

Dynamic Selection of Evolutionary Operators Based on Online Learning and Fitness Landscape Analysis

Pietro A. Consoli · Yi Mei · Leandro L.
Minku · Xin Yao

Received: date / Accepted: date

Abstract Self-adaptive mechanisms for the identification of the most suitable variation operator in Evolutionary Algorithms rely almost exclusively on the measurement of the fitness of the offspring, which may not be sufficient to assess the optimality of an operator (e.g., in a landscape with an high degree of neutrality). This paper proposes a novel Adaptive Operator Selection mechanism which uses a set of four Fitness Landscape Analysis techniques and an online learning algorithm, Dynamic Weighted Majority, to provide more detailed information about the search space in order to better determine the most suitable crossover operator. Experimental analysis on the Capacitated Arc Routing Problem (CARP) has demonstrated that different crossover operators behave differently during the search process, and selecting the proper one adaptively can lead to more promising results.

Keywords Adaptive Operator Selection · Online Learning · Capacitated Arc Routing Problem

P. A. Consoli
CERCIA, School of Computer Science, University of Birmingham, Birmingham, West Midlands, B15 2TT, United Kingdom
E-mail: p.a.consoli@cs.bham.ac.uk

Yi Mei
Victoria University of Wellington, Wellington 6140, New Zealand
E-mail: yi.mei@ecs.vuw.ac.nz

Leandro L. Minku
Department of Computer Science, University of Leicester, University Road, Leicester, LE1 7RH, United Kingdom
E-mail: leandro.minku@leicester.ac.uk

Xin Yao
CERCIA, School of Computer Science, University of Birmingham, Birmingham, West Midlands, B15 2TT, United Kingdom
E-mail: x.yao@cs.bham.ac.uk

1 Introduction

Parameter Setting is an important area of research in the Evolutionary Computation field. Since an a-priori identification of the optimal configuration of the parameters is always time-consuming and often impractical, one must employ a dynamic selection strategy of the optimal configuration which is performed while the search is being executed. In addition, a static set of parameters is not always the optimal choice for a large number of problems where self-adapting techniques have proven to be more effective [14].

The problem of identifying the most suitable variation operator among several, also known as Adaptive Operator Selection (AOS), can be divided into two sub-tasks: the Credit Assignment (CA) mechanism, used to evaluate the performance of the operators, and the Operator Selection (OS) Rule, necessary to determine the most suitable operator using the information provided by the CA mechanism. The majority of the Credit Assignment approaches in literature are based on the evaluation of the fitness of the offspring generated by the operator, which is compared either to the current best solution [11], to the median fitness [20] or to the parents' fitness [2]. A different strategy evaluating both fitness and diversity of the offspring was proposed in [28]. The reward has been mostly considered as the value assessed during the last evaluation (*Instantaneous* reward), as the average reward over a window of the last N evaluations (*Average* reward), and as the biggest improvement achieved over a window of the last N evaluations (*Extreme* reward)[15].

The use of alternative metrics has been recently considered in [36], where an evolvability metric replaces the evaluation of the fitness. A different approach for population-based meta-heuristics, proposed in [8], assesses the reward as the proportion of solutions generated by each operator which have been selected by the ranking phase of the evolutionary algorithm. The Credit Assignment mechanism is coupled with Operator Selection rules such as Probability Matching [16], Adaptive Pursuit[38] or Multi Armed Bandit solvers (MAB) [10]. Several improvements of the MAB strategy have been proposed, as in [3], [7], [22]. Reinforcement Learning has been also used in Parameter Setting [21], as in [13], where a Reinforcement Learning procedure is adopted to modify the parameters on-the-fly, or in [34] where the selection probability of the operators is adaptively changed using a Reinforcement Learning approach.

From the analysis of the existing literature, it is clear that almost all the existing CA strategies rely exclusively on the mere evaluation of the fitness of the offspring. However, the information provided by the fitness at a single generation may not be sufficient to assess the optimality of an operator (e.g. in a landscape with a high degree of neutrality). The purpose of our work is therefore to develop a new dynamic CA mechanism which considers a suite of measures, and that can be adopted also as an Operator Selection Rule. We consider the Memetic Algorithm with Extended Neighborhood Search (MAENS*) [8] algorithm as a case study and for comparison purposes. More specifically, we aim to answer the following research questions in our paper:

- **RQ1** : What kind of additional information we can provide to the Credit Assignment technique for a more “aware” calculation of the reward and does this information effectively help to improve the prediction ability of the algorithm?
- **RQ2** : What technique would be useful to handle such data and to select the most suitable operator in such a dynamic environment? Would the prediction ability of the technique be better than that of MAENS*? Would the use of this technique improve the optimization ability of MAENS*?

The contributions of our work include:

- An ensemble of four different online Fitness Landscape Analysis techniques, performed during the execution of the MAENS* algorithm in order to give a more accurate description of the current population (RQ1);
- A Credit Assignment technique based on the use of an online learning algorithm to predict the reward of the most suitable operator (RQ2);
- Two different Reward Measures are studied: one based on the survival ability of the offspring and another one based on the analysis of their diversity.

This work extends our previous work in [9] with new experiments and contributions. In particular: (a) we investigate the use of a novel Reward Measure called Diversity Based Reward (DBR); (b) we study the adoption of a different Operator Selection Rule, named Concurrent Strategy (CS); (c) we extend our analysis by testing our algorithms on a dataset of large CARP instances. The results of the experiments carried out show that the proposed approach is able to produce results with comparable solution quality to a state-of-the-art strategy and reveal how in some cases the presence of a set of measures have a beneficial effect on the optimization ability of the AOS.

The rest of the paper is organized as follows. Section 2 introduces the case scenario and the base MAENS* algorithm. Section 3 describes the novel Reward Measures and Operator Selection Rules investigated in this work. Section 4 describes the ensemble of Fitness Landscape Techniques used in conjunction with the CA mechanism of the MAENS* algorithm. Section 5 describes the online Learning algorithm that has been used and adapted for the CA system. Section 6 presents the proposed MAENS*-II algorithm. Section 7 describes the experiments that have been carried out to verify the assumptions of this research and their results. Finally, the last section includes the conclusions and some future work ideas.

2 Background

In order to investigate over our research questions, we consider the MAENS* algorithm [8] for the Capacitated Arc Routing Problem (CARP) [17], as the case study of this research, as it already utilizes an Adaptive Operator Selection scenario and provides a term of comparison with alternative techniques. The strong relationship between CARP and specific real world problems, such

as winter gritting, waste collection or postal service make this a problem of great interest for the scientific community, and a large number of heuristics, exact methods and meta-heuristics have been proposed for this problem and its many variants. Although the hyper-heuristic proposed in this work is applied to the Capacitated Arc Routing Problem, it would be possible to adapt it to different NP-Hard problems by replacing the low level heuristics and by identifying the best Fitness Landscape Analysis metrics that better describe the specific landscapes of the different NP-Hard problem.

2.1 MAENS*

MAENS*, the case study the for this research, extends the Memetic Algorithm named MAENS [37] introduced in 2009. MAENS is a memetic algorithm which makes use of a crossover operator, a local search combining three local move operators and a novel long move operator called MergeSplit, and a ranking selection procedure called Stochastic Ranking (SR)[33]. The major differences between MAENS and MAENS* are: (a) MAENS uses a single crossover operator, whereas MAENS* uses a set of crossover operators, (b) a dynamic MAB mechanism (dMAB) [15] is adopted as an AOS rule, (c) a novel CA mechanism assigns a reward to the operators which is proportional to the number of solutions generated by each operator that “survived” the ranking phase, named Proportional Reward, (d) the Stochastic Ranking is improved considering also the diversity of the solutions (dSR) using a (e) novel diversity measure for the CARP search space.

The dMAB [15] approach, adopted in this work, combines the UCB1 algorithm [1] with the Page-Hinckley (PH) statistical test [19] to detect changes in the environment. When the PH test is triggered, the MAB system is restarted and the information gathered in the previous generations is discarded. The MAENS* algorithm represents one case study of our research, as the presence of a suite of crossover operators allows the study of other AOS approaches.

2.2 Capacitated Arc Routing Problem

The Capacitated Arc Routing Problem (CARP) can be formally defined as the problem of minimizing the total service cost of a routing plan, given a set \mathbf{T} of tasks (which corresponds to a subset of the arcs of a graph) and a fleet of \mathbf{m} vehicles with capacity \mathbf{C} . Each task \mathbf{t} has a *service cost* \mathbf{sc} , a *demand* \mathbf{d} (the load of the vehicle necessary to service the task), a unique \mathbf{id} , a reference to its head and tail vertices, and must be served once and entirely within the same route \mathbf{R}_j . A CARP solution S can be represented as

$$S = \{\{t_0, t_k, \dots, t_l, t_0\}, \dots, \{t_0, t_p, \dots, t_q, t_0\}\}$$

which is a permutation of the whole set of tasks, divided into several routes \mathbf{R}_j . Each route must start and end in a specific vertex called *depot*. We use a

dummy task t_0 with null demand and service cost to show the start and the end of a route in the depot. The service cost of a single route is calculated by adding the service cost of all the tasks in the route plus the cost of the shortest path \mathbf{sp} between each task. The problem can be formally defined as follows:

$$\min \text{TC}(S) = \sum_{i=1}^{\text{length}(S)-1} (sc(t_i) + sp(t_i, t_{i+1})),$$

subject to the constraints

$$\text{load}(R_j) \leq C, \text{app}(t_i) = 1 \text{ and } \forall t_i \in T, m \leq n_{veh},$$

$$\text{load}(R_j) = \sum_{i=1}^{\text{length}(R_j)} d(t_{ij}),$$

where $\text{app}(t_i)$ gives the number of appearances of tasks t_i in the sequence of the tasks in S and n_{veh} is the number of available vehicles.

2.3 Recombination Operators for CARP

As explained in section 2.1, MAENS* uses a set of crossover operator, instead of a single crossover operator. This section describes the four crossover operators introduced in MAENS* to deal with the CARP problem.

2.3.1 Greedy Sequence Based Crossover (GSBX)

The GSBX operator, can be considered as a variant of the Sequence Based Crossover (SBX) operator. In this case, one route is extracted from each parent solution following a greedy strategy that influences the selection towards those routes whose vehicle is still not full (the total demand of the route is the smallest). Once selected, the two routes are recombined using a one point crossover mechanism. The new route replaces the original ones in the parent solutions to generate the offspring. Since problems caused by double servicing of tasks or tasks not being serviced might arise, the solution goes through a repair phase to guarantee its feasibility.

2.3.2 Greedy Route Crossover (GRX)

In GRX, an offspring is created by alternatively copying the routes of the parents into it. Routes are extracted from the parent solutions, giving higher priority to the routes with a higher quality measure. Tasks that have been inserted into the offspring are consequently removed from both parents, to avoid the double servicing of tasks. The procedure is repeated until the remaining routes in the parents have less than a certain amount of tasks. In that case, the remaining tasks are inserted in the existing routes or merged into new ones.

2.3.3 Pivot Based Crossover (PBX)

Two routes are randomly selected from the parent solutions. The PBX operator works by identifying, among the tasks belonging to such routes, the one that is most suitable to be placed in the middle of a route, which is named *pivot*, as it splits the route in two parts. The route is then rebuilt inserting the remaining tasks in the position that minimizes the total service cost of the route. Finally, the offspring is obtained by replacing the original route in one of the two parents. As in the case of GSBX, the solution goes through a repair phase to guarantee the feasibility of the solution.

2.3.4 Shortest Path Based Crossover (SPBX)

The SPBX operator works analogously as the PBX operator, except that in this case the *pivot* is represented by a path between two of the available tasks. The couple of pivoting tasks is selected as the one that serves the largest number of available tasks along their path and that minimizes the distance between the extremities of the path and the deposit.

3 Adaptive Operator Selection

As previously mentioned, AOS is conventionally composed of two different sub-tasks, the Credit Assignment and the Operator Selection. For the former part, we propose the use of two different Reward Measures, named Proportional Reward (PR) and Diversity Based Reward (DBR). For the latter, we study the performance of two different strategies: a simple Instantaneous Reward (IR) approach and a Concurrent Strategy (CS) based approach.

3.1 Credit Assignment

The choice of the proper Credit Assignment Strategy can be fundamental for the performance of the algorithm. As one objective of this work is to evaluate more than just the fitness of the individuals, we adopt two different strategies that involve the evaluation of different measurements. The first one, named Proportional Reward, was first used in [8] together with a Multi Armed Bandit approach. For the second case, we develop a novel measure based on the evaluation of the diversity of the offspring, named Diversity Based Reward.

3.1.1 Proportional Reward (PR)

PR[8] is a measure of the survival ability of the offspring generated by each crossover operator. We assign a reward r , where $r \in [0, 1]$ corresponds to the percentage of the solutions that have survived the selection phase of the algorithm, and are going to become the parent population for its generation.

```

1 find  $p_1(t)$  and  $n_1(t) \forall t$  in  $S_1$ ;
2 find  $p_2(t)$  and  $n_2(t) \forall t$  in  $S_2$ ;
3 for (each task  $t$ ) do
4   if (task  $t$  is served in both
      solutions) then
5     if (  $p_1(t) = p_2(t)$  ) then
6       add one;
7     end
8     if (  $n_1(t) = n_2(t)$  ) then
9       add one;
10    end
11  end
12 end
13 divide the obtained value by  $2N$ .

```

Fig. 1: Diversity measure for CARP solutions

The use of this technique is a way to entrust the algorithm itself for the evaluation of the offspring. In the case of MAENS*, the offspring able to survive the ranking phase are evaluated according to their fitness value, the amount of violation of the constraints and the average pairwise diversity from the other individuals of the population. The performance of the crossover operator is in this case evaluated at the end of the generation: rather than evaluating the individuals as soon as they are generated, the PR evaluates their performance in a longer period of time (e.g. an iteration). The PR can be formally represented with the following formula:

$$PR(i)^t = \frac{|x_i : x_i \in \text{parent}^{t+1}|}{|\text{parent}^{t+1}|}$$

where i refers to the i -th operator, x_i is an individual obtained through the use of operator i , t is the t -th generation and parent^{t+1} is the parent population at the $t + 1$ generation. If more than one operator is used during the same generation, the PR can be calculated in the following way:

$$PR(i)^t = \frac{|x_i : x_i \in \text{parent}^{t+1}|}{|\text{offspring}_i^t|}$$

where offspring_i^t is the set of individuals generated using the operator i during the t -th generation.

3.1.2 Diversity Based Reward (DBR)

In the case of the DBR, we propose an approach that is opposite to that of the PR, as we evaluate the crossover offspring as soon as they have been generated. As one purpose of the crossover operator is that of introducing diversity in the population through the exploration of new areas of the landscape, we adopt a measure of the diversity introduced by the offspring. In particular, for each operator, we want to measure how distant the offspring are from the parent

population, and how wide is the area explored. Therefore, we define a parent distance measure

$$P_d(x) = \frac{d(x, p_1) + d(x, p_2)}{2}$$

as the average distance from the offspring x to its parents p_1 and p_2 and we can consequently compute the average parent distance for operator i , $P_d(i)$, by averaging the $P_d(x)$ of all the offspring generated by such operator. To measure the distance between individuals, we adopt the distance measure for CARP developed in [8]. The pseudocode of the distance measure is shown in figure 1. Since a CARP solution is represented by a sequence of tasks t , split into different routes, we can define $p_i(t)$ and $n_i(t)$ as two functions that return respectively the previous and the next tasks of task t in the sequence of solution S_i . A task t has a perfect correspondence in both solutions if its previous and next tasks match. In the most extreme cases, for two solutions S_1 and S_2 , the value of the distance measure will be equal to 1 if $p_1(t) = p_2(t)$ and $n_1(t) = n_2(t), \forall t$, and will be equal to 0 if $(p_1(t) \neq p_2(t) \text{ and } n_1(t) \neq n_2(t), \forall t)$. In the former case, the two solutions are identical, as there is a full correspondence between the $p_i(t)$ and the $n_i(t)$ of both solutions for each task t , while in the latter case the two solutions are completely different. It is important to point out that while the order in which the tasks in each route are serviced is considered, the order of the routes is not. Therefore two solutions are still identical if they have perfectly corresponding routes even if permuted in a different order. In all the other cases, when the correspondence is partial, the diversity measure will consequently assume values within the range $[0, 1]$.

We also define the *coverage* measure of the operator i

$$C_m(i) = \frac{\sum_i \sum_j d(x_a, x_b)}{N_i^2}$$

as the pairwise average distance between any pair of individuals x_a and x_b that have been generated by it, where N_i is the number of individuals generated through the use of operator i . We can compute the DBR of the i -th operator in the following way:

$$DBR(i) = P_d(i) * C_m(i).$$

Similarly to Compass [28], this Credit Assignment technique considers the diversity of the offspring as a criterion to evaluate the performance of the operators. However, there are several differences between such approaches. First, the Compass approach addresses the evaluation of both the fitness and the diversity while DBR only considers the diversity, being focused on the evaluation of crossover operators exclusively. Secondly, Compass makes use of the Hamming distance entropy as in [24] to measure the population diversity, while DBR deals with both the average pairwise distance of the offspring as well as the distance from the parent population using the CARP based diversity measure shown in figure 1.

3.2 Operator Selection Rule (OSR)

The second step of the AOS process is the Operator Selection Rule. The OSR, given the information gained through the use of the Credit Assignment mechanism, needs to decide what is the most suitable operator and how to use it. A first problem in this context is that of balancing the exploration of all the operators against the exploitation of the most useful one. In other words, while using the operator that has performed the best so far, one wants to verify whether there is another operator that can do better. A second aspect is that of identifying changes during the execution. As the search goes on, the operator that has performed the best so far might not necessarily be the best one afterwards. It is therefore necessary to balance how much of the “history” relative to each operator one must consider to perform the selection.

In this work we consider two different approaches for the OSR, namely a single operator based approach named Instantaneous Reward and a Reinforcement Learning-inspired one called Concurrent Strategy.

3.2.1 Instantaneous Reward (IR)

In the IR approach, the offspring is produced through the use of only one crossover operator per generation. As offspring and parent populations are merged in a unique population, it is still possible to evaluate all the crossover operators who have generated a solution that is still present in the population. The operator to use in the next generation ($t + 1$) is consequently selected as the operator op_i that has obtained the largest reward in the current generation (t):

$$op_i^{t+1} = \max_i (RW(op_i)^t), op_i \in \text{operators}$$

given $RW()$ as a reward measure. Those operators having produced more “extreme” improvements (e.g. discovered new optima) with respect to the others, will have a more favourable evaluation that will last for more generations, even when they have not been selected for the current generation.

The information relative to the previous performances of the operator, except for the last iteration, is discarded. IR is therefore designed to be more sensitive to changes, having a bias on the performance of the operators during the previous generation. Finally, the adoption of such approach has the potential risk of eliminating completely an operator from the competition if none of its offspring are present in the current population.

3.2.2 Concurrent Strategy (CS)

One of the disadvantages of adopting the Instantaneous Reward strategy is that it is not possible to identify changes in the environment when only one operator is used. A different approach, therefore, is that of allowing the use of all the operators during all the generations. Such approach, named CS, aims to maximize the gain obtained by using the best performing operator,

and thus allowing the generation of a larger fraction of the offspring by it, while the remaining part is still generated by the other operators. The CS is similar in its behaviour to the Adaptive Pursuit (AP) approach [38] in its intent to maintain a minimum percentage of the solutions to be generated by the less performing operators. The formula to assign the Selection Rate to each operator i , is the following:

$$SR_i = SR_{\min} + (1 - n \times SR_{\min}) \frac{e^{RW(op_i)^t / \psi}}{\sum_{j=1}^n e^{RW(op_j)^t / \psi}}$$

where SR_{\min} is the minimum selection rate, n is the number of operators, $RW(op_i)$ is the reward calculated for the operator op_i during the generation t , and ψ is a control parameter that regulates how quickly the system reacts to the changes in the environment. In this case, $n = 4$ since four operators are available.

4 Online Fitness Landscape Analysis

The existing Fitness Landscape Analysis (FLA) techniques have been analysed with the purpose to identify those that can be used in the CARP context. Such selection has been driven by both the necessity to reduce the computational effort by exploiting some calculations that are already performed by the algorithm, and the necessity to identify measures able to “capture” different features of the landscape. We identified four FLA techniques, consisting of one evolvability measure, two neutrality measures and one fitness distribution measure, to describe different features of the landscape and without much increasing the computational effort. The computation of such techniques is based on the evaluation of the neighbourhood of each solution. Such neighbourhood is already generated through the initial iteration of the local search operator of the MAENS algorithm, by using of the three different move operators involved in this process (Single Insertion, Double Insertion, Swap Insertion). The FLA techniques are employed during each generation, and their results are used as input features of an online learning algorithm to predict the value of one of the two Reward Measures introduced in section 3.1, in order to create a more accurate and informative “snapshot” of the current population which eventually might lead to a better selection of the crossover operator. A final remark is necessary about the constraints handling and how it affects the fitness of the individuals. The landscape in which MAENS* operates is that of solutions which may potentially violate the capacity constraints of the vehicles. Therefore, we consider the following fitness function, adopted from [37]:

$$f(S) = TC(S) + \lambda * TV(S)$$

where λ is an adaptive parameter depending on the cost, on the violation and on the best feasible solution found so far, $TC(S)$ is the total cost of the solution and $TV(S)$ its total violation.

The rest of this section will introduce the four FLA techniques that have been considered in this work and how they are integrated in the MAENS* algorithm.

4.1 Accumulated Escape Probability

The Accumulated Escape Probability [25] is a metric that aims to measure the evolvability, which can be defined as the capacity of the solutions to evolve into better solutions. The Accumulated Escape Probability is obtained by averaging the mean escape rate [31] (the proportion of solutions with equal or better fitness in the neighbourhood) of each fitness level with the formula:

$$aep = \frac{\sum_{f_i \in F} P_j}{|F|}, \text{ where } F = f_0, f_1, \dots, f_L$$

where f_i is a fitness level (subset of all the solutions with fitness equal to a value f_i), P_j is the average Escape Rate of all samples belonging to the f_j fitness level and L is the number of possible fitness levels. Being the mean value of a set of probabilities, the aep will be 0 when the instance is hard and higher (up to 1) otherwise. The calculation of the aep requires the analysis of the neighbourhood of each solution in order to identify how many individuals have a equal or better fitness than the original individual. We analyse therefore the evolvability of the solutions which have been selected (with probability equal to 0.2) for the local search. Since the calculation of the neighbourhood of each solution corresponds to the first step of the local search, no significant additional cost is required to compute the aep.

4.2 Dispersion Metric

The analysis of the distribution of the solutions within the landscape can be sometimes used to understand more about the difficulty that a “jump” between fitness levels requires and to gain some information on the global structure of the landscape. In this context, the Dispersion Metric (dm) [26] is a technique to obtain information about the global structure of the landscape, by measuring the dispersion of good solutions. Ideally, if good solutions are very close we might have a single funnel structure. If, on the contrary, solutions get more distant when their fitness improves, the landscape might be more like a multi funnel structure. The analysis can be described as follows:

1. A sample S of solutions is taken from the search space;
2. the best S_{best} solutions are selected from S (using a threshold value);
3. the average pairwise distances in S ($\bar{d}(S)$) and in S_{best} ($\bar{d}(S_{best})$) are calculated using the CARP diversity measure shown in figure 1;
4. the dm is obtained as the difference between $\bar{d}(S_{best})$ and $\bar{d}(S)$.

The calculation of the pairwise distance between all the individuals of the sample is already performed during the diversity-based Stochastic Ranking of MAENS* by using the distance measure shown in figure 1, and therefore requires no additional cost. Thus, the dm can be computed on the set of all the $popsiz * offset$ individuals created during each generation of MAENS*. Finally, it is possible to rely on the ranking performed by the diversity-based Stochastic Ranking operator and choose these solutions as the subset of the best ones.

4.3 Average Neutrality Ratio and Δ -fitness

Neutrality is the study of the width, distribution and frequency of neutral structures within a landscape (e.g. plateaus, ridges). A set of several neutrality measures was defined in [39]. Among these, we select the following two:

1. average neutrality ratio (\bar{r}): can be obtained by averaging the neutrality ratio (e.g. the number of solutions with equal fitness) of each individual with respect to its neighbourhood;
2. average Δ -fitness of neutral network ($\Delta(\bar{f})$): can be defined as the average fitness gain after one mutation step of each individual belonging to a neutral network.

In the same fashion as in the case of the aep, the computational effort of this technique can be absorbed by the generation of the neighbourhood of the initial solution during the local search.

5 Online learning

The AOS model followed by MAENS* is that of the Multi Armed Bandit approach, where the UCB1 [1] algorithm is used to balance the exploration and exploitation of the crossover operators and the Page-Hinckley [19] test is used to detect when a different operator has become the most suitable.

In this work, we propose the adoption of a different model. The abrupt and scarcely predictable changes of the most suitable operator which might happen during the search show many similarities to the notion of concept drift [35][32] in machine learning. Thus, in such a context, we might adopt an online learning algorithm capable of (a) predicting a reward for each operator using the online Fitness Landscape Analysis measures and (b) tracking the changes of the environment, relying only on a limited number of training instances. We can define more formally the learning problem in the following way. At a given generation of the EA, we compute the FLA metrics ($fla_1, fla_2, fla_3, fla_4$) and the reward of each operator ($RW(op_i)$). Tuples ($fla_1, fla_2, fla_3, fla_4, RW(op_i)$) are then used as training examples for the online learning algorithm, where ($fla_1, fla_2, fla_3, fla_4$) are the input features and ($RW(op_i)$) is the target output.

We employ the Dynamic Weighted Majority (DWM) [23] algorithm as our

online learning algorithm, which has proved to be one of the most effective techniques in the task of tracking concept drifts. The DWM algorithm can be described as follows. A set of learners (called experts) are used to classify the incoming instances $\{\vec{x}, y\}$, where \vec{x} is the vector of n input features and y is the output feature. Each expert e_j has its own weight w_j , and operates a classification λ of the instance. The global prediction is identified as the prediction with the largest sum of weights. All the experts which have failed to classify correctly the instance have their weights reduced by a β factor. Moreover, for every p instances, all the experts with a weight below a certain threshold θ , are deleted and a new expert is created if the global prediction is wrong.

5.1 DWM for regression tasks

As the DWM algorithm was originally conceived for classification it is necessary to adapt and modify some of its mechanism for the regression task of predicting the reward of a given operator based on the FLA techniques. A comparison between the revised DWM algorithm for the regression task (rDWM) and the original DWM itself is given in figure 2. The modifications introduced are:

1. The global prediction σ_i is obtained by calculating the weighted average of all predictions (line 10);
2. we consider a prediction correct if its difference from the output feature is less than a threshold τ (lines 5-6);
3. a new expert is created if the difference between the global prediction and the output feature is less than a t factor (lines 17-18);
4. we introduce a window containing the last n instances wTS , which is used to train the new experts upon creation (line 2).

6 MAENS*-II

The revised version of the algorithm adopting the rDWM as an AOS mechanism, named MAENS*-II, is shown in figure 3, along with the original MAENS* algorithm. Further information about MAENS* can be found in [8]. A set of four (one for each crossover operator) rDWM instances are created upon initialization of the algorithm (line 2). During each generation, one new training example is created for each rDWM instance by using the current set of FLA metrics as input features, and the reward associated to the operator as the output feature (lines 10, 13-14) obtained with a given Credit Reward strategy. The set of four rDWM instances are then used to predict the reward of each operator (line 4). Finally, an Operator Selection Rule is adopted to choose the operators to use during each generation.

Three different versions of the MAENS*-II algorithm were implemented employing the two different techniques for the Operator Selection Rule introduced

(a) Original DWM.

```

1 for (each instance  $\{\vec{x}_i, y_i\}$ ) do
2   for (each expert  $e^j$ ) do
3      $\lambda^j = \text{classify}(e^j, \vec{x}_i)$ ;
4     if ( $|\lambda_i^j| \neq y_i$ ) then
5        $w_j = \beta * w_j$ ;
6     end
7   end
8   normalize weights;
9    $\sigma_i = \text{select class with largest sum of weights}$ ;
10  if ( $p \bmod i = 0$ ) then
11    for (each expert  $e^j$ ) do
12      if ( $w_j < \theta$ ) then
13        delete expert;
14      end
15    end
16    if ( $\sigma_i \neq y_i$ ) then
17      create new expert;
18    end
19  end
20  for (each expert  $e^j$ ) do
21    train( $e^j, \vec{x}_i$ );
22  end
23 end

```

(b) DWM for the regression task.

```

1 for (each instance  $\{\vec{x}_i, y_i\}$ ) do
2   update wTS( $\vec{x}_i$ );
3   for (each expert  $e^j$ ) do
4      $\lambda^j = \text{predict}(e^j, \vec{x}_i)$ ;
5     if ( $|\lambda_i^j - y_i| > \text{tau}$ ) then
6        $w_j = \beta * w_j$ ;
7     end
8   end
9   normalize weights;
10   $\sigma_i = \text{global prediction}$ ;
11  if ( $p \bmod i = 0$ ) then
12    for (each expert  $e^j$ ) do
13      if ( $w_j < \theta$ ) then
14        delete expert;
15      end
16    end
17    if ( $|\sigma_i - y_i| > t$ ) then
18      create new expert
19      and train using wTS;
20    end
21  end
22  for (each expert  $e^j$ ) do
23    train( $e^j, \vec{x}_i$ );
24  end

```

Fig. 2: Pseudocode of DWM (left side) and DWM for the regression task (right side). The novelties introduced in the latter version, discussed in section 5.1, are highlighted with a gray background color.

in section 3.2 as well as the two different Credit Assignment mechanism presented in section 3.1. All the experiments were performed using the *weka* [18] implementation of REPTrees as base learners. Table 1 summarizes a list of the different versions of the algorithm and a description of their components. It is worth noting that the combination of the DBR strategy and the Instantaneous Reward was not considered, as the strategy of measuring the reward of the crossover offspring and the use of only one operator during each generation would lead to the its exclusive use for the whole execution of the algorithm.

Table 1: List of Algorithms. Each row represents a different combination of one Operator Selection Rules and one Reward Measure strategies.

combination	Operator Selection Rule	Reward Measure
a	Concurrent Strategy	Diversity Based Reward
b	Concurrent Strategy	Proportional Reward
c	Instantaneous Reward	Proportional Reward

(a) MAENS* pseudocode	(b) MAENS*-II pseudocode
<pre> 1 initialize a population <i>pop</i> of <i>popsi</i> individuals; 2 while (<i>termination condition not met</i>) do 3 choose the crossover operator <i>op_i</i> using <i>dMAB</i> 4 generate a population <i>pop_x</i> of <i>popsi*offset</i> individuals, choosing the parents from <i>pop</i> \cup <i>pop_x</i>; 5 generate <i>pop_{ls}(i)</i> for each individual <i>pop_x(i)</i> with probability = 0.2; 6 if (<i>pop_{ls}(i)</i> is better than <i>pop_x(i)</i>) then 7 overwrite <i>pop_xi</i>; 8 end 9 use d-Stochastic Ranking and overwrite <i>pop</i>; 10 use the PR measure to calculate the reward <i>R_i</i> for each <i>op_i</i>; 11 end </pre>	<pre> 1 initialize a population <i>pop</i> of <i>popsi</i> individuals; 2 initialize four rDWM_i instances 3 initialize four rewards <i>rw_i</i> 4 while (<i>termination condition not met</i>) do 5 choose the crossover operator <i>op_i</i> with largest <i>rw_i</i> 6 generate a population <i>pop_x</i> of <i>popsi*offset</i> individuals, choosing the parents from <i>pop</i> \cup <i>pop_x</i>; 7 generate <i>pop_{ls}(i)</i> for each individual <i>pop_x(i)</i> with probability = 0.2; 8 if (<i>pop_{ls}(i)</i> is better than <i>pop_x(i)</i>) then 9 overwrite <i>pop_xi</i>; 10 end 11 calculate $Y^t = \{aep, \bar{r}, \Delta(\bar{f}), dm\}$ 12 use d-Stochastic Ranking and overwrite <i>pop</i>; 13 for each <i>op_i</i> do 14 <i>out_i</i> = <i>CreditAssignment</i>(<i>op_i</i>) 15 <i>rw_i</i> = rDWM_i(<i>Y^t</i>, <i>out_i</i>) 16 end 17 end </pre>

Fig. 3: Pseudocode of MAENS* (right side) and MAENS*-II(left side). The novelties introduced in the latter version, discussed in section 6, are highlighted with a gray background color.

6.1 Improvements on Local Search Efficiency

One of the most effective features of MAENS [37] is its Local Search, which, however, has a high computational cost - the algorithm spends around 95% of its runtime performing this operation. Although the proposed modifications to the original MAENS algorithm, as explained in section 4, cause no significant increase of the runtime of the algorithm, a fast implementation of MAENS local search is introduced, which helped reducing effectively the runtime without incurring into extra memory consumption. The approach is similar to the one introduced in [40] for the Vehicle Routing Problem, but without relying on the use of memory.

The approach can be summarized by the following points:

1. Every individual a in the neighbourhood of a solution x is represented as a move M , where M stores the information relative to the move operator op_i such that $op_i(x) = a$, the tasks involved in the move, and the variations in terms of fitness and violation of the constraints w.r.t. the values of the initial solution;
2. The set of moves representing the whole neighbourhood is split into $V = R * R$ subsets, where R is the number of routes of the initial solution. Each

subset contains the moves relative to the move operators applied to the tasks belonging to the routes R_i and R_j ;

3. During the first iteration of the local search, the whole neighbourhood of moves is produced. A storage array of size V is kept to store the best solution of each subset;
4. The best move in the neighbourhood is identified with a computational time of $O(M)$. If the best move belongs to the subset relative to the routes R_i and R_j , the positions in the storage relative to the combination of either routes are set to null;
5. In the following iterations, the local search produces only the moves involving either the route R_i or R_j or both. The positions of the storage relative to such moves are consequently updated.

After the first iteration, the number of subsets to be evaluated is therefore decreased from R^2 to $2R + 1$, resulting in a significant reduction in terms of size of the neighbourhood that is necessary to evaluate during each local search iteration. It is worth mentioning that the use of the *move* notation itself reduces the cost of evaluating the fitness and the violation of one individual from $O(n)$ to $O(k)$, where n is the number of tasks and k is equal either to 7 or 8 (depending on the move operator considered).

7 Experimental Studies

A set of experiments were designed to understand the behaviour of MAENS*-II. As a first step, an oracle based on the Proportional Reward was implemented with the purpose of analysing a set of CARP instances in order to obtain optimal crossover operator selection rates and to analyze them. The oracle can be briefly described as follows. Four different populations are obtained during each generation by using each crossover operator. All the individuals of the four generations are merged into a single population which is sorted using the MAENS* ranking operator. The Proportional Reward mechanism is therefore used to assess the best operator. The results achieved by the oracle show that the predictions operated by the dMAB are not optimal, as better results can be achieved. Besides, the results of the oracle should be considered “optimal” only when the Proportional Reward strategy is considered, because they might not necessarily be optimal when in presence of a set of multiple measures, as in the case of MAENS*-II, or when a different credit assignment strategy is considered.

Two different datasets are considered for the experiments. The first one, named Group A, is composed of instances taken from the known benchmark test sets egl [12], Beullen’s C , D , E , F [5] and val [4]. The second group (Group B) corresponds to the large scale CARP instances of the dataset EGL-G [6]. The characteristics relative to each instance, in terms of number of vertices, number of edges, number of required edges and best fitness value found in literature are included in table 3.

We provide an example of a solution for the $D07$ problem instance produced

by one of the variants of the MAENS*-II algorithm in figure 8, to clarify what kind of results this algorithm produces.

The set of parameters adopted in all the MAENS*-II algorithm variants, included in table 2, were identified by a series of test-and-trial attempts and might not correspond to the most optimal choice. With regards to the parameters that are common with the other algorithms (MAENS and MAENS*) we adopted the same set of values in order to exclude different results due to different parameter configurations. These parameters can be identified in table 2. All the final results were obtained by averaging the output of 30 independent runs.

7.1 Single Operator Scenario

In order to understand the improvement achievable by MAENS*-II, the algorithm was executed on the two benchmark sets considering each of the four available crossover operators. The results of such experiments for group A are included in table 5. For each single operator MAENS* version, the results show the average fitness over 30 independent runs, the standard deviation and the fitness of the best individual found. The last column, named *best*, shows the results of what an “optimal” adaptive operator selection would achieve (picking the best results out of the four achieved). In the second to last row (named #), the table provides the number of instances with statistically different results according to the results of a Wilcoxon Rank-Sum test with Holm-Bonferroni correction at the significance level of 0.05. The row at the bottom (named *W*) shows the number of comparison won against the other algorithms. A Wilcoxon Rank-Sum test was performed on the results achieved on every instance by each pair of algorithms, with Holm-Bonferroni correction to deal with the multiple comparisons. The results across all the problem instances were then compared using the Wilcoxon Signed-Rank test. Each problem instance with comparable results was treated as paired results and therefore omitted from the test. The results of such test, subject to the Holm-Bonferroni correction, are included at the bottom row (*pBest*) of table 5.

For Group A, it is possible to notice how there is a great number of instances for which the four versions achieve statistically different results. The only exception is represented by the comparison between the *PBX* and the *SPBX* based versions, for which there is a limited number of statistically different instances (7); in all the other cases, the statistically different instances are at least 24. The GSBX-based operator seem to be the one performing the worst, losing the comparison to most of the instances, while the other three operators achieve the best results in a similar number of times. The statistical difference between the results of the GSBX-version and the other three versions is also confirmed by the results of the Wilcoxon Signed-Rank test. None of the four single operator based versions of MAENS* algorithm is able to perform as good as the “optimal” results (in the *best* column), as testified by the results

of the Wilcoxon Signed-Rank test included in the *pBest* row in table 5.

The results of the comparison for the CARP instances belonging to the group B are included in table 6. The table shows the results of the four different versions of the algorithm, based on the use of one of the four crossover operators available. Analogously to the previous, the table presents the average fitness (*best*), the standard deviation (*std*) and the best result found (*best*) for each algorithm. The results show how the best results are achieved always by the *PBX* and *SPBX* based versions of the algorithm (providing better results on all statistically different instances against the other two versions). The *GRX* based version, in contrast, is the one that performs the worst (losing the comparison 5 times out of 6 against the *GSBX* based version and on all the statistically different instances for the other two versions).

In both datasets, *SPBX* and *PBX* operators appear to be the operators whose usage leads to the best results. This can be explained by the fact that such operators can introduce a fair amount of diversity in the offspring as one or more routes are built from scratch. On the other hand, they maintain the good traits of the parents copying the routes that are not affected by the recombination. The *GSBX* operator, in contrast, might not introduce much diversity in the offspring as the new routes are a combination of the subroutes of the parents. Therefore, despite being the least disruptive operator, on the long term it produces a minor contribution than *SPBX* and *PBX*. The *GRX* operator, on the other hand, has a larger disruptive capacity as only the best routes are preserved in the offspring. In the context of large instances with a great number of routes as in the case of dataset B, therefore, this operator might introduce an excessive level of exploration and consequently perform worse than the others.

A further experiment was conducted to analyse the behaviour of the four crossover operators. A population of 10000 solutions was generated using the initialization operator. Each operator was then used to generate a population of 10000 solutions, using a random parent selection mechanism. Table 4 reports the number of instances for which the operators achieved the worst results in terms of fitness, violation, average pairwise distance of the offspring population and average distance from the parents. This experiment has been repeated for both datasets.

The results show that in dataset A the operator *GSBX* achieves the worst results in the largest number of instances for each of the characteristics analysed. This is coherent with the results achieved by the four evolutionary algorithms. On the other end, for dataset B, *GRX* is the worst algorithm for both fitness and violation and the second worst for both the diversity measures, which reflects the behaviour of the algorithm. It is however worth specifying that these results refer to the behaviour of the operators with a population of low quality solutions (as they are generated through the use of the initialization operator) and might not necessarily reflect the behaviour of the crossover operators during the most advanced phases of the search.

Table 2: Parameters of the MAENS*-II algorithms. In the upper part of the table, we report the set of parameters and the respective values that are shared with the MAENS* algorithm. In the bottom part of the table, the new parameters introduced for the MAENS*-II algorithm.

MAENS* parameters		
Name	Description	Value
psize	population size	30
ubtrial	maximum attempts to generate a solution	50
opsize	size of the offspring during each generation	6*psize
parent selection	crossover parent selection strategy	random selection
P_{ls}	probability of performing the local search	0.2
pMS	routes selected during MergeSplit	2
G_{max}	maximum generations	500
SR_{r1}	probability of sorting solutions using diversity	0.25
SR_{r2}	probability of sorting solutions using fitness	0.70
MAENS*-II parameters		
Name	Description	Value
p	expert removal period	5
β	decrease factor for expert weights	0.75
τ	expert weight reduction threshold	0.05
θ	threshold for expert removal	0.05
t	threshold for expert creation	0.10
ψ	control parameter for concurrent strategy	0.002

7.2 Operator Selection Rules and Reward Measures: a Comparison

The performance of the algorithm using different Operator Selection Rules and Reward Measures is shown in tables 8 and 9, respectively for the groups A and B. We include the results of the three combinations introduced in table 1, along with the optimal result considering the best performance of the single-operator versions (in the last column, named *best*). In table 8, the results of the statistical tests show how the three versions of the algorithm achieve statistically different results only on a limited subset of the instances (at most 7 between MAENS*-IIa and MAENS*-IIc). The versions achieving the best results appear to be the ones adopting the Concurrent Strategy as an OSR (MAENS*-IIa and MAENS*-IIb). The two versions achieve extremely similar results (differing only in 3 instances), while the version using the Instantaneous Reward (MAENS*-IIc) differs from the other two respectively in 7 and 5 instances, and loses the comparison in the majority of the cases. In contrast, MAENS*-IIc is the variant that differs the least w.r.t. the best results achieved by the single-operator versions of the algorithm (only 6 statistically different instances, while the other two versions differ in 8 and 9 instances).

Although such results show small differences between the performances of the algorithms when adopting one OSR rather than the other, it is possible to see

Table 3: Characteristics of the instances of groups A (top part) and B (bottom part). For each instance the table provides informations such as the number of vertices of the graph ($|V|$), the number of required edges or tasks ($|R|$), the number of edges of the graph ($|E|$) and the best known solution in literature (BK).

instance	$ V $	$ R $	$ E $	BK	instance	$ V $	$ R $	$ E $	BK
C1	69	79	98	1590	e3-C	77	87	98	10163
C10	60	55	82	2190	e4-A	77	98	98	6408
C11	83	94	118	1725	e4-B	77	98	98	8884
C15	97	107	140	1765	e4-C	77	98	98	11427
C18	93	121	133	2315	S1-B	140	75	190	6384
C5	56	65	79	2410	S2-A	140	147	190	9824
C6	38	51	55	855	S2-B	140	147	190	12968
C9	76	97	117	1775	s2-C	140	147	190	16353
D1	69	79	98	725	s3-A	140	159	190	10143
D11	83	94	118	920	s3-B	140	159	190	13616
D21	60	76	84	695	s3-C	140	159	190	17100
D23	78	92	109	715	s4-A	140	190	190	12143
D7	54	52	70	735	s4-B	140	190	190	16093
D8	66	63	88	615	s4-C	140	190	190	20375
E1	73	85	105	1855	F1	73	85	105	1065
E11	80	94	113	1810	F11	80	94	113	1015
E12	74	67	103	1580	F12	74	67	103	900
E15	85	107	126	1555	F14	53	55	72	1025
E19	77	66	103	1400	F19	77	66	103	685
E21	57	72	82	1700	F24	97	86	142	975
E23	93	89	130	1395	F4	70	77	99	930
E5	68	61	94	2130	F7	73	50	94	1080
E9	93	103	141	2160	F9	93	103	141	1145
e1-B	77	51	98	4498	val4D	41	69	69	526
e2-B	77	72	98	6305	val5D	34	65	65	573
e3-B	77	87	98	7704	val8C	30	63	63	518
					val10D	50	97	97	525

instance	$ V $	$ R $	$ E $	BK	instance	$ V $	$ R $	$ E $	BK
EGL-G1-A	255	347	375	970495	EGL-G2-A	255	375	375	1061103
EGL-G1-B	255	347	375	1085097	EGL-G2-B	255	375	375	1173286
EGL-G1-C	255	347	375	1201030	EGL-G2-C	255	375	375	1295036
EGL-G1-D	255	347	375	1325317	EGL-G2-D	255	375	375	1430267
EGL-G1-E	255	347	375	1461469	EGL-G2-E	255	375	375	1557159

Table 4: Statistics relative to the performances of the four crossover operators for datasets A and B. The table shows the number of instances for which each of the four crossover operators (GSBX, GRX, PBX, SPBX) achieved the worst results in terms of fitness, violation, average pairwise diversity and distance from parents. The values are obtained averaging the statistics of four different populations of size 10000 generated using each crossover respectively from the same parent population of 10000 individuals.

Dataset A	GSBX	GRX	PBX	SPBX
Fitness	18	14	14	12
Violation	17	12	13	16
Diversity	17	16	11	16
Parent Distance	16	12	13	15

Dataset B	GSBX	GRX	PBX	SPBX
Fitness	3	3	3	1
Violation	3	6	1	0
Diversity	4	3	0	3
Parent Distance	0	3	4	3

how the Concurrent Strategy appears to perform slightly better. This might be explained by several factors. First, the use of more than one crossover operator might introduce higher diversity in the whole offspring population. Secondly, the capacity of monitoring and verifying the performance of all the crossover operators might be important to detect changes in the environment. With regards to the Reward Measure adopted, the two approaches achieved similar results. This could be interpreted by similar importance of the requirements that the two measures try to satisfy (diversity and survival ability of the offspring). The balance might be different when tackling larger CARP instances, as in the case of those in Group B, where the exploration ability of the operator might have a bigger impact on the performance of the algorithm. The results achieved by MAENS*-IIa and MAENS*-IIb on Group B are in-

Table 5: Experimental results on group A using alternatively each crossover operator. The first column shows the instance name (*inst*). For each operator (one among GSBX,GRX,PBX,spbx) the table includes the average fitness of the best solution (*avg*), the standard deviation (*std*), the best solution (*best*). Last column (*best*) shows the best *avg* result among the four crossover operators. In the rows at the bottom, the number of comparisons (#) against every operator with statistically different results according to the Wilcoxon Rank Sum Test with Holm-Bonferroni correction at the 0.05 significance level. The following row (*W*) shows the number of comparisons where the algorithm achieves a better average fitness. In the last rows (p and pBest), the p-values relative to the Wilcoxon Signed-Rank test with Holm-Bonferroni correction at the 0.05 significance level between each single based version and against the column of the best results.

inst	GSBX			GRX			PBX			SPBX			best
	avg	std	best	avg	std	best	avg	std	best	avg	std	best	avg
C01	4175.33	26.80	4150	4159.00	14.80	4150	4151.67	4.53	4150	4153.17	7.13	4150	4151.67
C05	5373.75	15.29	5365	5366.00	2.00	5365	5366.50	2.29	5365	5367.50	2.50	5365	5366.00
C06	2545.75	4.82	2535	2537.50	4.61	2535	2541.00	4.90	2535	2540.33	4.99	2535	2537.50
C09	5274.92	19.27	5260	5281.17	19.57	5260	5262.83	9.97	5260	5263.33	9.52	5260	5262.83
C10	4709.42	16.26	4700	4709.00	12.21	4700	4712.33	10.14	4700	4703.33	7.45	4700	4703.33
C11	4648.58	16.33	4640	4643.17	3.29	4640	4641.33	2.21	4640	4641.50	3.20	4630	4641.33
C15	4964.42	16.43	4940	4946.83	10.76	4940	4946.50	5.19	4940	4946.33	4.27	4940	4946.33
C18	5639.83	9.83	5620	5642.17	6.28	5620	5635.00	8.37	5620	5640.17	9.26	5620	5635.00
D01	3225.00	9.04	3215	3235.00	0.00	3235	3230.83	5.01	3215	3229.67	6.45	3215	3225.00
D07	3115.83	2.76	3115	3115.33	1.80	3115	3115.00	0.00	3115	3115.67	2.49	3115	3115.00
D08	3052.83	4.12	3045	3058.00	10.38	3045	3045.67	2.49	3045	3047.67	4.42	3045	3045.67
D11	3761.08	3.17	3760	3760.33	3.14	3755	3760.83	2.27	3755	3760.17	3.76	3745	3760.17
D21	3058.33	3.25	3050	3056.67	2.98	3050	3059.83	2.41	3050	3060.00	2.24	3055	3056.67
D23	3171.17	11.74	3140	3158.17	8.51	3145	3187.17	10.30	3155	3177.50	10.47	3145	3158.17
E01	4916.83	5.84	4910	4910.33	1.25	4910	4912.00	3.32	4910	4912.67	4.23	4910	4910.33
E05	4623.33	21.67	4585	4623.33	27.34	4585	4607.33	19.09	4585	4608.67	18.39	4585	4607.33
E09	5855.33	25.00	5820	5838.83	23.55	5810	5836.50	19.88	5815	5832.67	17.83	5810	5832.67
E11	4697.25	24.91	4660	4671.67	4.35	4670	4675.00	11.25	4670	4678.33	12.67	4670	4671.67
E12	4228.50	17.28	4190	4209.33	19.09	4190	4200.67	11.95	4190	4201.83	9.79	4195	4200.67
E15	4220.67	7.04	4205	4215.00	6.06	4210	4221.50	5.19	4210	4220.00	6.19	4210	4215.00
E19	3244.17	2.76	3235	3239.33	4.96	3235	3244.67	1.80	3235	3243.83	3.08	3235	3239.33
E21	3733.50	2.29	3730	3731.33	2.21	3730	3734.50	1.50	3730	3734.00	2.00	3730	3731.33
E23	3718.83	5.87	3715	3715.17	1.57	3710	3715.33	1.80	3710	3717.33	2.81	3710	3715.17
egl-e1-B	4512.47	12.04	4498	4504.87	10.94	4498	4503.03	8.78	4498	4501.77	8.41	4498	4501.77
egl-e2-B	6328.65	11.75	6317	6321.93	8.68	6317	6322.60	6.03	6317	6327.10	10.19	6317	6321.93
egl-e3-B	7792.07	15.71	7775	7780.90	9.45	7775	7784.97	8.83	7777	7782.57	4.96	7777	7780.90
egl-e3-C	10328.18	19.82	10292	10324.73	16.07	10305	10315.87	17.93	10292	10310.67	16.24	10292	10310.67
egl-e4-A	6464.97	5.39	6444	6464.23	4.26	6461	6463.47	3.00	6456	6463.90	1.83	6461	6463.47
egl-e4-B	9021.28	17.37	8988	9059.70	25.38	8988	9024.40	15.16	8998	9013.10	14.09	8988	9013.10
egl-e4-C	12032.60	1047.53	11559	11593.13	22.87	11554	11586.40	18.91	11539	11584.97	25.66	11543	11584.97
egl-s1-B	6415.70	21.28	6388	6405.50	19.41	6388	6393.17	2.48	6388	6399.43	13.95	6388	6393.17
egl-s2-A	9942.62	26.54	9895	9949.67	24.62	9890	9929.37	23.69	9889	9939.50	27.44	9889	9929.37
egl-s2-B	13201.76	35.16	13144	13244.63	90.51	13137	13163.57	30.84	13103	13181.60	29.65	13122	13163.57
egl-s2-C	16500.12	42.51	16430	16480.80	39.54	16430	16456.77	17.46	16430	16462.37	27.46	16430	16456.77
egl-s3-A	10298.59	31.43	10221	10305.10	45.56	10242	10284.60	25.31	10233	10300.73	25.47	10253	10284.60
egl-s3-B	13847.47	60.89	13713	13906.90	50.92	13771	13792.63	40.81	13713	13815.87	61.71	13707	13792.63
egl-s3-C	17317.86	38.59	17209	17290.47	41.79	17197	17292.90	39.24	17242	17287.70	37.00	17221	17287.70
egl-s4-A	12409.36	41.76	12296	12438.10	33.13	12389	12367.40	38.27	12315	12399.70	34.21	12316	12367.40
egl-s4-B	16448.97	43.34	16316	16499.40	45.08	16430	16384.27	43.71	16292	16405.10	40.22	16329	16384.27
egl-s4-C	25200.21	2151.47	20781	22201.00	143.44	21792	20801.63	103.60	20601	20796.70	112.85	20584	20796.70
F01	4046.81	3.04	4040	4047.00	3.32	4040	4047.50	3.10	4040	4046.67	3.50	4040	4046.67
F04	3499.47	5.10	3485	3500.17	3.76	3495	3498.83	5.43	3485	3500.50	5.06	3485	3498.83
F07	3347.73	33.33	3335	3338.33	17.95	3335	3338.33	17.95	3335	3355.00	40.00	3335	3338.33
F09	4750.71	9.95	4730	4743.00	7.48	4730	4751.67	12.20	4730	4753.83	10.70	4740	4743.00
F11	3850.38	13.65	3835	3845.83	11.91	3835	3839.83	6.89	3835	3844.83	10.99	3835	3839.83
F12	3416.14	28.91	3395	3450.33	21.21	3395	3423.67	22.10	3395	3410.17	16.71	3395	3410.17
F14	3339.29	11.03	3330	3349.33	18.25	3330	3349.00	6.88	3340	3359.33	14.13	3330	3339.29
F19	2553.97	9.38	2525	2531.17	10.46	2525	2565.33	9.21	2545	2564.83	12.88	2525	2531.17
F24	3234.38	11.30	3215	3230.17	13.01	3210	3245.50	7.89	3225	3244.17	6.20	3220	3230.17
val10D	531.47	1.79	528	532.73	1.26	530	530.73	1.18	528	530.50	1.52	528	530.50
val4D	532.03	3.01	530	531.63	2.44	530	530.27	0.85	530	530.00	0.00	530	530.00
val5D	583.37	1.83	579	586.80	3.53	578	581.87	2.38	575	583.60	2.42	579	581.87
val8C	524.37	1.96	521	526.80	1.08	525	524.50	1.84	521	525.00	1.83	521	524.37
	GRX	PBX	SPBX	GSBX	PBX	SPBX	GSBX	GRX	SPBX	GSBX	GRX	PBX	
#	26	31	24	26	26	27	31	26	7	24	27	7	
W	8	6	5	18	8	13	25	18	3	19	14	4	
p	0.126	0.0004	0.0021	0.126	0.1738	0.3524	0.0004	0.1738	-	0.0021	0.3524	-	
pBest	<0.001			0.002			0.005			0.007			

cluded in table 9. The two algorithms show a comparable result on 9 instances out of 10, with the only statistically different result according to the Wilcoxon Rank-Sum test being that of the instance EGL-G2-C, with a p-value of 0.0004. The similarity of the results achieved by the two different versions of the algorithm in both datasets can be explained by the fact that the use of the FLA metrics makes the algorithm more robust with respect to the Reward

Table 6: Experimental results on Group B using alternatively each crossover operator. The results of the four versions of the algorithm are split in two different rows. The first column shows the instance name (*inst*). For each operator (one among GSBX,GRX,PBX,spxb) the table includes the average fitness of the best solution (*avg*), the standard deviation (*std*), the best solution (*best*). In the rows at the bottom, the number of comparisons (#) against every operator with statistically different results according to the Wilcoxon Rank Sum Test with Holm-Bonferroni correction at the 0.05 significance level. The following row shows the number of comparisons where the algorithm achieves a better average fitness.

inst	GSBX			GRX		
	avg	std	best	avg	std	best
EGL-G1-A	983408.50	3866.28	974054	979095.97	4718.81	968747
EGL-G1-B	1091041.60	5539.83	1081857	1094405.97	7413.40	1079590
EGL-G1-C	1213921.77	5924.27	1198728	1212862.87	7906.85	1198393
EGL-G1-D	1337645.60	6620.72	1322885	1336158.93	6527.20	1326125
EGL-G1-E	1473433.10	5394.11	1461472	1472399.20	6338.61	1461155
EGL-G2-A	1107790.40	5057.72	1099946	1138050.43	10215.79	1110914
EGL-G2-B	1225440.43	5982.44	1214762	1255569.68	16865.83	1224099
EGL-G2-C	1359221.30	3798.77	1349981	1404364.32	10734.40	1369046
EGL-G2-D	1497934.97	6544.95	1486595	1563310.77	5689.30	1552126
EGL-G2-E	1641472.67	8022.86	1626564	1713877.53	8561.59	1687159
#	GRX	PBX	SPBX	GSBX	PBX	SPBX
W	6	7	7	6	8	7
	5	0	0	1	0	0

inst	PBX			SPBX		
	avg	std	best	avg	std	best
EGL-G1-A	980746.30	5681.29	970911	978411.33	4308.05	969682
EGL-G1-B	1087682.37	5542.29	1074857	1087255.93	4986.75	1079899
EGL-G1-C	1208196.23	5138.15	1198557	1210928.20	5972.77	1202072
EGL-G1-D	1330286.83	6554.80	1321271	1333503.57	6436.41	1324605
EGL-G1-E	1462940.17	5293.51	1452158	1467270.13	5003.86	1458893
EGL-G2-A	1104884.20	3252.86	1099756	1105976.47	3662.89	1098458
EGL-G2-B	1221379.87	3586.68	1213622	1220895.90	4180.51	1212440
EGL-G2-C	1351635.77	5294.39	1343015	1354111.13	5087.04	1343399
EGL-G2-D	1490662.50	4435.48	1484014	1492414.43	4495.21	1484208
EGL-G2-E	1635578.87	4851.51	1623322	1634917.43	5199.54	1623417
#	GSBX	GRX	SPBX	GSBX	GRX	PBX
W	7	8	1	7	7	1
	7	8	1	7	7	0

Measure considered. Further experiments with more aggressive Credit Assignment strategies might reveal more differences between the adoption of the two different Reward Measures. Finally, we provide a comparison with a version the algorithm selecting one crossover operator randomly during each generation. The results of such algorithm are included in table 7, in the rightmost column named *random*. At the bottom it is possible to see the number of statistically different instances according to the Wilcoxon Rank-Sum test with a level of significance of 0.05 ((line #) with respect to the three algorithms MAENS*-IIa, MAENS*-IIb and MAENS*-IIc, along with the number of times the random algorithm has won the comparison (line *W*). It is worth noting that the random algorithm achieves a fairly good performance, as it achieves statistically comparable results with the proposed techniques for most of the instances. This result could be interpreted as a probable sign of positive interaction between the crossover operators that have been considered in this case study.

7.3 Effectiveness of the FLA measures

An experiment was designed to understand whether the use of the online FLA techniques has a beneficial effect on both the optimization ability and the prediction capacity of the algorithm. Therefore, MAENS*-IIc was com-


Table 8: Experimental results on the instances of Group A for *MAENS*-IIa*, *MAENS*-IIb* and *MAENS*-IIc*. The first column shows the instance name (*inst*). For each version of the algorithm tested the table includes the average fitness of the best solution (*avg*), the standard deviation (*std*), the best solution (*best*). In the rows at the bottom, the number of comparisons against every operator with statistically different results according to the Wilcoxon Rank Sum Test with Holm-Bonferroni correction at the 0.05 significance level. The following row shows the number of comparisons where the algorithm achieves a better average fitness.

inst	MAENS*-IIa			MAENS*-IIb			MAENS*-IIc			best
	avg	std	best	avg	std	best	avg	std	best	avg
C01	4159.50	17.53	4150	4156.50	12.66	4150	4160.00	17.75	4150	4151.67
C05	5367.17	3.34	5365	5366.83	2.41	5365	5369.00	5.83	5365	5366.00
C06	2540.00	5.00	2535	2539.67	4.99	2535	2541.00	4.90	2535	2537.50
C09	5268.00	16.00	5260	5261.67	7.23	5260	5264.00	10.12	5260	5262.83
C10	4707.33	9.64	4700	4704.17	8.37	4700	4705.17	9.53	4700	4703.33
C11	4641.33	3.14	4630	4641.50	2.29	4640	4642.33	2.49	4640	4641.33
C15	4944.67	7.74	4940	4945.50	4.72	4940	4946.00	7.57	4940	4946.33
C18	5641.17	9.04	5620	5637.50	8.73	5620	5637.76	9.27	5620	5635.00
D01	3232.83	3.10	3225	3231.17	6.67	3215	3232.17	5.11	3215	3225.00
D07	3115.00	0.00	3115	3115.00	0.00	3115	3115.00	0.00	3115	3115.00
D08	3045.33	1.82	3045	3047.67	4.42	3045	3047.67	4.42	3045	3045.67
D11	3760.33	2.25	3755	3760.50	1.50	3760	3761.72	3.48	3755	3760.17
D21	3058.33	3.24	3050	3058.33	3.25	3050	3059.00	4.16	3050	3056.67
D23	3172.67	9.66	3150	3162.50	13.34	3135	3162.67	7.39	3150	3158.17
E01	4911.00	2.41	4910	4911.00	2.00	4910	4911.50	2.93	4910	4910.33
E05	4611.83	25.35	4585	4601.17	18.74	4585	4615.00	26.08	4585	4607.33
E09	5830.67	18.35	5810	5832.17	19.39	5810	5834.17	21.64	5810	5832.67
E11	4674.33	10.70	4670	4672.33	5.12	4670	4678.00	15.03	4660	4671.67
E12	4201.00	7.39	4195	4205.33	12.24	4195	4207.50	14.59	4180	4200.67
E15	4218.00	6.98	4210	4217.83	5.73	4205	4219.33	5.59	4210	4215.00
E19	3243.33	3.78	3235	3242.00	4.58	3235	3242.00	4.58	3235	3239.33
E21	3733.50	2.31	3730	3733.67	2.21	3730	3733.27	2.39	3730	3731.33
E23	3715.83	3.23	3710	3716.83	2.73	3715	3715.50	1.98	3710	3715.17
e1-B	4503.60	9.87	4498	4499.67	6.26	4498	4504.79	10.42	4498	4501.77
e2-B	6324.67	10.28	6317	6321.63	5.84	6317	6323.86	9.41	6317	6321.93
e3-B	7783.77	9.12	7775	7782.87	8.35	7777	7786.55	10.73	7777	7780.90
e3-C	10312.23	15.45	10292	10314.80	20.03	10292	10318.31	19.15	10292	10310.67
e4-A	6463.87	3.30	6454	6463.07	2.02	6461	6463.83	5.07	6446	6463.47
e4-B	9029.27	16.82	9000	9026.63	16.17	9000	9021.10	17.84	8990	9013.10
e4-C	11589.13	24.44	11540	11586.80	27.09	11536	11621.28	72.42	11555	11584.97
s1-B	6402.53	18.33	6388	6401.80	16.88	6388	6397.59	12.70	6388	6393.17
s2-A	9931.93	25.17	9889	9928.37	24.06	9889	9934.80	29.49	9881	9929.37
s2-B	13171.97	24.27	13122	13170.97	31.06	13107	13171.41	29.10	13123	13163.57
s2-C	16478.50	34.87	16425	16492.30	39.99	16442	16505.97	51.89	16434	16456.77
s3-A	10282.67	32.08	10221	10288.47	28.39	10221	10290.67	25.78	10251	10284.60
s3-B	13814.90	58.66	13722	13818.63	73.32	13717	13821.50	47.04	13747	13792.63
s3-C	17287.27	37.04	17205	17288.43	31.12	17223	17309.87	37.46	17221	17287.70
s4-A	12404.20	35.59	12301	12388.17	37.70	12304	12388.59	41.42	12316	12367.40
s4-B	16399.90	50.38	16305	16427.13	51.61	16278	16437.60	54.52	16281	16384.27
s4-C	20847.80	134.55	20603	20912.60	266.38	20565	21037.24	223.95	20648	20796.70
F01	4047.00	2.79	4040	4045.50	3.25	4040	4047.59	3.32	4040	4046.67
F04	3498.67	3.45	3485	3498.17	4.74	3485	3499.00	4.16	3485	3498.83
F07	3348.33	30.45	3335	3345.00	30.00	3335	3338.34	17.95	3335	3338.33
F09	4749.50	12.21	4730	4746.67	10.43	4730	4748.45	8.48	4730	4743.00
F11	3843.67	11.25	3835	3843.17	11.58	3835	3847.07	12.35	3835	3839.83
F12	3404.67	11.93	3395	3408.00	17.45	3395	3416.83	26.94	3395	3410.17
F14	3342.50	9.62	3330	3343.33	16.35	3330	3339.50	11.86	3330	3339.29
F19	2541.67	9.59	2525	2533.17	6.39	2525	2532.50	9.64	2525	2531.17
F24	3235.33	10.33	3215	3232.50	9.46	3215	3233.83	10.38	3210	3230.17
val10D	530.77	0.94	528	530.70	1.39	528	531.13	1.26	529	530.50
val4D	530.73	2.25	530	530.13	0.43	530	530.53	1.65	530	530.00
val5D	582.67	2.67	577	582.97	2.07	579	581.93	2.83	577	581.87
val8C	524.40	2.12	521	524.73	1.57	522	524.73	2.00	521	524.37
	b	c	best	a	c	best	a	b	best	
#	3	7	9	3	5	8	7	5	6	
W	1	5	0	2	5	0	2	0	0	


against the results achieved by MAENS*-IIc. The two algorithms produced statistically different results on 36 instances out of 53. MAENS*-IIrw achieved better results only on 3 instances, losing the comparison on 33. A Wilcoxon signed-rank test was consequently applied across the problem instances, which confirmed that the two algorithms produce significantly different results (respectively $W_{stat} = 26$ with $p < 0.05$ and $W_{stat} = 54.5$ sample size: 42). This can be interpreted as a signal that the rDWM is concretely affected by the FLA measures, which influence (in a beneficial way) the decisions made by the algorithm.

Table 9: Experimental results on the instances of Group B for MAENS*-IIa and MAENS*-IIb. The first column shows the instance name (*inst*). The second column shows the fitness of the best known (BK) solution for each instance[27]. For each version of the algorithm tested the table includes the average fitness of the best solution (*avg*), the standard deviation (*std*), the best solution (*best*)

inst	BK	MAENS*-IIa			MAENS*-IIb		
		avg	std	best	avg	std	best
EGL-G1-A	970495	978636.00	5267.70	964014	978127.07	5330.95	968157
EGL-G1-B	1085097	1086113.80	4709.52	1075069	1088504.40	5597.17	1076011
EGL-G1-C	1201030	1209512.20	5983.41	1197057	1208264.80	6225.46	1196975
EGL-G1-D	1325317	1331918.77	5702.86	1322682	1331367.00	5963.31	1318679
EGL-G1-E	1461469	1466771.97	6458.51	1451314	1465321.17	4890.69	1455995
EGL-G2-A	1061103	1107519.20	3744.93	1099674	1107461.13	2864.67	1101083
EGL-G2-B	1173286	1220912.47	4687.58	1213516	1220423.67	4751.97	1213237
EGL-G2-C	1295036	1356660.60	4883.30	1346969	1352307.90	5182.32	1338497
EGL-G2-D	1430267	1493163.27	5173.56	1482470	1494645.93	5294.38	1486269
EGL-G2-E	1557159	1635756.30	5750.90	1622468	1636974.10	6235.37	1626530



(a) Fitness distribution for *EGL-G2-A* instance (b) Fitness distribution for *egl-s1-b* instance



(c) Fitness distribution for *egl-s2-b* instance (d) Fitness distribution for *egl-e4-c* instance

Fig. 4: Boxplots relative to the four instances *EGL-G2-A* (a), *egl-s1-b* (b), *egl-s2-b* (c), *egl-e4-c* (d). The boxes refer to the first quartile, median and third quartile. Whiskers show minimum and maximum values over a sample size of 30 best fitness values relative to the independent runs of each algorithm reported in tables 5, 6, 7, 8, 9 and 10 .

7.4 Comparison with the state-of-the-art

The second research question in the introduction of this paper focuses on the performance of the proposed approach with respect to the existing ones. Therefore, the MAENS*-II variants that make use of the Proportional Re-

ward (a and b) were tested against the oracle. All the three variants were also compared against the results achieved by their base algorithm MAENS*. The results achieved by the oracle and by the MAENS* algorithm for Group A are included in table 7, in columns *MAENS** and *oracle*. In the bottom rows, the results of a Wilcoxon Rank-Sum test with Holm-Bonferroni correction at the 0.05 significance level show the number of instances with statistically different results. The results of the statistical test show how the number of statistically different results is small (4 for MAENS*-IIa and MAENS*-IIc and 2 for MAENS*-IIb). In these few instances, MAENS*-IIa and MAENS*-b perform better than MAENS*, while MAENS*-II wins the comparison in half of the instances (2 out of 4). The online learning system is therefore able to achieve results comparable to those achieved by the bandit solver.

The comparison with the oracle shows that MAENS*-IIa and MAENS*-IIb are able to achieve comparable results in most cases. In most of the instances with statistically different results, the oracle was able to perform better. It is worth noting that in a small number of instances the algorithm using the FLA measures was able to produce better results than the oracle. This is some evidence that, if the oracle represents a “lower bound” for the results that is possible to achieve using the Proportional Reward, the use of more than one measures (as in this case) can help the algorithm to achieve results beyond these bounds.

Finally, the results achieved by MAENS*-IIa and MAENS*-IIb, included in table 9, are compared against four state-of-the-art algorithms, whose results are included in table 10. We consider the results of MAENS* [8], of MAENS-RDG [29] and VND [30] and an algorithm combining Iterate Local Search and Variable Neighbourhood Descent [27].

It is possible to notice how MAENS*-IIa and MAENS*-IIb, as well as MAENS*, outperform all the other algorithms in terms of solution quality for the first 5 instances of group B 9. MAENS*-IIa, MAENS*-IIb and MAENS* produce a new best known solution for all of these instances, with MAENS*-IIa achieving the best ones on the first two instances (G1-A and G1-B), MAENS*-IIb on instances G1-C and G1-D and MAENS* finding the best one on the instance G1-E. In all these instances MAENS*-IIa and MAENS*-IIb achieve also the best average fitness in four cases. For the following 5 instances, the best results are achieved by either MAENS-RDG or VND. In all these cases, MAENS* is outperformed by both MAENS*-IIa and MAENS*-IIb. These results can be explained by the fact that their base algorithm, MAENS*, is already performing well for these instances. However, both variants MAENS*-IIa and MAENS*-IIb managed to outperform MAENS* in most of the instances. It is important to note that the runtime (not considered in this work) of these algorithms is not comparable to those of the decomposition based approaches, which manage to find these results in a fraction of the time required by MAENS*-IIa and MAENS*-IIb.

The behaviour of the algorithms can be analyzed also in terms of the fitness distribution of its solutions. Figure 4 shows the box plot relative to three representative instances belonging to Group A (*egl-e4-C*, *egl-s1-B*, *egl-s2-B*) and

Table 10: Comparison with some state-of-the art approaches for CARP. The first column shows the instance name (*inst*). For each algorithm the table includes the average fitness of the best solution (*avg*) and the best solution (*best*). The results are compared to those achieved by MAENS*-IIa and MAENS*-IIB included in table 9. The algorithms considered are MAENS*[8], MAENS-RDG[29], VND [30] and ILS-RVND[27]. No statistical test was carried out due to the partial availability of the results of the compared algorithms.

inst	MAENS*		MAENS-RDG		VNS		ILS-RVND	
	avg	best	avg	best	avg	best	avg	best
EGL-G1-A	977754.4	968897	1007223.0	1000575	10007393.0	997055	1014930.9	1004864
EGL-G1-B	1088706.5	1079793	1124751.0	1111971	1122077.0	1114120	1143221.7	1129937
EGL-G1-C	1209058.8	1195902	1251718.0	1243779	1253789.0	1243808	1270100.3	1262888
EGL-G1-D	1331595.8	1323397	1383619.0	1371443	1383997.0	1373480	1409811.9	1398958
EGL-G1-E	1469455.4	1449542	1524393.0	1512584	1525994.0	1517772	1556138.5	1543804
EGL-G2-A	1107363.0	1101559	1108916.0	1096027	1105870.0	1098454	1126561.0	1115339
EGL-G2-B	1223132.3	1213769	1222183.0	1213617	1220012.0	1211759	1237741.8	1226645
EGL-G2-C	1354725.3	1345587	1353118.0	1344148	1351845.0	1344184	1376931.6	1371004
EGL-G2-D	1495089.7	1486646	1489723.0	1481181	1489500.0	1481045	1520794.3	1509990
EGL-G2-E	1636140.6	1630656	1630132.0	1618955	1630048.0	1616119	1664230.2	1659217


one instance of Group B (*EGL-G2-A*). In the case of the *EGL-G1-A* instance, it is possible to notice how SPBX and GRX are the crossover operators whose usage leads to the distributions with the lowest median. The distribution of the three AOS considered in this case (MAENS*, MAENS*-IIa and MAENS*-IIB) are centered around the same median value, although MAENS*-IIa is capable of producing solutions of considerably better quality (bottom whisker) which translate into new minima for this instance. For the *egl-s1-B* instance (figure 4.b), the behaviour of the algorithms is quite similar, as in most of the cases the distributions lie around the same median. When considering the results of the versions of the algorithm using each crossover operator, GSBX, GRX and SPBX show a much wider distribution of their results although in the first two cases a large number of solutions are equal to the median value, while PBX results are much less spread. The different AOS strategies achieve overall comparable results. This instance represents an example of non optimal behaviour as none of the AOS strategies considered has managed to match that of the best crossover operator (GRX).

For *egl-s2-b* (figure 4.c), PBX is the operator that achieves the best results, while GRX performs the worst. MAENS*-IIB manages to achieve the same solution quality and similar median to PBX. This is also confirmed by the larger selection rate given to the PBX operator (figure 6.b).


In the case of *egl-e4-c* (figure 4.d), PBX and SPBX distributions have a similar median and similar quartiles performing the best among the four crossover operators. Among the AOS strategies, MAENS*-IIB solutions are distributed around a similar median but more spread.

7.5 Prediction Ability


In order to understand the behaviours of the algorithms, and to gain a deeper understanding of the selection mechanisms, we provide a comparison of the selection rates of the four different crossover operators, included in figures 5, 6 and 7. The plots refer to the selection rates relative to the instances *egl-s1-B*, *egl-s2-B* and *EGL-G2-A*. The y-axis in the figure refers to the Selection Rate




(a) MAENS*-IIa on s1-B




(b) MAENS*-IIb on s1-B



(c) MAENS*-IIc on s1-B




(d) MAENS* on s1-B




(e) oracle on s1-B

Fig. 5: Selection Rates for the instance *egl-s1-b*. Each graph shows the selection rates of the four different crossover operators (*GSBX*, *GRX*, *PBX*, *SPBX*) when using respectively MAENS*-IIa (fig. a), MAENS*-IIb (fig. b), MAENS*-IIc (fig. c), MAENS* (fig. d) and the oracle (fig. e)


(SR) of each crossover operators, where a SR of 0 means that the operator is not selected and a SR equal to 1 means that only that operator is selected. The x-axis corresponds to the average fitness of the population discretised into 50 intervals. We study, therefore, how the SR of the four operator changes while the search is carried out and the average fitness of the population decreases. In the first instance, *egl-s1-B* (figure 5), it is possible to notice three phases




(a) MAENS*-IIa on s2-B




(b) MAENS*-IIb on s2-B



(c) MAENS*-IIc on s2-B



(d) MAENS* on s2-B




(e) oracle on s2-B

Fig. 6: Selection Rates for the instance *egl-s2-b*. Each graph shows the selection rates of the four different crossover operators (*GSBX*, *GRX*, *PBX*, *SPBX*) when using respectively MAENS*-IIa (fig. a), MAENS*-IIb (fig. b), MAENS*-IIc (fig. c), MAENS* (fig. d) and the oracle (fig. e)


in the oracle prediction (figure 5.e). A first phase where the GRX operator is preferred over the others, an intermediate phase where the GRX and GSBX operators have nearly equal selection rates and a last phase characterized by a rise of the selection rate of the GRX operator which reaches 1 in the last moments of the search.

Both MAENS* (figure 5.d) and MAENS*-IIc (figure 5.c) award the GSBX



(a) MAENS*-IIa on G2-A

(b) MAENS*-IIb on G2-A



(c) MAENS* on G2-A

Fig. 7: Selection Rates for the three instances *EGL-G1-A*. Each row shows the selection rates of the four different crossover operators (*GSBX*, *GRX*, *PBX*, *SPBX*) when using respectively MAENS*-IIa (fig. a), MAENS*-IIb (fig. b) and MAENS* (fig. c). We do not include the results on *EGL-G2-A* for MAENS*-IIc as only versions MAENS*-IIa and MAENS*-IIb were tested on this dataset, as well as the oracle results, due to the extremely high computational cost required to perform this task

operator with the highest selection rate for the whole search, missing the prediction of the change in the environment made by the oracle. It is possible to see, however, how MAENS*-IIc increases the selection rate of GSBX more rapidly than MAENS*.

The SR of both MAENS*-IIa (figure 5.a) and MAENS*-IIb (figure 5.b) show different changes during the search, proving that the CS is more successful in predicting such events. In particular, MAENS*-IIb acknowledges the operators GSBX and PBX as the most useful ones during the search. It is worth remembering that MAENS*-IIa, makes use of a different Reward Measure and, therefore, is not comparable to the prediction made by the oracle. In this case, MAENS*-IIa, after an initial epoch of dominance of the operator *GRX*, shows an alternance of moments where the three operators *GRX*, *PBX* and *SPBX* show the highest selection rates.

On the second instance (figure 6), the oracle identifies a change in the environment halfway through the search (figure 6.e). The concept drift is not

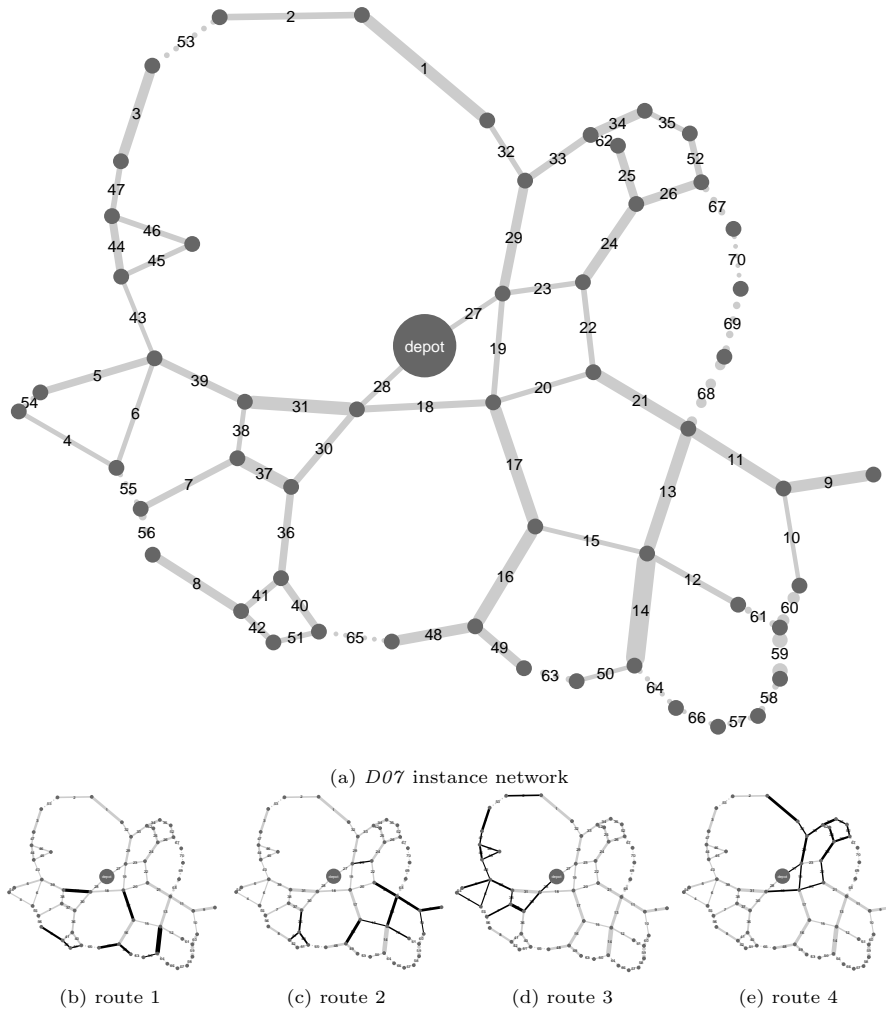


Fig. 8: The network relative to the *D07* instance (a) . The cost of serving each edge is proportional to the thickness of the line. Non required edges can be identified by dotted lines. In figures b, c, d and e the four routes that compose a solution for this problem instance, generated by the MAENS*-IIa algorithm. The edges served during each route have been highlighted in black.

detected by either MAENS* (figure 6.d) or MAENS*-IIc (figure 6.c), which, however shows an higher exploitation of the GSBX operator. MAENS*-IIb (figure 6.b) identifies the operators *GSBX* and *PBX* as the most successful ones; even in this case the change detected by the operator is not detected. As for the previous instance, MAENS*-IIa (figure 6.a) shows different moments where the three operators *GRX*, *PBX* and *SPBX* achieve the highest SR. The lowest SR for *GSBX* seems to indicate that this operator is probably the one that introduces the least diversity in the population.

For the large CARP instance EGL-G2-A, the behaviour of MAENS* (figure

7.c) shows a predominance of operator *GSBX* over the other ones. MAENS*-IIb (figure 7.b) shows a similar behaviour to that of MAENS*, identifying the *GSBX* operator as the one with the best performance during almost the whole search. Finally, MAENS*-IIa shows again an initial period of higher performance for the *GRX* operator, followed by an alternance of the *PBX* and *SPBX* operators (figure 7.a). The occurrence of this initial period of higher performance for the *GRX* operator seems to suggest that this operator is introducing the highest diversity in the initial part of the search, when the solutions are not extremely good.

The results of these experiments show that failing to detect a change in the environment does not necessarily translate into a worst performance of the algorithm and vice versa. This is confirmed by the fact that the algorithms produce good results despite the different selection rates. The relationship between the prediction ability of the algorithms and their results is therefore quite complex. There are several factors that influence its behaviour and that should be considered in order to fully grasp this mechanism, such as the interaction between the different operators, the performance of the single operators and the variation of the selection rates.

8 Conclusions and future work

In this work we proposed the adoption of a novel Adaptive Operator Selection scheme to identify the optimal crossover operator online. We consider the use of two different Reward Measure strategies, the Diversity Based Reward (DBR) and the Proportional Reward (PR), as well as two different Operator Selection Rule, namely the Instantaneous Reward (IR) and a Concurrent Approach (CA). The AOS proposed combines a set of four Fitness Landscape Analysis measures in conjunction with an online learning algorithm, to predict the most suitable crossover operator. We have chosen four FLA metrics to be used as inputs of our predictive model: accumulated escape probability, dispersion metric, average neutrality ratio and average delta fitness of neutral networks. These metrics have been chosen because (1) they can be computed without much increasing the computational effort and (2) they complement each other by capturing different features of the landscapes. Three versions of the MAENS*[8] algorithm were implemented and tested on two datasets of CARP instances. The results of such experiments were compared against those by state-of-the-art algorithms, and against an oracle. The results achieved by MAENS*-II show that this technique is able to compete with the state-of-the-art techniques and can, in some cases, exploit the multiple measures to outperform the state-of-the-art. In the dataset containing large CARP instances, MAENS*-II was able to outperform all the existing approaches in terms of average and best solution quality in half of the instances, and even discovered new lower bounds.

Our experiments seem to suggest a better performance of the Concurrent Strategy over the Instantaneous Reward, and a comparable performance of

the two Reward Measure Strategies.

This work leaves space for interesting directions that can be explored. First, the two Reward Measures might be combined to generate a novel measure that is able to predict better both the diversity and the survival ability of the offspring. Secondly, it would be interesting to test the behaviour of our algorithm when adopting an Average or Extreme Reward strategy and the use of different base learners. Adaptive Operator Selection might be extended to different cases. In particular, for MAENS, an AOS strategy can be adapted to choose among different parent selection strategies for the crossover operator, to analyze its impact on the offspring generation. Another direction is that of reducing the computational cost of MAENS*-II. Furthermore, due to the improved optimization ability provided by this approach, it would be interesting to test the use of MAENS*-II as the Single Objective routine for existing decomposition based approaches. Finally, our technique might be adopted to improve the performance of Evolutionary Algorithms for other combinatorial optimization problems.

Acknowledgements This work was supported by EPSRC (Grant Nos. EP/I010297/1 and EP/J017515/1). Xin Yao was supported by a Royal Society Wolfson Research Merit Award. The authors thank the editors and the anonymous reviewers for their constructive and insightful comments to improve the paper.

References

1. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Machine learning* **47**(2-3), 235–256 (2002)
2. Barbosa, H.J., Sá, A.: On adaptive operator probabilities in real coded genetic algorithms. In: *Workshop on Advances and Trends in Artificial Intelligence for Problem Solving* (2000)
3. Belluz, J., Gaudesi, M., Squillero, G., Tonda, A.: Operator selection using improved dynamic multi-armed bandit. In: *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference*, pp. 1311–1317. ACM (2015)
4. Benavent, E., Campos, V., Corberán, A., Mota, E.: The capacitated arc routing problem: lower bounds. *Networks* **22**(7), 669–690 (1992)
5. Beullens, P., Muyldermans, L., Cattrysse, D., Van Oudheusden, D.: A guided local search heuristic for the capacitated arc routing problem. *European Journal of Operational Research* **147**(3), 629–643 (2003)
6. Brandão, J., Eglese, R.: A deterministic tabu search algorithm for the capacitated arc routing problem. *Computers & Operations Research* **35**(4), 1112–1126 (2008)
7. Chen, W., Wang, Y., Yuan, Y.: Combinatorial multi-armed bandit: General framework and applications. In: *Proceedings of the 30th International Conference on Machine Learning*, pp. 151–159 (2013)
8. Consoli, P., Yao, X.: Diversity-driven selection of multiple crossover operators for the capacitated arc routing problem. In: C. Blum, G. Ochoa (eds.) *Evolutionary Computation in Combinatorial Optimisation - 14th European Conference, EvoCOP 2014, Granada, Spain, April 23-25, 2014, Revised Selected Papers*, no. 12 in *Lecture Notes in Computer Science*, pp. 97–108. Springer (2014)
9. Consoli, P.A., Minku, L.L., Yao, X.: Dynamic selection of evolutionary algorithm operators based on online learning and fitness landscape metrics. In: *Simulated Evolution and Learning*, pp. 359–370. Springer (2014)

10. DaCosta, L., Fialho, A., Schoenauer, M., Sebag, M.: Adaptive operator selection with dynamic multi-armed bandits. In: Proceedings of the 10th annual conference on Genetic and evolutionary computation, pp. 913–920. ACM (2008)
11. Davis, L.: Adapting operator probabilities in genetic algorithms. In: International Conference on Genetic Algorithms '89, pp. 61–69 (1989)
12. Eglese, R.W.: Routeing winter gritting vehicles. *Discrete applied mathematics* **48**(3), 231–244 (1994)
13. Eiben, A., Horvath, M., Kowalczyk, W., Schut, M.C.: Reinforcement learning for online control of evolutionary algorithms. In: Engineering Self-Organising Systems, pp. 151–160. Springer (2007)
14. Eiben, A.E., Hinterding, R., Michalewicz, Z.: Parameter control in evolutionary algorithms. *Evolutionary Computation*, IEEE Transactions on **3**(2), 124–141 (1999)
15. Fialho, Á., Da Costa, L., Schoenauer, M., Sebag, M.: Dynamic multi-armed bandits and extreme value-based rewards for adaptive operator selection in evolutionary algorithms. In: Learning and Intelligent Optimization, pp. 176–190. Springer (2009)
16. Goldberg, D.E.: Probability matching, the magnitude of reinforcement, and classifier system bidding. *Machine Learning* **5**(4), 407–425 (1990)
17. Golden, B.L., Wong, R.T.: Capacitated arc routing problems. *Networks* **11**(3), 305–315 (1981)
18. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: an update. *ACM SIGKDD explorations newsletter* **11**(1), 10–18 (2009)
19. Hinkley, D.V.: Inference about the change-point from cumulative sum tests. *Biometrika* **58**(3), 509–523 (1971)
20. Julstrom, B.A.: What have you done for me lately? adapting operator probabilities in a steady-state genetic algorithm. proceeding of: Proceedings of the 6th International Conference on Genetic Algorithms, Pittsburgh, PA, USA, July 15–19, 1995 (1995)
21. Karafotias, G., Hoogendoorn, M., Eiben, A.: Evaluating reward definitions for parameter control. In: Applications of Evolutionary Computation, pp. 667–680. Springer (2015)
22. Kim, M., McKay, R.I.B., Kim, D.K., Nguyen, X.H.: Evolutionary operator self-adaptation with diverse operators. In: Genetic Programming, pp. 230–241. Springer (2012)
23. Kolter, J.Z., Maloof, M.: Dynamic weighted majority: A new ensemble method for tracking concept drift. In: Data Mining, 2003. ICDM 2003. Third IEEE International Conference on, pp. 123–130. IEEE (2003)
24. Lardeux, F., Saubion, F., Hao, J.K.: Gasat: a genetic local search algorithm for the satisfiability problem. *Evolutionary Computation* **14**(2), 223–253 (2006)
25. Lu, G., Li, J., Yao, X.: Fitness-probability cloud and a measure of problem hardness for evolutionary algorithms. In: Evolutionary Computation in Combinatorial Optimization, pp. 108–117. Springer (2011)
26. Lunacek, M., Whitley, D.: The dispersion metric and the CMA evolution strategy. In: Proceedings of the 8th annual conference on Genetic and evolutionary computation, pp. 477–484. ACM (2006)
27. Martinelli, R., Poggi, M., Subramanian, A.: Improved bounds for large scale capacitated arc routing problem. *Computers & Operations Research* **40**(8), 2145–2160 (2013)
28. Maturana, J., Saubion, F.: A compass to guide genetic algorithms. In: Parallel Problem Solving from Nature–PPSN X, pp. 256–265. Springer (2008)
29. Mei, Y., Li, X., Yao, X.: Cooperative coevolution with route distance grouping for large-scale capacitated arc routing problems. *Evolutionary Computation*, IEEE Transactions on **18**(3), 435–449 (2014)
30. Mei, Y., Li, X., Yao, X.: Variable neighborhood decomposition for large scale capacitated arc routing problem. In: Evolutionary Computation (CEC), 2014 IEEE Congress on, pp. 1313–1320. IEEE (2014)
31. Merz, P.: Advanced fitness landscape analysis and the performance of memetic algorithms. *Evolutionary Computation* **12**(3), 303–325 (2004)
32. Minku, L.L., White, A.P., Yao, X.: The impact of diversity on online ensemble learning in the presence of concept drift. *Knowledge and Data Engineering*, IEEE Transactions on **22**(5), 730–742 (2010)

33. Runarsson, T.P., Yao, X.: Stochastic ranking for constrained evolutionary optimization. *Evolutionary Computation, IEEE Transactions on* **4**(3), 284–294 (2000)
34. Sakurai, Y., Takada, K., Kawabe, T., Tsuruta, S.: A method to control parameters of evolutionary algorithms by using reinforcement learning. In: *Signal-Image Technology and Internet-Based Systems (SITIS)*, 2010 Sixth International Conference on, pp. 74–79. IEEE (2010)
35. Schlimmer, J.C., Granger, R.H.: Beyond incremental processing: Tracking concept drift. In: *AAAI*, pp. 502–507 (1986)
36. Soria Alcaraz, J.A., Ochoa, G., Carpio, M., Puga, H.: Evolvability metrics in adaptive operator selection. In: *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pp. 1327–1334. ACM (2014)
37. Tang, K., Mei, Y., Yao, X.: Memetic algorithm with extended neighborhood search for capacitated arc routing problems. *Evolutionary Computation, IEEE Transactions on* **13**(5), 1151–1166 (2009)
38. Thierens, D.: An adaptive pursuit strategy for allocating operator probabilities. In: *Proceedings of the 2005 conference on Genetic and evolutionary computation*, pp. 1539–1546. ACM (2005)
39. Vanneschi, L., Pirola, Y., Collard, P.: A quantitative study of neutrality in gp boolean landscapes. In: *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pp. 895–902. ACM (2006)
40. Zachariadis, E.E., Kiranoudis, C.T.: A strategy for reducing the computational complexity of local search-based methods for the vehicle routing problem. *Computers & Operations Research* **37**(12), 2089–2105 (2010)