I. Triguero, S. García, F. Herrera, Differential Evolution for Optimiz-
1060 ing the Positioning of Prototypes in Nearest Neighbor Classification,
1061 *Pattern Recognition* 44 (4) (2011) 901–916.
- 1062 [53] I. Triguero, D. García-Gil, J. Maillo, J. Luengo, S. García, F. Herrera,
1063 Transforming Big Data into Smart Data: An Insight on the Use of the
1064 k-nearest Neighbors Algorithm to Obtain Quality Data, *WIREs Data*
1065 *Mining and Knowledge Discovery* 9 (2) (2019) 1289.
- 1066 [54] I. Triguero, S. González, J. M. Moyano, S. García, J. Alcalá-Fdez, J. Lu-
1067 engo, A. Fernández, M. J. del Jesús, L. Sánchez, F. Herrera, KEEL 3.0:
1068 An Open Source Software for Multi-Stage Analysis in Data Mining,
1069 *International Journal of Computational Intelligence Systems* 10 (2017)
1070 1238–1249, ISSN 1875-6883.
- 1071 [55] I. Triguero, D. Peralta, J. Bacardit, S. García, F. Herrera, MRPR: A
1072 MapReduce Solution for Prototype Reduction in Big Data Classification,
1073 *NeuroComputing* 150 (2015) 331–345, ISSN 0925-2312.

- 1074 [56] L.-Y. Tseng, C. Chen, Multiple trajectory Search for Large Scale Global
1075 Optimization, in: Proceedings of the IEEE Congress on Evolutionary
1076 Computation, 3052–3059, 2008.
- 1077 [57] F. Wilcoxon, Individual Comparisons by Ranking Methods, Biometrics
1078 Bulletin 1 (6) (1945) 80–83.
- 1079 [58] D. R. Wilson, T. R. Martinez, Reduction Techniques for Instance-Based
1080 Learning Algorithms, Machine Learning 38 (3) (2000) 257–286.
- 1081 [59] I. H. Witten, E. Frank, Data Mining: Practical Machine Learning Tools
1082 and Techniques, Elsevier, 2017.
- 1083 [60] X. Yao, Y. Liu, G. Lin, Evolutionary Programming Made Faster, IEEE
1084 Transactions on Evolutionary Computation 3 (2) (1999) 82–102.
- 1085 [61] A. A. Zaher, R. Berretta, N. Noman, P. Moscato, An Adaptive Memetic
1086 Algorithm for Feature Selection Using Proximity Graphs, Computa-
1087 tional Intelligence 35 (1) (2019) 156–183.
- 1088 [62] J. Zhang, A. C. Sanderson, JADE: Adaptive Differential Evolution with
1089 Optional External Archive, IEEE Transactions on Evolutionary Com-
1090 putation 13 (5) (2009) 945–958.
- 1091 [63] Z. Zhu, S. Jia, Z. Ji, Towards a Memetic Feature Selection Paradigm
1092 [Application Notes], IEEE Computational Intelligence Magazine 5 (2)
1093 (2010) 41–53.