Neil Urquhart*, Emma Hart, and William Hutcheson

Using MAP-Elites to support policy making around Workforce Scheduling and Routing

Abstract: Algorithms such as MAP-Elites provide a means of allowing users to explore a solution space by returning an archive of high-performing solutions. Such an archive, can allow the user an overview of the solution space which may be useful when formulating policy around the problem itself. The number of solutions that can potentially be returned by MAP-Elites is controlled by a parameter d that discretises the user-defined features into 'bins'. For a fixed evaluation budget, increasing the number of bins increases user-choice, but at the same time, may lead to a reduction in overall quality of solutions. We undertake a study of the application of Map-Elites to a Workforce Scheduling and Routing problem, using a set of realistic instances based in London.

1 Introduction and Motivation

Optimisation has traditionally focused upon providing specific solutions to a given problem. Typically the user supplies a problem instance (often an NP-Hard problem) and the optimisation algorithm presents them with a solution. Where the problem has a very clear single-objective, such as the Travelling Salesman Problem, the user may be satisfied with this. However, as we consider more complex real-world scheduling problems, optimisation methods that present a single-solution to the user become less attractive. Such real-world problems typically encompass multiple solution objectives and characteristics, it is no longer a case of searching for the one optimal solution and presenting that to the user. The adoption of many-objective optimisation methods for such problems [4] allowed the generation of a non-dominated front. By generating such a front, the user (typically a planner) is presented with a set of solutions that show the best possible trade-offs between objectives. In this paper we discuss the scenario of an an expert user being presented with a wider range of high-quality solutions using a *quality-diversity* (QD) algorithm [10]. This allows the user a greater understanding of the solution

*Corresponding author: Neil Urquhart, Emma Hart, William Hutcheson, School of Computing, Edinburgh Napier University, 10 Colinton Rd, Edinburgh.

space and the effects of constraining or relaxing characteristics. We propose, that this allows the user a greater degree of choice from a *policy making* perspective.

The Workforce Scheduling and Routing Problem (WSRP) is a commonly encountered real-world problem in which there is a requirement to allocate tasks to individuals within a workforce, and to provide a routing-plan between allocated tasks. As with many real-world problems, the ability to provide an end-user with a set of potential solutions is of great importance, enabling them to select between alternatives with respect to specific company priorities or objectives.

2 Background and Previous Work

The Workforce Routing and Scheduling Problem (WSRP) [3] is formulated around a set of mobile workers, who must travel between a set of locations as part of their work tasks. WSRP instances may cover a range of industrial areas including health and social home-care scheduling [6, 12], and technician scheduling [5]. An attempt to model the WSRP as a bi/multi-objective problem can be found in [2], the authors use cost and patient convenience as the twin objectives. In [6], a healthcare WSRP is formulated as a single objective problem that penalises violations of hard and soft constraints while trying to minimise working and travel-time. Within formulation a binary choice of transport mode (car or public transport) is allowed. In [2], the authors examine the trade off between the cost of solutions and client inconvenience within another home-care scheduling problem.

Approaches to finding solutions to the WSRP include meta-heuristics [1], and hyper-heuristics [5], the reader is directed to [3] for a detailed survey of methods used. Quality-Diversity (QD) algorithms [10] produce a set of high-quality solutions with respect to a set of user-defined features. The Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) was first introduced by Mouret *et al* [9] and provides a mechanism for illuminating search spaces by creating an archive of high-performing solutions mapped onto solution characteristics defined by the user. The majority of applications of illumination algorithms have been to *design* problems [9, 14]. There are very few applications of QD algorithms apparent in the combinatorial optimisation domain; to the best of our knowledge, the authors were the first to show that MAP-Elites could be successfully applied to combinatorial optimisation [12]. The algorithm was applied to produce multiple solutions to WSRP instances modelled using real data based upon the City of London, with four dimensions of user interest. The discretisation parameter d was arbitrarily chosen to give 160,000 cells in this work however, motivating the need for further exploration.

In summary we note that variants of the WSRP, including health-care applications, have been extensively investigated. A range of techniques has been used to produce solutions. Work to date has concentrated on the production of optimised solutions to provide solutions to specific problems. We consider the problem from a policy making perspective where the planner wishes to have a range of solutions that they can utilise to investigate the full range of possibilities when considering possible solutions to the problem. The ability to produce a range of solutions makes QD algorithms especially suited to policy type problems.

3 Methodology

3.1 WSRP Problem Description

The WSRP considered here is defined as follows; an organisation has to service a set of clients, who each require a single visit. Each visit v must be allocated to an employee, such that all client visits are made by an employee. Each visit v is located at g_v, where g represents an actual UK post-code, and has a visit length d_v and a time-window in which it must commence $\{e_v, l_v\}$. Visits are grouped into *journeys*, where each journey is allocated to an employee and contains a subset V_j of the V visits, starting and ending at a central office. In this formulation an unlimited number of employees are available.

Two modes of travel are available to employees, private transport (car) or public-transport. The overall goal of our WSRP is to minimise the total distance travelled across all journeys completed. Discussions with end-users [13] highlights four characteristics of solutions that are of interest:

– **Emissions** incurred by all employees on their journeys
– **Employee Cost** the cost (based on £/hour) of paying the workforce for the duration of the journeys and visits
– **Travel Cost** the cost of all of the travel activities undertaken by the workforce
– The % of Employees using **car travel** for their journeys

It should be noted that no preference is made by users' on the relative importance of these characteristics. For example, if a policy of low emissions is to be investigated, the planner can examine solutions with a range of emissions levels in order to ascertain the effects of requiring low emissions solutions.

We use a set of problem instances (introduced in [13]) based upon the city of London, divided into two problem sets, termed *Lon* (60 visits) and *BLon* (110 visits). For each of the problem sets, 5 instances are produced in which the visit duration 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 $<set>$- $numTimeWindows$. The 'rnd' instance (e.g. BLon-rnd) represents a randomly chosen mixture of time windows based on 1,2,4 and 8 hours.

Car journey distances and times are calculated based on Open StreetMap data[1], using the GraphHopper library[2]. Car emissions are calculated as 140 g/km based upon values presented in [8]. For public-transport, data is read from the Transport for London (TfL) API[3], emissions factors are based upon those published in [8].

3.2 The MAP-Elites Algorithm and Operators

We use a problem representation first described in [13], the genotype being a permutation of all v required visits. For each visit, the genotype also includes the preferred mode of transport. The tour contained within the genotype is decoded into individual employee shifts. Each shift represents one days work for an employee, starting from the central office, visiting a number of clients. The genome is decoded by examining each visit in order. The first visit in the genotype is allocated to the first shift, the travel mode associated with this visit in the genome is then allocated to the entire shift. The next visit in the permutation is added to the current shift, if *feasible*. Feasibility requires that the employee arrives from the previous visit (using the mode of transport allocated) within the time window associated with the visit. Subsequent visits are added to the shift until a hard constraint is violated, at which point the current shift is completed and a new shift is initiated.

The implementation of MAP-Elites used in this paper was used previously by [12] and was based on that given in [9]. The algorithm commences with an empty, N-dimensional map in which solutions $\mathcal{X}$ and their performances $\mathcal{P}$ are subsequently placed.

The total number of iterations to be executed by the algorithm is specified by I, the algorithm evaluates one solution per iteration. A new solution is generated either by random initialisation or through reproduction operators. While the number of solutions evaluated is less than G each solution is created randomly and added to the map in order to create an initial population within the map. Once G solutions have been evaluated, solutions are randomly selected from within map, a value

---

1 https://openstreetmap.org/

2 https://graphhopper.com/

3 https://api.tfl.gov.uk/

d	Cells	d	Cells	d	Cells	d	Cells
5	625	20	160000	35	1500625	50	5250000
10	10000	25	390625	40	2560000		
15	50625	30	810000	45	4100625		

Tab. 1: Numbers of cells in each map, based on the number of bins (d). The number of cells is calculated as $cells = d^n$, where $n = 4$.

of 500 is used for G within our experiments. The reproduction operators create a child solution by applying either crossover or cloning followed by a mutation. The mutation operator moves a randomly selected gene in the chromosome to a randomly selected point, or by changing the travel mode of a randomly selected gene. For each child solution, a *feature-descriptor* b is obtained by discretising the four features of interest associated with the solution. The upper and lower bounds required for discretisation are taken as the maximum and minimum values observed within these data sets in [13]. 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.

3.3 Methodology

We vary the number of bins (d) from 5 to 50 in steps of 5 (see table 1). The number of cells in the maps is $cells = d^n$ where n is the number of dimensions within the solution. MAP-Elites was executed on each problem instance 10 times for each bin configuration. The function evaluation budget is fixed at 5,000,000 evaluations all experiments. The fixed budget attempts to ensure that the computational effort applied remains constant as d is varied. As well as the MAP-Elites algorithm, we also compare results obtained with the Evolutionary Algorithm used in [12] - which contains a description of the algorithm. The EA uses the same representation and operators as our MAP-Elites algorithm and has the same evaluation budget.

4 Results

4.1 Effects of Bin Quantity on Solution Characteristics

If we consider the choice presented to the user, table 5 shows the average number of solutions returned in each map over 10 runs. We note two trends from table

Problem	d	Staff Cost	Travel Cost	CO2 Emissions	Car Use
	EA	974.67	85.79	163.75	0.25
	5	819.33	68.20	119.23	0.10
	10	780.33	69.75	137.04	0.08
	15	804.67	75.64	137.91	0
	20	800.00	78.03	136.99	0.00
lon-1	25	853.67	79.90	126.08	0.00
	30	878.00	81.32	113.48	0.00
	35	874.00	71.91	116.27	0.09
	40	900.00	70.37	117.42	0.09
	45	863.67	70.23	108.46	0.09
	50	904.67	73.53	102.95	0.00
	EA	1014.67	103.04	192.85	0.33
	5	864.67	74.91	132.60	0.14
	10	801.67	76.63	142.60	0.06
	15	846.00	86.15	146.23	0.00
	20	863.00	90.80	149.35	0.06
lon-2	25	872.67	91.89	137.46	0.05
	30	905.67	85.88	137.04	0.06
	35	931.00	86.47	129.63	0.06
	40	924.33	84.15	129.93	0.07
	45	927.00	80.30	131.53	0.07
	50	879.33	78.61	118.28	0.08
	EA	1276.00	116.74	194.59	0.33
	5	955.00	87.58	171.83	0.17
	10	879.67	88.14	148.25	0.05
	15	913.67	94.87	158.09	0.04
	20	899.67	94.82	160.57	0.04
lon-4	25	887.67	98.11	147.13	0.05
	30	922.33	92.54	150.66	0.05
	35	936.33	90.57	146.65	0.05
	40	916.33	80.82	147.72	0.06
	45	912.00	85.29	125.73	0.06
	50	881.00	79.01	125.38	0.06
	EA	1376.67	140.10	240.54	0.35
	5	1114.00	92.40	174.22	0.13
	10	946.67	97.88	166.72	0.05
	15	924.00	101.87	170.76	0.03
	20	967.67	103.09	165.29	0.04
lon-8	25	948.00	100.57	154.74	0.04
	30	985.33	95.20	153.65	0.05
	35	943.00	95.56	155.53	0.05
	40	926.00	83.73	137.80	0.05
	45	926.00	93.27	122.60	0.05
	50	1119.00	100.16	126.61	0.05

Tab. 2: The lowest values found for each problem characteristic. For the EA [12] this is the single solution with the lowest fitness returned. For the MAP-Elites results the lowest values found for each characteristic are collated and presented.

5, firstly that with the exception of the smallest map ($d = 5$) MAP-Elites offers more choice that the equivalent non-dominated front, secondly the choice available increases as d increases. In both cases we would expect to see these trends. The increase of choice compared to the non-dominated front is due to MAP Elites sampling the entire solution space (as divided up by d) as opposed to only returning non-dominated solutions.

Tables 2 and 3 show the best results obtained across problem characteristics. The EA despite having the same evaluation budget under performs in every objective, MAP-Elites not only provides greater choice, but also far greater optimisation that the EA.

A visual indication of the solutions found is shown in figure 9 which examines results achieved from the BLon data set. These figures chart the results obtained with differing numbers of bins (d) (note that each row represents the final archive

Problem	d	Staff Cost	Travel Cost	CO_2 Emissions	Car Use
	EA	2182.33	207.17	506.02	0.20
	5	1901.00	210.55	522.38	0.14
	10	1899.00	216.07	522.21	0.09
	15	1917.67	223.42	567.54	0.08
	20	1962.67	228.67	534.16	0.04
blon-1	25	2036.67	220.92	509.82	0.09
	30	2028.33	212.94	477.98	0.09
	35	2025.00	204.73	474.68	0.09
	40	2089.67	208.21	479.26	0.09
	45	2052.00	200.61	455.32	0.08
	50	2136.33	199.48	441.52	0.04
	EA	2385.67	243.55	581.16	0.32
	5	2163.67	228.47	529.49	0.17
	10	2089.67	227.08	556.65	0.15
	15	2064.33	252.47	548.47	0.11
	20	2076.00	249.59	575.19	0.06
blon-2	25	2122.00	236.21	547.34	0.07
	30	2115.67	226.39	509.25	0.12
	35	2101.67	226.71	475.37	0.11
	40	2066.00	219.49	499.98	0.09
	45	2086.67	213.41	474.89	0.07
	50	2148.67	196.31	447.46	0.11
	EA	2545.67	272.34	637.26	0.33
	5	2295.67	250.02	644.65	0.22
	10	2120.67	269.40	593.73	0.15
	15	2196.33	270.69	574.03	0.11
	20	2131.33	265.86	598.51	0.12
blon-4	25	2106.67	256.85	571.83	0.07
	30	2110.67	242.61	546.63	0.09
	35	2116.33	230.05	554.83	0.11
	40	2071.33	231.87	528.15	0.12
	45	2150.33	226.90	478.24	0.10
	50	2027.67	232.77	473.68	0.11
	EA	2772.00	311.52	637.50	0.38
	5	2424.33	254.95	584.78	0.17
	10	2323.00	283.87	594.38	0.15
	15	2204.33	276.81	610.69	0.11
	20	2180.00	275.85	609.11	0.09
blon-8	25	2176.33	261.57	568.33	0.10
	30	2086.67	245.02	536.85	0.11
	35	2134.67	245.19	516.28	0.13
	40	2201.33	233.73	492.97	0.09
	45	2390.33	236.58	459.60	0.13
	50	2259.67	242.10	476.96	0.09

Tab. 3: The lowest values found for each problem characteristic. For the EA [12] this is the single solution with the lowest fitness returned. For the MAP-Elites results the lowest values found for each characteristic are collated and presented.

	Avg				Best		
	5	50	% increase		5	50	% Increase
Lon-1	343.17	450.32	23.79%		169.72	264.4	35.81%
Lon-2	316.14	450.24	29.78%		189.68	292.77	35.21%
Lon-4	289.52	433.86	33.27%		201.27	307.19	34.48%
Lon-8	267.03	409.79	34.84%		209.85	307.17	31.68%
Lon-rnd	302.87	456.55	33.66%		202.48	313.57	35.43%
Blon-1	908.88	1179.35	22.93%		576.88	874.86	34.06%
Blon-2	843.4	1168.4	27.82%		595.33	919.96	35.29%
Blon-4	812.18	1132.02	28.25%		593.46	903.2	34.29%
Blon-8	739.43	1069.19	30.84%		575.02	893.62	35.65%
Blon-rnd	772.14	1141.09	32.33%		560.96	906.73	38.13%

Tab. 4: The increase (i.e. loss) in average and best fitness between the lowest and highest number of bins (d) used. Recall this is a minimisation problem so lower objective values are preferred
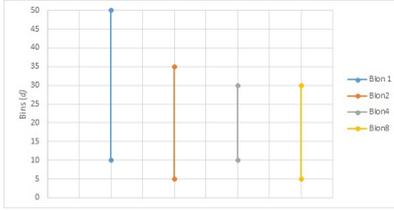
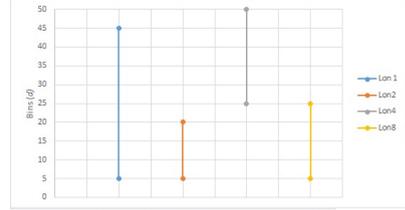**Fig. 1:** The best/worst average coverage obtained for the BLon datasets.

Fig. 2: The best/worst average coverage obtained for the Lon datasets.

of a specific run). We note that although the number of cells increases, the shape of the filled in area, largely remains the same. The resolution of the heat map increases, but the image remains largely unaltered. Those areas which are not filled in on the 5 bin heat maps are largely the same as those not filled in on the 50 bin heat maps. The lightest cells (green) represent those solutions with the least distance cost and we see largely the same distribution of colours. A similar pattern of results is noted for other solutions.

In terms of the best values (tables 2 and 3) we note the following trends:

– The lowest staff and travel costs are generally found with lower values of d
– The lowest emissions values are mostly found at small values of d, the exceptions being the more constrained Lon problems (4 and 8).
– Car use does not follow such a defined pattern, the lowest values being found at a range of d values depending on the problem being examined.

4.2 Case: Study - Exploring policy decisions

Let us assume that a planner is attempting to make some policy decisions regarding the WSRP instance and car use, policies which will be applied when solving future WSRP instances. It is beneficial to demonstrate the effect of changes in d, from the planners' perspective. If we assume that the planner is investigating the Blon-2 problem (the problem instance being picked at random) figures 3 and 6 show the Parallel Coordinates (PC) plots [7] relating to the solutions within each archive. The first thing that we note is the very obvious difference in the number of solutions produced (32 and 3095 respectively). The initial reaction may be that the plot for $d = 5$ is simpler to comprehend then that for $d = 50$. At this stage the planner wishes to examine possible policies to reduce the CO_2 produced. By having the
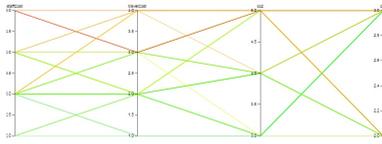
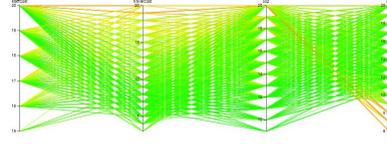**Fig. 3:** The solutions returned for Blon2 where $d = 5$

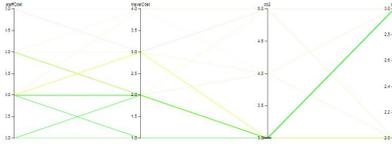**Fig. 6:** The solutions returned for Blon2 where $d = 50$

Fig. 4: As for figure 3 but with the lowest CO2 solutions highlighted.
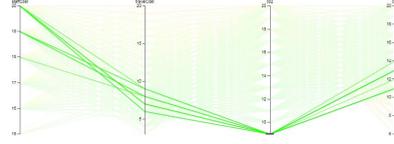

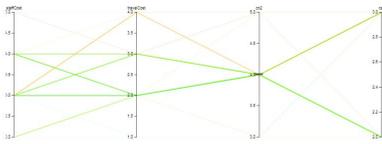**Fig. 7:** As for figure 6 but with the lowest CO2 solutions highlighted.

Fig. 5: As for figure 4 but highlighting possible compromise solutions.
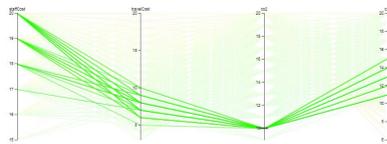

Fig. 8: As for figure 7 but highlighting possible compromise solutions.

PC plots within an interactive tool [4] it is possible to highlight those solutions with the lowest CO_2 values - figures 4 and 7. In the case of figure 4 sufficient solutions are highlighted to demonstrate trends associated with low CO_2 solutions such as:

1. Lower CO_2 equates to lower travel costs.
2. Lower CO_2 equates to higher staff costs
3. Lower CO_2 tends towards lower car use

Thus a policy of lowering car use, will lower (CO_2) (trend 3) is likely to result in higher wage costs (trend 2) and lower travel costs (trend 3).

When $d = 5$ (figure 4) it becomes difficult to draw such definite conclusions across a range of solutions. We can also look for compromise solutions, for instance, what advantages may be gained in other characteristics if a slightly higher CO_2 solution is accepted? Figure 8 shows that accepting a slightly higher CO_2 characteristic allows solutions that have lower staff and travel costs.

---

4 https://commute.napier.ac.uk/upload

This small case study demonstrates that where a higher value of d is used, the user may then have a larger archive which allows greater exploration of the solution space and in turn evaluation of polices on solution characteristics.

5 Conclusions

In this paper we set out establish the ability of MAP-Elites to support policy making, through presenting the end-user (planner) with a choice of solutions. We vary the range of choice through altering the values of d.

The potential for user-choice is reflected by the size of the map, which is determined by the number of bins. Table 4 shows a decrease in the performance of the algorithm (measured as best and average fitness) as the number of bins (d) increases. For these instances, a user interested in maximising objective performance can expect an average loss in performance of between 32% and 38% if the value of d is increased in order to maximise choice. On the other hand, the coverage results show that choice is maximised by setting d to higher values: an approximate gain in coverage of between 10 and 20% can be found by judicious choice of d depending on the problem when compared to setting it to 5 — the value which maximises the objective function.

Table 5 shows that when $d = 5$ there are far fewer solutions and far fewer opportunities for trading off between objectives. When $d = 50$ (figure **??**) we note the dramatic increase in solutions and the potential to to find trade-offs, as described in section 4.2. The increase in the quantity of solutions found (Table 5) represents an increase in the choice available to users selecting a solution.

It was noted that higher values of d considerably slowed the execution time of the algorithm, whilst implementation issues may be argued to be out of scope for this paper, this issue cannot be ignored. As a constant evaluation budget is utilised in all experiments, this increase is most likely due to the data structures used for holding the solutions archive. Further development work is required to deal with this scaling problem.

In summary we have provided evidence which quantifies the implications of altering the value of d from the perspective of the user. Larger values of d will result in increased choice (see section 4.2). It is possible that increasing the evaluation budget in line with increases in d will allow the trade-off between solution quality and choice to be lessened. Any increase in the evaluation budget must take into account the potential scaling issue noted above.
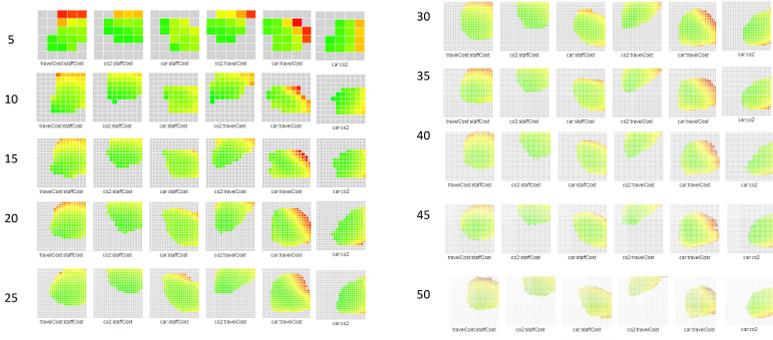
Fig. 9: Heat maps showing the solutions contained within the map. Each row contains a heat map for each pair of dimensions.

Problem	Front Size	d				
		5	10	15	20	25
Lon-1	229	69	432	1450	3344	6360
Lon-2	315	84	1080	4640	12374	25013
Lon-4	271	122	1501	6926	19325	38961
Lon-8	473	106	1490	6979	19144	39533
		30	35	40	45	50
Lon-1		10328	15524	21129	28759	38252
Lon-2		42039	65781	89835	118815	153787
Lon-4		65684	101561	140768	185533	236537
Lon-8		67495	103840	145798	194424	248200
		5	10	15	20	25
Blon-1	225	56	362	1271	3263	6258
Blon-2	418	54	565	2549	7180	15470
Blon-4	662	68	752	3305	9376	20473
Blon-8	639	72	859	3821	10993	22940
		30	35	40	45	50
Blon-1		11038	17539	24545	33898	44642
Blon-2		28173	44827	64970	88276	112549
Blon-4		35916	56615	81887	109628	141184
Blon-8		41354	63907	92374	123681	157905

Tab. 5: Average number of solutions produced for each problem instance, the front sizes [11] were obtained using a portfolio of multi-objective evolutionary algorithms to construct a non-dominated front.

References

[1] Stefan Bertels and Torsten Fafle. "A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem." In: *Computers & Operations Research* 33.10 (2006). (Visited on 07/10/2014).

[2] Kris Braekers et al. "A bi-objective home care scheduling problem: Analyzing the trade-off between costs and client inconvenience." In: *European Journal of Operational Research* 248.2 (2016), pp. 428 –443.

[3] J. Arturo Castillo-Salazar, Dario Landa-Silva, and Rong Qu. "A survey on workforce scheduling and routing problems." In: *Proceedings of the 9th international conference on the practice and theory of automated timetabling*. 2012, pp. 283–302.

[4] Kalyanmoy Deb et al. "A Fast Elitist Multi-Objective Genetic Algorithm: NSGA-II." In: *IEEE Transactions on Evolutionary Computation* 6 (2000), pp. 182–197.

[5] Emma Hart, Kevin Sim, and Neil Urquhart. "A Real-world Employee Scheduling and Routing Application." In: *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 2014, pp. 1239–1242.

[6] Gerhard Hiermann et al. "Metaheuristics for solving a multimodal home-healthcare scheduling problem." In: *Central European Journal of Operations Research* 23.1 (2013).

[7] A. Inselberg. "Multidimensional Detective." In: *Proceedings of the 1997 IEEE Symposium on Information Visualization (InfoVis '97)*. INFOVIS '97. Washington, DC, USA: IEEE Computer Society, 1997, pp. 100–. ISBN: 0-8186-8189-6. URL: http://dl.acm.org/citation.cfm?id=857188.857631.

[8] Transport for London. *Travel in London: Key trends and develpments*. Tech. rep. Transport for London, 2009.

[9] Jean-Baptiste Mouret and Jeff Clune. "Illuminating search spaces by mapping elites." In: *CoRR* (2015).

[10] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. "Quality Diversity: A New Frontier for Evolutionary Computation." In: *Frontiers in Robotics and AI* 3 (2016), p. 40.

[11] Neil Urquhart and Achille Fonzone. "Evolving Solution Choice and Decision Support for a Real-world Optimisation Problem." In: *Proceedings of the Genetic and Evolutionary Computation Conference*. GECCO '17. ACM, 2017, pp. 1264–1271.

[12] Neil Urquhart and Emma Hart. "Optimisation and Illumination of a Real-World Workforce Scheduling and Routing Application (WSRP) via Map-Elites." In: *Parallel Problem Solving from Nature – PPSN XV*. Cham: Springer International Publishing, 2018, pp. 488–499. ISBN: 978-3-319-99253-2.

[13] Neil B. Urquhart, Emma Hart, and Alistair Judson. "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*. GECCO Companion '15. Madrid, Spain: ACM, 2015, pp. 1503–1504. ISBN: 978-1-4503-3488-4. DOI: 10.1145/2739482.2764649. URL: http://doi.acm.org/10.1145/2739482.2764649.

[14]   Vassilis Vassiliades, Konstantinos Chatzilygeroudis, and Jean-Baptiste Mouret. "Using Centroidal Voronoi Tessellations to Scale Up the Multi-dimensional Archive of Phenotypic Elites Algorithm." In: (Aug. 2017), pp. 1–1.