

# Optimisation and Illumination of a Real-world Workforce Scheduling and Routing Application via Map-Elites (Final Draft)

Neil Urquhart and Emma Hart

School of Computing, Edinburgh Napier University, Scotland, UK  
n.urquhart,e.hart@napier.ac.uk

**Abstract.** Workforce Scheduling and Routing Problems (WSRP) are very common in many practical domains, and usually have a number of objectives. Illumination algorithms such as Map-Elites (ME) have recently gained traction in application to *design* problems, in providing multiple diverse solutions as well as illuminating the solution space in terms of user-defined characteristics, but typically require significant computational effort to produce the solution archive. We investigate whether ME can provide an effective approach to solving WSRP, a *repetitive* problem in which solutions have to be produced quickly and often. The goals of the paper are two-fold. The first is to evaluate whether ME can provide solutions of competitive quality to an Evolutionary Algorithm (EA) in terms of a single objective function, and the second to examine its ability to provide a repertoire of solutions that maximise user choice. We find that very small computational budgets favour the EA in terms of quality, but ME outperforms the EA at larger budgets, provides a more diverse array of solutions, and lends insight to the end-user.

## 1 Introduction

Workforce scheduling and routing problems (WSRP) [3] are challenging problems for organisations with staff working in areas including health care [2] and engineering [5]. Finding solutions is the responsibility of a planner within the organisation who will have an interest in the wider organisational policy decisions surrounding the solution. Such wider issues could include the implications of solutions with a lower environmental impact, the effects of switching to public transport, or the impact of changing the size of the workforce.

Multi-objective optimisation approaches are commonly used to find solutions, to WSRP instances, as they can provide a front of solutions that trade-off objectives [13]. However, fronts may only comprise a small section of the total solution space, and are difficult to visualise if there are many dimensions. Thus, it can be difficult for a planner to understand the range of solutions, why solutions were produced, and in particular to know whether other compromise solutions might exist.

A class of algorithms known as *illumination* algorithms have recently been introduced by Mouret *at al* [7], with a number of variants following, e.g. [10,8]. Fundamentally different to a traditional search algorithm, the approach provides a holistic view of how high-performing solutions are distributed throughout a solution space [7]. The method creates a map of high-performing solutions at each point in a space defined by dimensions of variation that are chosen by a user, according to characteristics of a solution that are of interest. The resulting map (a Multi-dimensional Archive of Phenotypic Elites -) enables the user to gain specific insight into how combinations of characteristics of solutions correlate with performance, hence providing insight as well as multiple potential solutions. As the approach encourages diversity, it has often been shown to be more capable of fully exploring a search-space, outperforming state-of-the-art search algorithms given a single-objective, and can be particularly helpful in overcoming deception [9]. We therefore hypothesise that an illumination algorithm might provide particular benefit to real-world problems such as WSRP, which contain multiple, and sometimes conflicting, objectives. However, in contrast to the majority of previous applications of (ME) Map-Elites however which fall mainly in the domain of design problems (e.g. designing robot morphology), WSRP is a repetitive problem, which requires solving new instances repeatedly and obtaining acceptable solutions in reasonable time. While investing effort into producing an archive of solutions can pay off in a design domain, it may prove prohibitive for repetitive problems. Therefore, in the context of a WSRP based on the city of London, using real geographical locations and real transport information, we consider the following questions:

1. How does the relative performance of ME compare to a standard Evolutionary Algorithm (EA) in terms of satisfying a single objective-function over a range of evaluation budgets?
2. Does MAP-Elites provide useful insights into problem characteristics from a real-world perspective through providing a range of diverse but high-quality solutions?

Using 10 realistic problem instances, we demonstrate that for a small fixed evaluation budget, MAP-Elites does not outperform an EA in terms of the objective function, but as the budget increases, it outperforms the EA on the majority of instances tested. Furthermore, even when it is outperformed by an EA in terms of the single objective, it can discover solutions that have better values for the individual characteristics. From a user-perspective, it may therefore present an acceptable trade-off between overall quality and insight.

## 2 Previous work

The WSRP was defined in [3] as a scenario that involves the mobilisation of personnel in order to perform work related activities at different locations. It has been tackled by a variety of methods including meta-heuristic [1] and hyper-heuristics [5]. It can involve consideration of many constraints and objectives,

for example transport modality, time-windows, transport cost, travel cost etc. and hence is often treated as multi-objective problem, e.g. [13]. The reader is referred to [3] for a detailed survey on previous approaches.

The Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) was first introduced by Mouret *et al* [7] and as discussed in the introduction, provides a mechanism for illuminating search spaces by creating an archive of high-performing solutions mapped onto solution characteristics defined by the user. To date, the majority of applications of illumination algorithms have been to *design* problems [7,16]. Another tranche of work focuses on *behaviour* evolution in robotics, for example Cully *et al* [4], who evolve a diverse set of behaviours for a single robot in a “pre-implementation” simulation phase: these are then used in future when the robot is in operation to guide intelligent choice of behaviour given changing environmental conditions.

To the best of our knowledge, an illumination algorithm has never been used to solve repetitive problems, i.e. problems faced in the real-world where acceptable solutions to problems have to be discovered in short time-frames, often many times a day. Typically these types of problems are combinatorial optimisation problems, e.g. scheduling, routing and packing, that often utilise indirect genotypic representations as a result of having to deal with multiple constraints. This contrasts to much of the existing work using MAP-Elites which uses a direct representation of design parameters (although the use of MAP-Elites with an indirect representation was discussed in [11]).

### 3 Methodology

We consider a WSRP characterised by time-windows, multiple transport modes and service times, variations of this scenario include the scheduling of health care workers as well as those providing other services such as environmental health inspections.

We assume an organisation has to service a set of clients, who each require a single visit. Each of the visits  $v$  must be allocated to an employee, such that all clients are serviced, and an unlimited number of employees are available. Each visit  $v$  is located at  $g_v$ , where  $g$  represents a real UK post-code, has a service time  $d_v$  and a time-window in which it must commence described by  $\{e_v, l_v\}$ , i.e. the earliest and latest time at which can start and finish. Visits are grouped into *journeys*, where each journey contains a subset  $V_j$  of the  $V$  visits and is allocated to an employee. Each journey  $j$  starts and ends at the central office. Two modes of travel are available to employees: the first mode uses private transport (car), the second makes uses of available public-transport, encouraging more sustainable travel. The overall goal is to minimise the total distance travelled across all journeys completed and forms the objective function for the problem. However, in addition, discussions with end-users [14] highlights four characteristics of solutions that are of interest:

- The total **emissions** incurred by all employees over all visits

- The total **employee cost** the total cost (based on £/hour) of paying the workforce
- The total **travel cost** the cost of all of the travel activities undertaken by the workforce
- The % of employees using **car travel**

We develop an algorithm based on Map-Elites to minimise the distance objective through projecting solutions onto a 4-dimensional map, with each axis representing one of the above characteristics. Solution quality is compared to an Evolutionary Algorithm that uses exactly the same distance function as an objective, and an identical representation, crossover and mutation operators.

Both the Map-Elites algorithm and the EA use an identical representation of the problem, previously described in [14]. The genotype defines a *grand-tour* [6], i.e. a single permutation of all  $v$  required visits. This is subsequently divided into individual feasible journeys using a *decoder*. The genotype also includes  $v$  additional genes that denote the model of transport to be used for the visit, i.e. public or private.

The decoder converts the single grand tour into a set of journeys to be undertaken by an employee. It examines each visit in the grand tour in order. Initially, the first visit in the grand tour specified by the genotype is allocated to the first journey. The travel mode (car or public transport) associated with this visit in the genome is then allocated to the journey: this travel mode is then adopted for the entire journey (regardless of the information associated with a visit in the genome). The decoder then examines the next visit in the grand tour: this is added to the current journey if it is *feasible*. Feasibility requires that the employee arrives from the previous visit using the mode of transport allocated to the journey within the time window associated with the visit. Note that a travel mode cannot be switched during a journey. Subsequent visits are added using the journey mode until a hard constraint is violated, at which point the current journey is completed and a new journey initiated.

### 3.1 The MAP-Elites Algorithm

The implementation of MAP-Elites used in this paper is given in Algorithm 1 and is taken directly from [7]. The algorithm starts with an empty,  $N$ -dimensional map of elites: {solutions  $\mathcal{X}$  and their performances  $\mathcal{P}$  }. An initialisation method generates  $G$  random-solutions; subsequent solutions are generated from elites stored in the archive. Following random selection of a solution (or solutions) from the archive, the *RandomVariation()* method applies either crossover followed by mutation, or just mutation, depending on the experiment. All operators utilised are borrowed from [14]. The *mutation* operator moves a randomly selected entry in the grand-tour to another randomly selected point in the tour. The *crossover* operator selects a random section of the tour from parent-1 and copies it to the new solution. The missing elements in the child are copied from parent-2 in the order that they appear in parent-2. For each solution  $x'$ , a *feature-descriptor*  $b$  is obtained by discretising the four features of interest associated with the

solution (Section 3) into 20 bins; for 4 dimensions this gives a map containing  $20^4 = 160,000$  cells. The upper and lower bounds required for discretisation are taken as the maximum and minimum values observed by [14] for each dimension during an extensive experimental investigation. A new solution is placed in the cell in the archive corresponding to  $b$  if its fitness  $p$  (calculated as total distance travelled) is better than the current solution stored, or the cell is currently empty.

---

**Algorithm 1** MAP-Elites Algorithm, taken directly from [7]

---

```

procedure MAP-ELITES ALGORITHM
  ( $\mathcal{P} \leftarrow \emptyset, \mathcal{X} \leftarrow \emptyset$ )
  for iter = 1  $\rightarrow$  I do
    if iter < G then
       $x' \leftarrow \text{randomSolution}()$ 
    else
       $x' \leftarrow \text{randomSelection}(\mathcal{X})$ 
       $x' \leftarrow \text{randomVariation}(\mathcal{X})$ 
    end if
     $b' \leftarrow \text{featureDescriptor}(x')$ 
     $p' \leftarrow \text{performance}(x')$ 
    if  $\mathcal{P}(b') = \emptyset$  or  $\mathcal{P}(b') < p'$  then
       $\mathcal{P}(b') \leftarrow p'$ 
       $\mathcal{X}(b') \leftarrow x'$ 
    end if
  end for
  return feature-performance map( $\mathcal{P}$  and  $\mathcal{X}$ )
end procedure

```

---

### 3.2 The Evolutionary Algorithm

The EA uses exactly the same representation and operators as the Map-Elites algorithm. The EA uses a population size of 100, with 40 children being created each generation. Each child is created by cloning from one parent or crossover using two parent. Parents are selected using a tournament of size 2. A mutation-rate of 0.7 is applied to each child. The children are added back into the population, replacing the loser of a tournament, providing the child represents an improvement over the loser. The parameters for the EA were derived from the authors' previous experience with similar algorithms applied to the same problem instances.

### 3.3 Problem Instances

We use a set of problem instances based upon the city of London, divided into two problem sets, termed London (60 visits) and BigLondon (110 visits). These instances were first introduced in [14]. Each visit represents a real post-code

within London. For each of the problem sets, 5 instances are produced in which the duration of each visit is fixed to 30 minutes. Visits are randomly allocated to one of  $t$  time-windows, where  $t \in \{1, 2, 4, 8\}$ . For  $t = 1$ , the time-window has a duration of 8 hours, for  $t = 2$ , the time-windows are “9am-1pm” and “1pm-5pm” etc. These instances are labelled using the scheme  $\langle set \rangle$ -  $numTimeWindows$ , i.e. *Lon-1* refers to an instance in the London with one time-window and *Blon-2* refers to an instance of the BigLondon problem with 2 time windows. The fifth instance represents a randomly chosen mixture of time windows based on 1,2,4 and 8 hours.

If a journey is undertaken by car, paths between visits and distance is calculated according to the real road-network using the GraphHopper library<sup>1</sup>. This relies on Open StreetMap data<sup>2</sup>. Car emissions are calculated as 140 g/km based upon values presented in [12]. For journeys by public-transport, data is read from the Transport for London (TfL) API<sup>3</sup> which provides information including times, modes and routes of travel by bus and train. Public transport emissions factors are based upon those published by TfL [12].

### 3.4 Experimental Parameters

The function evaluation budget is fixed in all experiments. We test two values: 1 million evaluations and 5 million, in each case the first 1000 evaluations are used for initialisation.. Each treatment is repeated 10 times on each instance. The best objective (distance) value is recorded for both treatments in each run. We apply Vargha and Delaney’s  $\hat{A}$  statistic [15] to assess difference between the algorithms. This is regarded as a robust test when assessing randomised algorithms. The test returns a statistic,  $\hat{A}$ , that takes values between 0 and 1; a value of 0.5 indicates that the two algorithms are stochastically equivalent, while values closer to 0 or 1 indicate an increasingly large stochastic difference between the algorithms. One of the most attractive properties of the Vargha-Delaney test is the simple interpretation of the  $\hat{A}$  statistic: for results from two algorithms, A and B, then is simply the expected probability that algorithm A produces a superior value to algorithm B. We follow the standard interpretation that a value in the range  $0.5 \pm 0.06$  indicates a small effect,  $0.5 \pm 0.14$  a medium effect and  $.5 \pm 0.21$  a large effect. In addition we use two metrics to further analyse Map-Elites that are now de-facto in the literature:

- *Coverage* represents the area of the feature-space covered by a single run of the algorithm, i.e. the number of cells filled. For a single run  $x$  of algorithm  $y$ ,  $coverage = noOfCellsFilled/C_{Max}$  where  $C_{Max}$  is the total number of cells filled by combining all runs of any algorithm on the problem under consideration.

<sup>1</sup> <https://graphhopper.com/>

<sup>2</sup> <https://openstreetmap.org/>

<sup>3</sup> <https://api.tfl.gov.uk/>

- *Precision* is also defined as *opt-in reliability*: if a cell is filled in a specific run, then the cell-precision is calculated as the inverse of the performance-value (distance) found in the that cell in that run, divided by the best-value ever obtained for cell in any run of any algorithm (as this is minimisation). Cell-precision is averaged over all *filled* cells in an archive to give a single precision value for a run.

From the perspective of a planner, this represents the choice of solutions available to them, while precision indicates whether a cell contains a solution that is likely to of potential use to the planner. The averaged precision for a run indicates the overall quality of the solutions produced.

	London Problems					Big London Problems				
	Lon-1	Lon-2	Lon-4	Lon-8	Lon-rnd	Blon-1	Blon-2	Blon-4	Blon8	Blon-rnd
ME(1M) vs EA(1M)	↔	↓↓↓	↓↓↓	↓↓↓	↓↓↓	↓↓↓	↓↓↓	↓↓↓	↓↓↓	↓↓↓
ME(5M) vs EA(5M)	↑↑↑	↑↑↑	↑↑↑	↑↑↑	↑↑↑	↓↓↓	↓↓	↑	↑	↓↓↓

Table 1: Comparison of Map-Elites (ME) to Evolutionary Algorithm (EA) at  $n$  million evaluations. Arrows show Vargha-Delaney A Test Effect Size and Direction

## 4 Results

The first research question aims to compare the performance of MAP-Elites and EA algorithms under different evaluation budgets to determine whether MAP-Elites might be useful in producing a set of acceptable solutions quickly. Two values are tested : the first is relatively small with 1 million evaluations (as in [14]); the second increases this to 5 million.

Figures 1(a,b) show the objective fitness values achieved by ME and the EA under both budgets on each of the problem instances. Table 1 shows effect size and direction according to the Vargha-Delaney metric.

We note firstly that for 1 million evaluations for both sets of problems, the EA outperforms Map-Elites: the median of the EA is lower than ME, and the effect size is large in each case. However, when the budget is increased to 5M, Map-Elites outperforms the EA on all of the smaller problems with a large effect size; it also outperforms the EA on two of the larger problems, although the effect size is small. In the remaining 3 cases, the EA still wins.

Note that the figures 1a and 1b only show performance in terms of distance and do not take into account the four characteristics which provide insight to the end-users. These values are given in table 2. Firstly we note that for the smaller Lon problems, the best-value for each characteristic is obtained from the MAP-Elites algorithm in call cases. This includes Lon-8 in which the best objective value for a solution is obtained by the EA, but the solution has poorer values

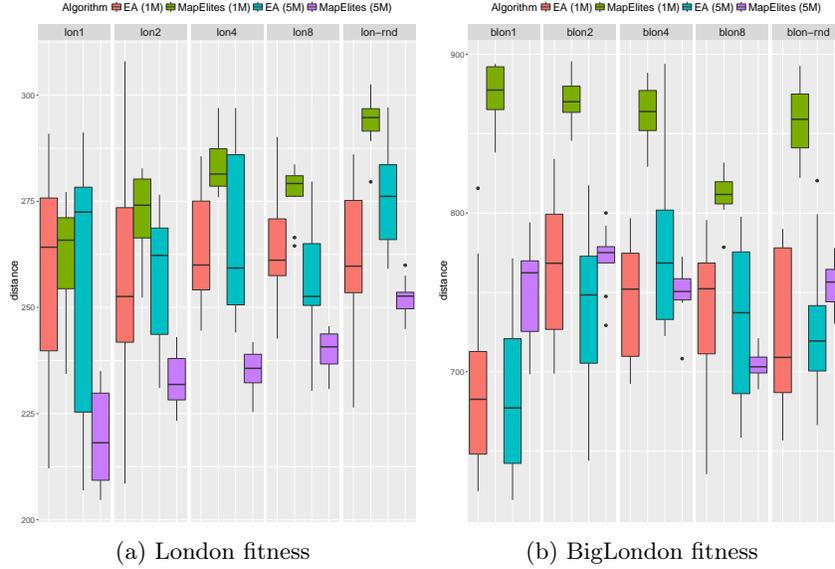


Fig. 1: Performance of MAP Elites and the EA with budgets of 1 million and 5 million evaluations.

for each of the 4 characteristics than the best solution obtained by MAP-Elites. Examining the results for the larger BLon problem demonstrates that MAP-Elites, despite a sub-optimal performance (w.r.t the objective function), can still find solutions that out perform the EA in terms of the individual characteristics. We also note that for the Lon problem ME performs equally well regardless of time window constraints.

	Dist	StaffCost	TravelCost	CO2	CarUse
<b>Lon-1</b>	<b>204.64</b> : 206.93	<b>841</b> : 974.67	<b>82.54</b> : 85.79	<b>133.83</b> : 163.75	<b>0</b> : 0.25
<b>Lon-2</b>	<b>223.3</b> : 231.02	<b>870.67</b> : 1014.67	<b>89.71</b> : 103.04	<b>148.94</b> : 192.85	<b>0.06</b> : 0.33
<b>Lon-4</b>	<b>225.37</b> : 244.09	<b>904.33</b> : 1276	<b>94.63</b> : 116.74	<b>158.77</b> : 194.59	<b>0.04</b> : 0.33
<b>Lon-8</b>	230.8 : <b>230.34</b>	<b>967.33</b> : 1376.67	<b>103.5</b> : 140.1	<b>159.07</b> : 240.54	<b>0.04</b> : 0.35
<b>Lon-Rnd</b>	<b>244.91</b> : 259.11	<b>944</b> : 1140.33	<b>99.48</b> : 107.4	<b>155.17</b> : 216.53	<b>0.04</b> : 0.33
<b>Blon-1</b>	698.48 : <b>619.15</b>	<b>1987</b> : 2182.33	222.63 : <b>207.17</b>	527.27 : <b>506.02</b>	<b>0.04</b> : 0.25
<b>Blon-2</b>	729.21 : <b>644.07</b>	<b>2107.67</b> : 2385.67	244.54 : <b>243.55</b>	584.99 : <b>581.16</b>	<b>0.07</b> : 0.32
<b>Blon-4</b>	<b>708.25</b> : 722.53	<b>2183.33</b> : 2545.67	<b>267.85</b> : 272.34	<b>584.19</b> : 637.26	<b>0.08</b> : 0.33
<b>Blon-8</b>	688.94 : <b>658.52</b>	<b>2209</b> : 2772	<b>272.22</b> : 311.52	<b>586.81</b> : 637.5	<b>0.08</b> : 0.38
<b>Blon-rnd</b>	730.3 : <b>666.29</b>	<b>2256</b> : 2717.67	<b>251.31</b> : 263.1	<b>580.16</b> : 602.47	<b>0.09</b> : 0.36

Table 2: Best results found for performance (distance) and the 4 solution characteristics (5 million evaluations). Values are shown for MAP Elites on the left and the EA on the right.

### 4.1 Coverage and Precision

The coverage metric evaluates the ability of an individual run of an algorithm to place individuals in each of the cells. Note that it is possible that some of the cells cannot be filled in because the characteristics of that instance do not allow a feasible solution in that area.

The coverage achieved is displayed in figures 2a and 2b. Observe that coverage of over 70% is common with MAP-Elites, but the EA gives very poor coverage as it converges to a single solution. In real-world terms, the EA leaves the user with little choice of solution and no insight into the problem.

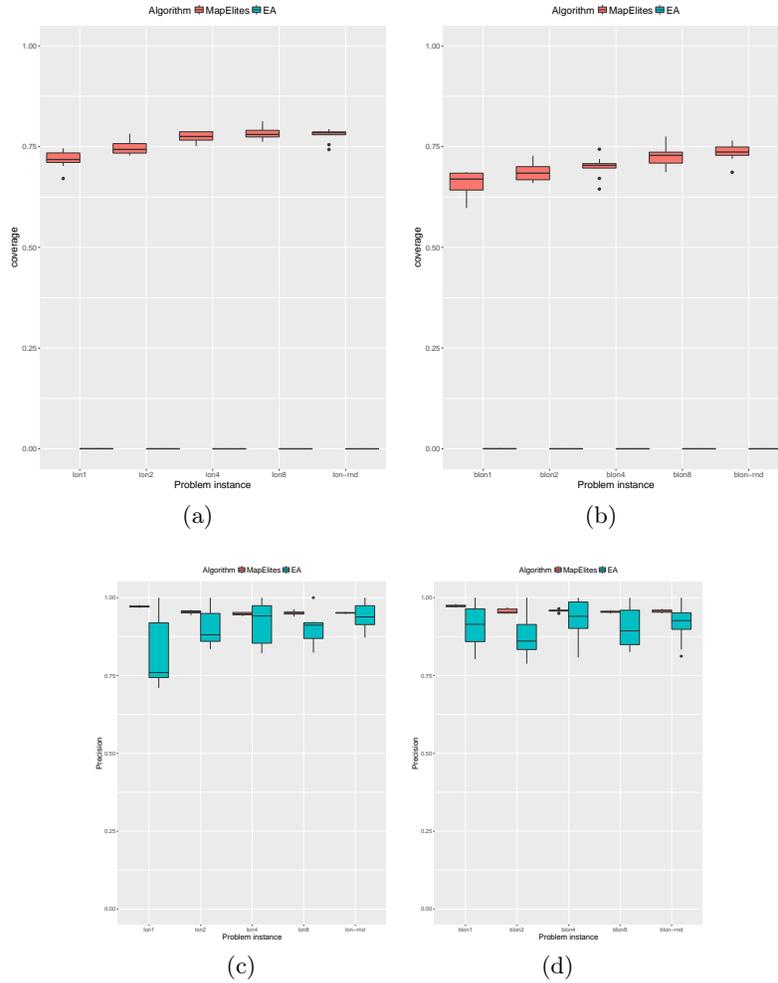


Fig. 2: Coverage and Precision for Map-Elites and the EA on both problem sets

Figures 2c and 2d show the precision achieved by MAP Elites and the EA. Note that the highest precision achieved by the EA outperforms ME, note that precision is calculated over only those cells that are filled. The EA allocates all of its evaluations to very few cells, and thus find good solutions for those cells. In contrast, MAP-Elites has to distribute the same budget of evaluations across a much larger number cells, making it hard to always find a high-performing solution in each cell. In addition, many of the low-precision scores for ME occur when one run does not find as high-performing a solution in a cell as another run of MAP-Elites. Running ME for more evaluations would likely improve precision (without danger of convergence due to its propensity to enforce diversity).

## 4.2 Gaining Insight into the Problem Domain

Figure 3 plots the cells, and the elite solutions contained, for each 2-dimensional pairing of the 4 dimensions. Although the archive could be drawn in 4-dimensions, discussion with users suggested that presenting 2-dimensional maps provides more insight. Within each plot, each cell that is occupied is coloured to represent the distance objective value of the elite solution - lowest (best) values being green, highest being red. Note that most of the cells have a solution within them. Where there is an area with no solutions it tends to be at a corner of the plot. For instance, there are a lack of solutions with low  $CO_2$  and high travel costs (figure 3e) or high car use and low  $CO_2$  (figure 3a). From a planning perspective, figure 3 indicates (1) combinations of objectives that have no feasible solutions, and (2) quality of feasible solutions.

MAP-Elites tends to cover a larger part of the solution space. A common trend is that the solutions that are better in terms of one or two of the four characteristics are not always solutions that exhibit the lowest distance objective. The map also quantifies trade-offs in objective value: for example, the extent to which increased car use increases  $CO_2$  compared to options that utilise more public transport. Another insight to be gained is the effects of higher public transport use (i.e. low car use) and staff cost: staff costs rise as public transport usage increases (figures 3c a). This is due to the longer journey times experienced with public transport leading to increased working hours for staff.

A planner with responsibility for determining policies regarding staff scheduling may make use of the diagrams in figure 3c to understand what solutions are possible given a specific priority. For instance, if it is determined that reducing  $CO_2$  is a priority then they can determine what possible trade-offs exist for low  $CO_2$  solutions. Where a balance is required (i.e. lowering  $CO_2$  but also keeping financial costs in check) MAP-Elites allows the planner to find compromise solutions that are not optimal in any single dimension, but may prove useful when meeting multiple organisational targets or aspirations.

## 5 Conclusions

In this paper we have applied MAP-Elites to a real world combinatorial optimisation problem domain — a workforce scheduling and routing problem. Unlike

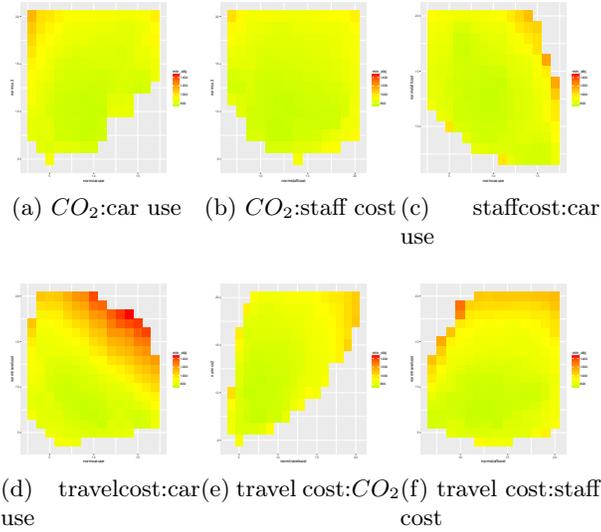


Fig. 3: Maps produced from a single run of the Blon-1 problem: (all possible pairings of the 4 characteristics)

previous applications of MAP-Elites that have tended to concentrate on design problems, WSRP is an example of a repetitive problem, requiring an optimisation algorithm to find acceptable solutions in a short period of time. In addition to an acceptable solution however, a user also requires choice, in being able to select potential solutions based on additional criteria of relevance to a particular company.

With reference to the research questions in section 1, we note that MAP-Elites tends to require a larger evaluation budget to produce results that are comparable with a straightforward EA for the problems tested. However, for small problems, affording a larger evaluation budget to Map-Elites enables it to discover improved solutions, compared to the EA. For larger problems, although our results show that MAP-Elites cannot outperform the EA in terms of objective performance, it does find solutions that outperform the EA in terms of the individual characteristics. It is likely that running MAP-Elites for longer would continue to improve its performance, without risking convergence. The increased cpu-time required for such a budget may be easily obtained through the use of multi-core desktop computers or cloud based resources in a practical setting. We also note that the illumination aspect of MAP-Elites may aid the ability of planners to understand the factors that lead to good solutions and subsequently influence policy planning/determine choices based on organisational values, and that this aspect is of considerable benefit. Illumination of the solution-space also provides additional insight to planners, who can gain understanding into the influence of different factors on the overall cost of a solution.

Future work will focus on further exploration of the relationship between objective quality and function evaluations, to gain insight into the anytime performance of Map-Elites, for use in a real-world setting. The granularity of the archive clearly influences performance and should be investigated by depth. Finally, an additional comparison to multi-objective approaches is also worth pursuing — while this may improve solution quality however it is unlikely to offer the same insight into the entire search-space.

## References

1. Bertels, S., Fafle, T.: A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Computers & Operations Research* **33**(10) (2006)
2. Braekers, K., Hartl, R.F., Parragh, S.N., Tricoire, F.: A bi-objective home care scheduling problem: Analyzing the trade-off between costs and client inconvenience. *European Journal of Operational Research* **248**(2), 428 – 443 (2016)
3. Castillo-Salazar, J.A., Landa-Silva, D., Qu, R.: A survey on workforce scheduling and routing problems. In: *Proceedings of the 9th international conference on the practice and theory of automated timetabling*. pp. 283–302 (2012)
4. Cully, A., Clune, J., Tarapore, D., Mouret, J.B.: Robots that can adapt like animals. *Nature* **521**(7553), 503 (2015)
5. Hart, E., Sim, K., Urquhart, N.: A real-world employee scheduling and routing application. In: *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*. pp. 1239–1242. ACM (2014)
6. Laporte, G., Toth, P.: *Vehicle routing: historical perspective and recent contributions*. EURO Journal on Transportation and Logistics (2013)
7. Mouret, J., Clune, J.: Illuminating search spaces by mapping elites. *CoRR* (2015)
8. Pugh, J.K., Soros, L.B., Stanley, K.O.: Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* **3**, 40 (2016)
9. Pugh, J.K., Soros, L.B., Stanley, K.O.: Searching for quality diversity when diversity is unaligned with quality. In: *International Conference on Parallel Problem Solving from Nature*. pp. 880–889. Springer (2016)
10. Smith, D., Tokarchuk, L., Wiggins, G.: Rapid phenotypic landscape exploration through hierarchical spatial partitioning. In: Handl, J., Hart, E., Lewis, P.R., López-Ibáñez, M., Ochoa, G., Paechter, B. (eds.) *Parallel Problem Solving from Nature – PPSN XIV*. pp. 911–920. Springer International Publishing, Cham (2016)
11. Tarapore, D., Clune, J., Cully, A., Mouret, J.B.: How do different encodings influence the performance of the map-elites algorithm? In: *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. pp. 173–180. GECCO '16, ACM, New York, NY, USA (2016)
12. TFL: *Travel in london: Key trends and developments*. Tech. rep., Transport for London (2009)
13. Urquhart, N., Fonzone, A.: Evolving solution choice and decision support for a real-world optimisation problem. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. pp. 1264–1271. GECCO '17, ACM (2017)
14. Urquhart, N.B., Hart, E., Judson, A.: Multi-modal employee routing with time windows in an urban environment. In: *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation*. pp. 1503–1504. ACM (2015)

15. Vargha, A., Delaney, H.D.: A critique and improvement of the cl common language effect size statistics of mcgraw and wong. *Journal of Educational and Behavioral Statistics* **25**(2), 101–132 (2000)
16. Vassiliades, V., Chatzilygeroudis, K., Mouret, J.B.: Using centroidal voronoi tessellations to scale up the multi-dimensional archive of phenotypic elites algorithm pp. 1–1 (08 2017)