Predicting Genetic Algorithm Performance on the Vehicle Routing Problem Using Information Theoretic Landscape Measures

Mario Ventresca[1], Beatrice Ombuki-Berman[2], and Andrew Runka[3]

[1] Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Canada
mario.ventresca@utoronto.ca

[2] Department of Computer Science, Brock University, St. Catharines, Canada
bombuki@brocku.ca

[3] Department of Computer Science, Carleton University, Ottawa, Canada
arunka@connect.carleton.ca

Abstract. In this paper we examine the predictability of genetic algorithm (GA) performance using information-theoretic fitness landscape measures. The outcome of a GA is largely based on the choice of search operator, problem representation and tunable parameters (crossover and mutation rates, etc). In particular, given a problem representation the choice of search operator will determine, along with the fitness function, the structure of the landscape that the GA will search upon. Statistical and information theoretic measures have been proposed that aim to quantify properties (ruggedness, smoothness, etc) of this landscape. In this paper we concentrate on the utility of information theoretic measures to predict algorithm output for various instances of the capacitated and time-windowed vehicle routing problem. Using a clustering-based approach we identify similar landscape structures within these problems and propose to compare GA results to these clusters using performance profiles. These results highlight the potential for predicting GA performance, and providing insight self-configurable search operator design.

1 Introduction

We study the well known $\mathcal{NP}$-hard [7] vehicle routing problem (VRP). Due to its wide applicability the VRP has been widely studied (for detailed reviews, see [3,17,10]). In this paper, we focus on the capacitated vehicle routing problem (CVRP) [22] and vehicle routing problems with time windows (VRPTW) [3]. A typical VRP aims to design least-cost routes from a central depot to a set of geographically dispersed points/customers with various demands. Each customer is to be serviced exactly once by only one vehicle, and each vehicle has a limited capacity. The Vehicle Routing Problem with Time Windows (VRPTW) is an extension of the VRP whereby a *time window* during which service must be completed is associated with each customer. A vehicle may arrive early, but

M. Middendorf and C. Blum (Eds.): EvoCOP 2013, LNCS 7832, pp. 214–225, 2013.

it must wait until the designated time window to open before service can commence. The objective of the VRPTW is to minimize the number of vehicles used and the total distance travelled to service the customers without violating the capacity and time window constraints.

The main question we shed light on in this paper is whether common information theoretic summary measures of the fitness landscape structure actually provide reliable feedback as to the relative difficulty of solving specific VRP problem instances. That is, whether problem instances can be grouped based on these measures, and whether these groups are indicative of the eventual solution quality one could observe after an evolutionary algorithm terminates. So, we are not aiming to predict the final objective value or behaviour of the algorithm over time, rather, whether there is a correspondence between the measures and performance. We address this question by considering a large subset of benchmark VRP and VRPTW instances and measure the influence that common mutation and crossover search operators have.

A fitness-distance based analysis of problem difficulty for the CVRP was conducted in [9], using a variety of distance measures. Their results indicate the existence of a possible "big valley" structure in the landscape that contains more than half of the sampled problem instances. They argue that this provides a plausible explanation for the success of some well known heuristics on the given problem instances. The analysis considered a much smaller set of problem instances than is provided in this study, as well as considering only the CVRP, whereas we consider VRPTW as well. Similar conclusions were found in [5,6]. The waste-collection vehicle routing problem with time windows was studied in [19]. Many other studies of fitness landscapes exist in the literature for a variety of problems [13,12,1,11,21,25,16,20,18]. To our best knowledge, cluster analysis of information theoretic measures and their relationship to observed GA performance is unique.

2 Fitness Landscapes

A fitness landscape can be defined as a tuple $\mathcal{L} = (\mathcal{S}, f, \mathcal{N})$ where $\mathcal{S}$ is the search space of feasible solutions, $f\colon \mathcal{S} \mapsto \mathbb{R}$ is a fitness function. The function $\mathcal{N}(s)$ assigns to every s a set of neighbour solutions. Traditionally, neighbours are solutions reachable through a single application of the search operator, but this need not be the case.

Without loss of generality, the following assumes a maximization problem with search space $\mathcal{S}$ where a solution $s \in \mathcal{S}$ is defined to be a *local maximum* if its fitness is greater than or equal to all of its neighbours, i.e., $f(s) \geq f(w)\ \forall\ w \in \mathcal{N}(s)$, where the *neighbourhood* $\mathcal{N}(s)$ is defined as the set of solutions reachable from s by a single application of the search operator being considered. If a relatively high number of local optima are present in the landscape, it is termed *rugged*. When few optima exist, the landscape could be either smooth or flat depending on the existence of large attractive basins.

A *basin of attraction* of a solution s_n is defined [8] as the set of vertices $B(s_n) = \{s_0 \in V | \exists s_1, ..., s_n \in V$ where $s_{i+1} \in \mathcal{N}(s_n)$ and $f(s_{i+1}) > f(s_i)\ \forall i,\ 0 \leq i \leq n\}$.

The size of a basin is generally considered to be defined as the number of solutions within it. Local optima with relatively small attractive basins can be considered *isolated* [8]. Larger basins of attraction typically imply a smoother landscape. Landscapes characterized by few local optima generally contain large amounts of *neutrality* [2]; the fitness of neighbouring solutions remains essentially equal. When existing in neutral epochs, the current set of solutions will randomly drift about these neutral networks.

Problem difficulty may be deduced from analyzing the characteristics defined above. For instance, a landscape having few isolated optima with a high degree of neutrality is likely going to be more difficult to search than a smooth landscape with a single global optima (i.e., a large hill) because on average searching the landscape provides little information indicating the location of peaks. Various measures have been proposed to ascertain properties of the search space, for example [26,14,24,8]. We focus on the information theoretic measures proposed in [24] and [23].

The Information Content (IC) measures the ruggedness with respect to the flat or neutral areas of the landscape. The degree of flatness sensitivity is based on an empirically decided parameter ε which is restricted to the range $[0, ..., L]$, where L is the maximum fitness difference along the random walk. Consequently, the analysis will be most sensitive when $\varepsilon = 0$. This measure is calculated a

$$H(\varepsilon) = -\sum_{p \neq q} Pr_{[pq]} \log_6 Pr_{[pq]} \tag{1}$$

where probabilities $Pr_{[pq]}$ represent the probabilities of possible fitness transitions from solution p to q while performing a random walk. Each $[pq]$ are elements of the string $S(\varepsilon) = s_1 s_2 s_3 s_n$, of symbols $s_i \in \{\bar{1}, 0, 1\}$, where each s_i is recursively obtained for a particular value of ε based on Equation (2), so $s_i = \Psi_f(i, \varepsilon)$. Thus,$\varepsilon$ can be said to represent an accuracy or sensitivity parameter of the analysis.

$$\Psi(i, \varepsilon) = \begin{cases} \bar{1}, & \text{if } f_i - f_{i-1} < -\varepsilon \\ 0, & \text{if } |f_i - f_{i-1}| \leq \varepsilon \\ 1, & \text{if } f_i - f_{i-1} > \varepsilon \end{cases} \tag{2}$$

The Partial Information Content (PIC) indicates the modality or number of local optima present on the landscape. The underlying idea is to filter out repeated symbols of $S(\varepsilon)$ in order to acquire an indication of the modality of the random walk. The formula for computing PIC is given in Equation (3), where n is the length of the original walk and μ is the length of the summarized string $S'(\varepsilon)$.

$$M(\varepsilon) = \frac{\mu}{n} \tag{3}$$

The value for $\mu = \Phi_s(1, 0, 0)$ is determined via the recursive function

$$\Phi_s(i, j, k) = \begin{cases} k, & \text{if } i > n \\ \Phi(i+1, i, k+1), & \text{if } j = 0 \text{and} s_i \neq 0 \\ \Phi(i+1, i, k+1), & \text{if } j > 0, s_i \neq 0, s_i \neq s_j \\ \Phi(i+1, j, k), & \text{otherwise.} \end{cases} \tag{4}$$

When the value of $M(\varepsilon) = 0$ it indicates that no slopes were present on the path of the random walk, meaning that the landscape is rather flat or smooth. Similarly, if $M(\varepsilon) = 1$ then the path is maximally multi-modal and likely very rugged. Furthermore, it is possible to calculate the expected number of optima of a random walk of length n via

$$\mathbb{E}[M(\varepsilon)] = \left\lfloor \frac{nM(\varepsilon)}{2} \right\rfloor, \tag{5}$$

although we do not consider the expected modality in this analysis.

The Density-Basin Information (DBI) measure (Equation 6) indicates the flat and smooth areas of the landscape as well as the density and isolation of peaks. It therefore provides an idea of the landscape structure around the optima.

$$h(\varepsilon) = -\sum_{p\in\{\bar{1},0,1\}} Pr_{[pp]} \log_3 Pr_{[pp]} \tag{6}$$

$Pr_{[pp]}$ represents the probability of sub-blocks $\bar{1}\bar{1}$, 00 and 11 of occurring. A high number of peaks within a small area results in a high DBI value. Conversely, if the peak is isolated the measure will yield a low value. Thus, this information gives an idea as to the size and nature of the basins of the landscape. Landscapes with a high DBI content should be easier for an evolutionary algorithm to attract to the area of fitter solutions.

3 Representation and Genetic Operators

We use 66 standard VRPTW and CVRP benchmark instances[1] and consider seven search operators, four mutation and three crossover [4,15]:

- **Swap:** swap two random elements.
- **Inversion:** reverse the order of a contiguous segment of elements (i.e., 2-opt).
- **Insertion:** move an element to a random index.
- **Displacement:** select and move a contiguous segment of elements.
- **PMX:** exchange contiguous segments of elements between parents.
- **UOX:** randomly select subset of elements from each parent, maintaining ordering.
- **CX:** include a random element from parent 1 (P1), then include the element in P1 found at the index of P2 corresponding to the previous included element. Repeat until an element is encountered that already exists in the child, then repeat using an unselected element of P2.

We represent solutions as an array of integers, where each integer appears only once, and corresponds to a stop of the vehicle (i.e., a city). Crossover and mutation operators are applied directly to the representation. Transcribing a solution representation into a valid solution is accomplished by linearly traversing

[1] Available at `web.cba.neu.edu/~msolomon/problems.htm` and `http://osiris.tuwien.ac.at/ wgarn/VehicleRouting/neo/Problem%20Instances/instances.html`.

the representation and adding each stop to a vehicle until its capacity limit is reached. The process repeats using a new vehicle until the representation is fully traversed.

4 Landscape Analysis Results

In this section we present the results of the landscape analysis. The required statistics and probabilities are gathered by taking 2,000 random walks each of length 10,000 steps. We consider the PIC, IC and DBI values as features for a given (instance, operator) pairing and perform a clustering on the scaled values. The optimal clustering model is determined according to the Bayesian Information Criteria (BIC) for expectation maximization initialized by hierarchical clustering for parametrized Gaussian mixture models. The implied landscape of each group can then be analyzed separately. Due to space limitations we provide summary results, but full statistics are available by contacting the authors.

Figure 1 presents the clustering results for the CVRP, presented with respect to the two main principal components. The left figure displays the 8 clusters found for the crossover operators, and the right figure shows the 11 mutation operator-based clusters. Each point represents a (problem instance, operator) pairing and is labelled according to the operator, as indicated in the figure caption. Immediately noticeable is that most clusters are composed of solely a single type of search operator, indicating similarly induced landscape structure for the corresponding problem instances. Additionally, the DBI value is found to explain very little of the total variance.

Table 1 shows the cluster means for each cluster in Figure 1. Class names have been determined by using the operator name (I)nsertion, (D)isplacement, (S)wap, in(V)ersion and sequence of the particular 1, 2 or 3-combination. For instance IV-1, is the first cluster that is represented entirely by insertion and inversion operators, where there are more insertion operators in the class. Cluster ALL contains a mix of all operators.

All 66 UOX results cluster together (UOX-1), indicating that the operator is invariant with respect to the chosen metrics to the problem being considered. The PMX operator is grouped into three clusters. PMX-1 and PMX-2 have very similar values; indeed they are adjacent in the clustering of Figure 1. The practical difference between these two classes is that PMX-1 will have a slightly more rugged landscape. Class PMX-3 has 4 elements, and their corresponding information theoretic measures indicate the landscapes contain a larger variety of shapes, but the overall landscape is slightly flatter with a relatively high degree of peak density. The CX operator has four very distinct landscape structures. CX-1 has a relatively large IC value, implying a landscape that may contain more ruggedness. The DBI measure shows that the density of these peaks is moderately high, compared to UOX-1. The majority (52/66) of PMX results have a landscape that seems less rugged. This leads to the hypothesis of UOX having the most desirable search space, followed by PMX and CX, respectively.

An important aspect is the separation of clusters by problem size, except for UOX-1, which seems problem-invariant. Clusters 1 and 4, represent the PMX

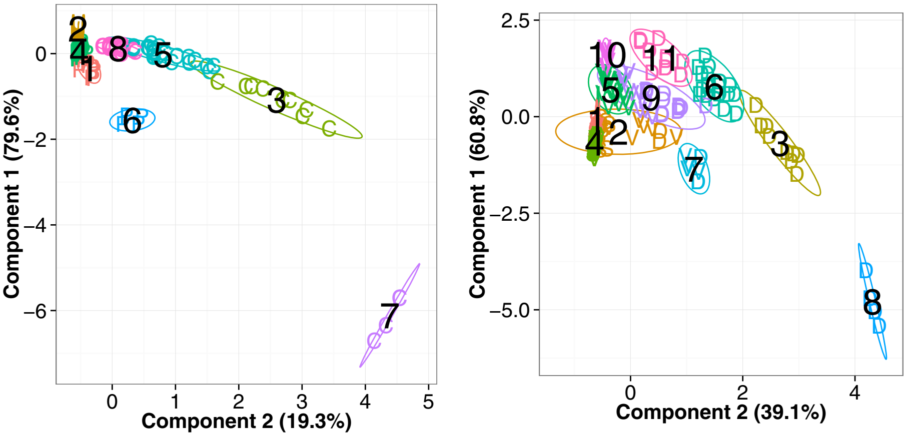

Fig. 1. Optimal clustering of the crossover (left) and mutation (right) problem-operator pairings for CVRP. The plot is shown with respect to the top 2 principal components. Using this approach, 8 distinct crossover and 11 mutation landscapes have been labelled, respectively. Points are placed using a multidimensional scaling with Euclidean distance metric and labelled as (C)X, (U)OX and (P)MX crossover. Mutations are labelled as (S)wap, in(V)ersion, (I)nsertion and (D)isplacement.

operator, where the separation of clusters corresponds nearly perfectly to a difference in problem size; cluster 1 contains C/R10X problem and cluster 4 contains mostly C/R20X problems, respectively. Similarly for clusters 3 and 5, but considering the CX operator. Given this information it can be deduced that the CX and PMX operators have similar landscape structures.

Figure 2 displays the clustering results for the VRPTW. The crossover landscapes form 5 clusters, each having nearly negligible covariance between the principal components. In the right diagram 10 mutation landscape clusters are identified. Both clusterings show very little overlap between elements and sufficient separation that subsequent comparisons could be more straightforward than for the CVRP.

Table 2 shows the cluster means for the VRPTW clustering results shown in Figure 2. For mutation-based clustering, clusters 4 and 9 are composed mostly of displacement operator results. Both of these search spaces contain a larger variety of shapes than the other 8 clusters, while maintaining a high degree of ruggedness as is evident from the PIC and DBI measures. Moreover, the attractive basins seem to be relatively small as well. In contrast, clusters containing the swap and inversion operators have indications of smoother landscapes (low IC and PCI accompanied by high DBI measures). The inversion operator has characteristics that typically result in landscapes that contain a slightly higher degree of ruggedness.

Table 1. The cluster means for CVRP. The first 8 clusters are crossover and the next 11 are mutation-based.

Cluster #	Class	IC	PIC	DBI	# elements
1	PMX-1	0.4196	0.5821	0.5936	10
2	UOX-1	0.4059	0.6273	0.5699	66
3	CX-1	0.6601	0.5465	0.5804	9
4	PMX-2	0.4074	0.6008	0.5851	52
5	CX-2	0.5148	0.5983	0.5716	27
6	PMX-3	0.4754	0.5234	0.6142	4
7	CX-3	0.8054	0.3009	0.6361	4
8	CX-4	0.4571	0.6051	0.5764	26
1	ISV-1	0.3940	0.5327	0.6183	53
2	ALL-1	0.4275	0.5241	0.6186	21
3	D-1	0.6699	0.5163	0.5979	9
4	SIV-1	0.3915	0.5166	0.6247	68
5	IV-1	0.4156	0.5535	0.6076	28
6	D-2	0.5712	0.5556	0.5876	19
7	VD-1	0.5406	0.4985	0.6183	7
8	D-3	0.8091	0.4066	0.6430	4
9	VD-2	0.4765	0.5480	0.6033	27
10	V-1	0.4095	0.5777	0.5969	13
11	D-4	0.4912	0.5753	0.5879	15

Table 2. The cluster means for VRPTW. The first 5 clusters are crossover and the next 10 are mutation-based.

Cluster #	Class	IC	PIC	DBI	size
1	PMX-1	0.4089	0.6338	0.5656	12
2	UOX-1	0.4073	0.6443	0.5590	21
3	CX-1	0.4647	0.6333	0.5581	12
4	PMX-2	0.4068	0.6152	0.5770	9
5	CX-2	0.4631	0.6180	0.5679	9
1	SI-1	0.4295	0.5795	0.5938	15
2	V-1	0.4230	0.6205	0.5721	8
3	I-1	0.4672	0.6038	0.5757	7
4	D-1	0.5058	0.6151	0.5628	12
5	S-1	0.4013	0.5473	0.6117	7
6	VS-1	0.4273	0.5969	0.5849	7
7	S-2	0.4004	0.5309	0.6186	3
8	VS-2	0.4153	0.5668	0.6016	7
9	DI-1	0.4989	0.5888	0.5795	12
10	V-2	0.4079	0.6136	0.5779	6

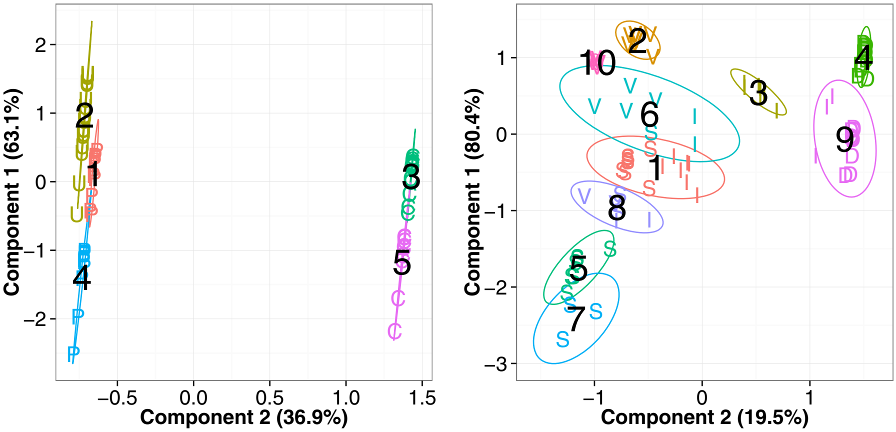

Fig. 2. Optimal clustering of the crossover (left) and mutation (right) problem-operator pairings for VRPTW. Using this approach, 5 distinct crossover and 10 mutation landscapes have been discovered, respectively. Points are placed using a multidimensional scaling with Euclidean distance metric and labelled as (C)X, (U)OX and (P)MX crossover. Mutations are labelled as (S)wap, in(V)ersion, (I)nsertion and (D)isplacement.

5 Genetic Algorithm Results

We employed a genetic algorithm that exclusively uses each of the seven search operators. The GA is run for 5000 generations with a population size of 200. Selection is according to a 3-tournament whereby the best of three randomly selected individuals is carried on to the next population (repeated until 200 individuals are selected). An elitism strategy is also incorporated; we maintain the top two best found solutions at each generation. The GA is run 30 times.

Given space limitations we forgo presenting full statistics about the obtained objective value, etc. Such results are obtainable from the authors. Our goal is to ascertain whether the information theoretic measures are useful indicators of problem difficulty. Since optimal objective values are not scaled to the same range, and in general each problem instance will have different possible evaluations. Instead, we examine how the different search operators compare relative to each other and create performance clusters based on these results. Subsequent comparison of the elements of these performance clusters and landscape clusters is then performed.

CVRP. In all cases the UOX crossover operator showed significantly better fitness across all problem instances when compared to the other crossover operators. Moreover, UOX typically yielded a more desirable outcome than all the

Table 3. Performance profile clusters for mutation operators on CVRP and VRPTW. A < indicates a statistically significant difference between the respective groups and a ∼ represents a non-significant result. The Welsh t-test was used at a 0.95 confidence level to ascertain significance. Left to right ordering of operators is according to mean value. The first 20 groups are CVRP and the next 11 are VRPTW.

Group	Relationship	Problems
1	inversion < insertion < swap ∼ displace	A-n53-k7, A-n54-k7, A-n55-k9, B-n66-k9, B-n67-k10, c50
2	inversion ∼ insertion ∼ swap < displace	A-n60-k9, B-n50-k8, B-n57-k9
3	insertion ∼ inversion < swap < displace	A-n62-k8, f134
4	inversion < displace < swap ∼ insertion	A-n63-k9, E-n76-k10
5	inversion < insertion ∼ swap < displace	A-n63-k10, A-n80-k10, B-n63-k10, c75, E-n101-k14, M-n151-k12, P-n55-k10
6	inversion < insertion < swap < displace	A-n64-k9, A-n69-k9, B-n68-k9, B-n78-k10, c100, c100b, E-n76-k7, E-n76-k8, E-n101-k8, M-n101-k10, M-n121-k7, C101, R101, RC101
7	inversion ∼ displace < insertion ∼ swap	A-n65-k9, B-n51-k7
8	inversion ∼ insertion < displace ∼ swap	B-n50-k7
9	inversion < insertion < displace ∼ swap	B-n52-k7, E-n51-k5, C201, R201
10	inversion ∼ insertion< swap ∼ displace	B-n56-k7, C101_50, R101_50
11	inversion ∼ insertion < swap < displace	c120, c150
12	swap < inversion ∼ insertion < displace	c199
13	inversion < insertion < displace < swap	f71, C201_50, R201_50, RC101_50, RC201_50, RC201
14	swap ∼ inversion ∼ insertion < displace	M-n200-k17
15	inversion < swap ∼ insertion <displace	P-n60-k10
16	insertion < inversion < swap ∼ displace	tai75a, tai75b, tai150c
17	insertion ∼ inversion < swap ∼ displace	tai75c
18	insertion ∼ inversion < displace ∼ swap	tai75d
19	insertion < inversion < swap <displace	tai100a, tai100b, tai100d, tai150b, tai150d
20	insertion < inversion ∼ swap < displace	tai150a
1	swap < insertion < inversion < displace	C101, C207
2	swap ∼ insertion ∼ inversion < displace	C102
3	inversion ∼ insertion ∼ swap < displace	C103, C203
4	inversion < insertion ∼ swap < displace	C104
5	swap ∼ insertion < inversion < displace	C105, C107
6	swap < insertion ∼ inversion < displace	C106, C206, C208
7	insertion ∼ swap ∼ inversion < displace	C108, C109, R103
8	swap < insertion < inversion ∼ displace	C201, C202, C205, R101
9	inversion < insertion < swap < displace	C204
10	swap ∼ insertion < inversion ∼ displace	R102
11	insertion < inversion < swap < displace	R104

mutation operators. The PMX operator consistently yielded more desirable results when compared to the CX operator. However, there is no significant trend of more desirable results of PMX when compared to the mutation operators.

Table 3 presents the results of comparing the final mean values for each of the 66 problem instances. The four mutation operators are compared using a Welch t-test at 0.95 confidence level, and the pairwise results of this test are reported, where a $<$ indicates statistical significance and $\sim$ notes the lack thereof, respectively. The relationship ordering was determined according to the mean values attained (not shown). Overlap with problem instances shown in the landscape analysis is large, yielding approximately 70% similarity. Merging clusters with single elements into an existing cluster increases the similarity to 85% similarity. Similar results were also observed for crossover landscapes, but were omitted due to space limitations.

The displacement operator typically occupies the lowest rank (in all but 6 groups), as would be expected considering the landscape analysis. Moreover, Class D-2 in Table 2 indicates a relatively easy search space. Investigating the particular instances for the associated problems A-n63-k9, A-n65-k9, B-n51-k7 and E-n76-k10 further supports the landscape analysis. The results for the displacement operator on these instances is significantly improved from other displacement results (using the relationship ranking as a measure), as the operator is the second rank for Group 4 and 7, respectively.

VRPTW. The results found when running a GA using the four mutation operators are given in Table 3, and grouped according to the statistical dominance relation described above. The swap, inversion and insertion operators tend to occupy the lowest rank (i.e., most desirable outcome), with the swap operator being most frequent. The results from the clustering of landscape measures had indicated that this result should occur.

Another prediction implied by the search space analyses is the deficiency of the displacement operator. For all the 11 rankings in Table 3 displacement is found to be least desirable. Table 3 shows the relatively large degree of displacement operator deficiency as the mean results can be observed to have large effect sizes (in the negative result sense).

We conducted a similar analysis for the three crossover operators (not shown), revealing the relative power of the UOX operator for these problems. Indeed, dominance was observed over all crossover and mutation operators; additionally, a very large effect size was evident. As was also discovered above for the CVRP, the PMX operator consistently, and statistically, dominates the results obtained by CX. In comparison to the fitness landscape clusters we see approximately 87% overlap of problem instances in the clusters.

6 Conclusion

The main question we aimed to address in this study was aimed at whether information theoretic landscape measures can actually be used to discriminate between problem instance difficulty for VRPTW and CVRP. To this end, numerous benchmark problem instances were considered and seven common search operators were examined. We found that the landscape measures can be clustered into groups that tend to contain mostly one type of search operator. This was true of both CVRP and VRPTW. In order to ascertain whether these clusters can be used to predict outcomes of a genetic algorithm we proposed the use of performance profiles that represent relative ordering of GA results. These profiles are also clustered according to the ordering they represent. We find significant overlap between the landscape and performance clusters. Further study may shed light on automatic search operator design and configuration.

More study of the performance profile approach, and other methods of clustering and comparing landscape and algorithm output may provide deeper insight into predictability of GAs. Future work also includes the examination of different problem representations, which have greatly impact the ability of an algorithm to obtain quality results and whether these results are limited to GAs. In the same vein, consideration of combinations of these, and more advanced, search operators may give some insight into practical implementations.

Acknowledgement. We thank the anonymous reviewers for their valuable insight and comments, and the Natural Sciences and Engineering Research Council of Canada for funding.

References

1. Alander, J.T., Zinchenko, L.A., Sorokin, S.N.: Analysis of fitness landscape properties for evolutionary antenna design. In: IEEE International Conference on Artificial Intelligence Systems, pp. 363–368 (2002)
2. Barnett, L.: Netcrawling-Optimal Evolutionary Search with Neutral Networks. In: Congress on Evolutionary Computation, pp. 30–37 (2001)
3. Braysy, O., Gendreau, M.: Vehicle routing problem with time windows, part ii: Metaheuristics. Transportation Science 39, 119–139 (2005)
4. Caramia, M., Onori, R.: Experimenting crossover operators to solve the vehicle routing problem with time windows by genetic algorithms. International Journal of Operational Research 3(5), 497–514 (2008)
5. Czech, Z.J.: Statistical measures of a fitness landscape for the vehicle routing problem. In: IEEE International Symposium on Parallel and Distributed Processing, pp. 1–8 (2008)
6. Czech, Z.J.: A parallel simulated annealing algorithm as a tool for fitness landscape exploration. In: Ros, A. (ed.) Parallel and Distributed Processing, pp. 247–271. InTech (2010)
7. Garey, M.R., Johnson, D.S.: Computers and Intractability, A Guide to the Theory of NP-Completeness. W. H. Freeman and Company (1979)
8. Jones, T.: Evolutionary Algorithms, Fitness Landscapes and Search. PhD thesis, University of New Mexico (1995)
9. Kubiak, M.: Distance measures and fitness-distance analysis for the capacitated vehicle routing problem. In: Doerner, K.F., Gendreau, M., Greistorfer, P., Gutjahr, W., Hartl, R.F., Reimann, M. (eds.) Metaheuristics. Operations Research Computer Science Interfaces, vol. 39, pp. 345–364. Springer (2007)

10. Laporte, G.: Fifty years of vehicle routing. Transportation Science 43, 408–416 (2009)
11. Mattfeld, D.C., Bierwirth, C., Kopfer, H.: A search space analysis of the job shop scheduling problem. Annals of Operations Research 86, 441–453 (1999)
12. Merz, P., Freisleben, B.: Memetic Algorithms and the Fitness Landscape of the Graph Bi-Partitioning Problem. In: Eiben, A.E., Bäck, T., Schoenauer, M., Schwefel, H.-P. (eds.) PPSN 1998. LNCS, vol. 1498, pp. 765–774. Springer, Heidelberg (1998)
13. Merz, P., Freisleben, B.: Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. IEEE Transactions on Evolutionary Computation 4(4), 337–352 (2000)
14. Naudts, B., Kallel, L.: A Comparison of Predictive Measures of Problem Difficulty in Evolutionary Algorithms. IEEE Transactions on Evolutionary Computation 4(1), 1–16 (2000)
15. Nazif, H., Lee, L.S.: Optimized crossover genetic algorithm for vehicle routing problem with time windows. American Journal of Applied Sciences 7(1), 95–101 (2010)
16. Ombuki-Berman, B., Ventresca, M.: Search difficulty of two-connected ring-based topological network designs. In: IEEE Symposium on Foundations of Computational Intelligence, pp. 267–274 (2007)
17. Potvin, J.: State-of-the art review evolutionary algorithms for vehicle routing. INFORMS Journal on Computing 21, 518–548 (2009)
18. Reeves, C.: Direct statistical estimation of GA landscape properties. In: Foundations of Genetic Algorithms 6, pp. 91–107 (2000)
19. Runka, A., Ombuki-Berman, B., Ventresca, M.: A search space analysis for the waste collection vehicle routing problem with time windows. In: Genetic and Evolutionary Computation Conference, pp. 1813–1814 (2009)
20. Schiavinotto, T., Stutzle, T.: A review of metrics on permutations for search landscape analysis. Computers and Operations Research 34(10), 3143–3153 (2007)
21. Tavares, J., Pereira, B., Costa, E.: Multidimensional knapsack problem: A fitness landscape analysis. IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cynernetics 38(3), 604–616 (2008)
22. Toth, P., Vigo, D.: The Vehicle Routing Problem. SIAM Monographs on Discrete Mathematics and Applications (2002)
23. Vassilev, V.K., Fogarty, T.C., Miller, J.F.: Information Characteristics and the Structure of Landscapes. Evolutionary Computation 8(1), 31–60 (2000)
24. Vassilev, V.K., Fogarty, T.C., Miller, J.F.: Fitness Landscapes: from Theory to Application. In: Advances in Evolutionary Computation: Theory and Applications, pp. 3–44. Springer (2003)
25. Ventresca, M., Ombuki-Berman, B.: Search space analysis of recurrent spiking and continuous-time neural networks. In: IEEE International Joint Conference on Neural Networks, pp. 8947–8954 (2006)
26. Weinberger, E.: Correlated and Uncorrelated Landscapes and How to Tell the Difference. Biological Cybernetics 63, 325–336 (1990)