Preference learning for efficient bundle selection in horizontal transport collaborations

Steffen Elting[1,*], Jan Fabian Ehmke[1], and Margaretha Gansterer[2]

[1]Department of Business Decisions and Analytics, University of Vienna, Austria,
steffen.elting@univie.ac.at, jan.ehmke@univie.ac.at
[2]Department of Production Management and Logistics, University of Klagenfurt, Austria,
margaretha.gansterer@aau.at
[*]Corresponding author

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

March 4, 2025

Keywords: Routing, Preference learning, Collaborative transportation planning, Combinatorial auction, Vehicle routing problem with time windows

1 INTRODUCTION

To reduce routing costs and emissions in the home delivery sector, transport service providers (carriers) can collaborate by reallocating their delivery orders among them, such that all collaborators gain some profit compared to the setting in which they operate in isolation (Cleophas et al., 2019). Due to a lack of mutual trust, carriers prefer to keep the level of information sharing to a minimum. Combinatorial auctions have proven to strike a balance between high collaboration gains and limited information sharing (Berger and Bierwirth, 2010, Gansterer and Hartl, 2018b) as they account for the substitution and complementary interactions between the delivery orders that shall be reallocated. A carrier's marginal routing cost of fulfilling a combination of orders is often lower than the sum of costs of fulfilling each order from that combination alone (Vetschera et al., 2024).

At the heart of such combinatorial auctions is the Winner Determination Problem (WDP) (or Combinatorial Auction Problem (CAP)). This optimization problem computes the order-to-carrier assignment with the lowest overall routing costs. With m orders to be reallocated, there are $2^m - 1$ non-empty combinations of orders, called bundles. To obtain the coefficients for the WDP, all bids, i.e., bundle valuations, must be available to the auctioneer. In the setting of freight transportation for Attended Home Deliveries (AHD), a carrier's valuation of a bundle is the marginal routing cost of adding the bundles' orders to its already existing route plan. Computing this marginal cost entails solving the Vehicle Routing Problem with Time Windows (VRPTW), for which even finding any feasible solution is NP-hard (Savelsbergh, 1985). Therefore, to obtain all coefficients for the WDP, every carrier has to solve $2^m - 1$ complex optimization problems, which renders any exact solution approach too demanding and time-consuming for practical applications.

A remedy to this situation is presented in the Bundle Selection Problem (BuSP) (Gansterer and Hartl, 2018a, Rüther and Rieck, 2022) which limits the number of bundles that carriers must bid on with as little decrease in collaboration gain as possible. The auctioneer must decide which bundles to query the carriers' valuations for, without exceeding a maximum number of value queries. To find good queries, it searches through query candidates using a substitute

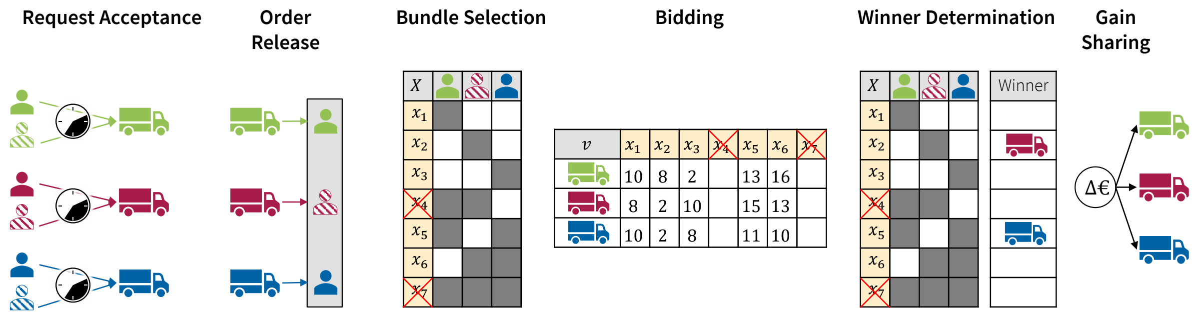

Figure 1 – *Example combinatorial auction-based order reassignment process using the one-shot approach for bundle selection. Of the 7 possible bundles that can be constructed from the 3 released orders, the auctioneer has not selected x_4 and x_7. Consequently, the three carriers do not report bids and thus, these bundles cannot be considered in the WDP.*

fitness function and a search algorithm. In this paper, we generalize existing approaches from the transportation literature (Gansterer and Hartl, 2018a, Rüther and Rieck, 2022), and apply methods from the field of preference learning (PL, Sandholm and Boutilier (2006), Fürnkranz and Hüllermeier (2010)) that are new for solving the BuSP in horizontal carrier collaboration. In particular, we consolidate the steps of **bundle selection** and **bidding** into a unified iterative process to increase the use of available information for decision-making. To that end, we demonstrate (1) the necessity of taking the WDP constraints into account already in the bundle selection stage and (2) that fitting parameterized bid estimators in a PL scheme can outperform the conventional "one-shot" approach to bundle selection.

2 METHODOLOGY

We consider an auction-based horizontal carrier collaboration that is based on the order reallocation scheme by Berger and Bierwirth (2010). Consider the example of an extended version of their 5-step in Figure 1. During **Request Acceptance**, a set of n carriers must attend to dynamically incoming customer requests by offering each customer a set of delivery time windows during which a delivery is feasible. Depending on the offered time windows and the customer's choice, a request is either discarded or accepted. After some cut-off time, no more requests are accepted, and in the **Order Release** phase, the carriers decide which of their accepted customers they want to release to the auction pool. As mentioned before, the number of order bundles grows exponentially with the size of the auction pool. This renders exhaustive bidding approaches impractical even for moderate collaborations. The auctioneer therefore **selects a subset of** L **bundles** to be biddable and **carriers must report bid values** for all selected bundles. We refer to this strictly sequential process of bundle selection and bidding as the one-shot approach. Using the carriers' bids, the auctioneer solves a **WDP** to find the optimal assignment of orders to carriers, making sure that (1) all orders are assigned to exactly one carrier, and (2) no carrier receives more than one of the available bundles to ensure integrity of the reported bids. Eventually, the cost savings that were generated through the re-assignment are shared among the participating carriers with an appropriate **Gain Sharing** mechanism.

Our main contribution lies in the integration of the Bundle Selection and Bidding steps using a PL mechanism. Rather than selecting all L bundles at once, we apply the following iterative algorithm: Initially, each carrier reports the marginal insertion cost for the bundle of orders that it released itself. This constitutes the base knowledge of the auctioneer, which it will use to approximate each carrier's preference function using a selected regression model. Then, as long as at least one bidder has answered fewer than L queries, the auctioneer will use the approximated preference functions to decide on a set of new value queries to ask each of the carriers. Bundles can be selected either completely independently, simply searching for the bundles that maximize each carrier's approximate valuation function. In our setting, the assumption of free disposal does not hold, therefore we also test two alternative selection approaches that consider combinations of bundles that exactly cover all orders in the auction pool. After receiving the carriers' responses (bids) to the selected queries (bundles), the auctioneer can update the regression models to find better queries in the next query round.

3 COMPUTATIONAL EXPERIMENTS

We simulate a 3-carrier collaboration in the city of Vienna: each carrier receives 50 customers in the request acceptance phase and releases 5 requests to the auction pool.

Our initial results show that neglecting the constraints of the downstream WDP in the query selection will decrease relative collaboration gains by 8 to 12 percentage points. Then, to evaluate the efficacy of the suggested machine-learning based PL mechanism, we first mimic the one-shot approach by using non-parametric models to predict a carrier's bundle valuation (Figure 2). The results are then compared to using parametric regression models like (regularized) linear regression (Figure 3), regression trees, and artificial neural networks. To evaluate, we compute the collaboration gain (CG) as the difference in total travel duration between the isolated and the collaborative planning scenario and divide this by the duration of the isolated scenario to obtain the relative CG (RCG). A random fitness function is used to establish a benchmark, so that the improvement compared to this benchmark can be computed (ΔRCG).

We limit the number of value queries L the auctioneer is allowed to ask to control the computational burden on the collaborating carriers. Multiple interesting patterns can be observed:

1. When (1) no (or only very loose) time window constraints are in place and (2) sufficient query-response pairs are available for fitting (i.e., L is large enough), then using the suggested iterative PL framework with parametric valuation models is superior to the one-shot approach

2. With short 1-hour time windows, any effort to fit any of the tested parametric models to the query-response pairs is futile, as using non-parametric models is always superior in terms of collaboration gains. The regression models we considered appear to be highly misspecified for the tight routing restrictions.

3. Linear regression models outperform regression trees and artificial neural networks. Despite the interactions between orders, approximating carrier valuations of order bundles with these purely additive models yields better or equally good collaboration gains. This is probably due to model mismatch, i.e. the inherent model bias.

4 DISCUSSION

Combinatorial auctions are a promising tool to enable horizontal carrier collaboration that is based on order reallocations. One of the impediments of these auctions is the need for carriers to report bids for each of the exponential number of order bundles. Limiting the number of biddable bundles reduces the collaboration gains. Our work focuses on minimizing this tradeoff by introducing methods from the field of preference learning to this application. Instead of selecting the biddable bundles all at once, we demonstrate how the auctioneer can accumulate information over multiple query iterations to find the bundles that maximize the downstream WDP while limiting the computational burden of the carriers. Our results clearly show that using the PL scheme to fit parametric bid estimators is effective if enough training data is available and the underlying routing problem does not use tight delivery time windows.

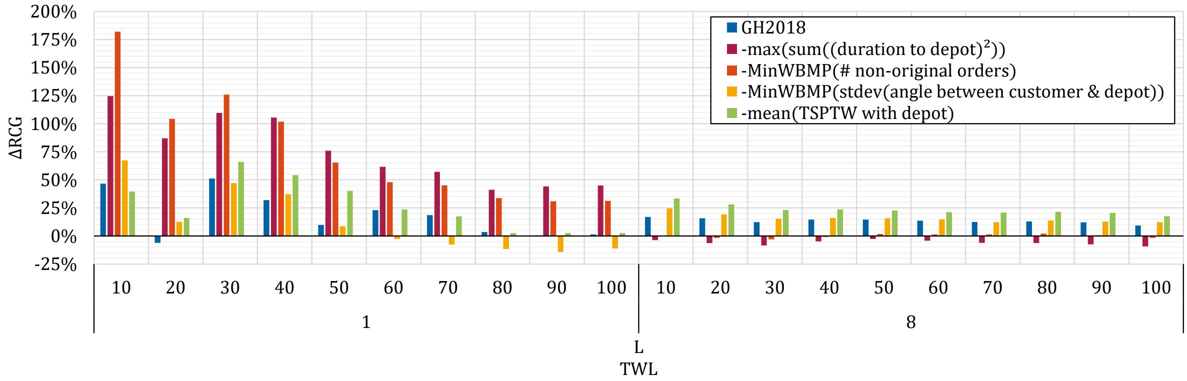

Figure 2 – *ΔRCG of non-parametric models. The x-axis is split into the 1-hour time window (left) and 8-hour time window setting (right), where 8 hours is equivalent to no time windows because we assume an 8-hour working day. The second x-axis level (L) are the number of queries the carriers had to answer.*

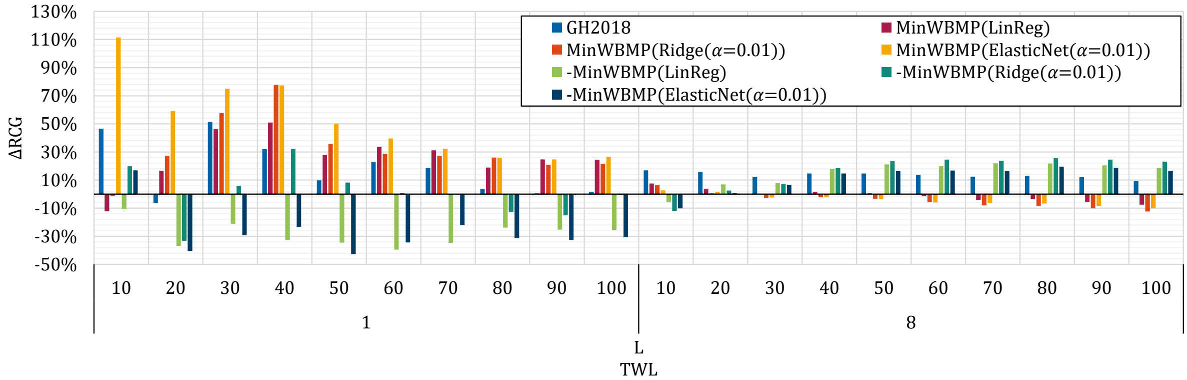

Figure 3 – *ΔRCG of parametric models. For 1-hour time windows, it is clearly better to stick to non-parametric alternatives. For 8-hour (no) time windows, the parametric regression models keep on increasing their lead over the random benchmark as their amount of training data (L) increases.*

References

Berger, S. and Bierwirth, C. (2010). Solutions to the request reassignment problem in collaborative carrier networks. *Transportation Research Part E: Logistics and Transportation Review*, 46(5):627–638.

Cleophas, C., Cottrill, C., Ehmke, J. F., and Tierney, K. (2019). Collaborative urban transportation: Recent advances in theory and practice. *European Journal of Operational Research*, 273(3):801–816.

Fürnkranz, J. and Hüllermeier, E., editors (2010). *Preference Learning*. Springer, Heidelberg; New York.

Gansterer, M. and Hartl, R. F. (2018a). Centralized bundle generation in auction-based collaborative transportation. *OR Spectrum : Quantitative Approaches In Management*, 40(3):613–635.

Gansterer, M. and Hartl, R. F. (2018b). Collaborative vehicle routing: A survey. *European Journal of Operational Research*, 268(1):1–12.

Rüther, C. and Rieck, J. (2022). Bundle selection approaches for collaborative practical-oriented Pickup and Delivery Problems. *EURO Journal on Transportation and Logistics*, 11:100087.

Sandholm, T. and Boutilier, C. (2006). Preference Elicitation in Combinatorial Auctions. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*. The MIT Press.

Savelsbergh, M. W. P. (1985). Local search in routing problems with time windows. *Annals of Operations Research*, 4(1):285–305.

Vetschera, R., Knyazev, D., and Rehsmann, D. (2024). Properties of bundle valuations in carrier collaboration. *Central European Journal of Operations Research*, 32(2):241–266.