

Preference Learning for Horizontal Collaboration in Transport Operations

Federico Toffano

Insight Centre for Data Analytics, School of Computer
Science and IT
University College Cork, Ireland
federico.toffano@insight-centre.org

Nic Wilson

Insight Centre for Data Analytics, School of Computer
Science and IT
University College Cork, Ireland
nic.wilson@insight-centre.org

ABSTRACT

We consider a system computing collaborative delivery plans to deliver a set of orders, by sharing trucks between the companies involved. Collaborative journeys are evaluated by the decision-makers considering multiple Key Performance Indicators (KPIs) such as cost, distance and time. Knowing the decision-makers' preferences over the KPIs would enable us to compute personalised solutions that increase the chance of finding attractive collaborations. With this work we address two problems. Firstly, we show how we can learn the decision-makers' preferences over the KPIs by observing their interactions with the system. Secondly, we show how to estimate the improvement or disimprovement in the KPIs arising from the collaboration.

KEYWORDS

Preference Learning; Collaborative Transport; Intelligent Logistics; Shapley Value

ACM Reference Format:

Federico Toffano and Nic Wilson. 2021. Preference Learning for Horizontal Collaboration in Transport Operations. In *Proc. of the 1st Workshop on Multi-Objective Decision Making (MODeM), Online, July 14-16, 2021, IFAAMAS*, 9 pages.

1 INTRODUCTION

Horizontal collaborations are inter-organisational relationships between two or more companies whose purpose is sharing resources to achieve a common objective. LOGISTAR (*Enhanced data management techniques for real time logistics planning and scheduling*) is an EU research project to develop a system to support horizontal collaborations of transport operations between competitors, and one of the main purposes of this project is the computation of collaborative delivery plans sharing trucks between the companies involved. Collaborative deliveries can be very beneficial, not only economically but also environmentally, since the optimisation of the truck-load reduces the number of journeys and the total CO₂ emissions. In fact, companies cannot always fill their trucks, especially during the return journey, when the truck could be completely empty. Therefore sharing the truck capacity with other companies can improve the efficiency of a network of transports with benefits for all the partners involved.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Proc. of the 1st Workshop on Multi-Objective Decision Making (MODeM), Hayes, Mannion, Vamplew (eds.), July 14-16, 2021, Online. © 2021 Copyright held by the owner/author(s).

The horizontal collaboration we focus on in this paper involves two companies, that we call P and N, with whom we have defined the problem requirements and who will be involved in the testing of our software. Decision makers from P and N aim to find collaborative delivery plans sharing some of their trucks. In this context, the collaborative journeys proposed by the LOGISTAR system are computed by an optimiser and evaluated by the decision-makers considering multiple Key Performance Indicators (KPIs) such as cost, distance and time. The optimiser needs then to consider such KPIs during the optimisation of a delivery plan for a set of input orders, and each decision-maker may have different preferences with respect to the KPIs. Thus a fair objective function for the computation of collaborative journeys should consider the preferences of all partners involved.

Our goal is to learn the preferences of the decision-makers by observing their decisions with respect to the collaborative journeys proposed by the system. We assume a parameterised preference model for each decision-maker, with the domain of the parameters being the set of all the feasible preferences. The preference model considered is the weighted sum of the KPIs evaluating the journeys, and the weights of the KPIs vector define the decision-maker preferences. The general idea is to estimate the preference model of the two decision-makers involved and merge them to create a common preference model used as the objective function for the optimiser computing the collaborative journeys. The method presented in our paper to compute the joint preference model applies only for two decision-makers; however, it could be easily extended to cover a generic number of decision-makers. Although the details on how the optimiser computes the journeys are out of the scope of this paper, it is important to note that the goal is to find a collaborative solution that performs better than non-collaborative solutions and not a collaborative solution that reduces the costs for all the decision-makers involved. The arrangement between the partners P and N is based on mutual trust, and it is not in the interest of either partner to undermine that trust by deliberately misleading the other party.

Our preference learning method is based on the observation of the acceptance or rejection of the collaborative journeys proposed by the system. In the user interface, each collaborative journey is associated with two KPIs vectors: (i) the KPIs of the collaborative journey computed by the optimiser; (ii) an estimated KPIs vector for a non-collaborative delivery of the orders planned in the collaborative journey. This enables the decision maker to consider potential advantages and disadvantages of the collaborative journey, in terms of the KPIs, by comparing the components of these KPIs vectors. The decision-maker is also asked to input a motivation for their

decisions from a predefined list, and we generate (soft) constraints on the preference model of the decision-maker if the motivation is related to the two KPIs vectors. The constraints generated from all the past decisions will then be used as input for a Support Vector Machine (SVM) algorithm to estimate the weights vector of the decision-maker.

The two main contributions of this work are the following. Firstly, we propose a method to learn the decision-makers' preferences by observing the motivation of the decisions taken in the system. In particular, we allow each decision-maker to express the importance of a specific KPI, and we show how to translate this preference information into soft constraints for the preference model. Secondly, we present a method based on the Shapley value to split the KPIs of a journey among the corresponding orders. This is a crucial step of our preference learning approach since we need to compare a collaborative KPIs vector associated with a single collaborative journey with a non-collaborative KPIs vector of the same orders composing the collaborative journey. However, the orders of a collaborative journey may be part of multiple non-collaborative journeys along with other orders. Thus, supposing, for example, two companies, we may very well not have two non-collaborative journeys that together deliver the same orders as the collaborative journey. Then we cannot directly compare the corresponding KPIs vectors. On the other hand, if we have a method that generates a KPIs vector for any given order we could evaluate any set of orders by combining the corresponding KPIs vectors.

The rest of the paper is organised as follows. Section 2 provides a brief literature review of relevant works. Section 3 defines the key concepts and goals of our work. Section 4 presents our preference learning approach based on the observation of the interaction of the decision-makers with the system. Section 5 defines how to estimate the KPIs of a generic set of orders using the KPIs of the journeys planned by the optimiser. Section 6 shows how to split the KPIs of a journey among the corresponding orders with a method based on the Shapley value. Section 7 presents some computational results showing the time performances of the methods proposed in this work. Section 8 summarises our work, including some related future works.

2 RELATED WORK

The problem of computing collaborative journeys delivering orders of two or more e-logistics operators is a collaborative Vehicle Routing Problem (VRP). See, e.g., [10] for a survey on this topic. To the best of our knowledge, there is no work on preference learning for multi-objective collaborative VRP. In general, collaborative VRPs are classified as *centralised* or *decentralised*. In centralised systems the total profit is maximised jointly (see, e.g., [5, 8, 32]); on the other hand, in a decentralised system, collaborations are evaluated separately by the partners sharing a limited amount of information (see, e.g., [4, 6, 23]). Our work could be of interest for both of these settings in a multi-objective context since we first learn the preferences of every single partner over a set of KPIs and afterwards we combine them to define a joint preference model. Thus, the preferences of every single partner could be used in a decentralised system to define personalised objective functions not shared among the collaborators; the joint preference model instead

could be used for the computation of a fair collaborative solution with a centralised system.

More generally, our work could be classified as a multi-objective multi-agent system (MOMAS) with a utility-based approach (see, e.g., [18] for a survey). One of the main goals of these systems is to define a utility function for every agent mapping a vector of values evaluating the objects to optimise to a scalar value. MOMAS are defined as *cooperative* (see, e.g., [3, 20]) if all the agents collaborate to improve the performance of the entire system, or *competitive* (see, e.g., [7, 17]) if any win for one agent implies a loss for another. In our context, we have a mix of these two characterisations since we assume that the LOGISTAR partners collaborate to improve the whole network of transportation; however, the preferences over the KPIs may be conflicting.

In the literature we can find several approaches to learn a utility function representing the preferences of the decision-makers. In particular, we are interested in preference model represented by the weighted sum of the objectives with the weights representing the possible preferences of the decision-makers. The classical preference elicitation approach arises from multiattribute utility theory [9, 19]. It aims to precisely specify such a utility function with a series of questions for the users to identify some key values. Alternative methods based on interviews are such as the Analytic Hierarchy Process [16, 21, 28, 31]. However, these approaches can be difficult and error-prone tasks. Artificial Intelligence (AI) methods based on comparisons of alternatives require instead much less cognitive effort. The idea is to show to the decision-maker two or more alternatives, asking them to select the most preferred alternative or to indicate a preference order over the whole set of alternatives. Certain qualitative approaches [13, 22, 25–27] reduce the uncertainty of the value function by reducing the set of feasible parameters by translating input preferences into hard constraints on the parameter space. These methods are efficient in terms of the number of queries for the decision-makers to obtain a good estimation of the preference model; however, inconsistent preference information with respect to the real (unknown) preference model may lead to a wrong estimation of the preference model. In our context, we assume that we will collect tens or hundreds of input preferences, increasing then the chance of collecting inconsistent data. Therefore, we adopted a machine learning approach based on SVM [12, 15] where the input preference information is translated into soft constraints. Alternatively, one could adopt a Bayesian approach [1, 11, 29, 30] that represents the user preferences with a probability distribution over the feasible parameters of the value function, and the input preference information is used to update the probability distribution.

3 FORMAL SETTINGS

Order: Let O be the input set of orders from either N or P . In our context, an *order* $\alpha \in O$ is an atomic element representing a set of items that has to be transported from a pickup location to a drop-off location.

Journey: We define a *journey* (of a truck) to be a non-empty set of orders $J \subseteq O$. A journey is defined to be *collaborative* if it contains at least one order of N and at least one order of P . Otherwise it is defined to be *non-collaborative*. Each journey J is evaluated with a

vector $K(J) = (K_1(J), \dots, K_5(J)) \in \mathbb{R}^5$ representing the following KPIs:

- Monetary cost, which includes the driver compensation and the fuel cost.
- Total journey distance in kilometres.
- Total journey time in hours.
- CO2 emissions of the journey.
- Empty running, i.e., distance travelled in kilometres without any load.

Preference model: We define

$$u_w(x) = w \cdot x = \sum_{i=1}^5 w_i x_i \quad (1)$$

to be the value function of a decision-maker parametrised with respect to the *weights vector* w , where x is a KPIs vector and $w \in \mathcal{U}$ and $\mathcal{U} = \{w \in \mathbb{R}^5 : w_i \geq 0, \sum_{i=1}^5 w_i = 1\}$.

The (unknown) weights vector w defines the trade-offs of a decision-maker with respect to the KPIs. An estimation of the weights vector w defines then a value function $u_w(x)$. The simple structure of this value function may not precisely describe the real user preferences. However, this was a deliberate choice aimed to have a simple objective function that can be used for a fast computation of personalised solutions with an optimiser. Note that the goal is to minimise the KPIs and thus to minimise the objective function. We assume that the KPIs are independent, so that none is a function of the others. (If this were the case, there would be an infinite number of utility functions representing the same preferences; we could solve this issue by eliminating appropriate KPIs.)

Goal 1: We want to learn two weights vectors w^n and w^p , one for N and one for P, and a common weights vector w^c which considers the preferences of both N and P.

Solution: We define a *solution* as a non-empty set of journeys. A solution is defined to be *collaborative* if it is optimised with an objective function considering the weights vector and the input orders of all the partners involved. A solution is defined to be *non-collaborative* if it is optimised considering the weights vector and the input orders of a single partner. The objective function of the optimiser is the weighted sum of the sum of the KPIs associated with the journeys planned in a solution. In our context, an optimiser computes a collaborative solution for N and P, using the weights vector w^c , and two non-collaborative solutions, one for N and one for P, using the weights vectors w^n and w^p respectively.

Goal 2: Let \mathcal{Z}^c be the set of collaborative journeys of the collaborative solution, i.e., the set of journeys of the collaborative solution containing at least one order from N and one order from P. The system will show to the decision-makers only the collaborative journeys \mathcal{Z}^c . For a collaborative journey $J^c \in \mathcal{Z}^c$, we want to compute an estimated non-collaborative KPIs vector $E(J^c) = (E_1(J^c), \dots, E_5(J^c)) \in \mathbb{R}^5$ to estimate the improvement or disimprovement in the KPIs arising from the collaborative journey J^c . The estimated non-collaborative KPIs vector $E(J^c)$ will then be shown to the decision-makers along with $K(J)$ to facilitate the decision, and the decision-makers' preference between $K(J)$ and $E(J^c)$ will be used to learn their weights vector. It might be argued

that the entire collaborative solution involving all the journeys should be evaluated with respect to a non-collaborative solution rather than evaluating every single journey. However, the industrial partners of Logistar wanted to have the possibility of evaluating each individual collaborative journey.

4 PREFERENCE LEARNING

The preference learning method considered in this work is based on the observations of the preferences of the decision-makers with respect to the KPIs vectors evaluating the collaborative journeys. For each collaborative journey $J^c \in \mathcal{Z}^c$ computed by the optimiser, the system shows to the decision-makers the details of the journey, a KPIs vector $K(J^c)$ evaluating the journey, and an estimated KPIs vector $E(J^c)$ evaluating a non collaborative delivery for the orders in J^c . The purpose of the latter KPIs vector is to evaluate the improvement or disimprovement in the KPIs arising from the collaboration. The preferences of the decision-makers arising from the comparison of $K(J^c)$ with $E(J^c)$ will be used to generate constraints for an SVM algorithm estimating the value functions of the two decision-makers.

Constraints generation: The two decision-makers N and P need to either accept or reject each collaborative journey, and also express a motivation for their decision from a predefined list (see Table 1). We want to distinguish three scenarios with the selected motivation:

- (1) The decision was made considering the KPIs.
- (2) The decision was made considering mainly a specific KPI.
- (3) The decision was made for other reasons.

Scenario (1). Suppose that a decision-maker accepts a collaborative journey selecting *better collaborative KPIs overall* as motivation. Since the motivation is related to the KPIs vector, we translate this input into the following (soft) constraints

$$w \cdot (K(J^c) - E(J^c)) \leq 0, \quad (2)$$

(using the notation introduced in Equation 1) where $K(J^c)$ is the KPIs vector associated with the collaborative journeys, $E(J^c)$ is the estimated non-collaborative KPIs vector, and w is the (unknown) weights vector of the decision-maker. If, instead, the decision-maker rejects the collaborative solution selecting *better non-collaborative KPIs overall* as motivation, then we learn the constraint $w \cdot (E(J^c) - K(J^c)) \leq 0$.

Scenario (2). In this case, we learn the same (soft) constraint as scenario (1) (since the decision was made considering the KPIs) and a further set of (soft) constraints to highlight the importance of the main KPI considered. Suppose that a decision-maker accepts a collaborative journey and that the main KPI considered has index i . The constraints we learn in this case are Equation 2 (scenario (1)), and also the constraint:

$$w_i(K_i(J^c) - E_i(J^c)) + w_j(K_j(J^c) - E_j(J^c)) \leq 0 \text{ for all } j \neq i. \quad (3)$$

If the decision-maker rejects the collaborative journey, we learn the same constraints but exchanging $K(J^c)$ and $E(J^c)$. Intuitively, the set of constraints given by Equation 3 means that the gain of the i -th KPI is more important than any other eventual loss in the remaining KPIs.

Scenario (3). In this case we do not learn any constraint since we assume that the decision was not related to the KPIs vector.

Table 1: Possible motivations shown to the user for their acceptance or rejection of a collaborative journey, and the corresponding scenarios described in Section 4.

Acceptance	Rejection	Scenario
Better collaborative KPIs overall	Better non-collaborative KPIs overall	(1)
Better collaborative cost	Better non-collaborative cost	(2)
Better collaborative distance	Better non-collaborative distance	(2)
Better collaborative time	Better non-collaborative time	(2)
Better collaborative CO2 emissions	Better non-collaborative CO2 emissions	(2)
Better collaborative empty running	Better non-collaborative empty running	(2)
Other reasons	Other reasons	(3)

Example 4.1. Suppose we have only three KPIs, e.g., cost in euro, distance in kilometres and time in hours, and suppose that a decision-maker accepts a collaborative journey selecting *better collaborative cost* as motivation (scenario(2)). Let $K(J^c) = (1000, 100, 2.5)$ and $E(J^c) = (1200, 80, 3)$. Then we learn:

- (a) $(w_1, w_2, w_3) \cdot (1000, 100, 2.5) \leq (w_1, w_2, w_3) \cdot (1200, 80, 3)$
- (b) $(w_1, w_2) \cdot (1000, 100) \leq (w_1, w_2) \cdot (1200, 80)$
- (c) $(w_1, w_3) \cdot (1000, 2.5) \leq (w_1, w_3) \cdot (1200, 3)$

Constraint (a) corresponds to Equation 2. Constraints (b) and (c) corresponds to Equation 3. Constraint (b) can be interpreted as: the collaborative KPIs are preferred even if we ignore the loss in time of $E(J^c)$ with respect to $K(J^c)$. Constraint (c) does not reduce the set of possible weights vectors since it is always verified.

Estimation of the weights vectors: To estimate the weights vector of a specific decision-maker, we use an SVM algorithm named Ranking SVM [12, 15]. Roughly speaking, the idea is to compute the hyperplane containing the origin that produces the largest margin (distance) from a set of preference points representing input preference constraints. In our context, the preference points are derived from the inequalities learned observing the decision-maker's preferences concerning the KPIs shown by the system, and the normal vector of the output hyperplane represents the estimated decision-maker's trade-offs between the KPIs.

Ranking SVM is defined by a quadratic optimisation problem with linear soft constraints. Formally, the optimisation problem to solve is the following:

$$\arg \min_{\omega, \xi_k} \frac{1}{2} \|\omega\|^2 + C \sum_{k=1}^N \xi_k \quad (4a)$$

subject to

$$1 + \omega \cdot \lambda_k - \xi_k \leq 0 \quad \forall k \quad (4b)$$

$$\xi_k \geq 0 \quad \forall k \quad (4c)$$

$$\omega_i \geq 0 \quad \forall i \in \{1, \dots, 5\}. \quad (4d)$$

where $\omega \in \mathbb{R}^5$ is the normal of the hyperplane estimated by SVM and N is the number of constraints. This method is very similar to that described by Equation 13 of [15]. The differences are in the normal ω of the hyperplane, which in our case is non-negative (Equation 4d) since we do not want to penalise any KPI, and in the sign of λ_k (Equation 4b) since our goal is to minimise the value function of the decision-maker (and not to maximise it). ξ_k is the slack variable used to assign the penalty to the k -th constraint when

it is violated. C is a constant scaling the penalty of inconsistent constraints; a larger value of C corresponds to assigning a higher penalty. λ_k is the coefficients vector of the k -th constraint derived from the input preference information of the decision-makers as described below.

To generate the values λ_k , we first need to normalise the values range of every component of all the KPIs vectors $K(J^c)$ and $E(J^c)$, dividing it by the corresponding maximum value. This step is required otherwise the maximisation of the margin performed by the SVM algorithm could focus mainly on the KPIs with a larger scale. We consider then a vector $M = (M_1, \dots, M_5) \in \mathbb{R}^5$ where M_i is the maximum value of the i -th KPI of a database of previous instances. Thereafter, we define the vector λ_k by scaling the i -th value of the k -th constraint learned (whose structure is defined by Equation 2 or Equation 3) by the reciprocal of M_i , for all $i \in \{1, \dots, 5\}$.

Example 4.2. Suppose we have shown to a decision-maker two collaborative journeys J' and J'' , from which we derived two constraints $w \cdot (K(J') - E(J')) \leq 0$ and $w \cdot (E(J'') - K(J'')) \leq 0$. Then the constraints for the SVM problem represented by Equation 4b are

$$1 + \omega \cdot \left(\frac{K_1(J') - E_1(J')}{M_1}, \dots, \frac{K_5(J') - E_5(J')}{M_5} \right) - \xi_1 \leq 0 \quad (5a)$$

and

$$1 + \omega \cdot \left(\frac{E_1(J'') - K_1(J'')}{M_1}, \dots, \frac{E_5(J'') - K_5(J'')}{M_5} \right) - \xi_2 \leq 0. \quad (5b)$$

Let $M^{-1} = (\frac{1}{M_1}, \dots, \frac{1}{M_5})$. The normal ω of the hyperplane computed by the SVM algorithm is related to the constraints defined on the rescaled KPIs. Thus the estimated weights vector w of a decision-maker defining their value function $u_w(x)$ for a KPIs vector x can be computed with the following transformation $f: \mathbb{R}^5 \rightarrow \mathbb{R}$:

$$w = f(\omega) = \frac{\omega \circ M^{-1}}{\omega \cdot M^{-1}} \quad (6)$$

where $\omega \circ M^{-1}$ is the pointwise product (i.e., the Hadamard product) between ω and M^{-1} . The normalisation $\frac{1}{\omega \cdot M^{-1}}$ of Equation 6 is required to obtain a weights vector w with $\sum_{i=1}^5 w_i = 1$, so that $w \in \mathcal{U}$.

The weights vectors w^n and w^p of the two decision-makers N and P will then be estimated by transforming with Equation 6 the two outputs ω^n and ω^p of our SVM algorithm executed considering the constraints derived by the preferences of N and P, respectively.

The joint weights vector w^c considering the preferences of two decision-makers is estimated by transforming with Equation 6 the midpoint of the segment connecting ω^n and ω^p :

$$w^c = f\left(\frac{\omega^n + \omega^p}{2}\right). \quad (7)$$

The resulting weights vector can be interpreted as a fair set of trade-offs among the KPIs considering the preferences of the two decision-makers involved. Thus the corresponding value function $u_{w^c}(\cdot)$ can be used as the objective function to compute collaborative solutions. Supposing we had $n > 2$ decision-makers, one could compute the joint weights vector as the transform f of center of gravity of the normal vectors ω^h of the hyper-planes h estimated for every decision-makers, i.e., $w^c = f\left(\frac{\sum_{h=1}^n \omega^h}{n}\right)$.

Without any preference information, i.e., when we do not have received any feedback from the decision-makers regarding the KPIs vectors, we compute the collaborative and non-collaborative solutions with a weights vector minimising the cost.

5 COMPUTATION OF $E(J^c)$

Our preference elicitation method learns the decision-makers' weights vector based on their preferences between the KPIs vector $K(J^c)$ evaluating a collaborative delivery J^c , and an estimated KPIs vector $E(J^c)$ evaluating a corresponding non-collaborative delivery plan for the orders in J^c . The idea is that the KPIs vector $E(J^c)$ will be considered by the decision-maker to evaluate the improvement or disimprovement in the KPIs arising from the collaboration. However, the computation of $E(J^c)$ is not straightforward since the non-collaborative deliveries for the orders in J^c planned in the non-collaborative solutions may be assigned to several non-collaborative journeys also delivering other orders not in J^c . Thus, since the KPIs vector computed by the optimiser evaluates journeys and not single orders, we may not have a set of non-collaborative KPIs vectors that can be directly used to evaluate a non-collaborative plan for the orders in J^c .

Example 5.1. Suppose we have two orders $N1$ and $N2$ from N , and two orders $P1$ and $P2$ from P . Suppose that the optimiser finds two distinct collaborative journeys, J_1^c delivering $N1$ and $P1$ with KPIs vector $K(J_1^c)$, and J_2^c delivering $N2$ and $P2$ with KPIs vector $K(J_2^c)$. The question is: how can we evaluate the value of the collaboration of each single journey? Consider, for example, the collaborative delivery J_1^c . We want to estimate a KPIs vector $E(J_1^c)$ for a non-collaborative delivery for the orders in J_1^c and evaluate the collaboration by comparing $K(J_1^c)$ with $E(J_1^c)$. Suppose that the computation of a non-collaborative solution leads to two non-collaborative KPIs evaluating a non collaborative delivery for $N1$ and $N2$, and a non-collaborative delivery for $P1$ and $P2$. The issue is that we are interested in a non-collaborative KPIs vector evaluating $N1$ and $P1$ to estimate the improvement or disimprovement in the KPIs arising from the collaborative journey J_1^c , but we do not have such a KPIs vector since these orders are delivered by two distinct non-collaborative journeys along with other orders, i.e., $N2$ and $P2$.

Our idea is to split the KPIs of the non-collaborative journeys among the corresponding orders to evaluate a non-collaborative plan for a generic subset of orders, even if they belong to different

non-collaborative journeys. In fact, having a KPIs vector for every order, we can estimate the KPIs vector of any subset of orders by summing up the individual KPIs vectors.

Recall that \mathcal{Z}^c is the set of collaborative journeys of the collaborative solution computed by the optimiser with objective function u_{w^c} . Our goal is then to estimate a non-collaborative KPIs vector $E(\alpha) \in \mathbb{R}^5$ for each $\alpha \in J^c$ and $J^c \in \mathcal{Z}^c$, and to compute $E(J^c)$ as:

$$E(J^c) = \sum_{\alpha \in J^c} E(\alpha). \quad (8)$$

Let \mathcal{Z}^{nc} be the union of the journeys of the two non-collaborative solutions for N and P computed by the optimiser based on objective functions u_{w^n} and u_{w^p} , respectively. Suppose that the delivery of every order planned with the collaborative journeys in \mathcal{Z}^c is also planned in a non-collaborative journey in \mathcal{Z}^{nc} . We then compute $E(\alpha)$ as follows

$$E(\alpha) = \mu(\alpha)K(J^{nc}(\alpha)), \quad (9)$$

where $J^{nc}(\alpha)$ is the unique non-collaborative journey in \mathcal{Z}^{nc} delivering α and with KPIs vector $K(J^{nc}(\alpha))$, and $\mu(\alpha)$ is a non-negative weight. We call $\mu(\alpha)$ the *Shapley weight* of α , and it is computed using a method based on the *Shapley value* described in Section 6. Roughly speaking, the weights computed by this method attempt to split fairly the cost of a journey J^{nc} among the corresponding orders, considering the extra kilometres required to deliver each single order. Thus we have that $\sum_{\alpha \in J^{nc}} \mu(\alpha) = 1$. Note that $\sum_{\alpha \in J^c} E(\alpha)$ might not be equal to $K(J^c)$ since J^c is a collaborative journey, and the values $E(\alpha)$ are computed splitting the corresponding non-collaborative journeys $J^{nc}(\alpha)$.

In our work we also had to address the following issue. The delivery of an order planned with a collaborative journey in \mathcal{Z}^c may not be planned in any of the non-collaborative journeys in \mathcal{Z}^{nc} . This can happen for several reasons; for example, there may not be enough trucks to deliver all the orders or there may not be a feasible solution matching all the delivery time windows of the input set of orders. Thus we may not be able to compute $E(J^c)$ with Equation 8. To overcome this issue we designed a simple fix described as follows. Let $J_x^c \subseteq J^c$ be the set of orders scheduled in the non-collaborative solution, i.e., the set of orders $\alpha \in J^c$ such that there exists $J^{nc} \in \mathcal{Z}^{nc}$ containing α . We compute the estimated KPIs vector $E(J^c)$ evaluating the non-collaborative deliveries for the orders in J^c as the rescaled sum of the estimated KPIs of the non-collaborative deliveries of the orders $\alpha \in J_x^c$:

$$E(J^c) = \frac{|J^c|}{|J_x^c|} \sum_{\alpha \in J_x^c} E(\alpha) = \frac{|J^c|}{|J_x^c|} \sum_{\alpha \in J_x^c} \mu(\alpha)K(J^{nc}(\alpha)). \quad (10)$$

The scale factor $\frac{|J^c|}{|J_x^c|}$ is used to estimate the KPIs of non-collaborative deliveries of $|J^c|$ orders having the estimation of only $|J_x^c|$ orders. We assume that a scenario with $J_x^c = \emptyset$ will not occur; however, in our implementation we also cover the case $J_x^c = \emptyset$ by computing $E(J^c)$ as the average KPIs vector of the non-collaborative journeys, i.e., $E(J^c) = \frac{1}{|\mathcal{Z}^{nc}|} \sum_{J^{nc} \in \mathcal{Z}^{nc}} K(J^{nc})$. This is just to cover all the possible scenarios, but an average of the KPIs of the non-collaborative journeys may not return realistic values for $E(J^c)$.

Example 5.2. Consider three input orders $P1$, $P2$ and $P3$ from P , and three input orders $N1$, $N2$ and $N3$ from N . Suppose that the optimiser computes a collaborative solution with the following

collaborative journeys evaluated with KPIs cost in euro, distance in kilometres and time in hours:

- Journey J_1^c with KPIs [900, 80, 2] delivering $P1$ and $N1$.
- Journey J_2^c with KPIs [800, 100, 3] delivering $P2$ and $N2$.
- Journey J_3^c with KPIs [1000, 115, 4] delivering $P3$ and $N3$.

Also, suppose that the optimiser computes the following non-collaborative solution:

- Journey J_1^{nc} with KPIs [900, 100, 2.5] delivering $P1$ and $P2$.
- Journey J_2^{nc} with KPIs [1000, 70, 2] delivering $N1$.
- Journey J_3^{nc} with KPIs [800, 100, 3] delivering $N2$ and $N3$.

Note that we are supposing that the optimiser didn't find a non-collaborative journey for order $P3$.

The first step of our procedure is to compute the Shapley weight of the orders of the non-collaborative journeys J_k^{nc} with $k \in \{1, 2, 3\}$. (In this example we do not show how to compute the Shapley weights; see Section 6 for an example). Then we use Equation 9 to estimate the KPIs vectors $E(\alpha)$ evaluating a non-collaborative delivery for the orders α planned in the non-collaborative solution:

- $\mu(P1) = 0.4$; $E(P1) = [0.4 \cdot 900, 0.4 \cdot 100, 0.4 \cdot 2.5] = [360, 40, 1]$
- $\mu(P2) = 0.6$; $E(P2) = [0.6 \cdot 900, 0.6 \cdot 100, 0.6 \cdot 2.5] = [540, 60, 1.5]$
- $\mu(N1) = 1$; $E(N1) = [1 \cdot 1000, 1 \cdot 70, 1 \cdot 2] = [1000, 70, 2]$
- $\mu(N2) = 0.4$; $E(N2) = [0.4 \cdot 800, 0.4 \cdot 100, 0.4 \cdot 3] = [320, 40, 1.2]$
- $\mu(N3) = 0.6$; $E(N3) = [0.6 \cdot 800, 0.6 \cdot 100, 0.6 \cdot 3] = [680, 60, 1.8]$

Finally, we estimate the KPIs vectors $E(J_k^c)$ for each collaborative journey J_k^c using Equation 10:

- $E(J_1^c) = [\frac{1}{3}(360+1000), \frac{1}{3}(40+60), \frac{1}{3}(1+1.5)] = [1360, 100, 2.5]$
- $E(J_2^c) = [\frac{1}{3}(540+320), \frac{1}{3}(40+60), \frac{1}{3}(1.5+1.2)] = [860, 100, 2.7]$
- $E(J_3^c) = [\frac{2}{3}(680), \frac{2}{3}(60), \frac{2}{3}(1.8)] = [1360, 120, 3.6]$

6 SPLITTING THE KPIS OF A JOURNEY AMONG THE CORRESPONDING ORDERS

A well-known method used to split the cost of a collaboration among a set n of participants *Shapley value* (see, e.g., [24]). This method is based on a cost function $v : 2^N \rightarrow \mathbb{R}$ with $v(\emptyset) = 0$, which defines the total cost of any collaboration of the set N of participants. Our intention is to divide the costs of a journey among the orders planned for delivery within the journey itself with the purpose of estimating a KPIs vector for each order. Thus, we interpret the set of orders defining a non-collaborative journey J^{nc} as the set of participants for the computation of the Shapley value. Regarding the cost function, we consider a function $TSP(t, I)$ which returns the total journey distance computed solving a travelling salesman problem with same start and end point t (which is the location of the truck assigned to J^{nc}), and with input locations defined as the pick up and drop locations of the orders $I \subseteq J^{nc}$. With the computation of $TSP(t, I)$, we consider precedence constraints for the pick up and drop location of each order in I . However, we ignore the capacity constraint of the corresponding assigned truck.

The Shapley value of an order $\alpha \in J^{nc}$ is then defined as:

$$\phi(\alpha) = \sum_{I \subseteq J^{nc} \setminus \{\alpha\}} \frac{|I|!(|J^{nc}| - |I| - 1)!}{|J^{nc}|!} (TSP(t, I \cup \{\alpha\}) - TSP(t, I)). \quad (11)$$

The Shapley value of an order $\alpha \in J^{nc}$ can be interpreted as the average increase of total distance caused by the inclusion of α in different journeys delivering all the possible permutations of the orders $J^{nc} \setminus \{\alpha\}$. A fundamental property derived from the theory behind the Shapley value is that $TSP(t, J^{nc}) = \sum_{\alpha \in J^{nc}} \phi(\alpha)$. Roughly speaking, the higher the Shapley value of an order is, the higher the extra distance that a truck has to travel to include the delivery of the order in the journey. We adopted this method since it seems reasonable to assign a higher fraction of KPIs to orders with higher Shapley value. The pseudocode of our implementation to compute $\phi(\alpha)$ for all the orders of a journey J^{nc} with starting point t is shown in Algorithm 1.

Once we have computed the Shapley value for every order α of a non-collaborative journey J^{nc} , we compute the corresponding Shapley weight $\mu(\alpha)$ as:

$$\mu(\alpha) = \frac{\phi(\alpha)}{\sum_{\beta \in J^{nc}} \phi(\beta)} \quad (12)$$

The Shapley weights are then used in Equation 9 to split the KPIs of a non-collaborative journey J^{nc} among the corresponding orders.

Algorithm 1 Shapley Value

```

1: procedure  $SV(J^{nc}, t)$ 
2:    $n \leftarrow |J^{nc}|$ 
3:    $S_0 \leftarrow \{\emptyset\}$ 
4:    $S_i \leftarrow \emptyset$  for each  $i \in [1, \dots, n-1]$ 
5:    $L(\emptyset) \leftarrow 0$ 
6:    $\phi(\alpha) \leftarrow 0$  for each  $\alpha \in J^{nc}$ 
7:   for  $m \in [0, \dots, n-1]$  do
8:      $c \leftarrow \frac{m!(n-m-1)!}{m!}$ 
9:     for  $I \subseteq J^{nc}$  with  $|I| = m+1$  do
10:       $S_{m+1} \leftarrow S_m \cup \{I\}$ 
11:      if  $\exists \alpha, \beta \in I$  with same pick up and drop location
12:      then
13:         $L(I) \leftarrow L(I \setminus \{\alpha\})$ 
14:      else
15:         $L(I) \leftarrow (TSP(t, I))$ 
16:      for  $I \in S_m$  do
17:        for  $\alpha \in J^{nc} \setminus I$  do
18:           $\phi(\alpha) \leftarrow \phi(\alpha) + c(L(I \cup \alpha) - L(I))$ 
19:   return  $\phi$ 

```

Example 6.1. Consider for example a non-collaborative journey composed by three orders $J^{nc} = \{\alpha_1, \alpha_2, \alpha_3\}$ with corresponding pick up and drop locations (p_i, d_i) shown in Figure 1. Suppose that t is the start and end location of the truck assigned to J^{nc} , and suppose that the distance between two locations of the map shown in Figure 1 is the corresponding Manhattan distance. We then get the following values of $TSP(t, I)$ for each subset I of J^{nc} :

$$\begin{aligned}
\text{TSP}(t, \emptyset) &= 0 \\
\text{TSP}(t, \{\alpha_1\}) &= 14 \\
\text{TSP}(t, \{\alpha_2\}) &= 14 \\
\text{TSP}(t, \{\alpha_3\}) &= 14 \\
\text{TSP}(t, \{\alpha_1, \alpha_2\}) &= 14 \\
\text{TSP}(t, \{\alpha_2, \alpha_3\}) &= 20 \\
\text{TSP}(t, \{\alpha_1, \alpha_3\}) &= 20 \\
\text{TSP}(t, \{\alpha_1, \alpha_2, \alpha_3\}) &= 20
\end{aligned} \tag{13}$$




Figure 1: Delivery map of three orders $J^{nc} = \{\alpha_1, \alpha_2, \alpha_3\}$ with start location t (circle) corresponding pick up (triangle) and drop (square) location (p_i, d_i)

Following Equation 11 we get:

$$\begin{aligned}
\phi(\alpha_1) &= \phi(\alpha_2) = \frac{1}{3} \text{TSP}(t, \{\alpha_1\}) \\
&+ \frac{1}{6} (\text{TSP}(t, \{\alpha_1, \alpha_2\}) - \text{TSP}(t, \{\alpha_2\})) \\
&+ \frac{1}{6} (\text{TSP}(t, \{\alpha_1, \alpha_3\}) - \text{TSP}(t, \{\alpha_3\})) \\
&+ \frac{1}{3} (\text{TSP}(t, \{\alpha_1, \alpha_2, \alpha_3\}) - \text{TSP}(t, \{\alpha_2, \alpha_3\})) \\
&= \frac{14}{3} + \frac{0}{6} + \frac{6}{6} + \frac{0}{3} = \frac{34}{6}.
\end{aligned} \tag{14}$$

$$\begin{aligned}
\phi(\alpha_3) &= \frac{1}{3} \text{TSP}(t, \{\alpha_3\}) \\
&+ \frac{1}{6} (\text{TSP}(t, \{\alpha_1, \alpha_3\}) - \text{TSP}(t, \{\alpha_1\})) \\
&+ \frac{1}{6} (\text{TSP}(t, \{\alpha_2, \alpha_3\}) - \text{TSP}(t, \{\alpha_2\})) \\
&+ \frac{1}{3} (\text{TSP}(t, \{\alpha_1, \alpha_2, \alpha_3\}) - \text{TSP}(t, \{\alpha_1, \alpha_2\})) \\
&= \frac{14}{3} + \frac{6}{6} + \frac{6}{6} + \frac{6}{3} = \frac{52}{6}.
\end{aligned} \tag{15}$$

Thus, following Equation 12 we get:

$$\mu_1 = \mu_2 = \frac{\phi(\alpha_1)}{\phi(\alpha_1) + \phi(\alpha_2) + \phi(\alpha_3)} = 0.283, \tag{16}$$

$$\mu_3 = \frac{\phi(\alpha_3)}{\phi(\alpha_1) + \phi(\alpha_2) + \phi(\alpha_3)} = 0.433. \tag{17}$$

Note that with a straightforward method to compute the weights such as $\mu_i = \frac{\text{TSP}(t, \{\alpha_i\})}{\sum_{j \in \{1, 2, 3\}} \text{TSP}(t, \{\alpha_j\})}$ we would have got the same weight for each order in this example.

The computational complexity of the Shapley value is exponential with respect to the number of orders composing a Journey. In fact, following Equation 11, for each non-collaborative journey J^{nc} we should solve $2^{|J^{nc}|}$ TSPs, one for each subset of J^{nc} .

In our use case we may have more than ten orders delivered with a single journey, and the computational time required for an exact computation of the Shapley value with this amount of orders is too slow (see Section 7). However, we have several orders with the same pick up and drop location. We thus compute the Shapley values considering orders with the same pick up and drop location as if they were one single order, and equally split the corresponding Shapley value among the orders grouped together. The resulting weights will be different, however dividing the KPIs between stretches of journeys and equally splitting the KPIs among orders with the same pick up and drop location sounds a reasonable approach which could also result in a more intuitive solution. The pseudocode of this variation is shown in Algorithm 2.

Example 6.2. Consider again the setup of Example 6.1 shown in Figure 1. If we group together the orders α_1 and α_2 and we consider them as if they were one unique order α_{12} , we obtain the following Shapley values:

$$\begin{aligned}
\phi(\alpha_{12}) &= \frac{1}{2} \text{TSP}(t, \{\alpha_{12}\}) \\
&+ \frac{1}{2} (\text{TSP}(t, \{\alpha_{12}, \alpha_3\}) - \text{TSP}(t, \{\alpha_3\})) \\
&= \frac{14}{2} + \frac{6}{2} = 10
\end{aligned} \tag{18}$$

$$\begin{aligned}
\phi(\alpha_3) &= \frac{1}{2} \text{TSP}(t, \{\alpha_3\}) \\
&+ \frac{1}{2} (\text{TSP}(t, \{\alpha_{12}, \alpha_3\}) - \text{TSP}(t, \{\alpha_{12}\})) \\
&= \frac{14}{2} + \frac{6}{2} = 10
\end{aligned} \tag{19}$$

If we equally split the Shapley value of the orders grouped together, we get $\phi(\alpha_1) = \phi(\alpha_2) = \frac{\phi(\alpha_{12})}{2} = 5$ and the Shapley weights are then:

$$\mu_1 = \mu_2 = \frac{\phi(\alpha_1)}{\phi(\alpha_1) + \phi(\alpha_2) + \phi(\alpha_3)} = 0.25, \tag{20}$$

$$\mu_3 = \frac{\phi(\alpha_3)}{\phi(\alpha_1) + \phi(\alpha_2) + \phi(\alpha_3)} = 0.5. \tag{21}$$

If also Algorithm 2 was too slow for a real application, one could consider the Monte Carlo approximation for the Shapley value presented in [2].

7 EXPERIMENTAL RESULTS

We briefly summarise some results of our experimental testing. All experiments were performed on a computer facilitated by a Core i5 2.70 GHz processor and 8 GB RAM. We used the Java library Jsprit to compute the TSP problems, and we tested our algorithms with randomly generated instances. Table 2 and Table 3 show the time performance of Algorithm 1 and Algorithm 2 with respect to the number of orders and the number of pairs of pick up and drop locations. As we can see, the execution time of Algorithm 1 grows exponentially with respect to the number of orders, and it is not very

Algorithm 2 Simplified Shapley Value

```

1: procedure  $SSV(J^{nc}, t)$ 
2:    $\phi(\alpha) \leftarrow 0$  for each  $\alpha \in J^{nc}$ 
3:    $H(\alpha) \leftarrow \emptyset$  for each  $\alpha \in J^{nc}$ 
4:    $J^{nc'} \leftarrow \emptyset$ 
5:   for  $\alpha \in J^{nc}$  do
6:     if  $\exists \beta \in J^{nc'}$  with same pick up and drop location of  $\alpha$ 
       then
7:        $H(\beta) \leftarrow H(\beta) \cup \alpha$ 
8:     else
9:        $J^{nc'} \leftarrow J^{nc'} \cup \alpha$ 
10:   $n \leftarrow |J^{nc'}|$ 
11:   $\mathcal{S}_0 \leftarrow \{\emptyset\}$ 
12:   $\mathcal{S}_i \leftarrow \emptyset$  for each  $i \in [1, \dots, n-1]$ 
13:   $L(\emptyset) \leftarrow 0$ 
14:  for  $m \in [0, \dots, n-1]$  do
15:     $c \leftarrow \frac{m!(n-m-1)!}{n!}$ 
16:    for  $I \subseteq J^{nc'}$  with  $|I| = m+1$  do
17:       $\mathcal{S}_{m+1} \leftarrow \mathcal{S}_{m+1} \cup \{I\}$ 
18:       $L(I) \leftarrow (\text{TSP}(t, I))$ 
19:    for  $I \in \mathcal{S}_m$  do
20:      for  $\alpha \in J^{nc'} \setminus I$  do
21:         $\phi(\alpha) \leftarrow \phi(\alpha) + c(L(I \cup \alpha) - L(I))$ 
22:  for  $\alpha \in J^{nc}$  do
23:    if  $H(\alpha) \neq \emptyset$  then
24:      for  $\beta \in H(\alpha)$  do
25:         $\phi(\beta) \leftarrow \frac{\phi(\alpha)}{|H(\alpha)|+1}$ 
26:     $\phi(\alpha) \leftarrow \frac{\phi(\alpha)}{|H(\alpha)|+1}$ 
return  $\phi$ 

```

sensitive with respect to the number of pairs of pick up and drop locations (which cannot be greater than the number of orders since each order is associated with only one pair of locations). On the other hand, the execution time of Algorithm 2 grows exponentially with respect to the number of pairs of pick up and drop locations, and it is not very sensitive with respect to the number of orders. This is due to the computational burden of the Shapley values, which is strictly related to the number of orders for Algorithm 1, and to the number of pairs of pick up and drop off locations for Algorithm 2. In our context we should not have more than six pairs of pick up and drop off locations for each journey, but we could have more than ten orders. Therefore Algorithm 2 seems to be the best choice since it would improve the time response of the system.

8 CONCLUSIONS

With this work we presented a preference learning method for collaborative Vehicle Routing Problem developed within the European project LOGISTAR. In particular, we have shown how we can learn the preferences of the decision-makers involved in collaborative journeys, by observing their interactions with the system, and translating input preference information into constraints for an SVM algorithm estimating a preference model. The preference model considered is the weighted sum of a set of KPIs evaluating the collaborative journeys, and it defines the trade-offs over the

Table 2: Execution time of Algorithm 1 varying the number of orders and the number of pick up and drop locations

N. orders	N. locations pairs	Time [s]
2	1	2.8
2	2	3.6
4	1	5.1
4	4	5.5
6	1	7.6
6	6	9.3
8	1	23.2
8	8	27.2
10	1	80.2
10	10	86.7

Table 3: Execution time of Algorithm 2 varying the number of pick up and drop locations and the number of orders.

N. locations pairs	N. Orders	Time [s]
2	2	3.2
2	20	3.6
4	4	4.8
4	40	5.2
6	6	10.1
6	60	11.8
8	8	22.3
8	80	24.3
10	10	83.9
10	100	85.2

KPIs. With our approach we learn a preference model for every partner involved, and also a preference model considering the preferences of all the partners. We also have shown a method based on the Shapley value to evaluate the improvement or disimprovement on the KPIs arising from the collaboration. This method computes an estimated non-collaborative KPIs vector which can be compared with the KPIs vector associated with a collaborative journey with the purpose of evaluating the collaboration.

Future work could involve alternative techniques to translate the motivations of the decisions into preference constraints, the evaluation of different preference learning approaches such as Bayesian methods or robust ordinal regression, and methods for the estimation of joint preference models with more than two partners.

ACKNOWLEDGMENTS

This material is based upon works supported by the Science Foundation Ireland under Grant No. 12/RC/2289-P2 which is co-funded under the European Regional Development Fund, by the LOGISTAR project [14], which is funded by the European Commission under the Horizon 2020 programme. We are grateful also to our various partners on the LOGISTAR project.

REFERENCES

- [1] Craig Boutilier. 2002. A POMDP formulation of preference elicitation problems. In *Proceedings of AAAI/IAAI* 239–246.
- [2] Javier Castro, Daniel Gómez, and Juan Tejada. 2009. Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research* 36, 5 (2009), 1726–1730.
- [3] Jen Jen Chung, Carrie Rebhuhn, Connor Yates, Geoffrey A Hollinger, and Kagan Tumer. 2019. A multiagent framework for learning dynamic traffic management strategies. *Autonomous Robots* 43, 6 (2019), 1375–1391.
- [4] Sascha Dahl and Ulrich Derigs. 2011. Cooperative planning in express carrier networks—An empirical study on the effectiveness of a real-time Decision Support System. *Decision Support Systems* 51, 3 (2011), 620–626.
- [5] Bo Dai and Haoxun Chen. 2012. Mathematical model and solution approach for carriers' collaborative transportation planning in less than truckload transportation. *International Journal of Advanced Operations Management* 4, 1-2 (2012), 62–84.
- [6] Dave de Jonge, Filippo Bistaffa, and Jordi Levy. 2021. A Heuristic Algorithm for Multi-Agent Vehicle Routing with Automated Negotiation. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 404–412.
- [7] Erella Eisenstadt, Amiram Moshaiov, and Gideon Avigad. 2015. Co-evolution of strategies for multi-objective games under postponed objective preferences. In *2015 IEEE conference on computational intelligence and games (CGI)*. IEEE, 461–468.
- [8] Elena Fernández, Mireia Roca-Riu, and M Grazia Speranza. 2018. The shared customer collaboration vehicle routing problem. *European journal of operational research* 265, 3 (2018), 1078–1093.
- [9] Peter C Fishburn. 1970. *Utility theory for decision making*. Technical Report. Research analysis corp McLean VA.
- [10] Margaretha Gansterer and Richard F Hartl. 2018. Collaborative vehicle routing: a survey. *European Journal of Operational Research* 268, 1 (2018), 1–12.
- [11] Shengbo Guo and Scott Sanner. 2010. Real-time multiattribute bayesian preference elicitation with pairwise comparison queries. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 289–296.
- [12] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. 133–142.
- [13] S. Kaddani, D. Vanderpooten, J. M. Vanpeperstraete, and H. Aissi. 2017. Weighted sum model with partial preference information: Application to multi-objective optimization. *European Journal of Operational Research* 260, 2 (2017), 665–679.
- [14] LOGISTAR. 2021. <https://logistar-project.eu/>.
- [15] Mojtaba Montazery and Nic Wilson. 2016. Learning User Preferences in Matching for Ridesharing. In *ICAART (2)*. 63–73.
- [16] Enrique Mu and Milagros Pereyra-Rojas. 2017. Understanding the analytic hierarchy process. In *Practical decision making*. Springer, 7–22.
- [17] Narges Norouzi, R Tavakkoli-Moghaddam, M Ghazanfari, M Alinaghian, and A Salamatbakhsh. 2012. A new multi-objective competitive open vehicle routing problem solved by particle swarm optimization. *Networks and Spatial Economics* 12, 4 (2012), 609–633.
- [18] Roxana Rădulescu, Patrick Mannion, Diederik M Roijers, and Ann Nowé. 2020. Multi-objective multi-agent decision making: a utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020), 1–52.
- [19] H. Raiffa. 1968. *Decision analysis*. Addison-Wesley.
- [20] Diederik M Roijers, Shimon Whiteson, Frans A Oliehoek, et al. 2014. Linear support for multi-objective coordination graphs. In *AAMAS'14: PROCEEDINGS OF THE 2014 INTERNATIONAL CONFERENCE ON AUTONOMOUS AGENTS & MULTIAGENT SYSTEMS*. 1297–1304.
- [21] Thomas L Saaty. 1990. How to make a decision: the analytic hierarchy process. *European journal of operational research* 48, 1 (1990), 9–26.
- [22] A. Salo and R. P. Hämäläinen. 2010. Preference Programming – Multicriteria Weighting Models under Incomplete Information. In *Proceedings Handbook of Multicriteria Analysis*. Springer Berlin Heidelberg, 167–187.
- [23] Christoph Schneeweiss. 2012. *Distributed decision making*. Springer Science & Business Media.
- [24] Lloyd S Shapley. 2016. *A value for n-person games*. Princeton University Press.
- [25] Stefano Teso, Andrea Passerini, and Paolo Viappiani. 2016. Constructive preference elicitation by setwise max-margin learning. *arXiv preprint arXiv:1604.06020* (2016).
- [26] Federico Toffano, Paolo Viappiani, and Nic Wilson. 2021. Efficient Exact Computation of Setwise Minimax Regret for Interactive Preference Elicitation. In *20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*.
- [27] Federico Toffano and Nic Wilson. 2020. Minimality and comparison of sets of multi-attribute vectors. In *Ebook Series: Frontiers in Artificial Intelligence and Applications*. IOS Publishing, 913–920.
- [28] Omkarprasad S Vaidya and Sushil Kumar. 2006. Analytic hierarchy process: An overview of applications. *European Journal of operational research* 169, 1 (2006), 1–29.
- [29] Ivan Vendrov, Tyler Lu, Qingqing Huang, and Craig Boutilier. 2020. Gradient-based optimization for Bayesian preference elicitation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 10292–10301.
- [30] Paolo Viappiani and Craig Boutilier. 2020. On the equivalence of optimal recommendation sets and myopically optimal query sets. *Artificial Intelligence* (2020), 103328.
- [31] Ling Wang, Jian Chu, and Jun Wu. 2007. Selection of optimum maintenance strategies based on a fuzzy analytic hierarchy process. *International journal of production economics* 107, 1 (2007), 151–163.
- [32] Wenyu Zhang, Zixuan Chen, Shuai Zhang, Weirui Wang, Shuiqing Yang, and Yishuai Cai. 2020. Composite multi-objective optimization on a new collaborative vehicle routing problem with shared carriers and depots. *Journal of Cleaner Production* 274 (2020), 122593.