

1 A regression analysis of the impact of  
2 routing and packing dependencies on the  
3 expected runtime

4 Mohamed El Yafrani<sup>1,4</sup>, Marcella Scoczynski<sup>2\*</sup>, Markus  
5 Wagner<sup>3</sup> and Peter Nielsen<sup>1</sup>

6 <sup>1</sup>Materials and Production, Aalborg University, Aalborg,  
7 Denmark.

8 <sup>2</sup>Department, Federal University of Technology - Parana, Ponta  
9 Grossa, Brazil.

10 <sup>3</sup>School of Computer Science, The University of Adelaide,  
11 Adelaide, Australia.

12 <sup>4</sup>Quantitative Assets, Norlys Energy Trading A/S, Aalborg,  
13 Denmark.

14 \*Corresponding author(s). E-mail(s): [marcella@utfpr.edu.br](mailto:marcella@utfpr.edu.br);  
15 Contributing authors: [mey@mp.aau.dk](mailto:mey@mp.aau.dk), [mey@net.com](mailto:mey@net.com);  
16 [markus.wagner@adelaide.edu.au](mailto:markus.wagner@adelaide.edu.au); [peter@mp.aau.dk](mailto:peter@mp.aau.dk);

17 **Abstract**

18 Problems with multiple interdependent components offer a better rep-  
19 resentation of the real-world situations where globally optimal solutions  
20 are preferred over optimal solutions for the individual components. One  
21 such model is the Travelling Thief Problem (TTP); while it may offer  
22 a better benchmarking alternative to the standard models, only one  
23 form of inter-component dependency is investigated. The goal of this  
24 paper is to study the impact of different models of dependency on  
25 the fitness landscape using performance prediction models (regression  
26 analysis). To conduct the analysis, we consider a generalised model of  
27 the TTP, where the dependencies between the two components of the  
28 problem are tunable through problem features. We use regression trees  
29 to predict the instance difficulty using an efficient Memetic Algorithm  
30 that is agnostic to the domain knowledge to avoid any bias. We re-  
31 port all the decision trees resulting from the regression model, which

2 *Impact of dependency features on the expected runtime*

is the core in understanding the relationship between the dependencies (represented by the features) and problem difficulty (represented by the runtime). The regression model was able to predict the expected runtime of the algorithm based on the problem features. Furthermore, the results show that the contribution of the item value drop dependency is significantly higher than the velocity change dependency.

**Keywords:** Runtime regression analysis, Evolutionary Algorithms, Travelling Thief Problem

40 **1 Introduction**

Tackling real-world problems can be very challenging compared to standard benchmarking integer programming models [17, 33, 34, 37]. Multiple aspects contribute to creating this gap, such as bad modelling practices (e.g., oversimplification, linearisation), intractability (high computational complexity of solution methods), hard and soft constraints (heavy constraints...), external factors (stochastic environment, uncertainty about data), and composition of interdependent (mutually dependent) sub-problems [4, 28].

In this paper, we are interested in the last aspect, namely, dependencies between the problem's components (sub-problems). This type of problems are referred to as problems with multiple interdependent components [2] or multi-hard problems [32], as the components are NP-hard when tackled separately. To the best of our knowledge, the first benchmark model covering this aspect is the Travelling Thief Problem (TTP), introduced by Bonyadi et al [2] as a combination of the Travelling Salesman Problem and the Knapsack Problem. The model was then simplified by Polyakovskiy et al [31], and a large dataset of instances was published. Furthermore, an extension has been proposed in Chand and Wagner [5] where multiple thieves are considered.

The standard TTP formulation in [31] considers a combination of the Travelling Salesman Problem (TSP) and the Knapsack Problem (KP). The problem considers a set of items scattered in different cities where a thief should visit each city once, picking some items on the way and returning to the first city, while trying to maximise the thief's gain. The dependency in this formulation is modelled in a number of manners, such as:

- Penalising the travelling time by tying the thief's velocity with the knapsack load (standard model [31]),
- Decreasing the value of items as the thief progresses in his journey,
- A combination of the above (bi-objective TTP model [2]).

Note that because the final aim is to cover some aspects from real-world problems, both dependency types were designed to reflect realistic situations. First, the load-velocity dependency can be adapted to reflect the relationship between the load and fuel consumption in the transportation of goods [16].

72 Second, the value drop dependency can have even more realistic and impor-  
73 tant applications within the shipping sector, such as the transportation of  
74 perishable goods [7, 14].

75 Since it was introduced, the TTP received the attention of researchers  
76 from the fields of evolutionary computation, metaheuristics and operations  
77 research, mainly due to the fact that the problem is easy to understand,  
78 yet challenging to solve. Several papers proposed heuristic solution meth-  
79 ods [9, 10, 26, 38], and fewer ones tried to analyse the problem empirically  
80 and theoretically [11, 42, 43]. These analyses focus on the impact of problem  
81 features, with little to no attention to the impact of the dependencies between  
82 the components. Furthermore, all of the above-mentioned works only consider  
83 the velocity change constraint, which is probably the result of the standard  
84 TTP library only supporting this dependency.

85 The gap we are trying to close with this work is the analysis of the impact  
86 of the above-mentioned dependency formulations and investigation of their  
87 impact on the difficulty of tackling the problem. To conduct the analysis,  
88 we consider cost models as a tool to empirically analyse the difficulty of the  
89 generalised TTP model, that embeds all these dependency models. The novelty  
90 here considers imposing that the value of a picked item drops by time, besides  
91 the evaluation of the difficulty of problem instances is done using a Memetic  
92 Algorithm based on MA2B [9]. The findings show that the item value drop  
93 dependency significantly impacts the difficulty of solving the TTP instances.  
94 More importantly, according to our analysis, its impact is stronger than the  
95 velocity change constraint.

96 In Section 2, we provide a background information with a brief literature  
97 review on the algorithms and analyses for the TTP and cost model-based  
98 fitness landscape analysis. Section 3 introduces the methodology adopted to  
99 analyse the problem, including the feature-based analysis and the adopted  
100 algorithm. In Section 4 describes the experiments and discusses the results.  
101 Finally, Section 5 summarises the findings and concludes the paper.

## 102 2 Background and Related Works

### 103 2.1 The Travelling Thief Problem

104 The standard TTP can be informally stated as follows: *Given are a set of cities*  
105 *and a set of items distributed among these cities. Each item is defined by its*  
106 *individual profit and weight. A thief must visit all the cities exactly once, pick*  
107 *some items while travelling, and return to the starting city. The knapsack has*  
108 *a limited capacity, which should not be exceeded. We also consider a knapsack*  
109 *renting rate (per time unit) which determines the amount that the thief must*  
110 *pay at the end of the journey.*

111 What makes the two components of the TTP interdependent is the veloc-  
112 ity of the thief, because it changes according to the weight of the knapsack.  
113 Specifically, the heavier the knapsack gets, the slower the thief becomes. The

4 *Impact of dependency features on the expected runtime*

114 objective is to maximise the total gain function defined as the total profit,  
115 minus the cost of the journey.

116 The interdependence, such as the one present in the TTP, can reflect the  
117 characteristics and the complexity of real-world problems [2]. For this reason,  
118 several authors have addressed the TTP by applying different methods.  
119 Polyakovskiy et al [31] presented the first heuristics for solving the TTP, generating  
120 a TSP tour using the classical Chained Lin–Kernighan heuristic [1]  
121 and with the fixed tour they applied some packing heuristics for improving the  
122 solution, such as Random Local Search (RLS) and (1+1)-EA.

123 As the problem is NP-hard and the objective function is non-linear, many  
124 researchers focused on (meta-)heuristic algorithms. The work proposed by  
125 Bonyadi et al [3] introduces a method named CoSolver, which is inspired by  
126 cooperative coevolution. The framework consists in splitting the problem into  
127 sub-problems and tackling them in a parallel and synchronous fashion. Mei  
128 et al [26] proposed a fast Memetic Algorithm called MATLS, which embeds  
129 multiple complexity reduction methods to solve large scale TTP instances.  
130 Faulkner et al [13] explored multiple operators for optimising the packing plan  
131 combining them in a number of simple and complex heuristics. The work in  
132 [9] proposed and compared two heuristic algorithms, a Memetic Algorithm  
133 (MA2B) and Simulated Annealing-based algorithm (CS2SA) which resulted  
134 in competitive performances for various problem sizes. Wagner [38] investigated  
135 the use of swarm intelligence approaches with two different TSP-specific  
136 local search operators and of “boosting” TTP solutions using TTP-specific  
137 local search. Two algorithms were proposed in [10] based on combining the  
138 2-OPT steepest ascent hill climbing algorithm for the TSP component and  
139 the simulated annealing metaheuristic for the KP component, named CS2SA\*  
140 and CS2SA-R. The obtained results showed that the proposed algorithms are  
141 competitive in many TTP instances.

142 In [39] the authors created a dataset with performance data of 21 TTP  
143 algorithms on the full original set of 9720 TTP instances. They also defined  
144 55 characteristics for TTP instances that can be used to select the best algorithm  
145 on a per-instance basis, and they used these algorithms and features to  
146 construct algorithm portfolios for TTP in order to outperform the single best  
147 algorithm.

148 Recently, the authors in [29] investigated the inter-dependency of the TSP  
149 and the KP by means of quality diversity (QD) approaches, conducting experimental  
150 studies that show the usefulness of using the QD approach applied  
151 to the TTP. Besides, the authors introduced a MAP-Elite based evolutionary  
152 algorithm called BMBEA, using well-known TSP and KP search operators.  
153 In the MAP-Elite solutions compete with each other to survive. Their results  
154 showed that QD approach can improve the best-known solution for a wide  
155 range of TTP instances.

## 156 2.2 Empirical algorithm analysis

157 Fitness landscapes represents the association between the search process and  
158 the fitness space [41]. A heuristic algorithm can be seen as a strategy for  
159 navigating the solution landscape structure in the search for an optimal so-  
160 lution. Thus, fitness landscape analysis is a set of tools and methods used  
161 to investigate the dynamics of heuristic search algorithms applied for specific  
162 optimisation problems [22].

163 Cost models are fitness landscapes methods that can help predicting the  
164 performance of algorithms by identifying features that make a problem more  
165 or less difficult to solve. These models can be expressed as linear, multiple re-  
166 gression models [25], decision trees [30], or other models of features and search  
167 cost; also, some models are more amendable to human interpretation than  
168 others. To aid interpretability, the features are extracted from the problem  
169 structure and the model can at times explain their influence in the difficulty  
170 level during the search [11].

171 Some authors have presented fitness landscape analysis for several prob-  
172 lems, as described in [44]. The work developed by [27] studied the relation  
173 between features of fitness landscapes and recombination and/or mutation  
174 operators for the Quadratic assignment problems. The authors in [35, 36] anal-  
175 ysed the fitness landscape of a dynamic optimisation problem, investigating  
176 the influence on the performance of the algorithm.

177 The work proposed by [18] designed a prediction model for the algorithmic  
178 performance of CMA-ES variants by using a random forest regression model  
179 based on exploratory features of fitness landscapes. Also, the authors in [24]  
180 presented a spatial-domain fitness landscape analysis framework to visualise  
181 the fitness landscapes regarding a specific combinatorial optimisation prob-  
182 lem and evaluate its properties. By extracting characteristics of combinatorial  
183 optimisation problems allowed study the behaviour of algorithms effectively.

184 Recently, the paper in [21] focused on the adaptability landscape features  
185 of optimisation problems by applying differential evolution algorithm. The  
186 authors presented a quantitative analysis of the fitness distance correlation  
187 information, evaluating the difficulty of solving the problem.

188 Some papers explored TTP using fitness landscape analysis. In [43] the  
189 authors considered local search operators and investigated the fitness landscape  
190 characteristics of some smaller instances of the TTP. The local search operators  
191 included 2-opt, Insertion, Bitflip and Exchange and metaheuristics included  
192 multi-start local search, iterated local search and genetic local search.

193 Another recent work in [12] investigated 3 dependency models of the TTP,  
194 in addition to a dependency-free model for comparison purposes. The authors  
195 used local optima networks as a fitness landscape analysis tool to study the  
196 difficulty of the search landscape for these dependency models. The results  
197 show that the dependency-free landscape is the most difficult to navigate based  
198 on the basin sizes and their correlation to the fitnesses. Such a result is diffi-  
199 cult to interpret as it is expected that the dependencies create more difficult

instances [4]. The authors speculated that this result is due to the join neighbourhood search algorithm, an algorithm that joins two neighbourhood sets into one, which creates an additional complexity even for the dependency-free model.

## 3 Proposed approach

### 3.1 Generalised TTP model

Herein, we provide the mathematical formulation of the generalised TTP, based on the standard TTP as introduced in Section 2, and the TTP2 formulation in [2]. The model embeds two types of dependency, and is conceived such that the strength of these dependencies is tunable through problem features.

#### *Input data*

We define the following problem input parameters:

- $N = \{1, \dots, n\}$  is the set of labels, representing the cities to be visited,
- $\{d_{ij}\}$  is the matrix of distances between these cities,
- $M = \{1, \dots, m\}$  is a set of labels corresponding to the items scattered among cities,
- $p_k$  and  $w_k$  represent the profit and weight of an item  $k \in M$ , respectively,
- $W$  is the knapsack capacity,
- $R$  is the knapsack renting rate, which is used to determine the cost of the journey,
- $v_{max}$  and  $v_{min}$  represent the maximum and minimum velocities, respectively.
- $A = \{A_1, \dots, A_m\}$  denotes the availability vector, such that  $A_k \in \{1, \dots, n\}$  contains the reference to the city that contains the item  $k$ .

#### *Decision variables*

A TTP solution is represented using two components: The first is the tour  $X = (x_1, \dots, x_n)$ , a vector containing the ordered list of cities, which is encoded as a permutation. The second is the picking plan  $Z = (z_1, \dots, z_m)$ , a binary vector representing the states of items, where 1 is associated to the packed items, and 0 to the unpacked ones.

#### *Interdependencies*

What makes the two components of the TTP interdependent is the velocity of the thief which changes according to the weight of the knapsack. Therefore, the velocity at city  $x_i$  is defined in Equation 1.

$$v_{x_i} = v_{max} - \left( \frac{v_{max} - v_{min}}{W} \right) \times w_{x_i} \quad (1)$$

where  $w_{x_i}$  the weight of the knapsack at city  $x_i$ .

In addition to the velocity change dependency in Equation 1, we add a second dependency on the value of items. This is achieved by imposing that

236 the value of a picked item  $k$  drops by time from its initial value  $p_k$  to  $p_k^{final}$   
 237 following Equation 2.

$$p_k^{final} = p_k \times \mathcal{D}_r^{\lceil \frac{T_k}{Q} \rceil} \quad (2)$$

238 where  $\mathcal{D}_r$  is the item value dropping rate,  $p_k$  is the initial value of item  $k$ ,  
 239  $T_k$  is the carrying time of item  $k$ , and  $Q$  is a constant calculated as shown in  
 240 Equation 3.

$$Q = T_{min} \frac{\ln(\mathcal{D}_r)}{\ln\left(\frac{P_{min}}{2 \times P_{max}}\right)} \quad (3)$$

241 Equation 3 seeks to satisfy  $P_{max} \times \mathcal{D}_r^{\frac{T_k}{Q}} = \frac{1}{2} \times P_{min}$ , where  $P_{max}$  and  $P_{min}$   
 242 represent the maximum and minimum item values, respectively, while  $T_{min}$  is  
 243 the minimum travelling time.

244 Furthermore, the dependencies in the Generalised TTP model are tunable  
 245 through the parameters  $v_{min} \in [0, v_{max}]$  and  $\mathcal{D}_r \in [0, 1]$ , where:


- 246 • The velocity drop dependency is controlled through  $v_{min}$ . When  $v_{min} =$   
 247  $v_{max}$ , this dependency is cancelled as the thief will be always travelling at  
 248 the maximum velocity  $v_{max}$  independently from the knapsack load.
- 249 • The item value drop dependency is controlled through  $\mathcal{D}_r$ . Similarly, this  
 250 dependency is deactivated by setting  $\mathcal{D}_r = 1$ .

251 Based on the above, the proposed model covers all dependency models  
 252 in [12], including the standard TTP.

253 There are two special cases to consider: (1) Setting  $v_{min} = 0$  can lead to a  
 254 velocity of 0. Based on Equation 1, we have  $v_x = v_{max}(1 - \frac{w_x}{W})$ , which leads  
 255 to the thief remaining stationary if the knapsack load reaches the maximum  
 256 capacity  $W$ . (2) Setting  $\mathcal{D}_r = 0$  makes all the picked items valueless in the end  
 257 as  $p_k^{final} = 0$ . Based on these remarks, the considered values for  $v_{min}$  and  $\mathcal{D}_r$   
 258 are always chosen strictly positive in our empirical study.

259 Note that it is difficult to predict how a change in dependency feature  
 260 values, e.g., lower values of  $\mathcal{D}_r$  and  $v_{min}$ , will impact the difficulty of the  
 261 problem instances. This partially depends on the algorithm used to tackle the  
 262 problem.

263 To illustrate how these interdependencies work, we consider a TTP instance  
 264 with 6 cities and 20 items, and a potential solution as shown in Figure 1. Note  
 265 that this is a simplified illustration as the distances, item values and weights  
 266 among other problem parameters are not shown for simplicity. Following the  
 267 velocity change dependency, assuming that  $v_{min} < v_{max}$ , the velocity will  
 268 start decreasing from city  $B$  where items 1 and 4 are picked, passing through  
 269 city  $C$  does not influence the velocity as no items are picked, and so on. Now  
 270 considering the item value drop dependency, the items 1 and 4 lose most of  
 271 their value as they are carried for almost the entire journey, while item 17 loses  
 272 the least of its initial value.



**Figure 1:** A simplified illustration of a TTP instance and solution. The circles represent the cities, labelled by letters from *A* to *F*. The rectangles represent the items associated to each city (except the first), labelled by numbers from 1 to 20. The dashed arrows form the route, and the highlighted rectangles represent the picked items.

### 273 *Objective function*

274 To focus on the dependency analysis, we consider a linear combination of the  
 275 total profit and cost of the journey as shown in Equation 4.

$$\text{Maximise } G(X, Z) = \sum_m p_m^{final} \times z_m - R \times \left( \sum_{i=1}^{n-1} \frac{d_{x_i, x_{i+1}}}{v_{x_i}} + \frac{d_{x_n, x_1}}{v_{x_n}} \right) \quad (4)$$

Subject to:

$$\sum_m w_m \times z_m \leq W \quad (5)$$

$$x_i \in \mathbb{Z}^+ \quad \forall i \in \{1, \dots, n\} \quad (6)$$

$$z_k \in \{0, 1\} \quad \forall k \in \{1, \dots, m\} \quad (7)$$

276 where  $t_{x_i, x_{i+1}} = \frac{d_{x_i, x_{i+1}}}{v_{x_i}}$  is the travel time from  $x_i$  to  $x_{i+1}$ .

277 As mentioned earlier, the introduction of the velocity change and item  
 278 value drop dependencies leads to a non-linear objective function. Note that  
 279 the non-linearity exists in the first term of the equation as well, since  $p_k^{final}$   
 280 has a non-linear formulation (Eq. 2) and depends on  $z$ . This increases the  
 281 difficulty of the problem as it cannot be solved using standard MILP methods,  
 282 and rather requires an ad-hod solution.



## 283 3.2 Memetic Algorithm

284 As a baseline for our analysis, we consider the memetic algorithm presented  
 285 in Algorithm 1<sup>1</sup>. The algorithm uses genetic operators (tournament selection  
 286 and crossover) combined with a hill climbing local search algorithm to evolve a  
 287 population of candidate solutions iteratively. The implementation is based on  
 288 the Memetic Algorithm MA2B initially designed for the standard TTP [9], but  
 289 differs from the original implementation by removing all domain knowledge  
 290 from the logic of the algorithm.

---

### Algorithm 1 Memetic Algorithm for the Generalised TTP

---

```

1:  $P \leftarrow \emptyset$ 
2:  $T \leftarrow 0$ 
3: for  $i \in \{0, \dots, N_{pop}\}$  do
4:    $S_i \leftarrow \{rand\_perm(n), rand\_bin(m)\}$  ▷ Random initialisation
5:    $G(S_i)$  ▷ Evaluate
6:    $T \leftarrow T + 1$ 
7:    $\{S_i, T\} \leftarrow hill\_climbing(S_i, T)$  ▷ Local search improvement
8:    $P \leftarrow P \cup S_i$ 
9: end for
10:  $P \leftarrow sort(P)$ 
11:  $S_{bsf} \leftarrow \{rand\_perm(n), rand\_bin(m)\}$ 
12: repeat
13:   if  $G(P_1) > G(S_{bsf})$  then
14:      $S_{bsf} \leftarrow P_1$ 
15:   end if
16:    $Q \leftarrow \emptyset$ 
17:   for  $i \in \{0, \dots, N_{offspring}\}$  do
18:      $\{p_1, p_2\} \leftarrow tournament(P, N_{tournament})$  ▷ Select parent solutions
19:      $S_{new} \leftarrow crossover(p_1, p_2)$  ▷ Crossover
20:      $G(S_{new})$ 
21:      $T \leftarrow T + 1$ 
22:     if  $random(0, 1) < R_{LS}$  then
23:        $\{S_{new}, T\} \leftarrow hill\_climbing(S_{new}, T)$  ▷ Local search improvement
24:     end if
25:     if  $S_{new} \notin Q$  then ▷ add to offspring population
26:        $Q \leftarrow Q \cup \{S_{new}\}$ 
27:     else ▷ Generate new solution to reduce premature convergence
28:        $S_{rand} \leftarrow \{rand\_perm(n), rand\_bin(m)\}$ 
29:        $\{S_{rand}, T\} \leftarrow hill\_climbing(S_{rand}, T)$ 
30:        $Q \leftarrow Q \cup \{S_{rand}\}$  ▷ add to offspring population
31:     end if
32:   end for
33:    $P' \leftarrow sort(P \cup Q)$ 
34:    $P \leftarrow \{P'_1, \dots, P'_{N_{pop}}\}$ 
35: until  $T \geq T_{max} \vee \frac{g^* - G(P_1)}{g^*} \leq \epsilon$ 

```

---

<sup>1</sup>The implementation of the memetic algorithm is done in Java based on the codes available at <https://github.com/yafrani/ttplab>

291 The motivation behind choosing a Memetic Algorithm is due to the fact  
 292 that it borrows aspects from classical local search heuristics and evolution-  
 293 ary algorithms (exploratory operators) [40]. Most of the other algorithms are  
 294 either based on (stochastic or deterministic) local search, or include domain  
 295 knowledge from TTP standard model, which is practically unusable for the  
 296 generalised formulation we proposed (due to the second dependency - item  
 297 value drop). These aspects combined result in an algorithm that can both ex-  
 298 plore and exploit the solution space. Furthermore, the other TTP heuristics  
 299 use problem knowledge as the main component of their logic [13, 26], making  
 300 them unsuitable for this analysis.

**Table 1:** Notations used in Algorithm 1. The remaining notations are aligned with the problem formulation in Sub-section 3.1

Notation	Description
$P$	Main population of solutions
$T$	Number of evaluations
$S_i$	An initial solution
$S_{bsf}$	Best-so-far solution
$R_{LS}$	Local search probability
$Q$	Offspring population
$p_1$ and $p_2$	Solutions selected to generate the offspring
$S_{new}$	Solutions generate using genetic operators and local search
$S_{rand}$	A randomly generated solution to avoid premature convergence
$g^*$	The optimal objective value

301 Note that, to simplify the pseudocode notations, we use the same notation  
 302 to evaluate a solution and to access its objective value. Therefore, only the  
 303 first call of evaluation function  $G(\cdot)$  is considered to increase the counter of the  
 304 number of evaluations  $T$ , which will be used to calculate the expected runtime  
 305 ( $ert$ ). Furthermore, the algorithm does not take into account the dependencies  
 306 as part of the domain knowledge used for the optimisation. This is done on  
 307 purpose to avoid giving the algorithm an unfair advantage for some specific  
 308 instance categories, and to obtain insights which can be used to improve the  
 309 efficiency the the algorithm.

310 Each solution in the population is initialised with a random permutation  
 311 for the tour and random binary vector for the picking plan (line 4. As we are  
 312 interested in the features making the problem difficult, random initialisation is  
 313 important to for the same reason stated above, i.e., ensuring a fair algorithm  
 314 and not favouring any specific instance categories.

315 A best improvement hill climbing procedure is then used to refine the  
 316 solutions by finding a local optima. The hill climbing algorithm uses two  
 317 neighbourhood searches for the tour and the picking plan sequentially. A sim-  
 318 plified pseudocode of the hill climbing procedure is shown in Algorithm 2. The  
 319  $\mathcal{N}_{2-OPT}(S)$  neighbourhood function returns a set of solutions where a new tour

320 is obtained by applying the 2-OPT operator [6] on the tour of  $S$ , while the pick-  
 321 ing plan is copied from  $S$  as is. The same logic is followed for the  $\mathcal{N}_{\text{bit-flip}}(S)$   
 322 function [8] where only the picking plan is updated. More details about these  
 323 operators can be found in [8].

---

**Algorithm 2** Best improvement hill climbing algorithm
 

---

```

1: function hill_climbing( $S, T$ )
2:   repeat                                     ▷ First neighbourhood search on the tour
3:     for  $S^* \in \mathcal{N}_{2\text{-OPT}}(S)$  do
4:       if  $G(S^*) > G(S)$  then
5:          $S \leftarrow S^*$ 
6:       end if
7:        $T \leftarrow T + 1$ 
8:     end for
9:   until  $S$  is not improved
10:  repeat                                       ▷ Second neighbourhood search on the picking plan
11:    for  $S^{**} \in \mathcal{N}_{\text{bit-flip}}(S)$  do
12:      if  $G(S^{**}) > G(S)$  then
13:         $S \leftarrow S^{**}$ 
14:      end if
15:       $T \leftarrow T + 1$ 
16:    end for
17:  until  $S$  is not improved
18:  return  $\{S, T\}$ 
19: end function

```

---

324 The population is then sorted and the best solution is identified (lines 10-  
 325 15). The tournament selection is applied to select 2 parent solutions, with a  
 326 tournament size of  $N_{\text{tournament}}$  to produce  $N_{\text{offspring}}$  (lines 18). This is fol-  
 327 lowed by the Maximal Preservative Crossover (MPX) operator, which returns  
 328 a new solution by combining the parents (line 19). The hill climbing func-  
 329 tion (Algorithm 2) is then applied to the new solution with a probability  $R_{LS}$   
 330 (lines 22-24). If the new solution does not already exists in the offspring pop-  
 331 ulation  $Q$ , it is included. Otherwise, a new solution is generated randomly  
 332 (lines 25-31).

333 Once the offspring is generated, it is combined with the population and the  
 334 best solutions are kept for further improvement (lines 33-34). The algorithm  
 335 stops when the maximum number of evaluations  $T_{\text{max}}$  is reached, or a near-  
 336 optimal solution (solution with a gap to optimal smaller than  $\epsilon$ ) is found  
 337 (line 35).

338 The parameters of Algorithm 1 are summarised in Table 2. The fine tuning  
 339 of the parameters is done taking into account the size of the instances used  
 340 for the analysis, with values chosen empirically based on a random sample  
 341 of instances (diversified based on the features in Section 4, and on previous  
 342 studies [9, 12]. This is ensure that the algorithm find a near-optimal solution  
 343 in most cases, which is important to conduct the analysis.

**Table 2:** Memetic Algorithm parameters

Parameter	Description	Value
$T_{max}$	Maximum number of evaluations	100000
$N_{pop}$	Population size	6
$N_{offspring}$	Offspring size	4
$N_{tournament}$	Tournament size	3
$R_{LS}$	Local search probability	0.1
$\epsilon$	$\epsilon$ -approximation	0.01

### 3.3 Estimation of the Expected Runtime (*ert*)

One way to measure the performance of an algorithm  $\mathcal{A}$  (search cost), is consider the expected number of function evaluations necessary to achieve an  $\epsilon$ -approximation. To achieve this, we save the number of function evaluations until an  $\epsilon$ -approximation is found which characterises a “success”. Otherwise, the search cost is set to a predetermined maximum  $T_{max}$ . This approach is similar to the one presented by Hansen et al [15], Liefoghe et al [23].

In this approach, we denote  $p_s \in [0, 1]$  as the probability of success of algorithm  $\mathcal{A}$ , and  $T_f$  as the random variable measuring the number of function evaluations for unsuccessful runs (failures). After  $(t - 1)$  failures, each one requiring  $T_f$  evaluations, and the final successful run of  $T_s$  evaluations, the total runtime is  $T = \sum_{i=1}^{t-1} T_f + T_s$ , where  $t$  is the random variable representing the number of runs.  $t$  follows a geometric distribution with parameter  $p_s$ . Equation 8 takes the expectation by considering independent runs for each instance, stopping at the first success:

$$\mathbb{E}[T] = (\mathbb{E}[t] - 1)\mathbb{E}[T_f] + \mathbb{E}[T_s] \quad (8)$$

Here, the estimated success rate ( $\hat{p}_s$ ) is computed by the ratio of successful runs over the total number of executions. The expectation of a geometric distribution for  $t$  with parameter  $p_s$  is equal to  $1/p_s$ . The expected runtime for unsuccessful runs  $\mathbb{E}[T_f]$  is set as a constant limit ( $T_{max}$ ) on the number of function evaluation calls, and the expected runtime for successful runs  $\mathbb{E}[T_s]$  is estimated as the average number of function evaluations performed by successful runs. Given these assumptions, *ert* can be expressed as an estimation of the expected runtime  $\mathbb{E}[T]$  as presented in Equation 9.

$$ert = \frac{1 - \hat{p}_s}{\hat{p}_s} T_{max} + \frac{1}{t_s} \sum_{i=1}^{t_s} T_i \quad (9)$$

where  $t_s$  is the number of successful runs,  $T_i$  is the number of evaluations for successful run  $i$ .

## 4 Experiments, Results and Discussion

### 4.1 Experimental setting

The experimental framework consists of generating enumerable instances based on the problem features of interest, then running the algorithm and modelling the expected runtime (*ert*)<sup>2</sup>.

The instance generator considers the following features:

- **Number of cities ( $n$ ):** Represents the number of cities to visit. This feature is not considered as an input to the regression model. Instead, the experiments are replicated for multiples values of  $n \in \{5, 6, 7, 8\}$ . Small values of  $n$  are chosen in order to be able to enumerate the solutions, which is used to derive the optimal values, as there is no exact methods able to solve the problem. Note that we will exclude some illustrations for  $n = 5, 6, 7$  for the sake of simplicity, but result summaries will be given for all the values of  $n$ .
- **Dropping rate ( $\mathcal{D}_r$ ):** Represents the dropping rate at which the value of an item decreases through time as shown in Equation 2.  $\mathcal{D}_r$  takes values from the set  $\{0.7, 0.75, 0.8, \dots, 1\}$ .
- **Minimum velocity ( $v_{min}$ )** is the minimum speed of the thief following Equation 1.  $v_{min}$  takes values from  $\{0.1, 0.2, \dots, 1\}$ .
- **Profit-weight correlation ( $\mathcal{T}$ ):** Defines the correlation between the weight ( $w_i$ ) and profit ( $p_i$ ) of each item. Three correlations have been defined in the TTP standard library [31]. We generate random weights and profits for each correlation type as follow:

- *Uncorrelated* (0):  $p_i \sim \mathcal{U}(10, 1000)$  and  $w_i \sim \mathcal{U}(10, 1000)$ ,
- *Uncorrelated with similar weight* (1):  $p_i \sim \mathcal{U}(10, 1000)$  and  $w_i \sim \mathcal{U}(1000, 1010)$ , and
- *Bounded strongly correlated* (2):  $p_i \sim \mathcal{U}(10, 1000)$  and  $w_i = p_i + 100$ .

As this feature can be considered an ordinal variable, numerical values (between parentheses) are assigned to represent the correlation strength, which will be useful for the regression analysis.

- **Knapsack capacity class ( $\mathcal{C}$ ):** Takes values from  $\{3, \dots, 10\}$ .  $\mathcal{C}$  is a factor occurring in the maximum weight of the knapsack which is given in Equation 10.

$$W = \frac{\mathcal{C}}{11} \sum_{x=2}^n \sum_{y=1}^{\mathcal{F}} w_{xy} \quad (10)$$

Note that values lower than 3 have been excluded as they can result in capacities smaller than the smallest item weight.

The features  $\mathcal{D}_r$  and  $v_{min}$  are considered as they control the strength of the dependencies. While the motivation for choosing  $\mathcal{T}$  and  $\mathcal{C}$  is to compare their impact on the expected runtime with that of the dependency features.

---

<sup>2</sup>We use Python3.10 with the statistical learning packages scikit-learn for the statistical analysis and regression models

403 Furthermore, the choice of  $\mathcal{T}$  and  $\mathcal{C}$ , instead of other features, is based on  
 404 previous analyses [11, 12].

405 10 instances are generated for each combination of feature values, resulting  
 406 in a total of  $|n| \times |\mathcal{D}_r| \times |v_{min}| \times |\mathcal{T}| \times |\mathcal{C}| \times 10 = 67,200$  instances. Then, 30  
 407 independent runs of Algorithm 1 are performed for each generated instance.  
 408 Afterwards, the *ert* is calculated for each instance based on Equation 9.

409 It is worth noting that the experiments result in 23 instances where the  
 410 algorithm fails to identify an optimal or near-optimal solution, i.e.,  $\hat{p}_s = 0$ ,  
 411 which results in a division-by-zero. To solve the issue, we set  $ert = \infty$ , and  
 412 consider these results as outliers for the analysis. Note that, based on a deeper  
 413 look into the individual results, we concluded that this happens for different  
 414 categories (combinations of features), i.e., all the categories are covered in the  
 415 analysis.



## 416 4.2 Results and analysis

### 417 4.2.1 Preliminary analysis of the runtime

418 The histograms in Figure 2 shows the distributions of  $\ln(ert)$ , where  $\ln(\cdot)$   
 419 denotes the natural logarithm function for the different values of  $n$ . The natural  
 420 logarithm is only used to better visualise the *ert* data. Indeed, *ert* follows a  
 421 distribution that is heavy-tailed (with a very high kurtosis  $kurt[ert] = 1390.38$   
 422 for  $n = 8$ ) and asymmetric (with a skewness  $\mu_3 = 32.51$  for  $n = 8$ ).

423 The nature of *ert* makes it virtually impossible to use standard linear re-  
 424 gression models. The most significant factor is the existence of outliers (large  
 425 *ert* values) – representing hard instances – which are difficult to properly in-  
 426 clude in the regression model. Note that the existence of heavy tails in the  
 427 dependent variable is not a problem in itself, except when it leads to heavy-  
 428 tailed residuals, which is the case here (based on experiments with linear  
 429 regression and regression tree models).

430 It is usually favourable for the observed residuals to be approximately  
 431 normally distributed. In order to obtain (approximately) normal residuals, a  
 432 transformation, such as the logarithm, square root or cubic root, should be  
 433 applied to the *ert*. While the mentioned transformations can be efficient in  
 434 taming outliers in many cases, the resulting residuals, even when using regres-  
 435 sion trees and other regression models, remain heavy-tailed due to the high  
 436 variation in the *ert* values. A better alternative in this case is to use the re-  
 437 ciprocals,  $\frac{1}{ert}$ , which results in a non-normal, but smaller-tailed distribution of  
 438 the residuals for a regression tree as shown in Figure 3. Furthermore, in order  
 439 to preserve the *ert* orders and improve the illustrations, we apply the trans-  
 440 formation  $1 - \frac{\min(ert)}{ert}$  instead, which does not impact the regression analysis  
 441 results. We will refer to the resulting values as the *runtime scores*.

**Figure 2:** Histogram of expected runtimes in logarithmic scale**Figure 3:** Histogram of the runtime scores (left) and residuals using a regression tree (right) for  $n = 8$ 

It is worth noting that  $\frac{1}{ert}$  should not be interpreted as the expected rate. Indeed, as the reciprocal is a convex function, we have:

$$E[\text{rate}] = E\left[\frac{1}{\text{runtime}}\right] \leq \frac{1}{E[\text{runtime}]} = \frac{1}{ert}$$

442 based on Jensen's inequality [19]. Furthermore, as the distribution of *ert* is  
 443 unknown and difficult to fit, it is hard to deduce  $E[rate]$  using the law of the  
 444 unconscious statistician. Therefore,  $\frac{1}{ert}$  can be simply interpreted as the *ert*  
 445 reciprocal.

446 As linear models were not efficient in capturing the variability of the data,  
 447 we utilise regression trees for two main reasons: (1) they are efficient in han-  
 448 dling larger data with outliers, (2) they can produce explainable outcomes,  
 449 which is important to understand the features' impact on the expected runtime  
 450 and identify what makes some TTP instances harder to solve. In other words,  
 451 regression trees offer a better trade-off between efficiency and explainability  
 452 compared to other alternatives.

### 453 4.2.2 Regression analysis

454 In order to analyse the impact of the problem features on the difficulty of  
 455 tackling the instances, one must map the problem features to the expected run-  
 456 time to identify near-optimal solutions. This can be achieved using non-linear  
 457 regression analysis. Specifically, we consider regression trees using the Mean  
 458 Squared Error (MSE) to measure the quality of split, and two parameters to  
 459 control the model's complexity. The first is the maximum depth, which rep-  
 460 resents the maximum number of node splits the regression tree makes before  
 461 returning a prediction. The second is the minimum sample size, which repre-  
 462 sents the minimum number of samples at each leaf. We denote by  $RT_{d_{max}, s_{min}}^n$   
 463 the regression tree <sup>3</sup> obtained for problem size  $n$ , by using the parameter val-  
 464 ues  $d_{max}$  as the maximum depth and  $s_{min}$  as the minimum sample size, during  
 465 the training phase.


466 Model tuning is achieved using a 5-fold cross-validation, and based on a  
 467 grid search with different values of  $d_{max}$  and  $s_{min}$ . Figure 4 shows a heat-map  
 468 of the trade-off between the regression tree complexity and the model fitness  
 469 for  $n = 8$ . A maximum  $R^2$  of 0.53 can be reached using the regression model  
 470  $RT_{9,40}^8$ , i.e., the corresponding models can explain 53% of the variation in ex-  
 471 pected runtime that is predictable from the problem features. The remaining  
 472 47% of the variation can be attributed to many other aspects, the most likely  
 473 is the stochasticity of the considered Memetic Algorithm, which leads to high  
 474 variation in the *ert* within some instance categories (combination of features),  
 475 and high variation in the *number of evaluations* for the same instances. An-  
 476 other, less likely scenario is that there are combinations of features that the  
 477 regression tree was not able to identify. This is believed to have minimum im-  
 478 pact as other machine learning algorithms (random forest and neural networks)  
 479 were investigated and could not result in a better coefficient of determination.

480 Naturally, more complex trees can result in higher  $R^2$  values. Nevertheless,  
 481 this comes with the cost of losing the ability to interpret the model. For the  
 482 purpose of this analysis, the goal is to understand the impact of the feature  
 483 combinations on the expected runtime. Therefore, we favour small explainable

---

<sup>3</sup>The implementation is done in Python 3.8.10 using scikit-learn 1.1.0





**Figure 4:** Regression tree complexity vs. coefficient of determination ( $R^2$ ) for  $n = 8$ . The complexity is represented by the maximum depth, shown in the y-axis, and the minimum sample size per leaf, shown in the x-axis; while  $R^2$  values are illustrated in the colour bar

484 trees, even if they explain a lower percentage of variability between the features  
 485 and the expected runtime.

486 Based on the above, we consider two regression models for each  $n$ .  $RT_{6,280}^5$ ,  
 487  $RT_{6,20}^6$ ,  $RT_{7,20}^7$ , and  $RT_{9,40}^8$  are the models with the highest accuracy; while  
 488  $RT_{3,380}^5$ ,  $RT_{3,380}^6$ ,  $RT_{3,380}^7$ , and  $RT_{3,380}^8$  represent a good trade-off between the  
 489 tree complexity and quality-of-fit for all  $n$  values. The evaluation metrics of  
 490 the two resulting models are shown in Table 3 for each value of  $n$ , where the  
 491 first model corresponds to the explainable one, and the second corresponds to  
 492 the one with the highest coefficient of determination.

**Table 3:** Evaluation metrics of the obtained regression trees. MAE is the mean absolute error, MSE is the mean squared error, RMSE is the root mean squared error, and  $R^2$  is the coefficient of determination

n	5		6		7		8	
<b>Model</b>	$RT_{3,380}^5$	$RT_{6,280}^5$	$RT_{3,380}^6$	$RT_{6,20}^6$	$RT_{3,380}^7$	$RT_{7,20}^7$	$RT_{3,380}^8$	$RT_{9,40}^8$
<b>MAE</b>	0.08	0.08	0.09	0.09	0.1	0.1	0.09	0.08
<b>MSE</b>	0.12	0.12	0.12	0.12	0.13	0.12	0.11	0.11
<b>RMSE</b>	0.02	0.01	0.02	0.01	0.02	0.02	0.01	0.01
<b><math>R^2</math></b>	0.37	0.4	0.45	0.48	0.4	0.45	0.5	0.55


493 In general, Table 3 shows that the loss in quality between the simple and  
 494 complex models, in terms of error metrics and  $R^2$ , is minimal and can be  
 495 neglected due to the gain in explainability. Nevertheless, a difference in the

496 performances can be noticed between the different models for each  $n$  value.  
 497 Looking back at the different *ert* distributions, this can be explained by the  
 498 fact that as  $n$  grows, the variance in *ert* decreases.

499 On the one hand, the complex regression trees ( $RT_{6,280}^5$ ,  $RT_{6,20}^6$ ,  $RT_{7,20}^7$  and  
 500  $RT_{9,40}^8$ ) can provide useful global insights on the importance of features and  
 501 to what extent we can explain the variability in the runtime variable, but it  
 502 could be less practical to extract which combination of feature values lead to a  
 503 specific expected runtime. On the other hand, the explainable regression trees  
 504 ( $RT_{3,380}^5$ ,  $RT_{3,380}^6$ ,  $RT_{3,380}^7$  and  $RT_{3,380}^8$ ) result in a competitive quality while  
 505 having the additional benefit of being explainable, allowing us to draw useful  
 506 conclusions on what makes instances easier or harder to solve.

**Table 4:** Feature importance based on all regression tree models

<b>n</b>	<b>5</b>		<b>6</b>		<b>7</b>		<b>8</b>	
<b>Model</b>	$RT_{3,380}^5$	$RT_{6,280}^5$	$RT_{3,380}^6$	$RT_{6,20}^6$	$RT_{3,380}^7$	$RT_{7,20}^7$	$RT_{3,380}^8$	$RT_{9,40}^8$
$\mathcal{T}$	0.13904	0.13409	0.13814	0.13800	0.17611	0.18711	0.27582	0.27084
$\mathcal{C}$	0.79609	0.79498	0.72660	0.67902	0.60393	0.54678	0.47955	0.44658
$\mathcal{D}_r$	0.05560	0.05555	0.12460	0.13655	0.20094	0.19975	0.22137	0.22135
$v_{min}$	0.00927	0.01537	0.01066	0.04643	0.01902	0.06636	0.02327	0.06123

**Figure 5:** Feature importance versus the problem size ( $n$ )

507 Table 4 provides a macroscopic idea on the impact of the problem features  
 508 using the Gini-Simpson Index [20]. The feature importance values vary between  
 509 the two models, but are aligned in sense that they rank features similarly. For  
 510 this holistic investigation, we focus on the models having the better quality-  
 511 of-fit.

512 The results show that the capacity class ( $\mathcal{C}$ ) is the most influential feature  
 513 contributing to the difficulty of instances. This is followed by the dropping rate  
 514 ( $\mathcal{D}_r$ ) and profit-weight correlation ( $\mathcal{T}$ ), while the minimum velocity ( $V_{min}$ ) has  
 515 a much lower impact index. Comparing the two dependency features shows  
 516 that items losing value with time is highly important to predict the efficiency  
 517 of the algorithm in tackling the problem. However, the change of velocity of the  
 518 thief based on the knapsack load seems to have minimum impact on predicting  
 519 the performance of the algorithm when tackling the problem.

520 Figure 5 shows the evolution of the feature importance indices in terms of  
 521  $n$ . The dependency features tend to have a higher importance index as the  
 522 problem size increases, this phenomenon is stronger for  $\mathcal{D}_r$  compared to  $v_{min}$ .  
 523 The same can be said about the feature  $\mathcal{T}$ . On the other hand, the importance  
 524 index for  $\mathcal{C}$  is decreasing significantly.

525 While the above analysis can help in identifying important features (and  
 526 potentially eliminating less important ones), a more in-depth analysis should  
 527 be done at a microscopic scale, which can be achieved using  $RT_{3,380}^n$ .

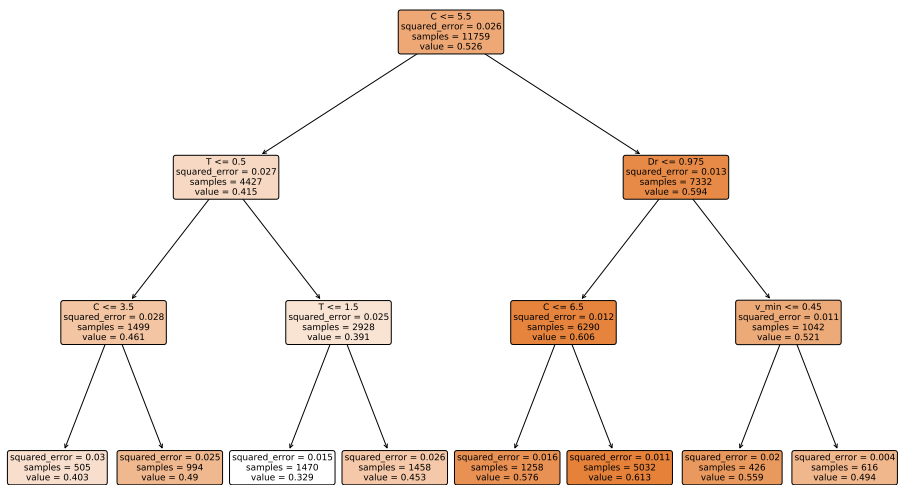
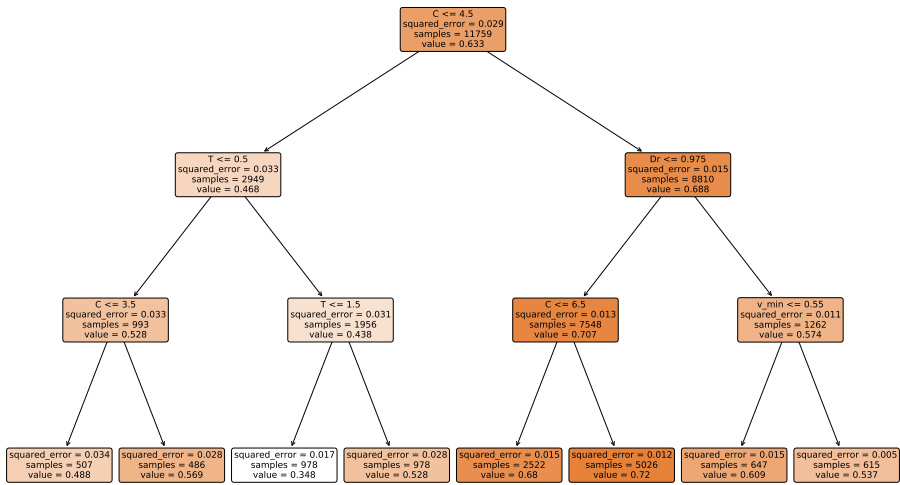
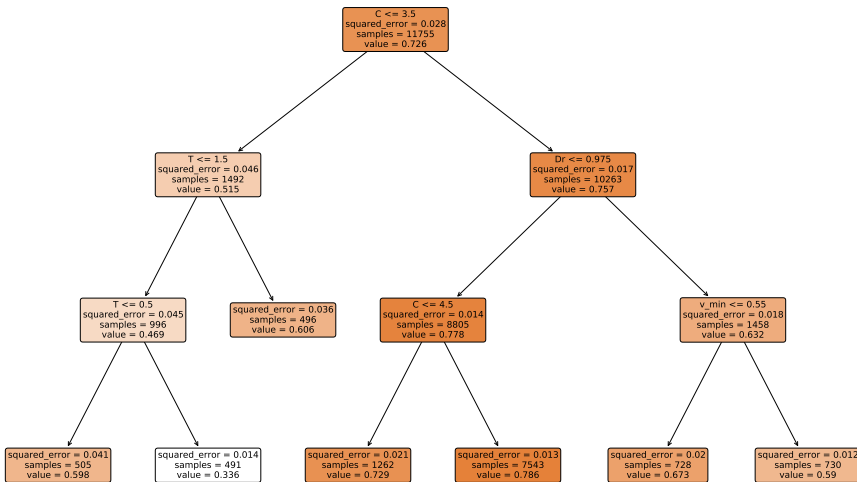


Figure 6: Regression tree  $RT_{3,380}^5$ . Colour brightness represents the difficulty of instance sets, where darker colour represents harder instances.

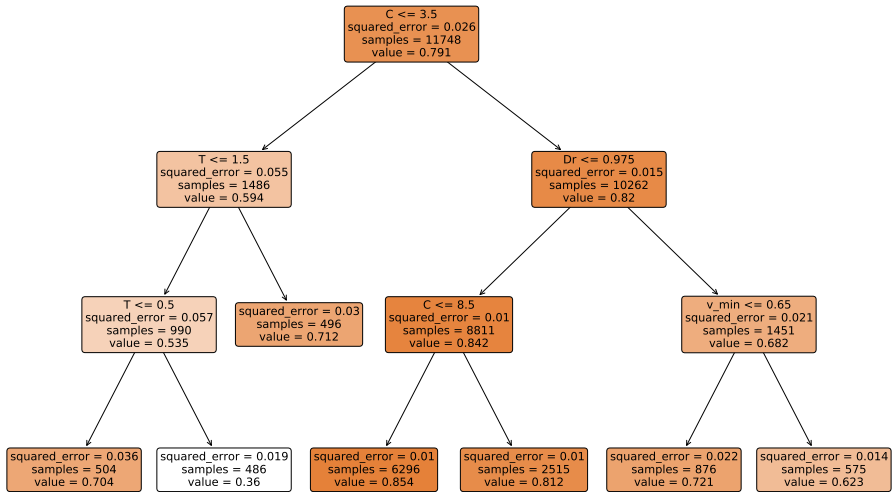
528 Figures 6, 7, 8, and 9 show the resulting regression tree models  
 529  $\{RT_{3,380}^n\}, n \in \{5, \dots, 8\}$ , which will help us compare the regression logic across  
 530 different instance sizes. Additionally, in Appendix A we reproduce the regres-  
 531 sion models for  $n = 6, 7$  to show that the results can be generalised, i.e., they  
 532 are not dependent on the generated instances, but on the problem features.



**Figure 7:** Regression tree  $RT_{3,380}^6$ . Colour brightness represents the difficulty of instance sets, where darker colour represents harder instances.



**Figure 8:** Regression tree  $RT_{3,380}^7$ . Colour brightness represents the difficulty of instance sets, where darker colour represents harder instances.



**Figure 9:** Regression tree  $RT_{3,380}^S$ . Colour brightness represents the difficulty of instance sets, where darker colour represents harder instances.

533 Based on this, we observe that the regression models follow a similar logic  
 534 which can be clearly seen in the branching nodes. Small deviations exist in the  
 535 branching nodes and predicted values. Looking at all the four models, we can  
 536 report the following general observations and explanations:

- 537 • Setting  $\mathcal{D}_r = 1$  results in the standard TTP model as the dropping rate  
 538 dependency is deactivated, making the minimum velocity feature the main  
 539 separator as it represents the velocity drop dependency. Furthermore, the  
 540 lower minimum velocities ( $v_{min}$ ) result in slightly harder instances.
- 541 •  $\mathcal{D}_r$  has a clear impact on the difficulty of instances. Lower dropping rate  
 542 values result in more difficult instances.
- 543 • Larger knapsack capacities ( $\mathcal{C}$ ) result in harder instances, as clearly seen  
 544 in branches  $B_2$ ,  $B_4$ ,  $B_5$  and  $B_6$ . This is in contradiction with the find-  
 545 ings in [11, 12] where the difficulty of instances is associated with the size  
 546 of basins of attraction, in the context of a local search algorithm applied  
 547 to the standard TTP. Further experiments considering only standard TTP  
 548 instances ( $\mathcal{D}_r = 1$ ) show a similar behaviour, confirming that larger knap-  
 549 sack capacities increase the difficulty of the instances. This is suggesting  
 550 that the contradiction is due to the nature of the search algorithm, not to  
 551 the complexity added by considering the dropping rate dependency. Indeed,  
 552 the local optima network analysis is only suitable for embedded neighbour-  
 553 hood search algorithms; and the conclusions cannot be extrapolated to more  
 554 sophisticated algorithms such as the considered evolutionary algorithms.

- 555 • For smaller capacities, the profit-value correlation ( $\mathcal{T}$ ) has a significant impact  
556 on the problem difficulty. The hardest instances are the ones where  
557 the profit and weight are bounded and strongly correlated ( $T = 2$ ) and the  
558 ones with no correlation ( $T = 0$ ), while the instances with similar weights  
559 ( $T = 1$ ) are the easiest to tackle.

560 It is interesting to see how the item value drop dependency results in significant  
561 gaps in the performance of the algorithm. When it is deactivated (standard TTP), the instances can be solved relatively fast, and the velocity  
562 change dependency has only a small impact on the runtimes. When it is activated,  
563 the impact of the velocity change dependency is dominated by the capacity and correlation features.  
565

566 A possible explanation of the stronger impact of  $\mathcal{D}_r$  compared to  $v_{min}$  can  
567 be attributed to the formulation of dependencies. Specifically, in the dropping  
568 rate dependency (see Equation 2), the updated item values are dependent on  
569 both the previous item value and the time the item has been carried. On the  
570 other hand, in the velocity change dependency (see Equation 1, the thief's  
571 updated velocity only depend on the knapsack load, and its previous velocity  
572 is completely omitted.

573 These findings show that the impact of different dependency relationships  
574 can differ significantly when formulating a problem with multiple components.  
575 In our case, the analysis shows that adding the item value drop dependency  
576 makes the TTP a much more challenging problem compared to the standard  
577 model, leading to harder instances, which is reflected by the expected runtime.  
578 Therefore, this aspect should be given more attention when investigated the  
579 TTP in particular, or other capacitated routing problems in general, especially  
580 because the impact of these features evolves based on the problem size.

581 Recognising that different dependencies can have a significantly different  
582 impact on the problem's difficulty is important. However, one can go beyond  
583 and seeks ways to use these results to improve the way these problems are  
584 tackled.

585 As it is possible to estimate how much time would be needed to find a  
586 good approximation given a specific problem instance, based only on known  
587 features. One way to directly use the results reported here is to make an informed  
588 and automated decision on the maximum number of iterations (or time  
589 budget) needed to achieve near-optimality, given a specific problem instance.  
590 It is also possible to revisit the model and simplify it to reduce the impact  
591 of particular dependencies. This should be done carefully as it can lead to  
592 undesirable outcomes due to model oversimplification. Another possibility is  
593 to use these results as domain knowledge included in the solution methods.  
594 Such an approach can be clustering instances into categories that can be tackled  
595 with different approaches based on their difficulty, or tuning the heuristic  
596 parameters based on these instances clusters.

## 5 Conclusion

In this paper, we investigated the difficulty of the Travelling Thief Problem by considering two types of dependency on the fitness landscape using performance prediction models (regression analysis). We introduce a tunable model allowing to control the velocity change and item value drop dependencies using associated problem features. The impact of these features on the expected runtime is then evaluated to understand how these features impact the search landscape for a Memetic Algorithm.

This study allowed us to better understand what makes instances harder, and gave us new insights on the impact of the problem features. The analysis shows that the inclusion of the item value drop dependency leads to harder instances, which is reflected by the expected runtime. The results also showed how the impact of these features evolves based on the problem size. In particular, the dropping rate tends to be more important as the problem grows and its impact is stronger than the velocity change constraint.

A continuation of this study is to map these features to the choice of algorithm parameter. Besides, as future work, we intend to explore other fitness landscape techniques such as fitness cloud, auto-correlation, time to local optimum, distance to global optimum, information analysis: entropy, information stability, partial information content and density basin information.

## Declarations

### Compliance with Ethical Standards

None of the authors of this paper have a financial or personal relationship with other people or organisations that could inappropriately influence or bias the content of the paper. This paper does not contain any studies with human participants or animals performed by any of the authors. This manuscript is the authors' original work and has not been published nor has it been submitted simultaneously elsewhere. All authors have checked the manuscript and have agreed to the submission.

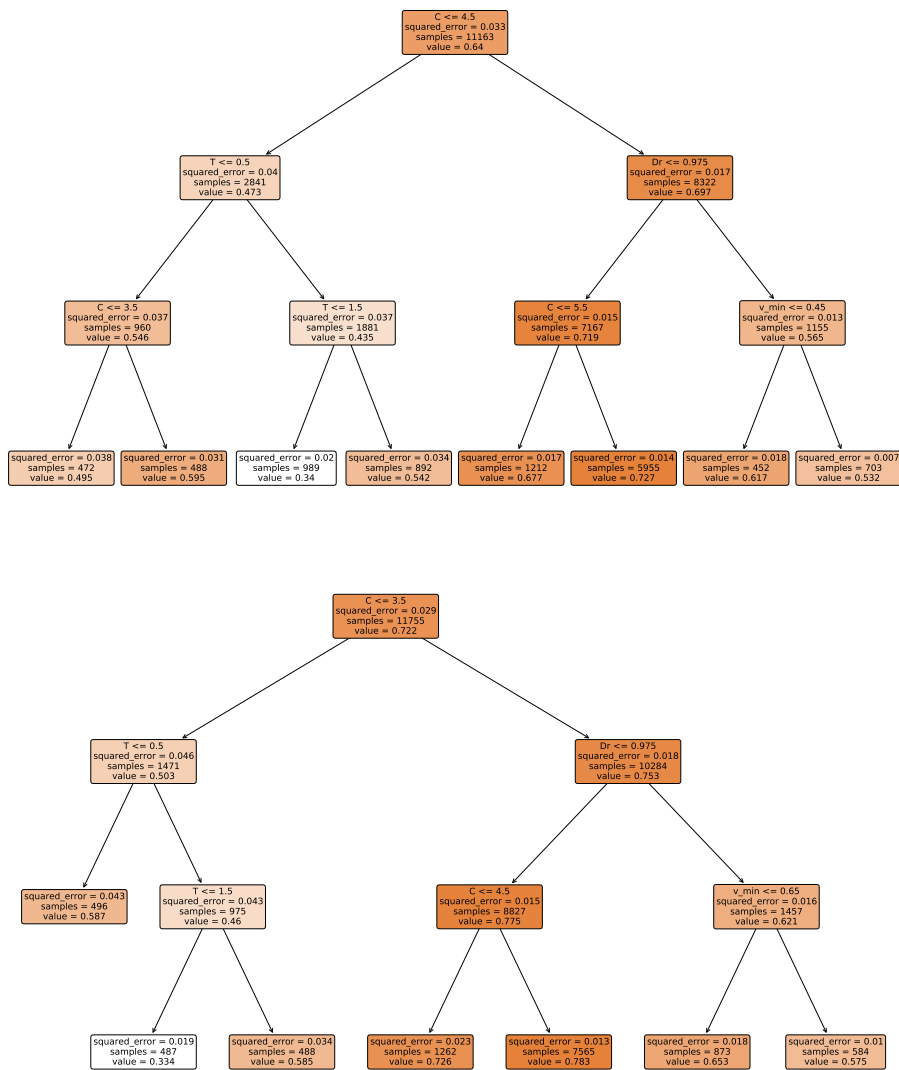
### Funding and Competing interests

No funding was received to assist with the preparation of this manuscript. Besides, the authors have no competing interests to declare that are relevant to the content of this article.

## Appendix A Additional results

In this appendix, we show the results for additional experiments on different sets of instances for  $n = 6$  and  $n = 7$ . The same process defined in the earlier sections was used to create them, but just different seeds of the random number general were used. The goal is to show that different sets of instances generate roughly the same regression model, i.e., regression trees with a similar logic as shown in Figure A1.

637 When comparing these two with the corresponding trees in Figures 7 and 8,  
 638 we can see that the conditions at the inner nodes are almost always identical  
 639 (i.e. for 10 of 13 inner nodes) or very similar, and the respective errors (shown  
 640 in Table A1) and sample numbers are very close matches, too.



**Figure A1:** Regression tree  $RT_{3,380}^{6*}$  for a new dataset with  $n = 6$  (top) and  $RT_{3,380}^{7*}$  with  $n = 7$  (bottom).



**Table A1:** Evaluation metrics of the regression trees.

<b>Model</b>	$RT_{3,380}^{6*}$	$RT_{3,380}^{7*}$
<b>MAE</b>	0.1	0.1
<b>MSE</b>	0.13	0.13
<b>RMSE</b>	0.02	0.02
$R^2$	0.45	0.4

641 Hence, we conclude that even though the methodology is based on ran-  
642 domly created instances and even though it employs a memetic algorithm as a  
643 randomised search heuristic, the achievable insights at a high level (i.e. when  
644 reasoning about the effects of dependencies) are unaffected.

## 645 References

- 646 [1] Applegate D, Cook W, Rohe A (2003) Chained lin-kernighan for large  
647 traveling salesman problems. *INFORMS Journal on Computing* 15(1):82–  
648 92
- 649 [2] Bonyadi MR, Michalewicz Z, Barone L (2013) The travelling thief prob-  
650 lem: The first step in the transition from theoretical problems to realistic  
651 problems. In: 2013 IEEE Congress on Evolutionary Computation, IEEE
- 652 [3] Bonyadi MR, Michalewicz Z, Przybyłek MR, et al (2014) Socially inspired  
653 algorithms for the travelling thief problem. In: Proceedings of the 2014  
654 Annual Conference on Genetic and Evolutionary Computation
- 655 [4] Bonyadi MR, Michalewicz Z, Wagner M, et al (2019) Evolutionary compu-  
656 tation for multicomponent problems: opportunities and future directions.  
657 In: Optimization in Industry. Springer
- 658 [5] Chand S, Wagner M (2016) Fast heuristics for the multiple travel-  
659 ing thieves problem. In: Proceedings of the Genetic and Evolutionary  
660 Computation Conference 2016, pp 293–300
- 661 [6] Croes GA (1958) A method for solving traveling-salesman problems.  
662 *Operations research* 6(6):791–812
- 663 [7] Darestani SA, Hemmati M (2019) Robust optimization of a bi-objective  
664 closed-loop supply chain network for perishable goods considering queue  
665 system. *Computers & industrial engineering* 136:277–292
- 666 [8] Eiben AE, Smith JE (2015) Introduction to evolutionary computing.  
667 Springer

- 668 [9] El Yafrani M, Ahiod B (2016) Population-based vs. single-solution heuristics for the travelling thief problem. In: Proceedings of the Genetic and  
669 Evolutionary Computation Conference 2016  
670
- 671 [10] El Yafrani M, Ahiod B (2018) Efficiently solving the traveling thief problem using hill climbing and simulated annealing. *Information Sciences*  
672 432  
673
- 674 [11] El Yafrani M, Martins MS, Krari ME, et al (2018) A fitness landscape analysis of the travelling thief problem. In: Proceedings of the Genetic  
675 and Evolutionary Computation Conference  
676
- 677 [12] El Yafrani M, Scoczynski M, Delgado M, et al (2022) On the fitness landscapes of interdependency models in the travelling thief problem. In: Proceedings of the Genetic and Evolutionary Computation Conference  
678 Companion  
679  
680
- 681 [13] Faulkner H, Polyakovskiy S, Schultz T, et al (2015) Approximate approaches to the traveling thief problem. In: Proceedings of the 2015  
682 Annual Conference on Genetic and Evolutionary Computation  
683
- 684 [14] Ghomi SF, Asgarian B (2019) Development of metaheuristics to solve a transportation inventory location routing problem considering lost sale for perishable goods. *Journal of Modelling in Management*  
685  
686
- 687 [15] Hansen N, Auger A, Ros R, et al (2010) Comparing results of 31 algorithms from the black-box optimization benchmarking bbob-2009. In: Proceedings of the 12th annual conference companion on Genetic and  
688 evolutionary computation, pp 1689–1696  
689  
690
- 691 [16] Işıklı E, Aydın N, Bilgili L, et al (2020) Estimating fuel consumption in maritime transport. *Journal of Cleaner Production* 275:124,142  
692
- 693 [17] Janardhanan M, Li Z, Nielsen P (2019) Model and migrating birds optimization algorithm for two-sided assembly line worker assignment and balancing problem. *Soft Computing* 23(21):11,263–11,276. <https://doi.org/10.1007/s00500-018-03684-8>  
694  
695  
696
- 697 [18] Janković A, Doerr C (2019) Adaptive landscape analysis. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp 2032–2035  
698  
699
- 700 [19] Jensen JLWV (1906) Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta mathematica* 30(1):175–193  
701
- 702 [20] Jost L (2006) Entropy and diversity. *Oikos* 113(2):363–375

- 703 [21] Li W, Meng X, Huang Y (2021) Fitness distance correlation and mixed  
704 search strategy for differential evolution. *Neurocomputing* 458:514–525
- 705 [22] Liefvooghe A, Verel S, Aguirre H, et al (2013) What makes an instance  
706 difficult for black-box 0–1 evolutionary multiobjective optimizers? In:  
707 International Conference on Artificial Evolution (Evolution Artificielle),  
708 Springer, pp 3–15
- 709 [23] Liefvooghe A, Verel S, Daolio F, et al (2015) A feature-based performance  
710 analysis in evolutionary multiobjective optimization. In: International  
711 Conference on Evolutionary Multi-Criterion Optimization, Springer, pp  
712 95–109
- 713 [24] Lu H, Zhou R, Fei Z, et al (2019) Spatial-domain fitness landscape analysis  
714 for combinatorial optimization. *Information Sciences* 472:126–144
- 715 [25] Martins MS, El Yafrani M, Santana R, et al (2018) On the performance  
716 of multi-objective estimation of distribution algorithms for combinatorial  
717 problems. In: 2018 IEEE congress on evolutionary computation (CEC),  
718 IEEE, pp 1–8
- 719 [26] Mei Y, Li X, Yao X (2014) Improving efficiency of heuristics for the large  
720 scale traveling thief problem. In: Asia-Pacific Conference on Simulated  
721 Evolution and Learning, Springer
- 722 [27] Merz P, Freisleben B (2000) Fitness landscape analysis and memetic al-  
723 gorithms for the quadratic assignment problem. *IEEE transactions on*  
724 *evolutionary computation* 4(4):337–352
- 725 [28] Michalewicz Z, Fogel DB (2013) *How to solve it: modern heuristics*.  
726 Springer Science & Business Media
- 727 [29] Nikfarjam A, Neumann A, Neumann F (2022) On the use of quality di-  
728 versity algorithms for the traveling thief problem. In: Proceedings of the  
729 Genetic and Evolutionary Computation Conference, pp 260–268
- 730 [30] Ochoa G, Malan K (2019) Recent advances in fitness landscape analysis.  
731 In: Proceedings of the Genetic and Evolutionary Computation Conference  
732 Companion, pp 1077–1094
- 733 [31] Polyakovskiy S, Bonyadi MR, Wagner M, et al (2014) A comprehensive  
734 benchmark set and heuristics for the traveling thief problem. In: Pro-  
735 ceedings of the 2014 Annual Conference on Genetic and Evolutionary  
736 Computation
- 737 [32] Przybyłek MR, Wierzbicki A, Michalewicz Z (2018) Decomposition  
738 algorithms for a multi-hard problem. *Evolutionary computation* 26(3)

- 739 [33] Rahman H, Nielsen I (2019) Scheduling automated transport vehicles for  
740 material distribution systems. *Applied Soft Computing Journal* 82. <https://doi.org/10.1016/j.asoc.2019.105552>  
741
- 742 [34] Rahman H, Janardhanan M, Nielsen P (2020) An integrated ap-  
743 proach for line balancing and agv scheduling towards smart assembly  
744 systems. *Assembly Automation* 40(2):219–234. [https://doi.org/10.1108/](https://doi.org/10.1108/AA-03-2019-0057)  
745 [AA-03-2019-0057](https://doi.org/10.1108/AA-03-2019-0057)
- 746 [35] Richter H (2008) Coupled map lattices as spatio-temporal fitness func-  
747 tions: Landscape measures and evolutionary optimization. *Physica D:*  
748 *Nonlinear Phenomena* 237(2):167–186
- 749 [36] Richter H (2013) Dynamic fitness landscape analysis. In: *Evolutionary*  
750 *computation for dynamic optimization problems*, Springer, pp 269–297
- 751 [37] Sung I, Nam H, Lee T (2013) Scheduling algorithms for mobile harbor: An  
752 extended m-parallel machine problem. *International Journal of Industrial*  
753 *Engineering : Theory Applications and Practice* 20(1-2):211–224
- 754 [38] Wagner M (2016) Stealing items more efficiently with ants: a swarm  
755 intelligence approach to the travelling thief problem. In: *International*  
756 *Conference on Swarm Intelligence*, Springer
- 757 [39] Wagner M, Lindauer M, Mısıř M, et al (2018) A case study of algorithm  
758 selection for the traveling thief problem. *Journal of Heuristics* 24(3):295–  
759 320
- 760 [40] Wang Jj, Wang L (2022) A cooperative memetic algorithm with feed-  
761 back for the energy-aware distributed flow-shops with flexible assembly  
762 scheduling. *Computers & Industrial Engineering* 168:108,126
- 763 [41] Watson JP (2010) An introduction to fitness landscape analysis and cost  
764 models for local search. In: *Handbook of metaheuristics*. Springer
- 765 [42] Wu J, Polyakovskiy S, Neumann F (2016) On the impact of the renting  
766 rate for the unconstrained nonlinear knapsack problem. In: *Genetic and*  
767 *Evolutionary Computation Conference*, ACM
- 768 [43] Wuijts RH, Thierens D (2019) Investigation of the traveling thief problem.  
769 In: *Proceedings of the Genetic and Evolutionary Computation Conference*
- 770 [44] Zou F, Chen D, Liu H, et al (2022) A survey of fitness landscape analysis  
771 for optimization. *Neurocomputing* 503:129–139