
TOWARDS DECISION SUPPORT IN DYNAMIC BI-OBJECTIVE VEHICLE ROUTING

Jakob Bossek

Optimisation and Logistics
The University of Adelaide
Adelaide, Australia
jakob.bossek@adelaide.edu.au

Christian Grimme

Information Systems and Statistics
University of Münster
Münster, Germany
christian.grimme@wi.uni-muenster.de

Günter Rudolph

Dept. of Computer Science
TU Dortmund University
Dortmund, Germany
guenter.rudolph@tu-dortmund.de

Heike Trautmann

Information Systems and Statistics
University of Münster
Münster, Germany
heike.trautmann@wi.uni-muenster.de

ABSTRACT

We consider a dynamic bi-objective vehicle routing problem, where a subset of customers ask for service over time. Therein, the distance traveled by a single vehicle and the number of unserved dynamic requests is minimized by a dynamic evolutionary multi-objective algorithm (DEMOA), which operates on discrete time windows (eras). A decision is made at each era by a decision-maker, thus any decision depends on irreversible decisions made in foregoing eras. To understand effects of sequences of decision-making and interactions/dependencies between decisions made, we conduct a series of experiments. More precisely, we fix a set of decision-maker preferences D and the number of eras n_t and analyze all $|D|^{n_t}$ combinations of decision-maker options. We find that for random uniform instances (a) the final selected solutions mainly depend on the final decision and not on the decision history, (b) solutions are quite robust with respect to the number of unvisited dynamic customers, and (c) solutions of the dynamic approach can even dominate solutions obtained by a clairvoyant EMOA. In contrast, for instances with clustered customers, we observe a strong dependency on decision-making history as well as more variance in solution diversity.

Keywords Transportation · vehicle routing · decision making · multi-objective optimization · combinatorial optimization · orienteering · dynamic optimization

1 Introduction

In industry, especially logistics [1] and 24/7 production [2], many decisions must be made under uncertain knowledge about future events. Production planning problems may include tasks appearing over time, which is often expressed via release dates [3], while logistic problems like the traveling salesperson problem (TSP) may be augmented with dynamically changing traffic [4] or – more abstract – with dynamically changing distance matrices [5].

Here, we consider the scenario that repeated decision requests are not independent from the previous ones; rather, every decision reduces the number of possible actions/decisions in future affecting the quality of the best possible solution that can be achieved.

If optimization methods are used to support the decision process, an optimizer could iterate for short time intervals (eras). After each era, a human *decision maker* (DM) may decide whether the proposed solution is executed or not. This situation is studied in the field of *interactive optimization* [6] (headword: human-in-the-loop).

In analyzing the quality of optimization algorithms it is quite common to replace the human DM by a software agent (automatic decision maker) [7] to speed up experiments or to make it amenable to a theoretical analysis. The software agent may obey some simple guidelines or even a complex set of user-supplied preferences to arrive at a decision. The situation becomes even more complex in case of multiple objectives [8]; if the multi-objective optimizer uses the a-posteriori approach, the (automatic) DM must additionally decide, which non-dominated solution should be picked from the Pareto-front.

Here, we also take the approach of replacing the human by an automatic decision maker in case of a dynamic bi-objective vehicle routing problem, where the goal is to minimize both the distance traveled by a single vehicle and at the same time minimize the number of unvisited customers which ask for services over time. This special variant of the TSP problem includes the additional problem of subset selection of serviced customers and is similar to the so-called TSP with profits [9]. However, the problem considered here comprises a dynamically growing set of customers, who request service over time.

It is important to note, that the focus of this work is not on the performance of the optimization algorithm but (1) on the impacts of specific (automatic) decision making rules on the final solution and (2) the visualization of subsequent decisions and solutions as a preliminary step towards an (interactive) decision support system. The dynamically generated solutions are compared to the solution of the so-called clairvoyant optimizer which has complete knowledge about the future, i.e., request times of dynamic customers. In our study we run experiments for all possible combinations of a fixed set of automatic DM decision rules to scrutinize as many aspects of the decision rule as possible.

The work is organized as follows: Section 2 details the dynamic multi-objective problem before describing the dynamic multi-objective evolutionary optimization algorithm supporting the decision maker in Section 3. The experimental setup (including the automatic decision maker) and results are described in Section 4. The conclusion and the prospects of this work towards inclusion in interactive decision support systems are presented in Section 5.

2 Problem Description

We consider a dynamic vehicle routing problem for which the overall goal is to have a single service vehicle visiting customer locations from the set C of all customer locations.

The dynamic character of this problem originates in the time-dependent appearance of customers from C . The set of customers $C \setminus \{1, N\} = C^m \cup C^o$ is divided into two disjoint subsets: Mandatory customers C^m are known at time $t = 0$ and must be visited by the vehicle while dynamic customers $c \in C^o$ ask for service at request time $r_t(c)$, $t > 0$ as time passes by. They can either be visited or not. In a real-world context, we may imagine the vehicle as a customer service vehicle with fixed orders and spare time to handle dynamically emerging service requests.

We assume that the vehicle leaves w.l.o.g. at a start depot $1 \in C$ and ends at an end depot $N \in C$.¹ The optimization task is to (1) minimize the overall tour length and at the same time (2) minimize the number of unserved dynamic customers. Undoubtedly, the goals conflict with each other and we are faced with a complex dynamic combinatorial multi-objective optimization problem (MOP), for which we strive to find a set of (near) optimal compromises. Here, we adopt the notion of *Pareto-dominance* for a definition of optimal compromise solutions: for two vehicle tours x and y we say that x dominates y , if x is not worse in any objective and strictly better in at least one objective [10]. The set of all non-dominated solutions is termed the *Pareto-set*, its image in the objective space is called the *Pareto-front*. Hence, in our scenario for each point in time $t \geq 0$ a bi-objective problem needs to be solved and the problem can be fully described by the sequence of all trade-off solution sets. Since time is continuous, this approach is infeasible in practice. A common approach is to discretize the time horizon, i. e., the time interval in which dynamic customers pose requests, into a number of phases n_t , so-called *eras*, of length $\Delta \in \mathbb{R}_{\geq 0}$ (see, e. g. [11]).

At the beginning of each era j time $t = (j - 1) \cdot \Delta$ has already passed and we may consider the set $C^m \cup C_{<t}^o$, with $C_{<t}^o$ being the set of dynamic customers, which asked for service before time t , as a static MOP. Tackling this static MOP with the algorithm of our choice results in an approximation of the Pareto-set. Finally, a decision maker (DM) is given the resulting set of trade-off solutions in each era j and has to decide on how to guide the vehicle on the road until the beginning of the next era where then, more knowledge about dynamic customers becomes available. A crucial aspect here is that in each era $j > 0$ time already passed and consequently the vehicle might already have visited a subset of mandatory and/or optional customer locations. These decisions are irreversible and (1) may have a strong impact on the achievable solution quality (this is because a part of the solution space may become infeasible) and (2) exhibit a strong dependence on the decisions made by the DM in foregoing eras.

¹This assumption is more general than starting and ending at the same depot, although a circular tour may be the normal case in real-world scenarios.

Static formulations of bi-objective vehicle routing problems or so-called traveling salesperson problems with profits [12] have been addressed by several authors so far, e. g. proposing exact ε -constraint methods [13], approximations schemes [14] or meta-heuristics [9, 15]. Additionally, dynamic decision making gained some attention in the context of vehicle routing problems in general [16, 17]. However, work on the intersection, i. e. dynamic multi-objective vehicle routing problems is still rare. Braekers et al. [18] show in their extensive literature review that less than 3% of the literature between 2009 and 2015 address dynamic aspects. According to them, mentionable research includes work by [19], [20], [21], [22], and [23]. This list may be extended by mostly evolutionary approaches of [24], [25]. Although of dynamic nature, the problem formulations have very diverse characteristics like moving service time windows, multiple vehicles, or changing structures of the network. Own work addressed a clairvoyant and non-dynamic variant of the here discussed problem with an evolutionary multi-objective algorithm (EMOA) [26]. We enhanced the EMOA in a follow-up work by local search integration into the evolutionary search process [27]. This clairvoyant approach is considered here as reference approach. A sophisticated dynamic variant was presented in [28].

3 The Dynamic Multi-Objective Evolutionary Algorithm

We adopt the DEMOA introduced in [28]. Note that the focus of this work is on the influence of subsequent decision-making. Therefore, and due to space limitations, we omit most implementation details and present the working principles. For detailed pseudo-code we refer the interested reader to [28]. Also the implementation is available in a public GitHub repository². We advise the reader to consult Fig. 1 for visual support while reading the following text.

The input for the DEMOA is a problem instance $C = \{1, N\} \cup C^m \cup C^o$, a time resolution Δ , a number of eras n_t and a population size μ . The optimization process starts at time $t = 0$ and the algorithm treats the problem as a sequence of n_t static MOPs (see Section 2). Note however, that the first era is a special case, since $C_{\leq t}^o = \emptyset$, i. e., no dynamic requests arrived so far, and there is no possibility to vary the second objective. Hence, in the 1st era, a single-objective Hamiltonian path problem (HPP) on the set $\{1, N\} \cup C^m$ has to be solved. An approximate solution is calculated with the state-of-the-art solver EAX [29] for the symmetric Travelling-Salesperson-Problem (TSP) after reducing the HPP to a symmetric TSP problem by a sequence of transformations [30]. Note that in the first era the decision-maker has no choice as there is just a single solution (see era 1 in Fig. 1). In subsequent eras $j = 2, \dots, n_t$ time $t = (j - 1) \cdot \Delta$ already passed, $C_{\leq t}^o$ is non-empty and as a consequence the problem turns into a true multi-objective problem. Here, the DEMOA calls a static EMOA whose internals are discussed in the following. The EMOA initializes a multi-set P of μ candidate solutions. Each candidate solution $x \in P$ is fully described by three vectors of length $N - 2$. A binary vector $x.b = (b_2, \dots, b_{N-1})$ indicates which customers are to be visited by the service vehicle (note that the depots 1 and N need to be visited in any case and are thus not encoded).³ Another vector $x.t$ holds a permutation of $C^m \cup C^o = \{2, \dots, N - 1\}$, i.e., the actual tour where during fitness evaluation only those entries i with $b_i = 1$ are considered. Finally, the vector $x.p = (p_2, \dots, p_{N-1}) \in [0, 1]^{N-2}$ stores per-customer mutation probabilities. If $x.p_i = 0$, the corresponding customer is fixed and not affected by mutation. While in the second era individuals are generated at random (fixing mandatory customers by setting $x.b_i = 1$ and $x.p_i = 0$ for $i \in C^m$), in eras $j \geq 3$ more effort is put into the initialization to transfer as much information from the solution set of the preceding era $j - 1$ as possible. The challenge here is that once era $j \geq 3$ starts, the vehicle may already have visited dynamic customers with request times $r_t(i) \leq (j - 2) \cdot \Delta$ (this is illustrated by means of example in Fig. 1 last column. Here, bold edges show the fixed, already driven initial tour) given by the decision at the end of the previous era (Fig. 1). As a consequence, those customers cannot be inactive and hence need to be treated as mandatory customers by the EMOA in all upcoming eras. Moreover, the initial tour, i.e., the part of the tour that has already been driven by the vehicle, needs to be identical for all feasible individuals. Here, the EMOA relies on a sequence of repairing mechanisms.

Given the population P the algorithm continues by adopting a $(\mu + \lambda)$ -strategy with NSGA-II [31] survival selection. Variation is based on feasibility-preserving mutation. Here, each bit $x.b_i$ is flipped independently with probability p_i . Subsequently, swap-mutation alters the permutation string $x.t$: with probability $p_{\text{swap}} \in (0, 1)$ a sequence of σ_{swap} exchanges is performed. In addition, every k generations the population is boosted towards shorter tours by applying EAX local-search where the EMOA accounts for the fact that certain nodes have already been visited. Once the stopping condition has been triggered, e.g., a maximum number of generations has been reached, the solution set is presented to a decision-maker who has to decide on exactly the solution which determines the adaptation of the ongoing vehicle route and which serves as a template for the initialization of the population in the next era.

²<https://github.com/jakobbossek/dynvrp>

³We want to point out that the current implementation knows the total number N of customers in advance for legacy reasons. However, it only operates on those customers, who asked for service before time t . Clearly, it is straight forward to adapt the implementation into a true black-box scenario, where the number of dynamic requests is not known a-priori.

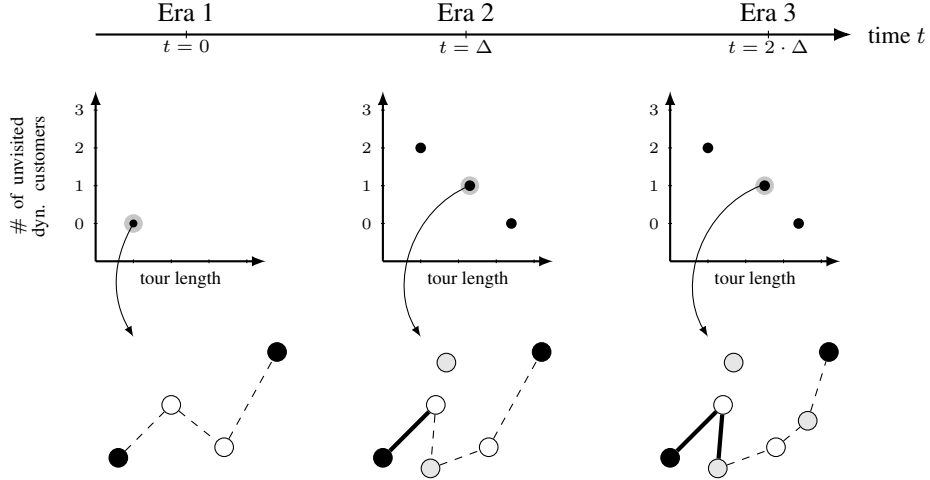


Figure 1: Exemplary illustration of three eras of the DEMOA. The scatterplots show the Pareto-front approximations at eras $j = 1, 2, 3$. Here, solutions selected by the DM are highlighted (\bullet). Below, the selected tours are visualized (depots \bullet , mandatory customers \circ and dynamic customers \odot). The thick solid path highlights the fixed partial tour already driven by the vehicle at the beginning of the corresponding era. Note that in this example solutions with three unvisited customers are infeasible in the 3rd era since one dynamic customer is already served.

4 Computational Experiments

4.1 Experimental Setup

In order to gain insights into the decision making process we conducted a two-stage study. In a first series of experiments we perform a systematic study of decision making strategies. Subsequent experiments focus on a selected sample of decision making strategies on a broader set of instances in order to confirm the lessons learned.

For the exhaustive experiments we selected three structurally different instances from the pool of instances introduced in [26]: one instance with customer locations spread uniformly at random in the Euclidean plane and two clustered instances with two and three groups of instances respectively. All instances share $|C^m| = 25$ mandatory customers (including depots) and $|C^o| = 75$ dynamic customers with $N = 100$ customers in total. For details on the generation process we refer to [26]. We fixed the number of eras $n_t = 7$ and considered three different ranking based rules for the decision maker in each era. For ranking, the solutions of the approximation set $P_i = \{p_1, \dots, p_\mu\}$ obtained in era i are sorted in ascending order of tour length and therefore in descending order of the number of unvisited customers. Let $p_{(1)} < p_{(2)} < \dots < p_{(\mu)}$ denote this order. The d -rank decision maker ($d \in [0, 1]$) then decides for the solution $p_{(k)}$ with $k = \lceil d \cdot \mu \rceil$. Note that small values of d favor solutions with short tours whereas values closer to 1 put a higher emphasis on keeping the number of unvisited dynamic customers low. Our setup considers $d \in D = \{0.25, 0.5, 0.75\}$ in each era⁴; we use a tuple notation for sequences of decision maker decisions, termed *decision paths* in the following, e. g., $(0.25, 0.25, 0.25, 0.25, 0.75, 0.75, 0.75)$ describes the decision path where the DM puts more emphasis on short tours in the first four eras but decides to cover more dynamic requests in the last three eras.

Table 1: Dynamic EMOA parameterization.

Parameter	Setting
Generations per era	65 000
μ, λ	100
p_{swap}	0.6
σ_{swap}	$N/10 = 10$
Local search at generations	initial, half-time, last
Time limit for local search	1s

⁴We do not consider the extremal values 0 and 1 for two reasons: (1) to avoid a combinatorial explosion of possible configurations for exhaustive evaluation; (2) extremal decisions are usually unrealistic as they either imply to ignore all optional customers (0) or to accept every optional customer (1).

Table 2: Mean values and standard deviations of tour lengths and the number of unvisited customers of the solutions in the last era split by the instance type and final decision in the last era.

Type	Last decision	Tour length		#Dyn. customers	
		Mean	Std	Mean	Std
2 clusters	0.25	997.7	59.94	22.495	5.7486
	0.50	1070.3	57.16	15.645	4.3148
	0.75	1174.9	65.14	8.105	2.4266
uniform	0.25	1227.6	28.61	37.654	0.7662
	0.50	1416.4	38.62	24.854	0.9167
	0.75	1657.1	42.32	12.069	0.8743

We run experiments for all $|D|^{n_t} = 2187$ decision paths, i. e., we cover all kinds of decision scenarios. At each stage, five independent runs were performed for each of the three instances resulting in a set of 32 805 experiments in total.

The results of this “exhaustive” experimentation served as starting point for subsequent experiments on a broader set of benchmark instances; all 75 instances with $N = 100$ introduced in [26]. However, due to combinatorial explosion only a small subset of four decision paths (with different outcomes in the last era) were considered here for an extensive analysis (for details on the selection process, see Section 4.2). For each combination of problem instance and decision path we run the DEMOA 25 times independently in this series of experiments. The parameter configuration of the DEMOA follows the suggestions in [28] and is listed in Table 1.

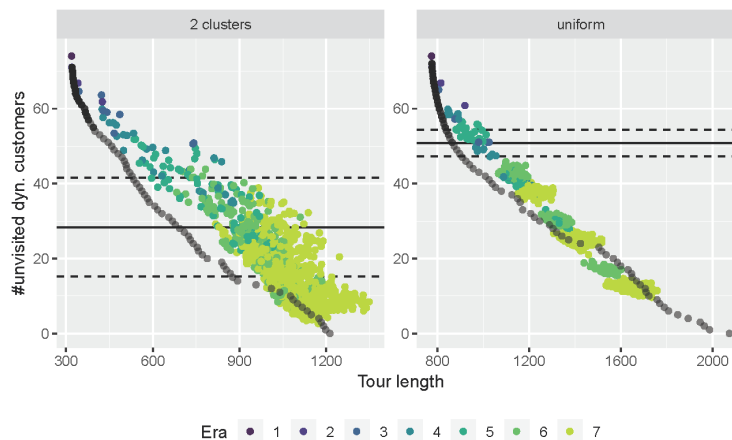


Figure 2: Union of all final decisions made across the complete set of all considered decision maker strategies. Solid and dashed lines indicate the mean upper bound \pm three times the standard deviation of the number of unvisited customers in the last era (based on all decision maker strategies). Small black dots represent the Pareto-front approximation of the clairvoyant EMOA.

4.2 Results

Next, we investigate the influence of considered graph topologies as well as the implications of final and intermediary decision making onto the solution development over time. Therefore, we perform a step-wise narrowing of perspective to focus on interesting insights for our considered instances, topologies, and decision strategies in the context of our exhaustive experimental results.

4.2.1 General observations

In Figure 2 we provide a first overview of the results for all decisions in each era and for all $|D|^{n_t}$ decision paths for the uniform instance and the instance with two clusters⁵. Note that in order to compare our results to the clairvoyant

⁵Results for the third instance with three clusters are omitted here since these are very similar to those of the 2-cluster instance.

EMOA approach – either visually or by performance metrics – we transformed the results of all eras to the a-posteriori solution space, which covers the whole potential of arbitrary decision paths.⁶

From this high-level perspective, we can identify an interesting property of solution distribution, which not only depends on the applied decision strategies but is strongly related to the considered topology of the instances: For clustered topologies, tour length tends to expose larger variability in later eras, and – on a first glance – no matter which combination of decisions is used, while in uniform topologies distinct clusters of solutions represent the different weighting of decision strategies.

4.2.2 Topology and final decision making influence

The observed results can be explained by a deeper analysis of the experiments. Here, we focus on two representative topologies consisting of two clusters and a uniformly distributed set of customers, respectively. For both topologies, each instance, and all decision maker strategies, we determine the upper bound of unvisited customers for the last era and compute the mean upper bound as well as standard deviation. More detailed values for mean and standard deviation of the upper bound of both objectives – and split according to the final preference – are available in Table 2. Each upper bound of unvisited customers is determined by the already visited customers on the traveled partial tour, which is the result of decisions made during previous eras.

For clustered topologies (ref. to Figure 2, left) we find a low mean upper bound and large standard deviation, while for the uniform topology the mean upper bound of unvisited customers is rather high with little variance. Consequently on average, for clustered topologies, the decision maker preference at the onset of the last era allows only little flexibility for final solutions. This leads to the stronger focus on the lower right area of objective space. For uniform instances, the on average larger upper bound of unvisited customers leads to a larger and less flexible range for diverse forming of (intermediary) solutions over all eras and finally to more distinct clustering solutions in objective space.

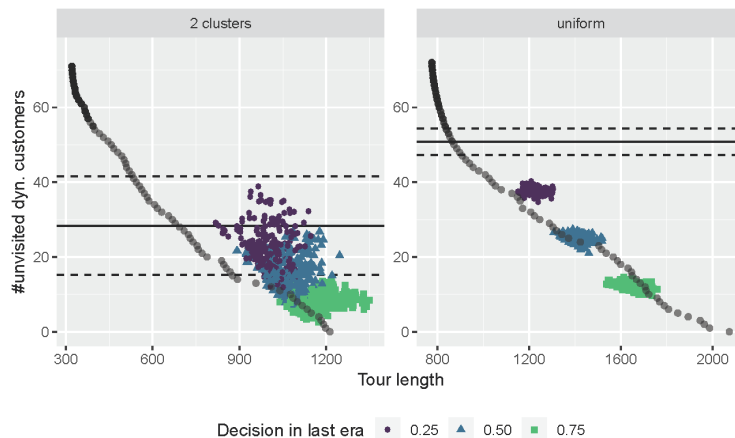


Figure 3: Union of all decisions made in last era colored and shaped by the decision maker preference in the last era. Small black dots represent the Pareto-front approximation of the clairvoyant EMOA.

The coloring of final solutions with respect to the final decision maker preference in Figure 3 provides additional insights into the partitioning of the exhaustively generated solutions. In fact, we find that in both cases the last decision preference has significant influence on the solution position. For the uniform instances, however, the preferences are more distinct due to more certain planning flexibility in the final era.

From Figure 3 and Table 2 we also conclude that solutions for uniform topologies expose less variance in quality and converge closer to the a-posteriori solution fronts determined by application of the clairvoyant EMOA. As a consequence, more solutions of the dynamic approach individually outperform solutions on the a-posteriori front.⁷ This effect is rooted in decreasing complexity of the tour planning component of the bi-objective problem under the successive dynamic decision making [28]. Due to decision making over time eras, partial tours are already completed

⁶We explain this transformation in more in detail here: in the first era, we have zero dynamic requests and consequently zero unvisited dynamic customers. However, in the a-posteriori solution space this solution corresponds to $|C^o|$ unvisited dynamic customers. Therefore, in order to make solutions comparable, a linear transformation of the second objective to the clairvoyant EMOA solution space is required.

⁷The a-posteriori front used here was achieved as non-dominated set of the union of ten clairvoyant EMOA runs.

such that the combinatorial decision space shrinks to the still available customers leaving the tour planning problem with less degrees of freedom. Clearly, this observation holds for clustered instances, too. Here, the mean upper bound is even lower. However, in the clustered setting the service vehicle might need to travel back and forth between clusters in order to fulfill the decision maker preferences which oftentimes might lead to enlarged tour length in particular in late eras and a high preference on the second objective. We will catch up on this important aspect later on.

4.2.3 Intermediate decision making

While so far we analyzed the final decision maker preference for some instances, the following stage of investigation is focused on the influence of intermediary decisions. A subset of decision paths, which led to the non-dominated as well as completely dominated solutions in the last era over all considered topologies is selected. For this selection, we detail the effects of decision steps that yielded very good and very bad results, compared on final solutions. For the following discussion, we investigate the results up to specific solution phenotypes, i. e., the development of specific tours over time. We present detailed results for two exemplary but representative out of 75 instances.

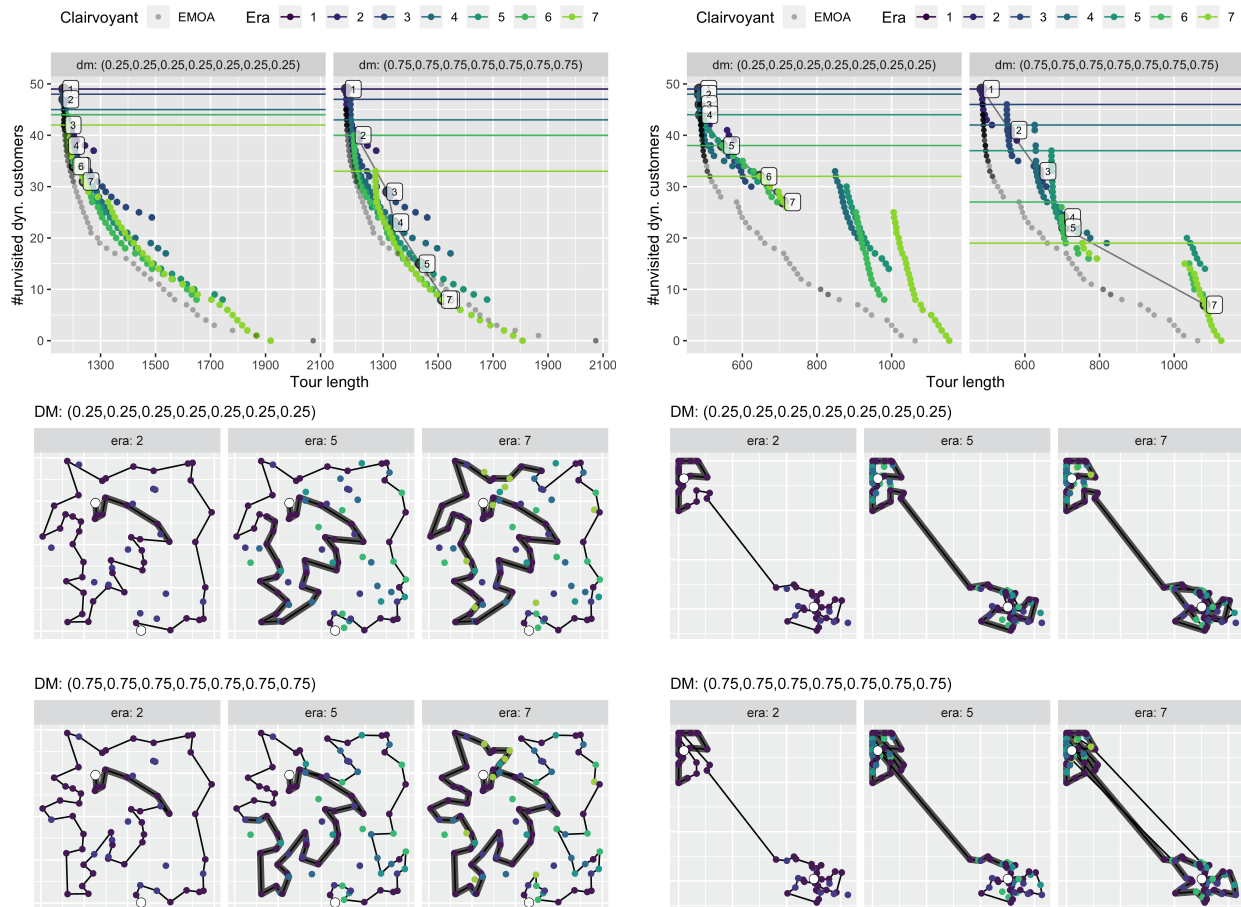


Figure 4: Pareto-front approximations for eras 2, 5, 7 and two selected decision maker strategies (top row plots) for an exemplary uniform instance (left side) and an instance with two clusters (right side) with 50 dynamic customers. Colored horizontal lines in the top plots show the upper bound for unvisited customers in the respective era. Numeric labels $i \in \{1, \dots, 7 = n_t\}$ indicate the decision made by the DM in the corresponding era. In the rows below phenotypes of the decisions are plotted for the respective era aus DM strategies.

Figure 4 provides detailed insight into the development of solutions under different DM strategies. In order to visualize the effect of the permutation of decisions inside a strategy (and also due to space limitations), we restrict ourselves to one uniform topology and one topology with two-clusters again and show results of single representative runs. At the top of the figure we show the non-dominated solutions of all eras of these topologies regarding four strategies that follow (a) only 0.25 preferences, (b) first four times 0.25 and then three times 0.75 preferences, (c) the inverse strategy

to (b), and (d) only 0.75 decisions. Below the non-dominated solutions, we visualize the development of exemplary tours of the solutions. We omit era 1, where the vehicle has not traveled yet and also omit some intermediate tours to show a second example tour. For each tour, the decision path via intermediate solutions is included into the respective top figure and annotated with the era number.

For the clustered topology, we find a strategy-sensitive behavior that is related to when (in which era) preferences are used. The overall observation is that preferences, which do not put a strong focus on minimizing the number of unvisited customers (represented by a sequence of only 0.25 preference) lead to rather short tours (according to the second objective). In these tours, the vehicle transfers to the other cluster only once. Introducing a strong preference for visiting all customers (represented by a sequence of only 0.75 preferences) forces the vehicle to transfer multiple times between clusters, see Figure 4 bottom right plot. This behavioral changes are also observable for planned tours, when preferences mix, e.g., when a strong preference for visiting many customers is only present at the beginning or the end of the strategy sequence. In many sequences with changing decision preferences (not shown here as figure), we observe that planned transfers in early eras vanish in following eras (due to 0.25 preferences later). From this behavior we conclude, that intermediary preference ordering can have decisive influence on the solution generation process for clustered instances. With respect to minimization of unvisited customers and dynamic appearance of customers, decision maker preferences have different degrees of greediness: a 0.25 preference is far less greedy than a 0.75 strategy and often allows the vehicle during tour planning to remain in the current cluster, as far fewer customers need to be served.

Introducing a more greedy strategy often forces a vehicle transfer to serve the preferred amount of customers. The observations however show, that flexible re-planning is still possible as long as the partial tour has not been realized.

Considering the exemplary but representative results from Figure 4, we can conclude, that strategy preferences are less important compared to the clustered case. As mandatory and dynamic customers are uniformly distributed, planned tours do not need to be changed extensively in order to visit or ignore a customer. When we switch preferences from less greedy to more greedy, customers “on the way” can be included. The same way, dynamic customers can be excluded again from a planned tour, often without significant changes in the overall tour length.

The described effects for clustered and uniform topologies can also be observed in the resulting non-dominated fronts for the eras. For clustered topologies, the front usually exposes a gap, which corresponds to the additional traveled distance in size. It appears, that solutions cannot be realized without transferring the vehicle multiple times. Such case usually does not happen for uniform topologies such that the approximated efficient front does not expose a gap.

4.2.4 Performance measurement

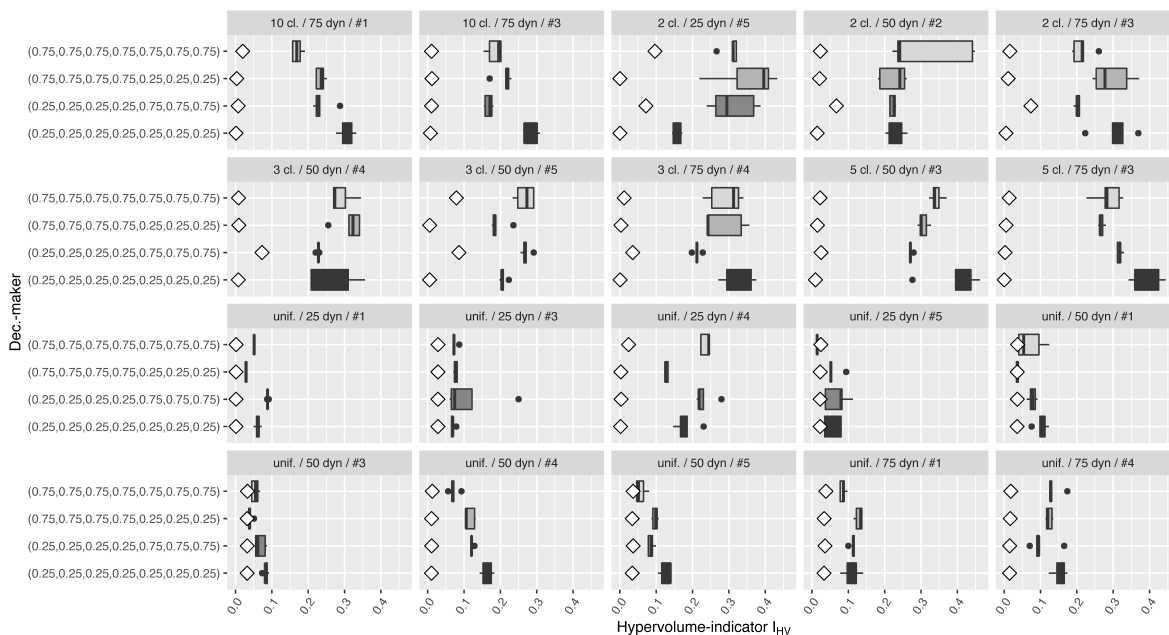


Figure 5: Distribution of hypervolume-indicator I_{HV} (lower is better) for 20 out of 75 randomly sampled instances. The white diamonds indicate the HV-indicator for the union of 10 runs of the clairvoyant EMOA, i. e., the baseline.

In order to support our observations from the previous paragraphs we continue with indicator-based performance assessment of the DEMOA in comparison to the approximation sets calculated by the clairvoyant EMOA. We aim to quantify the quality of the overall final approximation set in the last era. We use the hypervolume indicator $I_{HV}(P, R)$ [32] to measure the space enclosed by a reference set R (non-dominated set of the union of clairvoyant EMOA approximation and all front approximations for the problem instance obtained by the DEMOA) and the DEMOA approximation P . We restrict our analysis to the DEMOA approximation sets of the final era only. We take account for the upper bound that restricts the possible number of unvisited dynamic customers in the last era as follows: only solutions of the clairvoyant EMOA whose second objective is lower or equal to the maximum upper bound in the last era for each instance over all 25 performed runs are taken into consideration. We want to stress that this comparison – and the one in the next paragraph – is obviously highly unfair, i. e., (1) the clairvoyant EMOA has a clear advantage over the dynamic approach due to its a-priori knowledge of request times and (2) the solutions of the clairvoyant might not even be feasible anymore in the last era. Hence, we do not expect the DEMOA to beat the clairvoyant EMOA by any means. Instead, our goal is to learn how close we can approach the clairvoyant solutions with the dynamic approach. Figure 5 shows the distributions of the HV-indicator split by instance and the four DM-strategies discussed before. We show results for a random sample of 10 uniform and 10 clustered instances. The plots confirm our previous observations: in the case of customers distributed uniformly at random in the Euclidean plane the final approximation sets are close to the reference set. In contrast, for clustered topologies the situation is different. Here, as the vehicle possibly needs to transfer between clusters multiple times, the oracle-perspective of the clairvoyant EMOA is much more advantageous and has a much larger impact. In other words, the HV-indicator is less close to the one of the clairvoyant EMOA.

5 Conclusions

For bi-objective vehicle routing, problem dynamics have to be efficiently addressed while suitable decision maker strategies accounting for the trade-off of minimizing overall tour length and maximizing the number of served customers are required simultaneously. We build upon previous work which provides a sophisticated dynamic EMOA hybridized with local search and specifically investigate the influence of respective decision maker preferences and strategies.

As vehicle tours for a given problem instance evolve over the focused time horizon, decision maker preferences regarding both objective functions may change in the course of the day. We assume that the decisions for possibly altering a predefined tour based on new customer requests have to be made at predefined time intervals which of course subsequently impacts optimization algorithm behavior and thus also influences solution selection decisions which have to be made at later stages.

In systematic experiments, we investigated the influence of decision paths, i.e. sequences of (possibly different) decision maker preferences and solution selections. We present a decision support system enhanced by informative figures visualizing the vehicle tour over time and the characteristics of the candidate trade-off solutions at the points of required decisions.

We confirm the reasonable suspicion that decision making is sensitive to the underlying problem topology. For clustered topologies, intermediate decisions should be considered carefully, as too greedy approaches can lead to multiple vehicle transfers between clusters and massively deteriorate solution quality. For uniform instances, sensitivity is low and the last decision for the optimal trade-off solution is of major importance for final tour quality. Consequently, it is important for the decision maker to estimate the customer location topology for adjusting the greediness of decision making.

Future work directions are manifold with the most promising being listed below:

- The problem can be extended to a more realistic scenario which includes multiple vehicles or vehicle loading and unloading during service. Thereby, properties and challenges of the traveling thief problem [33, 34] are included into consideration.
- We see much room for algorithmic improvements. The insights gained in this paper suggest that biased mutation (e. g., activating customers in the current cluster with higher probability) may have beneficial effects on the solution quality of the DEMOA. Furthermore, utilizing probabilistic models to predict upcoming customer requests can be leveraged to achieve more thoughtful algorithmic tour planning.
- Last but not least the major goal is to refine the presented approach in terms of providing tool-support for informative interactive decision making in this highly dynamic environment. Moreover, we will include predefined agent-based decision maker paths into the algorithm which adapt to problem topology characteristics via automatically extracting problem features and which can be adjusted if needed.

References

- [1] Kannan Govindana, Mohammad Fattahib, and Esmaei Keyvanshokoohc. Supply chain network design under uncertainty: A comprehensive review and future research directions. *European Journal of Operational Research*, 263(1):108–141, 2017.
- [2] J. Mula, R. Poler, J.P. garcia Sabater, and F.C. Lario. Models for production planning under uncertainty: A review. *International Journal of Production Economics*, 103(1):271–285, 2006.
- [3] M. L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, New York, 4th edition, 2012.
- [4] T. Cheong and C. C. White. Dynamic traveling salesman problem: Value of real-time traffic information. *IEEE Transactions on Intelligent Transportation Systems*, 13(2):619–630, June 2012.
- [5] Xue-Song Yan, Li-Shan Kang, Zhi-Hui Cai, and Hui Li. An approach to dynamic traveling salesman problem. In *Proceedings of 2004 International Conference on Machine Learning and Cybernetics (IEEE Cat. No.04EX826)*, volume 4, pages 2418–2420 vol.4, Aug 2004.
- [6] David Meignan, Sigrid Knust, Jean-Marc Frayret, Gilles Pesant, and Nicolas Gaud. A review and taxonomy of interactive optimization methods in operations research. *ACM Transactions on Interactive Intelligent Systems*, 5(3), 2015.
- [7] Cristóbal Barba-González, Vesa Ojalehto, José García-Nieto, Antonio J. Nebro, Kaisa Miettinen, and José Francisco Aldana Montes. Artificial decision maker driven by PSO: an approach for testing reference point based interactive methods. In *Parallel Problem Solving from Nature - PPSN XV - 15th International Conference, Proceedings Part I*, pages 274–285, 2018.
- [8] Jürgen Branke, Kalyanmoy Deb, Kaisa Miettinen, and Roman Slowinski, editors. *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Springer, 2008.
- [9] Mansoureh Aghabeig and Andrzej Jaszkiwicz. Experimental analysis of design elements of scalarizing functions-based multiobjective evolutionary algorithms. *Soft Computing*, 2018.
- [10] Carlos A Coello Coello, Gary B Lamont, and David A Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic and Evolutionary Computation)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [11] C. Raquel and X. Yao. Dynamic multi-objective optimization: A survey of the state-of-the-art. In S. Yang and X. Yao, editors, *Evolutionary Computation for Dynamic Optimization Problems*, pages 85–106. Springer, Berlin Heidelberg, 2013.
- [12] Dominique Feillet, Pierre Dejax, and Michel Gendreau. Traveling salesman problems with profits. *Transportation Science*, 39(2):188–205, 2005.
- [13] J.-F. Berube, M. Gendreau, and J.-Y. Potvin. An exact [epsilon]-constraint method for bi-objective combinatorial optimization problems: Application to the traveling salesman problem with profits. *European Journal of Operational Research*, 194(1):39–50, 2009.
- [14] C. Filippi and E. Stevanato. Approximation schemes for bi-objective combinatorial optimization and their application to the TSP with profits. *Computers & Operations Research*, 40(10):2418–2428, 2013.
- [15] N. Jozefowicz, F. Glover, and M. Laguna. Multi-objective meta-heuristics for the traveling salesman problem with profits. *Journal of Mathematical Modelling and Algorithms*, 7(2):177–195, 2008.
- [16] V. Pillac, M. Gendreau, C. Guéret, and A. L. Medaglia. A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1):1–11, 2013.
- [17] S. Meisel. *Anticipatory Optimization for Dynamic Decision Making*, volume 51 of *Operations Research/Computer Science Interfaces Series*. Springer New York, 2011.
- [18] Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuyse. The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313, 2016.
- [19] Min Wen, Jean-François Cordeau, Gilbert Laporte, and Jesper Larsen. The dynamic multi-period vehicle routing problem. *Comput. Oper. Res.*, 37(9):1615–1623, 2010.
- [20] Sandro Lorini, Jean-Yves Potvin, and Nicolas Zufferey. Online vehicle routing and scheduling with dynamic travel times. *Comput. Oper. Res.*, 38(7):1086–1090, 2011.
- [21] Mostepha R. Khouadjia, Briseida Sarasola, Enrique Alba, Laetitia Jourdan, and El-Ghazali Talbi. A comparative study between dynamic adapted pso and vns for the vehicle routing problem with dynamic requests. *Appl. Soft Comput.*, 12(4):1426–1439, 2012.

- [22] Lianxi Hong. An improved lns algorithm for real-time vehicle routing problem with time windows. *Comput. Oper. Res.*, 39(2):151–163, 2012.
- [23] Mohamed Barkaoui and Michel Gendreau. An adaptive evolutionary approach for real-time vehicle routing and dispatching. *Comput. Oper. Res.*, 40(7):1766–1776, 2013.
- [24] O. E. Nahum and Y. Hadas. A framework for solving real-time multi-objective VRP. In J. Zak, Y. Hadas, and R. Rossi, editors, *Advanced Concepts, Methodologies and Technologies for Transportation and Logistics*, volume 572 of *Advances in Intelligent Systems and Computing*, pages 103–120. Springer, 2018.
- [25] S. F. Ghannadpour, S. Noori, and R. Tavakkoli-Moghaddam. A multi-objective vehicle routing and scheduling problem with uncertainty in customers’ request and priority. *Journal of Combinatorial Optimization*, 28:414–446, 2014.
- [26] S. Meisel, C. Grimme, J. Bossek, M. Wölck, G. Rudolph, and H. Trautmann. Evaluation of a Multi-Objective EA on Benchmark Instances for Dynamic Routing of a Vehicle. In *Proc. of the Genetic and Evolutionary Computation Conference, GECCO ’15*, pages 425–432, New York, NY, USA, 2015. ACM.
- [27] J. Bossek, C. Grimme, S. Meisel, G. Rudolph, and H. Trautmann. Local Search Effects in Bi-objective Orienteering. In *Proc. of the Genetic and Evolutionary Computation Conference, GECCO ’18*, pages 585–592, New York, NY, USA, 2018. ACM.
- [28] J. Bossek, C. Grimme, S. Meisel, G. Rudolph, and H. Trautmann. Bi-Objective Orienteering: Towards a Dynamic Multi-Objective Evolutionary Algorithm. In *Evolutionary Multi-Criterion Optimization*, volume 11411 of *Lecture Notes in Computer Science (LNCS)*, pages 1–12. Springer, 2019.
- [29] Y. Nagata and S. Kobayashi. A powerful genetic algorithm using edge assembly crossover for the traveling salesman problem. *INFORMS Journal on Computing*, 25(2):346–363, 2013.
- [30] Roy Jonker and Ton Volgenant. Transforming asymmetric into symmetric traveling salesman problems. *Oper. Res. Lett.*, 2(4):161–163, November 1983.
- [31] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [32] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evol. Comput.*, 8(2):173–195, 2000.
- [33] M.R. Bonyadi, Z. Michalewicz, and L. Barone. The travelling thief problem: the first step in the transition from theoretical problems to realistic problems. In *IEEE Congress on Evolutionary Computation*, pages 1037–1044. IEEE, 2013.
- [34] Julian Blank, Kalyanmoy Deb, and Sanaz Mostaghim. Solving the bi-objective traveling thief problem with multi-objective evolutionary algorithms. In *9th International Conference on Evolutionary Multi-Criterion Optimization - Volume 10173*, EMO 2017, pages 46–60, Berlin, Heidelberg, 2017. Springer-Verlag.