

Exact matching of attractive shared rides (ExMAS) for system-wide strategic evaluations

Kucharski, Rafal; Cats, Oded

DOI

[10.1016/j.trb.2020.06.006](https://doi.org/10.1016/j.trb.2020.06.006)

Publication date

2020

Document Version

Final published version

Published in

Transportation Research Part B: Methodological

Citation (APA)

Kucharski, R., & Cats, O. (2020). Exact matching of attractive shared rides (ExMAS) for system-wide strategic evaluations. *Transportation Research Part B: Methodological*, 139, 285-310.
<https://doi.org/10.1016/j.trb.2020.06.006>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Exact matching of attractive shared rides (ExMAS) for system-wide strategic evaluations

Rafał Kucharski*, Oded Cats

Department of Transport & Planning, Delft University of Technology, the Netherlands



ARTICLE INFO

Article history:

Received 29 November 2019

Revised 9 June 2020

Accepted 15 June 2020

Keywords:

Ride-hailing

Mobility on demand

Ride-sharing

Shareability

Public transport

ABSTRACT

The premise of ride-sharing is that service providers can offer a discount, so that travellers are compensated for prolonged travel times and induced discomfort, while still increasing their revenues. While recently proposed real-time solutions support online operations, algorithms to perform strategic system-wide evaluations are crucially needed. We propose an exact, replicable and demand-, rather than supply-driven algorithm for matching trips into shared rides. We leverage on delimiting our search for attractive shared rides only, which, coupled with a directed shareability multi-graph representation and efficient graph searches with predetermined node sequence, narrows the (otherwise exploding) search-space effectively enough to derive an exact solution. The proposed utility-based formulation paves the way for model integration in travel demand models, allowing for a cross-scenario sensitivity analysis, including pricing strategies and regulation policies. We apply the proposed algorithm in a series of experiments for the case of Amsterdam, where we perform a system-wide analysis of the ride-sharing performance in terms of both algorithm computations of shareability under alternative demand, network and service settings as well as behavioural parameters. In the case of Amsterdam, 3000 travellers offered a 30% discount form 1900 rides achieving an average occupancy of 1.67 and yielding a 30% vehicle-hours reduction at the cost of halving service provider revenues and a 17% increase in passenger-hours. Benchmarking against time-window constrained approaches reveals that our algorithm reduces the search-space more effectively, while yielding solutions that are substantially more attractive for travellers.

© 2020 The Authors. Published by Elsevier Ltd.
This is an open access article under the CC BY license.
(<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Mobility platforms dynamically match supply and demand in a two-sided market, enabling the provision of flexible on-demand transport services. Nowadays, thanks to the steadily growing market share of ride-hailing services and their instant operations, shared rides are becoming increasingly available and attractive for both sides of this two-sided market (Li et al., 2019b). The premise of ride-sharing is that travellers can travel for a reduced fee, while service providers can increase their revenues (Furuhata et al., 2013). Furthermore, shared rides are believed to align well with policy objectives, such as

* Corresponding author.

E-mail address: r.m.kucharski@tudelft.nl (R. Kucharski).




Fig. 1. Road network graph of Amsterdam used in experiments, with sample rides: a) non-shared, private ride marked blue in the North, b) ride shared by two travellers marked green at the South, and c) ride shared by three travellers marked brown in the Central Amsterdam. Stars denote origins, triangles destinations and grey bold lines travellers' shortest paths respectively.

increasing accessibility and reducing congestion. Notwithstanding, the potential of ride-sharing is yet to be realised in large-scale operations.

Travellers, requesting rides from mobility platforms, are offered a shared ride, where a vehicle offers to pick them up at their origin and drop them off at their destination at specific times. We illustrate this in Fig. 1, where private and shared rides in Amsterdam, the Netherlands, are marked. Both pick-up and travel times may deviate from the desired or minimal ones, since the vehicle needs to meet the requirements of all pooled travel requests. Service provider needs to compensate this inconvenience by means of offering a lower fare, compared to the private alternative. This can still be viable for the service provider, if it can now better utilise its capacity and charge several users for a ride section, for which a single driver commission has to be paid.

While offering potential benefits, shared rides can also be challenging for both service provider and service users. The mobility platform needs to bundle trips into attractive shared rides. Whilst the problem of identifying feasible shared rides suffers from the curse of dimensionality and is hardly solvable for real-size demand levels. Here, we exploit the fundamental prerequisite for the provision of shared rides, namely identifying rides that can be shared, while offering all travellers a service that is perceived as more attractive than the non-shared, private, alternative. The trade-off exercised between private and shared rides involves (i) the discounted price, (ii) travel time delays and, (iii) on-board discomfort. Notably, we demonstrate that in such context, the search-space is reduced effectively enough to propose exact method, since the problem, as we formulate it, is not NP-hard anymore.

While recently proposed real-time solutions (Alonso-Mora et al., 2017; Simonetto et al., 2019) effectively support online operations, algorithms to perform offline demand-driven comparative analyses are still missing. We propose an algorithm for obtaining exact solutions to efficiently assess a large number of city-wide strategic scenarios. The method results in a replicable solution, which allows performing a cross-scenario sensitivity analysis when a high quality solution is crucial for assessing the impact of offering shared rides through two-sided mobility platforms. The proposed method offers a demand-, rather than supply-driven analysis of the modal split, market competition and system-wide travellers welfare. Our solution approach can thus be embedded in the analysis of pricing strategies, regulation policies or competition settings including supply-demand feedback loops and emerging market equilibria. Furthermore, the proposed utility-based formulation paves the way for model integration in travel demand models, allowing for example to feed travel attributes back to the modal choice level, utilising the latest findings from travel behaviour research (e.g. Alonso-Gonzalez et al. (2020b,c)). Tools and methods to properly forecast ride-sharing platform market shares under various circumstances are of interest to competi-

tors (including public transport) and policy makers. Better understanding of system-wide properties is crucial to counteract potentially negative impacts of emerging modes on equity and accessibility of public transport.

1.1. Related work

Wang and Yang (2019) provide a recent and exhaustive overview of research topics in the domain of ride-sourcing systems including shared rides, while Mourad et al. (2019) review previously proposed matching algorithms for shared rides. The problem of matching trips into shared rides can be decomposed into the following sub-problems:

- determining if two or more rides may be shared,
- identifying all feasible shared rides for a set of trip requests,
- assigning trips to shared rides,
- realistic system representation (in terms of demand, fleet and platform operations),
- quantifying and assessing results using a set of performance indicators.

In the following, we discuss related work while considering how these aspects were addressed.

The shared rides problem can be classified based on its formulation, depending on system settings, objectives and the assumptions made. Trip requests may be treated as static, i.e. known in advance (oracle case of Santi et al. (2014)), or dynamic, i.e. announced as an incoming stream and not known before a given announcement time (Agatz et al., 2011). The travel request-vehicle assignment may be static, i.e. the full schedule is determined before the passenger enters the vehicle (Cordeau, 2006), or dynamic, i.e. new requests may be assigned en-route, not known to passenger before s/he enters the vehicle (Alonso-Mora et al., 2017). The fleet may be explicitly considered, i.e. any given ride is served by a certain vehicle that arrives at a certain time (Alonso-Mora et al., 2017; Simonetto et al., 2019), or assumed instantly available upon request at the origin location (Santi et al., 2014; Tachet et al., 2017). Vehicle capacity, the most important determinant of computational complexity, can be limited to two (Santi et al., 2014), four, eight or ten passengers (Alonso-Mora et al., 2017). Solution approaches may be designed for online application, hence focusing on fast and efficient solutions (Simonetto et al., 2019), or offline where focus is on network-wide impact on performance and mobility (Tachet et al., 2017). The search space may be thus either exhaustively explored (which was considered hitherto infeasible due to the curse of dimensionality apart from small instances like in Cordeau (2006)), decomposed into smaller spatially partitioned problems (Pelzer et al., 2015), temporally partitioned problems (Simonetto et al., 2019), explored via heuristics (Alonso-Mora et al., 2017), or narrowed due to assumed capacity limitations (Santi et al., 2014). Accordingly, solutions sought may be exact (Cordeau, 2006) or heuristic-based (Simonetto et al., 2019).

The objective of the trip-ride assignment has been formulated in terms of the cost for the: a) operator (vehicle-hours, revenues, served trips) (Hosni et al., 2014), b) traveller (generalised travel costs, travel times, delays, fares) (Santi et al., 2014), c) a mixture of both (Cáp and Alonso Mora, 2018), or d) system performance (Santos and Xavier, 2013). The constraints imposed on vehicles may include the maximal detour and the profitability of shared rides, whereas constraints imposed on travel requests may involve departure time window and maximal detour or maximal walking distance. Constraints were assumed fixed in all past work, with the exception of (de Ruijter et al., 2020) who expressed it using a compensatory function of trip attributes. The travel demand given as input, may be generated based on observed ride-hailing trips (Tachet et al., 2017), observed ride-sharing trips (Li et al., 2019b) or sampled from synthetic travel demand data (Hosni et al., 2014). Demand in most of studies is considered as exogenous, i.e. trip requests are fixed and not sensitive to sharing incentives and performance. Endogenous demand, where modal competition is reproduced to some extent, both between private (non-shared) and shared ride, as well as in relation to other modes (public transport and private car) is scarce. Demand is typically microscopically represented, with exception of matching algorithm by Friedrich et al. (2018), which can be integrated in existing macroscopic travel demand models. Finally, the most important network property, travel time, may be fixed and static, fixed and dynamic (Alonso-Mora et al., 2017) or flow-dependent (i.e. sensitive to flow changes induced by ride-sharing) (Narayan et al., 2019).

Most of the analytical experiments and simulation studies concluded that ride-sharing can yield remarkable efficiency gains. Findings suggested a fleet size reduction of 85% (Alonso-Mora et al., 2017), 47% or 35% mileage reductions (Qian et al., 2017) and (Friedrich et al., 2018), 42% reduction in the number of rides (Pelzer et al., 2015) or 19% cost reductions (Santos and Xavier, 2013).

The approach taken in all of the aforementioned research was to identify a set of shared rides, that satisfy pre-specified constraints related to the compatibility of trip combinations. The latter includes imposing thresholds, related to maximum waiting time or maximum detour that were set as fixed constraints for each passenger trip. The share of travel requests that could not be served given these constraints was used as one of the performance indicators. Travellers were considered captive users of the ride-sharing service and insensitive to the level of service offered. This resulted in generating unattractive shared rides, while significantly increasing the search space of shared rides. The search space of practical-size problems was never explored exhaustively, due to the deployment of over-optimistic assumptions on feasibility, leading to a combinatorically exploding search space. Furthermore, such assumptions may result with passenger-vehicle assignments that yield travel delays that may not be acceptable in reality (e.g. 7 minutes detour and 7 minutes extra waiting time in New York City reported by Simonetto et al. (2019)). While empirical studies on the performance of ride-sharing services are still scarce, there is evidence to suggest that the vast majority of passengers opt for the private ride alternative (e.g. 75% in New York City, (LLC, 2019); 93% in Chengdu, (Li et al., 2019b)) and even those that do choose for the shared option

only seldom end up actually sharing (part of) their ride with at least one other traveller (e.g. 18% of those requested in Toronto (Big Data Innovation Team, 2019)). There is, therefore, a need to better understand the viability of the ride-sharing alternative by analysing the relation between the discount offered for shared rides, the share of passengers that choose for a shared ride, the level of service delivered and service profitability.

The combinatorial problem of identifying shared trips becomes quickly difficult with an increasing number of travel requests and sharing degree (i.e. maximal number of passengers sharing a ride). Santi et al. (2014) argue that the combinatorial explosion of the matching problem will limit the potential of ride-sharing in practice. As argued by Mourad et al. (2019), exact solutions have been so far available only for oversimplified instances and there is a need for improved heuristics, in terms of quality, as well as computation times. Previous attempts to address this problem have considerably simplified the problem to reduce its complexity by: considering a small number of travel requests (Hosni et al., 2014), strongly restricting the number of trips that can be matched (Santi et al., 2014), or proposing heuristics (Simonetto et al., 2019; Alonso-Mora et al., 2017). Alternatively, the potential for trip sharing was explored using generic and aggregate spatial and temporal trip similarity characteristics to obtain rough approximations (Bicocchi and Mamei, 2014).

Within the rapidly growing literature on shared rides we build in this study on the concept of shareability graphs introduced by (Santi et al., 2014), further extended for multiple rides and fleet considerations in (Alonso-Mora et al., 2017) and the inclusion of ride-sharing utility in (de Ruijter et al., 2020). Santi et al. (2014) introduced the so-called shareability network, where arcs connect pairwise shareable trips. Trips are shareable in Santi et al. (2014) if: (i) requested within the same time-window and (ii) satisfy pre-set delay thresholds. Santi et al. (2014) cast the trip sharing problem as a maximal matching on shareability graph, with exact solution for pairwise shareable trips (with the objective of either maximise the number of shared trips or minimise their travel time in a weighted matching version). For rides consisting of three passengers they cast shareability graph as a hypergraph and propose heuristics, while arguing that higher degrees are computationally infeasible. Alonso-Mora et al. (2017) extend the concept of shareability proposed by (Santi et al., 2014) allowing for more than two people to share a ride (up to 10 per vehicle) and explicitly considering fleet operations. They divide the time-window of Santi et al. (2014), into separate parameters for maximal waiting and travel delay, which may compensate for each other. The shareability graph of Santi et al. (2014) is further exploited and rides of degree greater than two are searched as graph cliques, which Alonso-Mora et al. (2017) formally introduce as ride candidates. Ride candidates are feasible only if there is a vehicle that can serve it and if a sequence of pickup and drop-off points may be visited, while feasible for all travellers.

Past research focused on real-time operations and the need to develop heuristics to ensure their applicability. In most cases, real-time fleet operations of service provider are explicitly considered, while demand is represented as largely insensitive to system performance and only conditional on fleet availability. Moreover, the profitability of shared rides for service provider, expressed as a function of offered discount, was also not considered, presumably due to models' insensitivity to users' choice between private and shared rides. Conversely, this work pertains to offline strategic applications. In order to devise business strategies and policies, transport authorities and service providers need to assess a larger number of scenarios at the city-wide level. Scenarios of interest pertain, for example, to pricing strategies, regulation policies (Li et al., 2019a) or competition settings (Pandey et al., 2019; Kalczyński and Miklas-Kalczyńska, 2019). This calls for the development of an exact algorithm, that will enable offline city-wide evaluations, where the quality of the solution needs to support business or policy decisions, and the planner is interested in exploiting a range of potential specifications of system parameters. The algorithm needs hence to also be efficient, in particular given the number of requests that need to be handled and the resulting solution space.

1.2. Study objectives and approach

The prime objective of this paper is to propose a scalable algorithm, capable of handling large-scale travel demand to efficiently identify attractive shared rides of any degree within acceptable computation time for the purpose of assessing strategic scenarios. To this end, we propose an offline, demand-oriented solution to solve the ride-sharing trip-vehicle assignment. The proposed method of Exact Matching of Attractive Shared rides (ExMAS) matches trips into rides, that are attractive for all passengers' involved and (optionally) profitable for the service provider.

We argue that travel demand characteristics are such, that the number of attractive shared rides implodes, rather than explodes, with their degree, permitting the development and deployment of an exhaustive algorithm to enumerate all feasible and attractive shared rides. Without relying on heuristics, we represent the compatibility and sequencing of travel requests using a multi-graph representation, allowing the identification of shared rides of higher degree by performing efficient graph searches. We address the complexity of the trip-sharing problem by narrowing the searches to attractive rides only and then leverage on a directed multi-graph, to incrementally search for rides of higher degrees. In our hierarchical approach trips are gradually matched into rides of increasing degree. This is key to efficiently reduce the search space before the degree makes computations intractable. This is achieved by means of (i) hierarchically exploring the search space ensuring that any explored ride of n degree is composed of feasible shareable $n - 1$ degree rides and (ii) maintaining the order of visited origins and destinations, to predetermine the sequence and avoid combinatorial searches. This way, while searching for a ride shared by four passengers, we do not process all four-element subsets of trips, iterating through all feasible permutations of origins and destinations. Instead, we first identify triples, composed of trips which are all pairwise shareable with a potential fourth trip and evaluate already predetermined sequence of origins and destinations, reducing thus the number of computationally expensive calculations to a few precise and targeted ones. This approach makes an, otherwise infeasible, problem solvable within acceptable computation time.

The approach proposed in this study allows us to analyse overall passenger welfare changes under alternative demand, network and service settings as well as behavioural parameters. We perform a system-wide analysis of the ride-sharing performance, in terms of shareability of real-size urban demand patterns as well as indicators to quantify its impact on passengers, operator and the system as a whole. We complement system performance outputs with more detailed shareability indicators, computed for the total travel demand (and not limited to taxi demand only which arguably has specific properties), sampled from the actual Amsterdam demand patterns.

Trip utility formulas are incorporated in our approach, assuming that passengers consider a shared ride attractive if, and only if, its utility is greater than the utility of travel alternatives. An alternative may be any other mode, but hereby we consider that passengers are making a choice between private and shared on-demand services. In this study we assume that, like in other travel choice situations, travellers trade-off different attributes and their choice can be represented as a probabilistic outcome of a compensatory function composed of a set of travel attributes. We hereby adopt the approach proposed in [de Ruijter et al. \(2020\)](#). This implies that, instead of specifying time-windows constraints that must be satisfied for a ride to be shared, we formulate the willingness to share as a compensatory cost function at the individual passenger level considering trade-offs between delays caused by detours, travel discomfort related to sharing a vehicle and a fare discount associated with a shared ride. Recent behavioural research indicates that travellers preferences in the context of ride-sharing can indeed be well explained using compensatory discrete choice models ([Al-Ayyash et al., 2016](#); [Alonso-Gonzalez et al., 2020c](#); [Frei et al., 2017](#)). We, therefore, adopt the well-established utility maximisation modelling framework to allow travellers to choose between individual and shared rides. In order to examine the full sharing potential, we assume that the supply can scale accordingly and do not specify an a-priori limit on the fleet size and capacity. We propose an exact and complete solution technique that is applicable to the large-scale problems.

Benchmarking our results with the previously proposed non-compensatory time-windows approach reveals that: (i) on one hand, fixed time windows do not allow identifying attractive rides which do not fit the strict non-compensatory windows; and (ii) on the other hand, a significant share of the rides fitting within the windows are not perceived as attractive. Consequently, in our experiments ExMAS outperforms previous time-windows approaches in terms of: (i) improved performance for travellers (in terms of passenger hours reductions); and (ii) effectively reducing the search space to allow for exact searches. Consequently, computation times remain feasible for offline, city-wide analyses. Moreover, further decomposing the problem into smaller batches (as typically applied in online solutions) yields computation times which are promising for real-time operations (below 10s for demand levels around 5 000 trips/hour).

The remainder of the paper is organised as follows. In the next section, we introduce a method to determine attractive shared rides, search for them and finally optimally assign passenger trips to shared rides. [Section 3](#) details an extensive set of experiments applied for the case of Amsterdam. We discuss the implications of our findings, study limitations and directions for further research in [Section 4](#).

2. Method - Exact matching of attractive shared rides (ExMAS)

In the following, we first formulate the sharing utility in [Section 2.1](#), we then introduce the trip-ride assignment problem in [Section 2.2](#). In [Section 2.3](#) we detail algorithm to identify attractive shared rides of incrementally increasing degree. We first propose pair-wise searches ([Section 2.3.1](#)), on which we build a directed multi-graph in [Section 2.3.2](#), based on which we finally introduce complete searches for shared rides of higher degree in [Section 2.3.3](#). Altogether this sequence of steps constitutes the Exact Matching of Shared Rides (ExMAS) algorithm, formalised in [Section 2.3.4](#) and available online under public repository (github.com/rafalkucharskipk/ExMAS).

2.1. Attractive shared rides

We consider passenger trip requests $\mathbf{Q} = \{Q_1, Q_2, \dots, Q_n\}$, with single trip i defined as a combination of its origin o_i , destination d_i and desired pick-up time t_i^p :

$$Q_i = (o_i, d_i, t_i^p) \quad (1)$$

Travel times, t_i , are calculated on the road network graph $G(N, A)$, where A is the set of route-able arcs, and N is the set of nodes, pertaining to pick-up and drop-off locations $o_i, d_i \in N$. We use efficient look-up tables to query network travel times $t(o, d, t_o)$ between origin node o , destination node d at departure time t_o . Fare is calculated based on distance $l_i(o_i, d_i)$ with a given service fare λ^{ns} (€ /km). Here we consider that travellers, currently using a non-shared on-demand service (i.e. private ride), will opt for shared rides if they find them superior, i.e. more attractive. If passenger i decides to take a private ride, he is picked-up from his origin o_i at time t_i^p and reaches his destination d_i at time t_i paying a fee of $\lambda^{ns}l_i$.

Travel time and fare, coupled with value-of-time β^t and cost-sensitivity β^c allow to express the (dis)utility of the non-shared trip as:

$$U_i^{ns} = \beta^c \lambda^{ns} l_i + \beta^t t_i \quad (2)$$

Let us now consider a passenger being offered a shared ride r , such that he is picked-up from origin at time $\hat{t}_i^p$ and arrives at his destination at $\hat{t}_i$. For a typical shared ride the travel time is longer ($\hat{t}_i \geq t_i$) than for non-shared private ride and pick-up

time may be different from requested ($\hat{t}_i^p \neq t_i^p$). Which yields a different utility term, that we express as follows:

$$U_{i,r}^s = \beta^c \lambda^s l_i + \beta^t \beta^s (\hat{t}_i + \beta^d (\hat{t}_i^p - t_i^p)) \quad (3)$$

The utility of a shared ride involves a (possibly different) value of time, further penalised with β^s to account for the discomfort of a shared ride. The trip time has now two components: travel time $\hat{t}_i$ (possibly longer than for non-shared ride) and a delay ($\hat{t}_i^p - t_i^p$). Pick-up delay is weighted with an additional multiplier β^d , which may be different for early and late departures depending on the trip purpose. Such longer and less comfortable trips may be attractive only if the service provider offers a lower service fee rate λ^s , such that it compensates for the loss in utility caused by sharing.

To identify when sharing is attractive, we express the *sharing utility* as the difference between the shared and non-shared utility terms:

$$U_{i,r} = U_{i,r}^s - U_i^{ns} = \beta^c \lambda l_i + \beta^t (t_i - \beta^s (\hat{t}_i + \beta^d (\hat{t}_i^p - t_i^p))) + \varepsilon \quad (4)$$

where $\lambda = \lambda^s - \lambda^{ns}$ is the discount offered for sharing a ride, i.e. the discount in per-kilometre fare, compared to a non-shared ride. This formulation allows to express the utility of sharing with: one service design parameter λ , set of behavioural parameters (value-of-time β^t , cost-sensitivity β^c , delay sensitivity β^d , sharing discomfort β^s), and three variables (travel distance, travel time and pick-up delay). Mind, that both travel time and costs are perceived negatively (through negative signs of both β^c and β^t). The random term ε allows to consider both deterministic (if $\varepsilon = 0$) and probabilistic systems (where ε represents some random distribution). In the following we adopt a deterministic binary choice model.

We assert that a shared ride r is attractive for i if $U_{i,r}$ is positive, denoted in what follows with a dummy binary variable:

$$b_{i,r} = \begin{cases} 1 & \text{if } U_{i,r} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

In the remainder, we consider a system where the discount parameter $\lambda < 0$ (i.e. $\lambda^s < \lambda^{ns}$) allows to (at least partially) compensate for the downsides of a shared ride, namely: longer travel time, departure delay and the discomfort associated with sharing a trip. Since we are interested in the relative utility of shared vs. private trips, we set the pickup-up delay to null for private trips and for shared trips it represents the delay stemming from sharing. Other delay components (primarily those related to the service operations) are assumed constant over alternatives and thus cancel out each other in the utility formula (Eq. 4).

In above setting we consider the competition between non-shared and shared ride-hailing, which, without loss of generality, may be reformulated into competition with private car (in which case travel costs are different yet in-vehicle times remain roughly unchanged) or public transport (in which case λ^s would presumably be above λ^{ns}). Mind, that in the experiments for readability we use the more intuitive relative discount, i.e. $\lambda = -(\lambda^s - \lambda^{ns})/\lambda^{ns}$. So that $\lambda = 0.2$ in experiments implies that the shared ride fare is 20% lower than the private alternative.

The formulation proposed above is generic and can be extended to accommodate various utility specifications. In the following, we include in the utility function variables and parameters that are either known, or can be estimated. For example, recent studies provide estimates for the value-of-time for urban pooled on-demand services based on stated preferences data collected in choice experiments (Alonso-Gonzalez et al., 2020a). As new insights are gained from travel behaviour studies concerning: the value-of-reliability (Alonso-Gonzalez et al., 2020a), taste variation among riders (Lavieri and Bhat, 2019) and attitudinal classes (Alonso-Gonzalez et al., 2020b), the utility formulation may be specified accordingly and further extended. Furthermore, the increasing deployment of ride-pooling on-demand transport services will allow future research to estimate users' willingness-to-share (or in reversed terms, willingness-to-pay for a private ride) as well as the disutility induced by delay from revealed preferences data.

Importantly, our formulation does not require defining a (possibly arbitrary) time window for trip departure, since we derive the delay directly from utilities. To this end, we introduce the so-called trip flexibility parameter δ_i , expressing the maximal extent of trip travel time, for which sharing is still attractive (and beyond which it becomes unattractive). We read it directly from the utility formula (Eq. 4), assuming that the travel time t_i is extended by δ_i such that utility is zero, which yields:

$$\delta_i = \left(\frac{1}{\beta^s} - 1 \right) t_i + \frac{\beta^c}{\beta^t \beta^s} \lambda l_i \quad (6)$$

Solution of the above poses an upper limit on travel time extension, beyond which a shared trip becomes unattractive, i.e. break-even point at the individual level. We find this formulation efficient to determine rides, for which $\hat{t}_i \leq t_i + \delta_i$ and $|\hat{t}_i^p - t_i^p| \leq \delta_i / \beta^d$ holds true, as potentially attractive. Such a formulation allows to extend the static time window, used in previous studies, and directly identify bounds for attractive shared ride.

Let us define a ride r serving $\mathbf{Q}_r \subseteq \mathbf{Q}$ trips through:

$$r = (\mathbf{Q}_r, \mathbf{O}_r, \mathbf{D}_r, t_r^p) \quad (7)$$

where $\mathbf{Q}_r$ is a sequence of served trips, $\mathbf{O}_r$ and $\mathbf{D}_r$ are ordered sequences of served trips' origins and destinations, respectively, and the t_r^p denotes departure time from the first origin. We consider only rides for which origins always precede

destinations, i.e. no destination is served before the last origin. By generating sequences of origins and sequences of destinations we get the schedule of a ride, $N_r = (o_1, o_2, \dots, o_n, d_1, d_2, \dots, d_n)$. The number of trips served by a ride is called the *degree*, or *order* of a shared ride and is denoted by $m_r = \|\mathbf{Q}_r\|$.

Ride r serves a sequence of nodes n in time instances $\hat{t}_{r,n}$, calculated recursively starting at time instant t_1^p , shifted with δ_r , and moving forward for each served node with the travel time t between consecutive nodes, plus an additional service time t^s for each pick-up and drop-off:

$$\hat{t}_{r,n} = \begin{cases} t_n^p + \delta_r & \text{if } n = N_{r,1} \\ \hat{t}_{r,n-1} + t(N_{r,n-1}, N_{r,n}, \hat{t}_{r,n-1}) + t^s & \text{otherwise} \end{cases} \quad (8)$$

The total travel time of a ride is $t_r = \hat{t}_{r,N} - \hat{t}_{r,1}$. Travel time and departure delay for each served trip are respectively:

$$\hat{t}_{r,i} = \hat{t}_{r,d_i} - \hat{t}_{r,o_i} \quad (9)$$

$$\hat{t}_{r,i}^p = \hat{t}_{r,o_i} - t_i^p. \quad (10)$$

From the above, we can calculate ride utility for each of the trips sharing it, $U_{i,r}$. We are interested whether a longer travel time $\hat{t}_{r,i}$ and a pick-up delay $\hat{t}_{r,i}^p$ are compensated by the price discount λl_i in Eq. 4. We say that a ride is attractive only if it is attractive (positive sharing utility) for all the trips sharing the respective ride, which leads to following ride utility U_r formulation:

$$U_r = \sum_{i \in \mathbf{Q}_r} U_{i,r} \cdot b_r = \sum_{i \in \mathbf{Q}_r} U_{i,r} \cdot \prod_{i \in \mathbf{Q}_r} b_{i,r}, \quad (11)$$

where b_r is the ride feasibility, computed as the product of trip feasibility $b_{i,r}$ (Eq. 5). The utility of a ride depends, apart from consecutive travel times between served nodes, on the departure time from the first node, which may be shifted in order to maximise the trip utility. To determine the optimal shift, yielding the maximal utility for sharing passengers, we first update the trip flexibility with the actual travel time extension:

$$\delta_{i,r} = \delta_i - (\hat{t}_{r,i} - t_i), \quad (12)$$

which constitutes bounds for the acceptable departure time for each trip. Within the bounds, we identify the optimal departure time t_r^{*p} , for which U_r is maximised. Ride utility, expressed as a function of departure time, is a piece-wise linear function, where optimum can be found analytically. Here, we formulate an analytical solution only for trips of second degree, where the search space is still sufficiently large to call for the formulation of an analytical approach. For trips of higher degree, we find that $\delta_{i,r}$ narrows the search space effectively. So that, it can be exhaustively searched for minimum and analytical solutions are obsolete:

$$t_r^p = \operatorname{argmax}_{t_r^{*p}} U_r(t_r^{*p}) \quad (13)$$

2.1.1. Scheduling horizon

We explicitly account for scheduling constraints. Specifically, for a ride to be feasible it is not sufficient that it is attractive for all the trips, but it must also be feasibly scheduled in time. We introduce a so-called scheduling horizon T , meaning that trips are requested T seconds before the desired departure time t_i^p . This allows to introduce an additional criterion for shareability:

$$b_r > 0 \iff |t_i^p - t_j^p| < T \quad \forall i, j \in \mathbf{Q}_r \quad (14)$$

We treat T as an exogenous parameter in our model. Large values of T , imply that the full shareability potential is revealed, sometimes referred to as 'oracle' settings (Santi et al., 2014) since it assumes that everything is known well in advance. In contrast, when T is small the setting is closer to real-time system operations where trips are typically requested as late as possible, which significantly limits the shareability, as we demonstrate in our analysis.

2.1.2. Profitable shared ride

In the following, we focus on the profitability of the system under fixed demand for on-demand services. The key control variable for service provider is the discount rate offered for shared rides. If the latter is too low it will not encourage travellers to share and impede the sharing potential. Conversely, if set too high, service provider may find it hard to compensate for its lost revenues with detours, which, with high discount rate, are more acceptable for travellers. This stresses the need to develop a model, that accounts for these complex inter-dependencies and feedback loops, to analyse service performance.

A shared ride generates revenues and costs. Without loss of generality, we may assume revenues to be the sum of fares paid by travellers, while the costs, to be proportional to the ride length. The revenues generated by a shared ride are the sum of trips lengths l_i multiplied by the reduced trip fare λ^s while the costs are proportional to ride length l_r . The profit is the difference between the two. If, however, the trips are to be served by a non-shared private ride, revenues are the sum




Fig. 2. Visualisation of trip-ride incidence matrix $I_{m,r}$ for 50 trips and 320 rides. Rows denote trips and columns denote rides of increasing degree (number of trips per ride). Starting with a diagonal part of single rides, followed with FIFO pairwise shared rides, LIFO pairwise shared rides and finally triplets. Each trip may be assigned to only one ride, despite being shareable with multiple others trips via multiple rides.

of trip lengths multiplied by λ^{ns} while costs are proportional to the sum of trip lengths l_i . We are interested in determining whether a shared ride yields a higher profit for service provider than serving it by offering private rides:

$$\sum_{i \in \mathbf{Q}_r} \lambda^s l_i - l_r > \sum_{i \in \mathbf{Q}_r} \lambda^{ns} l_i - \sum_{i \in \mathbf{Q}_r} l_i \quad (15)$$

Which yields that the ride discount λ_r can be expressed as:

$$\lambda_r = 1 - \frac{l_r}{\sum_{i \in \mathbf{Q}_r} l_i} \quad (16)$$

This allows for a clear comparison with the discount λ offered to passenger, to determine whether the ride is profitable

$$b_r > 0 \iff \lambda_r \geq \lambda \quad (17)$$

Depending on the problem setting, one may wish to use the above condition directly, to determine if a shared ride is attractive from the perspective of all involved passengers as well as profitable. Keeping in mind that profitable rides of high degree may be composed of non-profitable rides of lower degree, the condition should be applied only after the search space has been explored and not as part of the hierarchical exploration, as introduced in the following sections, otherwise potentially feasible rides of higher degree will be left unidentified.

Note, that the above condition implies that each of the shared rides directly yields a profit, rather than the whole service. The proposed algorithm allows for specifying both myopic, as well as global profitability conditions, as illustrated in our experiments. To assess service profitability we introduce a term for the total cost reduction Λ . Calculated using a formula similar to (Eq. 17), yet aggregated for all the rides:

$$\Lambda = 1 - \frac{\sum_{r \in \mathbf{R}} l_r}{\sum_{r \in \mathbf{R}} \sum_{i \in \mathbf{Q}_r} l_i} = 1 - \frac{L_r}{L_q} \quad (18)$$

where we compute the cost reduction through the ratio of total rides length (L_r) to total trips length (L_q). This formula may be further used to determine service profitability at the aggregate level. In particular, if computed only for shared rides it indicates shared rides profitability Λ_r .

2.2. Trip-ride assignment problem

In the typical, urban demand pattern a single trip may be assigned to more than one ride, so that the number of feasible rides is greater than the number of trips (as illustrated on Fig. 2). To address this, we formulate a so-called trip-ride problem, formulated as an assignment problem, where each trip i is unilaterally assigned to an attractive ride r and the assignment yields minimal costs. It is then formulated as a problem of identifying a binary vector x_r of length equal to the number of rides $\|\mathbf{R}\|$, i.e. an assignment variable equal to one if a ride is selected and zero otherwise (Eq. 19c). The costs C_R are then expressed as the cost of each ride c_r , multiplied by the assignment variable x_r and aggregated for all rides (Eq. 19a).

In particular, we consider two cost functions: service provider perspective (minimising the sum of ride time t_r) and passenger perspective (maximise the sum of utilities $U_{r,i}$), which can be extended to any ride cost formula.

The assignment shall meet the constraint of assigning each trip to exactly one ride, obtained through the row-wise sum of assignment variable x_r and trip-ride incident matrix $I_{i,r}$. The latter is a binary matrix, which each entry is one if ride r serves trip i and zero otherwise (Eq. 19b), as illustrated in Fig. 2. Eventually, a solution to the problem (Eq. 19a) is the set of rides $\mathbf{R}^* \subseteq \mathbf{R}$ such that $x_r = 1 \forall r \in \mathbf{R}^*$. We express the shareability problem as the following binary program:

$$\min \quad C_R(x_r) = \sum_{r \in \mathbf{R}} c_r x_r \quad (19a)$$

$$\text{subject to } \sum_{i \in \mathbf{Q}} I_{i,r} x_r = 1, \quad (19b)$$

$$x_r \in \{0, 1\}. \quad (19c)$$

Other authors addressed this problem by solving the maximal matching problem on the shareability graph (Santi et al., 2014), which is suitable only for trips of a degree not greater than two, where trip-ride incidence matrix $I_{i,r}$ may be easily converted into a graph. If we consider rides of three or more trips, such mapping is not obvious, and requires a hypergraph representation (Santi et al., 2014) for which most of classic graph algorithms do not apply. Therefore, in the above, we adopt a linear problem formulation of (Alonso-Mora et al., 2017), which is applicable for any sharing degree. Notably, since rides identified with our algorithm are strictly attractive for travellers, we have now more flexibility in formulating objective functions to optimise for system externalities or platform objectives, since we can guarantee that they will be composed of attractive shared rides.

2.3. Algorithm to identify attractive rides

Despite its clear form, the trip-ride assignment problem described above is hardly solvable in practice. The main difficulty is to enumerate all the feasible attractive rides in which trips may be shared. This refers first to finding all the subsets $\mathbf{Q}_r \subseteq \mathbf{Q}$ (there are $2^n - 1$ subsets for n trips), secondly generating all the ordered permutations of origins and destinations of each identified subset ($2 \cdot n!$ for trip of n -degree), and finally evaluating the utility of each ride U_r (which requires querying travel time look-up table for each trip segment (Eq. 8)) and determining the optimal departure time to maximise ride utility (Eq. 13). We address this complexity by developing an original ExMAS algorithm as detailed in the following sections. The algorithm hierarchically searches within the space of feasible rides $\mathbf{R}$ and efficiently narrows it, before entering into costly computations. We initialise a set $\mathbf{R}_1$ of non-shared single rides (which are also single-element subsets of trip set $\mathbf{Q}$), with a null utility and a cost that equals the direct travel time t_i . We then incrementally increase the degree of searched rides. Starting with matching trip pairs into rides of second degree, from which we build a directed multi-graph, where we search rides of higher degree.

2.3.1. Attractive rides of second degree.

We start with identifying trip pairs, which can be attractively shared by forming a ride of second degree. Two trips, i and j , may be shared in six ways (as depicted in Fig. 3). Two of which are non-shared (a,b) on Fig. 3), for which the first trip ends before the second one starts and are not considered further here. Two for which origin i is visited prior to origin j (c,d). The former (c) is further called FIFO (since the first to be picked-up is also the first passenger to be dropped-off) and the latter (d) called LIFO (since the last passenger picked-up is the first one dropped-off). In the proposed algorithm the oriented trip pairs are explored, thus we first identify rides where origin i precedes origin j , while the opposite (e, f on Fig. 3) are explored through inverse searches j, i .

We identify two sets of pairwise shareable trips:




Fig. 3. Six possible ways to combine two trips i and j : non-shared sequences (a,b), FIFO and LIFO sequences with origin i visited first (c,d), FIFO and LIFO sequences with origin j visited first (e,f). Ones identified within the algorithm are a FIFO sequence where i precedes j for origin and destination (c) and LIFO sequence (d). First two are out of scope, since they are not shared, while last two are explored through the inverse search.




Fig. 4. Travel times in pairwise shared FIFO trip (left) and LIFO trip (right).

(a) $\mathbf{R}_2$ pairwise FIFO trips (20), for which $i < j$ for both origins and destination (21):

$$\mathbf{R}_2 = \{r_{ij} : \mathbf{Q}_r = (i, j), \mathbf{O}_r = (o_i, o_j), \mathbf{D}_r = (d_i, d_j)\} \quad (20)$$

$$o_i < o_j \wedge d_i < d_j \quad \forall r_{i,j} \in \mathbf{R}_2 \quad (21)$$

(b) $\mathbf{R}_{2'}$ pairwise LIFO trips (22), for which d_j is visited prior to d_i (23)

$$\mathbf{R}_{2'} = \{r_{ij'} : \mathbf{Q}_r = (i, j), \mathbf{O}_r = (o_i, o_j), \mathbf{D}_r = (d_j, d_i)\} \quad (22)$$

$$o_i < o_j \wedge d_j < d_i \quad \forall r_{i,j} \in \mathbf{R}_{2'} \quad (23)$$

For n trips we have n^2 trip pairs forming a potential shared rides to explore, which even for a medium trip set size results in a very large search space. To estimate the sharing utility of second degree shared ride, one needs to determine travel times for each trip segment $t(o_i, o_j)$, $t(o_j, d_i)$ and $t(d_j, d_i)$ (as depicted in Fig. 4) and determine the optimal departure time t_i^p .

We propose the following, sequential reduction, aiming to efficiently reducing the dimensionality of the search space. All filters below are derived from the trip utility formula (Eq. 4) and ordered to effectively reduce the size of the solution space under consideration, before executing costly travel time queries, needed for evaluating utilities. The resulting set of attractive pairwise rides r and r' is obtained if all the above conditions are met:

a) Scheduling horizon filter:

$$b_{i,r} = 1 \iff |t_j^p - t_i^p| < T \quad (24)$$

First we check, whether two trips are requested within the same scheduling horizon (Eq. 14).

b) Departure compatibility filter:

$$b_{i,r} = 1 \iff (t_j^p + \delta_j) \geq (t_i^p - \delta_i) \quad (25)$$

$$b_{i,r} = 1 \iff (t_j^p - \delta_j) \leq (t_i^p + \delta_i) \quad (26)$$

Followed with filtering trips, for which origins cannot be reached, even if travel times are zero. We employ the trip scheduling flexibility term (δ of Eq. 6), rather than fixed time-windows.

Conditions 24–26 allow to reduce the sample size at least by an order of magnitude in realistic settings. Notably, they can be evaluated using solely trip data t_i and δ_i (of size n , rather than trip-pair data of size n^2).

c) Origin compatibility filter:

$$b_{i,r} = 1 \iff (t_i^p + t(o_i, o_j) + \delta_i) \geq (t_j^p - \delta_j) \quad (27)$$

$$b_{i,r} = 1 \iff (t_i^p + t(o_i, o_j) - \delta_i) \leq (t_j^p + \delta_j) \quad (28)$$

We now include the travel time between o_i and o_j in Eqs. 25 and 26, to guarantee that origin j is reachable from i within departure flexibility windows of both trips. Mind, that travel times are first queried in (Eq. 27), when the reduced search space is already reduced significantly.

d) Attractiveness for i :

$$b_{i,r} = 1 \iff U_{i,r} > 0 \quad (29)$$

After making sure trip origins are reachable, we can now determine the sharing utility for trip i directly from Eq. 4. To this end, we query for the travel time on the second segment of the trip: $t(o_j, d_i)$.

e) LIFO attractiveness for i :

$$b_{i,r'} = 1 \iff U_{i,r} > 0 \wedge U_{i,r'} > 0 \quad (30)$$

At this stage we may exploit the utility dominance of FIFO over LIFO ($U_{i,r} \geq U_{i,r'}$) and use the already decimated trip candidates set to determine LIFO rides r' . We generate r' from r by switching the order of visited destination nodes and updating travel times t_i , t_j , according to Fig. 4. Travel time is now longer for i and shorter for j . Using Eq. 30 we check, whether the utility for i , which was positive when $\hat{t}_i$ was $t_{o_i, o_j} + t_{o_j, d_i}$, remains positive when $\hat{t}_i$ is now $t_{o_i, o_j} + t_{o_j, d_j} + t_{d_i, d_j}$.

f) Attractiveness for j :

$$b_{j,r} = 1 \iff U_{j,r} > 0 \quad (31)$$

$$b_{j,r'} = 1 \iff U_{j,r'} > 0 \quad (32)$$

The final condition for both FIFO and LIFO rides is that the utility (Eq. 4) is positive also for the second trip j . Now, we query for the last segment of the ride: $t(d_i, d_j)$ for FIFO and $t(d_j, d_i)$ for LIFO.

To determine the sharing utility with Eqs. 29 and 31, the optimal departure time t_r^p needs to be determined. Mind that, thanks to the filters Eqs. 25 and (26), the search space is already reduced to reachable origins. If destination j is reachable from o_i and is time-wise acceptable for both i and j , then we can distribute the total departure delay $(t_i^p + t(o_i, o_j) - t_j^p)$ evenly between sharing passengers. Otherwise, one of the boundaries needs to be exploited, namely the smaller among δ_i and δ_j . Which leads to the following analytical formulation:

$$\delta_r = \min(0.5(t_i^p + t(o_i, o_j) - t_j^p), \delta_i, \delta_j) \quad (33)$$

We summarise the above sequence with the following algorithm, to identify attractive shared rides of second degree (Algorithm 1). The sequencing of 'if' conditions aims at minimising the number of pairs in the trip set before querying the

Algorithm 1: Identify pairwise shareable FIFO and LIFO trips.

PairwiseShareable (R_1, T)

inputs :
 R_1

* single trips

 T

* network travel times

output:
 R_2

* pairwise shareable trips

foreach $i \in Q$ **do**

// iterate through all trip pairs

foreach $j \in Q$ **do**

// and identify pairs with:

if $|t_j^p - t_i^p| < T$ **then**

// Scheduling horizon filter (eq. 24)

if $(t_j^p + \delta_j) \geq (t_i^p - \delta_i)$ **then**

// Departure

if $(t_j^p - \delta_j) \leq (t_i^p + t_i + \delta_i)$ **then**

// compatibility filter (eq. 25, 26)

query $t(o_i, o_j)$
if $(t_i^p + t(o_i, o_j) + \delta_j) \geq (t_i^p - \delta_i)$ **then**
if $(t_i^p + t(o_i, o_j) + \delta_j) \leq (t_i^p + \delta_i)$ **then**

// Origin comp. filters (eq. 27, 28)

query $t(o_j, d_i)$
 $\hat{t}_r^p = \min(0.5(t_i^p + t(o_i, o_j) - t_j^p), \delta_i, \delta_j)$

* Optimal departure time (eq. 33)

 $r = ((i, j), (o_i, o_j), (d_i, d_j), \hat{t}_r^p)$
if $U_{i,r} > 0$ **then**

// positive U for i (eq. 29)

query $t(d_i, d_j)$
if $U_{j,r} > 0$ **then**

// positive U for j (eq. 31)

 $R_2 \leftarrow R_2 \cup r$

* success for FIFO

 $r' = ((i, j), (o_i, o_j), (d_j, d_i), \hat{t}_r^p)$

* potential LIFO candidate

query $t(d_j, d_i)$
if $U_{j,r'} > 0$ **then**

// positive U for j (eq. 30)

if $U_{i,r'} > 0$ **then**

// positive U for i
 $R_2' \leftarrow R_2' \cup r'$

* success for LIFO (eq. 32)

Result: $R_2 = R_2 \cup R_2'$

travel times and determining optimal departure time (which is most computationally expensive). This procedure proves to be highly useful: first, by cutting the computation times by an order of magnitude and second, by enabling the computation of pair-wise shareability of large trip sets. The efficiency of this approach is illustrated for a set of 1000 trips in Fig. 5 where 3600 shareable trip pairs are identified from a search space of one million trip pairs. Importantly, the algorithm is not




Fig. 5. Pairwise trip shareability matrix, where search space for shareable pairs i, j (marked in black) is gradually reduced with consecutive filters. The initial set of 1000000 pairs is reduced to 135000 with (Eqs. 24, 25 and 26), then to 66000 with (Eqs. 27, 28). Next, only 14000 pairs have a positive utility for i (Eq. 29) and finally just 3600 pairs happen to be attractively shareable also for j (Eq. 31). Computation time for one million pairs is around 15s.




Fig. 6. Shareability graph with multiple directed edges per pair of trips (FIFO and LIFO). Dashed edges illustrate two possible extensions of a ride of third degree into fourth degree.

heuristic and provides the exact solution within a reasonable amount of time. It is implemented using efficient SQL table joins and queries.

2.3.2. Shareability graph

Starting with the set of feasible rides of second degree $\mathbf{R}$, we incrementally explore rides of higher orders. First, we create a shareability graph $\mathbf{S} = G(\mathbf{Q}, \mathbf{R})$ where nodes are trips $Q \in \mathbf{Q}$ and arcs are pair-wise shareable rides $r \in \mathbf{R}_2$, as depicted in Fig. 6. Contrary to (Santi et al., 2014), our graph is directed, since arcs denote pairwise shareability, where the order of origins and destinations matters and is already determined. Moreover, multiple edges are allowed per node pair, one for FIFO and another for LIFO pair-wise shareability. In particular, two nodes (trips) may be connected with four edges (FIFO and LIFO rides per direction i, j and j, i).

Fig. 6 illustrates the proposed graph with eight trips, connected by twelve edges. Isolated node (Q_1) cannot be shared with any other trip. Trip Q_2 can be shared with Q_3 both in FIFO and LIFO sequence, yet only when Q_2 is picked-up first. Q_4 can be shared with Q_5 in both orders and with Q_2 as a second trip only. Trips Q_5 , Q_6 and Q_7 form potentially shareable triplet ($r_{5,6,7}$), where the order of origins is predetermined by the edges creating the cycle: $\mathbf{O}_{5,6,7} = (o_5, o_6, o_7)$, $\mathbf{D}_{5,6,7} = (d_5, d_6, d_7)$.

Directed multi-graph is found highly efficient in searching for higher order rides. We introduce the following theorems for extending a generic ride r with a trip q , which we utilise in identifying rides of higher order:

Theorem 1. Attractive shared ride r of an order greater than one, can be extended into an attractive ride of a higher degree r' by adding trip q if and only if all trips $i \in \mathbf{Q}_r$ are directly connected with node q through a direct edge $r_{i,q}$ in graph $\mathbf{S}$. Set of edges, such that each trip in r is connected with q with exactly one edge, is called an extension of r into q , denoted $E(r, q)$ (also known as a star in graph theory):

$$b_{r'} = 1 \iff b_r \cdot \prod_{r_{i,q} \in E(r,q)} b_{i,q} = 1$$

Theorem 1 can be derived directly from the fact that the utility of each trip in the extended ride $U_{r'}$ (Eq. 4) is lower or equal to the utility of this trip in the non-extended ride U_r , since both travel time and waiting time need to be either greater or equal for the extended ride. If so, r' is feasible only if all $r \in E(r, q)$ are feasible, i.e. they have been identified as rides of second degree and are therefore edges in $\mathbf{S}$.

Theorem 2. For each extension $E(r, q)$ of ride r to ride r' with trip q , there exists at most one feasible order of origins and destinations for ride r' , that is unambiguously read from types of edges $E(r, q)$, connecting trips of ride r with trip q .

Theorem 2 is derived as follows: 1) Thanks to the directed nature of the graph $\mathbf{S}$, the order in which origins are visited in the currently explored extension $E(r, q)$ is predetermined. By definition $o_i < o_j \quad \forall r_{i,j} \in \mathbf{R}$ for both FIFO and LIFO edges (Eqs. 21, 23). In particular, $o_i < o_q \quad \forall r_{i,q} \in E(r, q)$, thus origin d_q can be visited only after all origins are visited and before the first destination is visited. 2) Analogically, the order of destinations is also predetermined, albeit the destination may be placed at any position. To determine it, let us divide $E(r, q)$ into FIFO edges $E^+ = E(r, q) : r \in \mathbf{R}_2$ and LIFO edges $E^- = E(r, q) : r \in \mathbf{R}_2'$. From Eq. 21 we have lower bounds $d_r < d_q \quad \forall r \in E^+$, and from Eq. 23 we can establish upper bounds $d_q < d_r \quad \forall r \in E^-$. The combination of which yields two bounds between which destination d_q may be inserted: $\max(d_{E^+}) < d_q < \min(d_{E^-})$. Being either infeasible (when lower bound $\max(d_{E^+})$ is greater than upper bound $\min(d_{E^-})$) or pointing exactly to one feasible position (since in that case $\max(d_{E^+})$ always directly succeeds $\min(d_{E^-})$).

We illustrate this with the example shown in Fig. 6. Feasible ride r , serving trips Q_5 , Q_6 and Q_7 , may be potentially extended to ride r' serving also ride Q_8 , since there is a star-graph (dashed edges) connecting all the trips from r with Q_8 , allowing for a so called extension $E(r, q)$ of ride r with trip q (see Theorem 1). In fact, there are two such star graphs, since trips Q_7 and Q_8 are connected by two edges. Only one of them, however, yields a feasible ride (see Theorem 2). For both candidate rides of fourth degree origin o_8 will be visited fourth within ride r' . If we consider an extension composed of LIFO edges only, then we identify a feasible sequence of destinations, where d_8 is the first visited destination, since all the edges of extension $E(r, q)$ are LIFO type, i.e. $d_8 < d_i \quad \forall r_{i,q} \in E(r, q)$. If we, however, consider a second extension where Q_7 and Q_8 are connected with FIFO edge, we end up with an infeasible destination sequence: $d_8 < d_5 < d_7 < d_8$, i.e. $\max(d_{E^+}) > \min(d_{E^-})$ in Theorem 2.




Fig. 7. Ride utility U_r , as a function of departure time t_r^p for candidate rides of fourth degree. Bold lines denote feasible rides (i.e. utility is positive for all trips at some departure time), and light lines denote infeasible trips. A period of indifference can be observed for all rides. We select the earliest time instant of that period allowing for the longest acceptable delays in the actual real-time operations.

2.3.3. Matching trips into higher order rides

We utilise the above theorems to propose a method to identify all feasible and attractive extensions of ride r . We first identify nodes (trips) $\mathbf{S}_r$ pair-wise shareable with all trips of ride r . They are found through the intersection of trips S_i^+ shareable with each trip ride:

$$\mathbf{S}_r = \prod_{i \in \mathbf{Q}} S_i^+ \quad (34)$$

Then, for each candidate trip $r \in \mathbf{S}_r$, we enumerate star graphs (extensions) $E(q, r)$ connecting trips of ride r with trip q . Since graph $\mathbf{S}$ is a multi-graph with two possible arcs connecting an ordered pair of nodes, for trip of degree n there are 2^n possible extensions $E(q, r)$ if the graph is complete. However, the graph is likely to be sparse (as we show below) and sparsity increases for higher degrees. Moreover, for a generic ride of degree n not more than $n + 1$ destination insertion points are possible. For each extension $E(q, r)$ we determine an insertion point for the destination and its feasibility (using Theorem 2):

$$\max(d_{E^+}) < d_q < \min(d_{E^-}) \quad (35)$$

When we identify feasible ride r' , we need to determine its optimal departure time $t_{r'}^p$ with Eq. 13. We leverage here on the previously defined trip flexibility δ_i as the maximal departure delay for all trips of share rides. Thanks to this, and the fact that the shareability graph is sparse, so that the number of feasible rides strongly decreases with the degree, we can narrow the search space efficiently to allow performing an exhaustive search (with Eq. 13). Fig. 7 depicts the utility as a function of departure time, for a sample of rides of fourth degree. In case, when a departure time that yields a positive utility can not be found within the boundaries, we can already conclude that the ride is infeasible and terminate. Otherwise, we evaluate the utility of each trip within the new ride r' (and terminate as soon as a non-positive utility is found).

We summarise the above in Algorithm 2, which, for any generic ride r of degree n , finds all feasible rides of degree $n + 1$ extending it. Note that the already ordered sequence of trips in ride r is extended so that the new origin is always inserted at the last position, other sequences (if feasible) are explored as extensions of different rides. The proposed graph methods can be seen as an extension of ones proposed by (Alonso-Mora et al., 2017), yet suited for a directed graph and hierarchical searches, with the optimal departure time and predetermined sequence allowing to avoid heuristics of travelling salesman problem for rides of degree greater than three. Importantly, Algorithm 2 does not require querying travel times, since all travel times of ride r' are already stored in graph $\mathbf{S}$.

2.3.4. Implementation

Below we introduce the complete ExMAS algorithm (Algorithm 3), to match trips into attractive rides. The proposed algorithm first identifies attractive rides of second degree $\mathbf{R}_2$ (with Algorithm 1), based on which a directed multi-graph $\mathbf{S}$ is built. Then it iterates with incrementally increasing degree of searched rides until all feasible rides are explored. For each degree n , we loop over all identified rides of degree $n - 1$ and extend them to feasible rides (with Algorithm 2). When extensions of degree $n - 1$ rides are all identified, the algorithm proceeds to loop over rides of degree n and extends them to rides of degree $n + 1$. This continues until no more extensions are found.

Once all feasible rides are identified, we solve the trip-ride assignment problem, defining the ride cost c_r to find the optimal solution (with Eq. 19a). Optionally, before solving, we may apply a filter to maintain only profitable rides (by employing Eq. 17). Finally, we obtain $\mathbf{R}^*$, the set of rides serving all the trips with minimal cost c_r .

Algorithm 2: Extend ride r .

```

ExtendRide ( $r, S$ )
  inputs :
     $r$  // extended ride
     $S$  // shareability graph
  output:
     $R$  // set of feasible rides obtained as an extension of  $r$ 
     $S_r = \prod_{i \in Q_r} S_i^+$  // find trips shareable with all trips of ride  $r$ 
    foreach  $q \in S_r$  do // trip candidates
      foreach  $E(q, r) \in E(q, r)$  do // extensions of  $r$  with  $q$ 
         $E = E^+ \cup E^-$  // divide into FIFO and LIFO edges
        if  $\max(d_{E^+}) < \min(d_{E^-})$  then // if there is feasible order of destinations (eq. 35)
           $r' = r \cup q$  // extend ride  $r$ 
           $O_{r'} \leftarrow O_r \cup (o_q)$  // insert origin at the last position
           $D_{r'} \leftarrow \text{insert } d_q \text{ to } D_r \text{ at position } \min(E^-) - (\max(E^+))/2$  // insert destination
           $t_r^p = \arg\max_{t_r^{*p}} U_r(t_r^{*p})$  // determine optimal departure time
          foreach  $i \in Q_{r'}$  do // trips of a new ride  $r'$ 
            if  $U_{i,r'} \leq 0$  then // check if sharing utility is positive
              break // early exit
           $R += r'$  // success
  Result:  $R$ 

```

Table 1

Search space and its reduction for a sample of 3000 trips in Amsterdam. Out of 4.65×10^{20} theoretically feasible trips shared by five passengers only 123 have to be actually examined and 76 of which are identified as attractive and 44 of those are assigned as part of the optimal solution. Out of 6.47×10^{11} trips of third degree only 1807 need to be explored, 243 of which are found feasible and up to 160 are selected within the trip-ride assignment. For pair-wise computations, set of efficient filters (Algorithm 1) managed to gradually narrow to 5270 attractive pairwise shareable rides, forming arcs in our shareability graph.

	degree:	1	2	3	4	5	6	7
search	theoretical	3.00×10^3	3.60×10^7	6.47×10^{11}	1.55×10^{16}	4.65×10^{20}	1.67×10^{25}	7.01×10^{29}
space:	explored	3000	8,997,000	1807	226	123	24	0
	attractive	3000	5270	243	130	76	8	0
	assigned	1422	435	160	44	8	2	0

A key feature of the algorithm is that the search space, which theoretically explodes combinatorially, becomes smaller with each additional degree thanks to an efficient reduction. This can be seen in Table 1 where 3000 trips are matched into rides in realistic setting and parameterization (the specification thereof is detailed in the next section). Out of 4.65×10^{20} theoretically feasible trips shared by five passengers, only 123 needed to be explored, 76 were identified as attractive and 8 were assigned in the optimal solution. Out of 6.47×10^{11} trips of third degree, only 1807 needed to be explored, 243 were found feasible and finally 160 were selected within the trip-ride assignment.

The algorithm is implemented in Python 3.7 using pandas library for data management and SQL-like queries and networkx for graph operations. Computation times are reported in Section 3.2.9. The code is open-source and available on a public repository.

2.4. Shareability indicators

We analyse ride-sharing with the following set of system-wide indicators, computed for the set of rides R^* resulting from algorithm 3:

T_q total passenger-hours travelled, $T_q = \sum_{i \in Q} t_{i,r}$,
 T_r total vehicle-hours travelled, $T_r = \sum_{i \in R^*} t_r$,
 L_r total vehicle-kilometres travelled, $L_r = \sum_{i \in R^*} l_r$,
 U_r total passenger utility, $U = \sum_{i \in Q} U_{i,r}$,
 R number of rides, $R = \|R^*\|$,
 I total revenue, $I = \sum_{i \in Q} \lambda_i \cdot l_i$,
 O average occupancy, $O = T_q/T_r$.

We compare the values obtained for shared rides to the ones without sharing and report the relative difference that stems from sharing, e.g. $\Delta T_q = (T_q^s - T_q^{ns})/T_q^{ns}$.

Algorithm 3: ExMAS Exactly identify and Match Attractive Shared-rides for tripset Q on graph G .**SharedRides** (T, Q, N)**inputs :** Q

* trip requests

 T

* network travel times

 N

* maximal ride degree

output: R^*

* shared rides set

 $R_1 \leftarrow Q$

* Initialise with single rides

 $R_2 \leftarrow \text{PairwiseShareable}(R_1, T)$

* compute pairwise shareability with Algorithm 1

 $S \leftarrow G(Q, R_2)$

* create directed shareability graph

 $n \leftarrow 2$ **while** $R_n \neq \emptyset$ **or** $n \leq N$ **do**

// until set of extensible rides is empty

foreach $r \in R_n$ **do**

// for each extensible ride

 $R_{n+1} \leftarrow \text{ExtendRide}(r, S, T)$ * identify extensions of ride r with Algorithm 2 $n \leftarrow n + 1$ $R^* \leftarrow (r : r \in R, \lambda_r > \lambda)$

* limit rides to profitable only

 $c_r = C(r : r \in R)$

* determine ride costs

 $R^* \leftarrow \text{argmin}_{C_R}(x_r)$

* solve the shareability assignment problem with eq. 19a

Result: R^*

We hypothesise that ride-sharing will, compared to private rides, results in a decrease in-vehicle-hours T_r and an increase in passenger utility U_r albeit with an increase in passenger hours T_q . The first two relations are positive and are the primary benefits of sharing, while the latter is an inevitable cost compromised by passengers, in case the discount for sharing surpasses the loss in travel time. Results in the next section are discussed mainly with respect to these three key indicators. Note that the details of service operations are neglected in this study. Hence, only revenue-hours (with passenger on-board) are used when computing vehicle hours and kilometres. However, to obtain an indication for the potential fleet size savings, we use the number of rides per time unit for both non-shared and shared scenarios. We use the maximal number of simultaneous rides as a proxy for the fleet size, and the relative difference as a proxy for potential fleet savings (see work of Vazifteh et al. (2018) for exact fleet size calculations in a similar setting).

3. Application

We apply our algorithm to understand the benefits and costs, magnitudes, sensitivity and consequences of ride-sharing in a realistic urban context. We use the actual travel demand pattern in Amsterdam, the Netherlands, with realistic model parameter settings, to simulate ride-sharing in various system settings, behavioural parameters and demand levels.

Specifically, we aim to identify shared rides that are attractive, rather than those merely satisfying time windows, by properly setting the utility of non-shared and shared ride respectively. Not all of the behavioural parameters included in the utility formula (Eq. 4) are well established, yet most of them can be specified using reasonable assumptions, as already discussed in the Introduction. While the value-of-time, β^t , is known for the Dutch context (Kouwenhoven et al., 2014), willingness-to-share, β^s , remains hitherto largely unknown. Based on a recent stated preference choice experiments studies (Alonso-Gonzalez et al., 2020c; 2020a) we can obtain initial estimates, yet the actual values and shape of β^s remains unknown. Similarly, while the relative discomfort of waiting time (related to travel time) in the context of ride-hailing services is unknown and presumably depends on the context (trip purpose, time of day), we argue that well established findings from public transport generalised travel costs may be transferred and used in this context.

System settings are both observable and controllable, so we either use actual values or explore their impact by specifying them within reasonable ranges. The price of ride-hailing λ^{ns} as well as the sharing discount λ are known up-front when using a ride-hailing app. We keep the former fixed and set the discount as an exogenous variable subject to investigation in our experiments. Similarly, we examine the impact of the scheduling horizon T in our scenarios, as it is directly controlled by the operator and latent for passengers. Mind, that here we decided to use a more intuitive, relative, notion of discount, with $\lambda = 0.2$ meaning per-kilometre shared ride fare 20% cheaper than the corresponding private ride. Formulas remain consistent with the method introduced in the previous section, though for sake of readability we use now $\lambda = -(\lambda^s - \lambda^{ns})/\lambda^{ns}$.

Travel demand is also typically known, either from observations or from calibrated models, usually allowing to understand demand with information on origin, destination and departure time, which is sufficient to quantify potential demand for shared rides Q .

Table 2

Matching 3000 trips into 1888 rides in Amsterdam. Passengers are offered a shared discount λ of 0.3, matching objective function is minimising vehicle-hours T_r . Values obtained for non-shared and shared scenario and relative differences.

	Vehicle hours T_r	Passenger hours T_q	Utility U	Number of rides R	Fleet size F	Revenues I	Occupancy O
non-shared	1,280,334	1,280,334	-19860	3000	474	1,920,501	1
shared	896,224	1,499,395	-18644	1888	336	1,575,386	1.673
rel. diff (%)	-30.0	17.1	6.1	-37.1	-29.1	-18.0	

3.1. Experimental set-up

We run experiments on an a travel dataset available for Amsterdam from a nation-wide activity-based model (Arentze and Timmermans, 2004). The dataset contains the origin, destination and departure time for 241 thousand trips conducted within the boundaries of Amsterdam during a representative working day. We sample from this complete, multi-modal dataset, rather than from taxi demand only. Tripset size ranged from low demand levels to identify critical mass, to levels exceeding the current demand level for ride-hailing, taxis and shared-taxis typically observed in the urban context (average of 0.5% reported for US cities in Conway et al. (2018)), to identify the potential of those services and test the scalability of the proposed model.

We employ a detailed road network graph of Amsterdam available from OSM (Fig. 1), with over 26 thousand edges and 11 thousand nodes to estimate travel times (we assume a constant speed of 29km/h, inline with empirical observations).

In the following, we first discuss how sharing impacts central KPIs, followed with a discussion on how they change with the system settings, demand levels and travellers' behaviour. Then, we compare system optimal, versus user optimal matching. We discuss how vehicle capacity allows to improve occupancy, further enriched with discussion on profitability, analysed both at the system and individual levels. We then discuss how benefits of sharing are distributed across trips of various length and degree. Concluded with benchmarking ExMAS against previous, time-window based solutions and testing its' scalability, both for off-line and online applications.

Throughout the experiments we assume and kept constant the following parameters: value of time $\beta^t = -12.6\text{€}/\text{h}$, cost-sensitivity $\beta^c = -1$ (with negative signs to reproduce disutility of cost and time components), delay penalty $\beta^d = 1.5$, non-shared ride fare $\lambda^{ns} = 1.5\text{€}/\text{km}$, service time (extra time for each pick-up and drop-off) $t_s = 30$. Unless explicitly stated otherwise, we assume a network-wide speed of 29km/h and a willingness-to-share of $\beta^s = 1.3$. We solve the trip-ride assignment problem with the objective either to minimise vehicle-hours, or maximise travellers' utility. By default, we allow also for unprofitable rides in the solution. In all of the experiments conducted overall demand for on-demand services is exogenous and constant. The only passenger behaviour that we consider is their decision to opt for a shared ride (expressed as a binary choice using the deterministic sharing utility of Eq. 4). Other effects on mode choice, trip generation and other choice dimensions are not considered.

3.2. Results

First, we detail the results of matching 3000 trips in Table 2 with three main indicators computed for non-shared rides, shared rides and their relative difference. Travellers are offered a shared discount λ of 0.3 per-kilometre fare for shared ride 30% lower than for private ride), the planning horizon T is infinite, the matching objective function in trip-ride assignment is vehicle hours T_r . Thanks to sharing, the number of vehicle rides can be reduced by 37%, resulting in a reduction of approximately 30% in vehicle hours by attaining an occupancy of 1.67 instead of 1 for the non-shared scenario. On the other hand, since trips in shared rides are longer, passenger-hours T_q rise by 17.1%. Notwithstanding, thanks to the reduced price, total passenger utility is 6.1% higher.

3.2.1. Ride-sharing in various system settings

To assess, how stable these results are, we experiment with a range of system settings, varying one at a time, while keeping the others fixed to the base level. We are particularly interested in how the three key indicators (relative increase in total passenger utility ΔU_r , decrease in vehicle-hours ΔT_r and increase in passenger-hours ΔT_q) change under various conditions. We plot these key indicators as a function of:

- demand level (number of trips) $Q \in (100, 200, \dots, 5000)$, with $\lambda = 0.25$ and infinite planning horizon on Fig. 8,
- shared ride discount $\lambda \in (0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4)$ with demand $Q = 3000$ and infinite planning horizon on Fig. 9,
- scheduling horizon $T \in (5, 10, 30, 60, 120, 300, 600, 1200, \infty)$ with demand $Q = 3000$ and $\lambda = 0.25$ on Fig. 10

We can see that all of the above strongly impact the obtained indicators. The benefits, associated with sharing, increase with demand level Q (Fig. 8), sharing discount λ (Fig. 9) and scheduling horizon T .

Vehicle-hours are reduced by approx. 11% if 1000 trips are matched and drop by 22% when trip set size is 5000 (Fig. 8). Passengers also benefit more with higher demand levels, albeit exhibiting a weaker trend. Results stabilise for demand levels




Fig. 8. Gains from trip sharing expressed in terms of total utility improvement ΔU_r and vehicle hours reduction ΔT_r , with sharing costs reflected in increased passenger travel times ΔT_q . Passengers are offered a 25% discount and demand increases to 5000 trips per hour. Passengers' extra travel time increases with demand growth, yet stabilises around 10%. Similarly, passenger utility reaches a limit of 5%. Service cost reduction, however, steadily decreases, stabilising at around 20% reductions when the number of travellers reaches 3000.




Fig. 9. Gains from trip sharing expressed in terms of total utility improvement ΔU_r and vehicle hours reduction ΔT_r , with sharing costs reflected in increased passenger travel times ΔT_q . 3000 passengers are offered various sharing discounts λ . Low discounts hinder the potential of trip sharing, while high discounts unlock economies-of-scale. Turning point in vehicle-hours for higher discounts can be explained with profitability analyses of subsequent sections.

exceeding 3000 trips. This is a profound finding, strongly suggesting that the benefits of sharing may be obtained only after a certain critical mass of passengers is reached, above which results stabilise again.

Unsurprisingly, low sharing discounts constrain the benefits that can be generated by sharing (Fig. 9), with a threshold of around 10%, below which sharing is hardly attractive for travellers. Above 10%, however, both passenger utility and vehicle-hours improve with increasing discount rates. Hence, a certain discount rate is needed to unlock the efficiency gains made possible by sharing. Notably, at the same time passenger-hours increase, since travellers compromise their travel time to receive fare discounts. A turning-point is visible when increasing the discount from 35% to 40%. This is caused by exceeding the profitability level, as discussed in further sections.

We observe that the scheduling window T has a substantial impact on the obtained results (Fig. 10). While a scheduling horizon of more than 12 minutes already allows to fully exploit the trip-sharing potential, shorter horizons hinder it. Notably, while benefits for the service provider stabilise after 12 minutes, already 5 minutes seem to be enough for passengers to optimally share their trips, beyond which only limited gains can be made. Horizons shorter than 5 minutes, which seem to be the actual horizons used in practice, yield significantly lower benefits.




Fig. 10. Gains from trip sharing expressed in terms of total utility improvement ΔU_r and vehicle hours reduction ΔT_r , with sharing costs reflected in increased passenger travel times ΔT_q . 3000 passengers matched with varying scheduling horizon T (x-axis, in minutes). Short horizons do not allow reaching high benefits of sharing, specifically when horizons are shorter than 5 minutes benefits are low.




Fig. 11. Three indicators of trip sharing: a) vehicle hours reduction ΔT_r , b) passenger utility improvement ΔU_r and c) increased passenger travel times ΔT_q analysed against sensitivity to the willingness-to-share and the value-of-time. 2000 trips are matched into rides where a discount of 25% is offered, with a willingness-to-share varying from -1.05 to -1.5 and value-of-time ranging between -9 and -18 €/hour. Shareability benefits are highly sensitive to settings, the highest benefits are obtained for a low willingness-to-share (sharing penalty β^s) and a low value-of-time β^t . Otherwise, when the penalty is high and the value-of-time is high, sharing benefits diminish.

3.2.2. Ride-sharing with varying travellers' behaviour

To better understand the impact of traveller behaviour on trip shareability, we experiment with a range of value-of-time β^t and willingness-to-share β^s values. Unsurprisingly, their impact on the results is substantial. Essentially, the extent to which trips are shared and the associated benefits increase for lower values-of-time (since passengers are then ready to trade more time for the lower price of shared ride) and higher willingness-to-share. In Fig. 11 we alter the travel time multiplier β^t from -1.05 to -1.5 and double the value-of-time from -9 to -18 €/h, which results in dramatic changes in the KPIs, revealing strong diagonal patterns. For instance, vehicle-hours can be either barely altered (when value-of-time and travel time multiplier are high) or decrease by 25% (otherwise).

3.2.3. Trip-ride assignment optimal for traveller and system operator.

We compare two approaches in the trip-ride assignment: service provider perspective (minimising vehicle-hours T_r) and passenger perspective (maximise utility U_r), each of which yielding significantly different results. Table 3 reports the benefits of sharing 3000 trips with a discount of 25%, obtained under different cost formulations for the trip-ride matching (Eq. 19a). While costs in terms of additional travel time are similar when maximising utility and minimising ride travel time (around 10%), the benefits obtained are different. Maximising passenger utility yields a 7% improvement for passengers and 10% reduction in vehicle-hours, while minimising service time reduces it by 20% and yields a 4.4% improvement in utility terms. These, somehow orthogonal, results are both plausible, for different parties, and a practical solution may be attained by a combination of the two criteria in the objective function term with pre-defined weights.

Apart from the aggregate performance results, cost function impacts also the composition of ride degrees. When we match 3000 trips, the number of rides decreases as the discount offered goes up, with a diminishing share of single trips

Table 3

Benefits of sharing 3000 trips with 30% discount, obtained with objective to maximise passenger utility U_r or minimise vehicle hours $T_r \cdot X$.

c_r	ΔT_r	ΔT_q	ΔU_r
U_r	-0.115	0.104	0.070
T_r	-0.201	0.117	0.044




Fig. 12. Composition of rides degree (number of passengers per ride) for 3000 trips with varying discount for the objective of either maximising passenger utility U_r (left), or minimising vehicle hours T_r (right).

included in the solution (Fig. 12 left). Shared trips, however, consist mostly of second degree rides. Nevertheless, if we minimise vehicle-hours instead of maximising passenger utility, the number of single rides remains high, and so does the number of rides (2000 rides instead of 1800 for 0.4 discount). In this case (Fig. 12 right), however, for discounts greater than 25% attractive rides of higher degrees are found in the solution. Surprisingly, at the same time the share of single rides remains high. In conclusion, the compositions obtained for the two different objectives are substantially different, in terms of the number of rides as well as the presence of rides of higher degree rides and the share of single rides.

3.2.4. Occupation rate and vehicle capacity

As indicated above, various compositions of rides may be obtained for various settings. We further investigate the impact of capacity on the solution. We match 3000 trips with a discount of 30% and gradually increase the vehicle capacity (as an input for Algorithm 3). In particular we are interested in finding, whether extending capacity beyond a standard private car (three travellers plus driver) may be justified. In Table 4, we report the results of assigning 3000 trips into rides of various degrees with varying maximal capacity for two different objectives in the trip-ride assignment. When trips are matched to rides with the objective of maximising passenger utility, shared rides never reach a degree of four. Even when they are found, they are dominated in the solution by lower degree trips. Yet, trips exceeding a degree of four are found in the solution when we minimise vehicle hours, yet the number of these rides remains low, and their impact on the efficiency is minor. When we limit the capacity to three, vehicle hours are reduced by 24.3% with an occupancy of 1.55, which drops only to 26.3% and an occupancy of 1.57 for higher degrees, albeit at the cost of lesser traveller utility improvements (Table 4).


3.2.5. Ride-sharing profitability

We now turn to the service provider perspective, to analyse their benefits from offering shared rides. Since in our approach vehicle fleet is not explicitly considered, we use aggregate measures to assess supply-side profitability. Nonetheless, we are able to derive meaningful insights, first by tracing how the discount offered is compensated by vehicle kilometres reductions (Fig. 13), followed by revenue and fleet size reductions to better assess potential business models (Fig. 14). To conclude, we assess the case where profitability is considered at the individual (driver) level rather than centrally (Table 5).

Table 4

Trip composition for different maximum vehicle capacities. Results of assigning 3000 trips to rides optimal for travellers and vehicles with key indicators. In the top rows we show how the search space size increases with ride degree: we report the number of rides theoretically possible, actually explored and attractive for all passengers. Underneath it we report rides actually found in the optimal solution along with resulting KPIs. First for the solution optimal for travellers with a capacity rising from two to four, and then for the solution optimal for service provider, with a capacity reaching 9 and rides degree reaching 6.

	degree:	1	2	3	4	5	6	7				
search space:	theoretical	3.00×10^3	3.60×10^7	6.47×10^{11}	1.55×10^{16}	4.65×10^{20}	1.67×10^{25}	7.01×10^{29}				
	explored	3000	8,997,000	1807	226	123	24	0				
	attractive	3000	5270	243	130	76	8	0				
objective	capacity	1	2	3	4	5	6	7	O	ΔT_r	ΔT_q	ΔU_r
U_r travellers	2	726	1137	0	0	0	0	0	1.161	-0.051	0.101	0.125
	3	752	1118	4	0	0	0	0	1.164	-0.056	0.099	0.127
	4	699	1134	11	0	0	0	0	1.172	-0.061	0.100	0.127
T_r vehicles	2	1320	840	0	0	0	0	0	1.381	-0.209	0.091	0.073
	3	1317	495	231	0	0	0	0	1.505	-0.243	0.139	0.064
	4	1306	485	160	61	0	0	0	1.550	-0.251	0.160	0.060
	5	1330	459	159	55	11	0	0	1.567	-0.252	0.156	0.058
	6	1308	456	166	58	10	0	0	1.565	-0.257	0.163	0.060
	7	1326	468	162	49	10	1	0	1.575	-0.263	0.160	0.059
	8	1338	423	194	47	8	1	0	1.571	-0.264	0.157	0.061
	9	1422	435	160	44	8	2	0	1.542	-0.264	0.153	0.058

**Fig. 13.** Shared trips vehicle-hours reduction Λ_r as a function of discounts λ obtained for various demand Q .**Table 5**

Sharing 3000 trips with a discount of 25% while allowing only profitable rides or allowing all rides.

	ΔI	ΔL_r	ΔF	O
greedy, local	-0.202	-0.178	-0.202	1.307
system, global	-0.428	-0.264	-0.244	1.533

We analyse vehicle kilometres reductions for shared rides as expressed with Λ_r (Eq. 18), as a function of demand levels and discount offered at Fig. 13. The diagonal dashed line indicates a naive pricing level, i.e. when the service provider offers a discount equal to his costs reduction. For all the demand levels we see that the cost reductions are above the discount offered for some range of discounts. As discussed above, discounts below 5% bring no sharing benefits, thus it is not profitable to offer them. Discounts of 0.1–0.25 are firmly above the diagonal for all demand levels. As expected, higher demand levels bring a higher profitability. Nevertheless, regardless of the demand level, the system becomes unprofitable again for higher discounts offered. This suggests that the potential of trip-sharing is captured with discount levels of 10 to 30% and offering higher discounts cannot be compensated at such demand levels.




Fig. 14. Service operator benefits (reduced fleet size ΔF and vehicle-hours ΔT_r) along with revenues reductions ΔI (due to the discount offered for shared rides).




Fig. 15. Trip length distribution for various degrees of shared rides. We can observe that shared rides are typically longer, or in other words, short trips are unlikely to be shared. Rides shared by three travellers or more are typically long, since only a high discount can compensate for a longer travel and waiting time of a ride shared by many travellers.

For richer insights on the impact of shared rides on profitability, we further analyse this by, apart from previously analysed reduction in vehicle-hours, also investigating the fleet-size reduction (albeit a very rough approximation in this study) and total revenues (sum of fares paid with discount for shared rides and without for non-shared rides) for various demand levels with a 0.3 discount (Fig. 14). We see gradually decreasing revenues ΔI of up to 40% lower for 3000 trips, which may be compensated by up to 25% lower fleet size and 20% lower vehicle-hours, depending on the business model.

To understand whether shared rides profitability shall be considered based on local greedy considerations or global system considerations, we examine how the above results change, if we allow only for individually profitable rides in the assignment, i.e. all rides for which the yielded vehicle-hours reduction does not compensate for the respective discount are rejected (Eq. 17). The revenues are harmed by making such local greedy decisions (42% instead of 20%). Moreover, the obtained distance savings (ΔL_r) are also lower (18% instead of 26%). The fact that benefits for the service provider in terms of vehicle-hours are greater if profitability is not guaranteed at the ride-level, seems to advocate that the service provider shall evaluate sharing from the system perspective to maximise his benefits and occasionally accept trips which, if assessed in isolation, may seem unprofitable.

3.2.6. Distribution of sharing benefits among travellers

In addition to evaluating system-wide results, our model also enables the investigation of the distributional effects of service provision. We plot (Fig. 15) the trip length distribution for various degrees of shared rides, to explore whether the extent to which trips are shared is related to trip length. There is clearly a strong relation and average trip length increases with its degree. This is expected, since for short (and hence cheap) trips the discount hardly can compensate for the detours. Consequently, most of the short trips are not shared. For a similar reason most of the long trips are shared, since it is more likely to be possible to compensate for a detour with a higher fare reduction. This relation is also manifested by higher order trips: trips shared by four travellers are longer than trips shared by three passengers, since it is both easier to find four matching trips when they are long and they are also more likely to be attractive if they are long. This finding may shed




Fig. 16. The distribution of relative benefits of sharing among passengers of different degrees for shared rides.




Fig. 17. The distribution of relative benefits of two passengers sharing a FIFO (upper part) or a LIFO (lower part) ride depending on the order in which they are picked up. An unequal distribution is clearly visible for LIFO rides.

light on the future perspectives of using shared rides for first-last mile services, where trips are typically short (albeit often sharing the same origin and/or destination).

We also analyse the distribution of benefits stemming from sharing a trip among travellers. To this end, we plot the distribution of relative benefits of sharing among passengers of different kinds and degrees for shared rides. Fig. 16 reveals high variability (and thus inequality in the distribution) of benefits, with many trips of little relative improvement in their utility. Notwithstanding, the distribution is not characterised by long tails, hence sharing is not driven by few individuals, which are much better off, but rather by most individuals, being slightly better off through sharing their trip (with many by no more than 3%).

Unsurprisingly, the average utility when three travellers share a ride is higher than for four travellers. Though, pair-wise shared trips reveal two peaks in the distribution, which we further investigate using Fig. 17. When two trips are shared in a FIFO ride, the benefits are roughly equal among travellers; whereas when they share in a LIFO order, clearly the second traveller benefits more than the first one.

3.2.7. Ride-sharing with varying road congestion

We comment on the strong impact of travel times on the shareability results. Namely, the 14% vehicle-hour reduction for the current average speed in Amsterdam (29 km/h) increases to 16.5%, if the speed reaches 36 km/h. Conversely, it decreases to 12.5%, if the speed drops to 23 km/h. This suggests that, in less congested road networks, with a better performance as measured in terms of average speed, shared rides may bring higher benefits. If we also consider that increasing the number of shared trips may lead to a reduction in traffic flows on road networks (yielding lower travel times), this poses an interesting fixed-point problem (albeit of a limited effect as long as the modal share of shared rides remains low).

3.2.8. Benchmarking with time-window constrained approach

We benchmark our algorithm against previous approaches, where shared rides were assumed feasible if delay (extra waiting time at pick-up) and detour (extra travel time) were under a certain value Δ . They were explored for 2, 5 and 7 minutes in (Alonso-Mora et al., 2017) and up to 5 minutes in (Santi et al., 2014). To allow for model comparison, we adjust our algorithm so that a shared ride is assumed to be feasible as long as it fits within a predefined time window. We compare the results for 2000 trips in Amsterdam. We find that the 85th percentile of delay and detour for the results obtained with our algorithm are 300 and 212 seconds, respectively. We use those in the following comparative analysis.

Table 6

Differences in the results of the proposed algorithm and previous time-window based approaches. 2000 passengers are offered a shared discount λ of 0.3, matching objective function is minimising total ride distance L_r .

	Vehicle hours T_r	Passenger hours T_q	Utility U
windows	868,559	828,102	-16874
utilities	972,569	643,277	-12348
rel. diff (%)	10.7	-28.7	-36.7




Fig. 18. Distributions of attractiveness of shared rides identified with the proposed algorithm compared to those identified by the time-window based approach.

	0.1	0.15	0.2	0.25	0.3	0.35
500	2.2s (0.50)	2.5s (0.53)	2.6s (0.58)	3.3s (0.62)	3.9s (0.80)	4.5s (1.05)
1000	6.5s (1.02)	8.3s (1.09)	11.9s (1.32)	12.9s (1.56)	15.5s (2.12)	22.7s (3.43)
1500	14.0s (1.53)	20.3s (1.69)	24.2s (2.23)	33.2s (3.02)	43.2s (4.79)	49.1s (6.39)
2000	27.5s (2.06)	38.4s (2.38)	47.3s (3.15)	65.1s (4.69)	78.0s (7.23)	102.3s (12.13)
2500	46.1s (2.57)	65.9s (3.08)	85.6s (4.57)	109.1s (7.17)	147.6s (13.53)	199.2s (22.25)
3000	67.3s (3.11)	92.3s (3.96)	117.0s (5.75)	147.6s (9.96)	191.9s (15.99)	284.2s (31.43)
4000	131.2s (4.22)	181.4s (6.12)	225.0s (9.15)	292.2s (17.72)	432.8s (36.31)	667.9s (63.10)
5000	258.6s (5.47)	331.4s (7.43)	376.0s (13.80)	490.6s (26.14)	707.8s (51.59)	1223.7s (115.54)
6000	345.8s (6.38)	515.9s (10.10)	567.5s (20.70)	813.8s (39.73)	1166.6s (78.74)	2021.8s (165.27)
7000	523.8s (14.26)	766.6s (12.15)	879.1s (24.30)	1282.2s (63.72)	2304.9s (133.51)	4673.7s (359.44)

Fig. 19. Computation times in seconds and the number of feasible rides identified in the search space (thousands in brackets) for varying demand level (row) and a shared ride discount (column).

Evidently, the two algorithms yield substantially different outcomes (Table 6). In particular, the fixed time-windows approach induces higher travel impedance for passengers when compared with the utility-based algorithm proposed in this study. Consequently, the latter leads to higher service occupancy and requires a smaller fleet size. Notwithstanding, the time-window approach leads to fewer vehicle hours.

The approaches also differ in the number of shared rides that are identified. While the utility-based approach identified 19 000 rides of second degree, only 10 000 are identified when applying the time-windows approach. Conversely, while our algorithm identified only 1 000 shared rides of fourth degree, filtering only based on time-window constraints yielded 170 000 such rides. This suggests that the narrowing of an otherwise exploding search-space is obtained only when shared rides are attractive. In total, our algorithm identified 27 000 attractive shared rides, compared to 231 000 which are found fitting by the time-windows approach.

A significant share of the rides of higher degree generated by the time-windows approach are less attractive to travellers. In our algorithm, the utility of shared rides in general decreases with its degree (compare with Fig. 16) yet remains strictly positive, whereas in the window-based approach high degree rides, despite fitting in the windows, become highly unattractive to travellers (see Fig. 18). As long as the objective in the matching (Eq. 19a) is to minimise passenger hours, the results

Table 7

Computation time to match trips pooled into five minute batches for various demand levels (for a single batch of requests, rather than a full trips reported in Fig. 19.

demand Q [trips/h]	600	1200	2400	4800	7200	9600	12,000
time [s]	0.4	0.8	2.3	7.3	19.5	38.64	55.1

remain comparable, yet if we minimise for vehicle hours, the results of window-based approach may become unbearable for travellers.

In conclusion, the two different approaches led to the following twofold differences: (i) on one hand, fixed time windows do not allow identifying attractive trips, which do not fit the strict non-compensatory windows, and; (ii) on the other hand, a significant share of the trips fitting within the windows are not attractive.

3.2.9. Scalability

Finally, in this section we report the computation times of the proposed algorithm and the corresponding number of identified feasible rides. For all experiments we used a 2.3 GHz, 8GB RAM MacBook Pro. The implementation is single threaded, however graph searches of sect. 2.3.2 can be easily distributed over multiple cores for efficiency. Fig. 19 displays the algorithm computation times and the number of identified shared rides (thousands, in brackets) for different combinations of demand levels and discount rate for shared rides. Computation times remain feasible for high demand levels and high discount rates. In particular, the case of 7 000 trips per hour and a discount rate of 35% still yields a significant search-space reduction (compare with theoretical sizes reported in Table 1) and can be computed in time reasonable for off-line analyses.

To demonstrate the applicability of our approach in real-time settings we mimic the batch approach and in Table 7, we report time to match trips requested within a five minute interval as a function of the hourly demand. Computation time is less than 10 seconds on a standard laptop (even when the code is single-threaded) for number of travel requests around 5 000 trips per hour and approaches one minute for very high demand levels (above 10 000 trips per hour).

4. Discussion and conclusion

Mobility on demand can, potentially, contribute to reducing congestion only if rides are shared. The approach taken in this study is to identify a combination of trips, which are not only mutually compatible (i.e. feasible matches), but also where each individual trip remains attractive, in comparison to the private ride alternative. We propose a generic utility formulation to explicitly consider trade-offs between fares, delays and on-board discomfort by incorporating value-of-time and willingness-to-share parameters. Consequently, the solution space is substantially reduced allowing for the reformulation of the ride-sharing problem, so that an exact analytical method becomes feasible also for high-degree shared rides. For instance, the theoretical number of rides of fourth degree within the travel demand dataset of 3000 trips considered in our application is enormous, 1.55×10^{16} , whereas the number of searches needed with our method is 226 which is 10^{14} smaller. Consequently, our method allows analysing system-wide performance in terms of, among others, the share of on-demand trips that opt for sharing, overall passenger utility, vehicle hours and service occupancy.

We conducted an extensive set of experiments using travel demand data from Amsterdam and the results thereof reveal complex patterns and relations in the ride-sharing system. As expected, with sharing, vehicle-hours are reduced, at the cost of an increase in passenger hours. Notwithstanding, overall passenger utility increases due to the discounts offered, compensating for the longer travel times and lesser comfort. Consequently, service provider's revenues decrease, although these may be compensated by the reduced operational costs. Ultimately, whether this is profitable will depend on the underlying business model and the ability to attract demand beyond the current market share of on-demand services. The relative reduction in vehicle-kilometres strongly increases for higher demand levels, from 5% to 20% when demand grows from 500 to 3000 vehicles per hour in Amsterdam, suggesting that ride-sharing service properties are characterised by economies-of-scale. Service providers may adopt a myopic or a global but more risky strategy with regard to which rides may be shared. Our findings indicate that if the service provider requires that profitability is guaranteed at the individual ride level, the share of ride-sharing is significantly diminished, leading to the risk of impeding the attainment of economies-of-scale.

The results of assignment are substantially different depending on whether we design the trip-vehicle assignment for travellers versus if we optimise for service provider. Differences are in terms of the number of shared rides, their composition and the benefits for the respective parties. In particular, when designing the service to maximise users' utility, the matching yields 7% improvement in their utility and an 11.5% improvement for service provider. When matching is designed from service provider's perspective, the results yield a 20% improvement for service provider and an improvement of only 4.4% for travellers. Hence, both service users and service providers can simultaneously benefit from various settings of the ride-sharing, although the magnitude of which depends on the matching objective formulation and exhibits a trade-off between their interests. Future research may investigate the impact of different business models and regulatory conditions on a multi-criteria matching formulation. Furthermore, while our approach allows quantifying a series of key performance indicators relevant for service provider, coupling the method proposed in this work, with a microscopic representation of

fleet operations, including fleet rebalancing will allow to more precisely assess the fleet-related savings stemming from ride-sharing operations.

Our model also allows testing the impacts of two service design parameters: planning horizon and fee discounts. Introducing long planning horizons allows to exploit the full potential of sharing, albeit at the cost of making the system less user-friendly due to reduced spontaneity. Our results show, that the ride-sharing service is profitable within a certain range of discounts. Yet it becomes ineffective and unsustainable when the discount is outside of this range: if it is too low then travellers will not find it attractive to share a ride, if it is too high then cost savings cannot meet the reduced revenues. Interestingly, our results indicate that the upper limit for fee discounts increases with increasing demand. This stems from the underlying positive feedback loop, when high demand unleashes the sharing potential and allows offering higher discounts (35% instead of 25% when demand increases to 3000 from 1500 trips per hour). This relation shall be further exploited, calling for the formulation of an equilibrium state discount.

The compensatory utility term for assessing shared ride attractiveness allows testing how system performance depends on travellers' value-of-time and willingness-to-share. The share of ride-sharing is highest when travellers have a low value-of-time and a high willingness-to-share (i.e. low β^s multiplier). Future research on the behavioural determinants of travellers' choices, in the context of mobility on demand, will enable the empirical underpinning of the utility formulation adopted in this study. In particular, it can be expanded to a probabilistic choice model, including heterogeneity with respect to the value-of-time and willingness-to-share.

The prospects of ride-sharing also depend on network congestion, with two potentially counteracting effects. On one hand, if the network is congested the attractiveness of shared rides is hampered, due to the accumulation of delays along detours and thus reduce the share of ride-sharing. On the other hand, if the magnitude of sharing is sufficiently high, then it may contribute to relieving network congestion, thereby creating a positive feedback loop. Future research may investigate the conditions under which such a, potentially virtuous, cycle is created by considering modal shifts beyond on-demand services, as well as service reliability aspects of ride-sharing services.

When analysing the performance of ride-sharing systems it is important to consider the distributional aspects, since sharing and its benefits are not uniformly distributed across the trips and travellers. Short trips are less likely to be shared and shared trips are likely to be long. Notably, trips of a degree greater than two are feasible only if they are long enough, so that the fee discount in absolute terms becomes sufficiently large to compensate for the detours induced. Benefits are not uniquely distributed among sharing passengers, especially for LIFO trips, since the second passenger benefits at the cost of the first passenger. This calls for the future development of discriminatory pricing mechanisms, where the discounts are not uniform but rather dependent on the number of co-riders. Including customised approaches, which compensate for the general loss in riders' utility, based on the detours inflicted, the individual value-of-time and willingness-to-share values, as part of a yield management strategy.

The scalability and transferability of the proposed method facilitates variety of its future applications and extensions. It allows performing a cross-scenario sensitivity analysis of two-sided mobility platforms, with modal split, market competition and system-wide travellers welfare analysed from the demand perspective. It is designed to be seamlessly embedded in the analysis of pricing strategies, regulation policies or competition settings, including supply-demand feedback loops and emerging market equilibria. Thanks to utility-based formulation of sharing, the latest findings from travel behaviour research can be utilised and introduced into the models. In particular, proposed deterministic formulation can be extended to cover travellers heterogeneity and non-deterministic behaviour accordingly. Travel demand in ExMAS can become an endogenous variable. Thus enabling the analysis of the impacts of public transport service characteristics on the performance of ride-sharing and the potential positioning of the latter, between individual ride-hailing and public transport, in terms of both price and level-of-service. This will allow studying to what extent and under what circumstances, the ride-sharing service competes or complements public transport. ExMAS offers a comparative analysis tool for a wide range of cities, to establish how the shareability potential varies as a function of network and demand properties and their temporal and spatial distributions. Furthermore, the method can be used to analyse the feedback between service performance and service equity and coverage, assessing market entry strategies and on-demand platform competition. Proposed approach may be of particular interest to policy makers, who can now better understand system-wide implications of sharing rides and counteract to, potentially negative, impacts of emerging modes on equity and accessibility. While the method proposed in this study was primarily developed for offline matching applications, ExMAS also paves the way for improving online operations by means of (i) leveraging on utility-based filtering for attractive rides (instead of oversimplified non-compensatory thresholds) and (ii) improved graph searches. Finally, the method is available under public repository² and may be utilised in future research.

Declaration of Competing Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

² github.com/rafalkucharskipk/ExMAS

Acknowledgements

This research was supported by the CriticalMaaS project (grant number 804469), which is financed by the [European Research Council](#) and the Amsterdam Institute of Advanced Metropolitan Solutions.

References

- Agatz, N., Erera, A.L., Savelsbergh, M.W., Wang, X., 2011. Dynamic ride-sharing: a simulation study in metro atlanta. *Procedia - Social and Behavioral Sciences* 17, 532–550. doi:[10.1016/j.sbspro.2011.04.530](#). Papers selected for the 19th International Symposium on Transportation and Traffic Theory.
- Al-Ayyash, Z., Abou-Zeid, M., Kaysi, I., 2016. Modeling the demand for a shared-ride taxi service: an application to an organization-based context. *Transp. Policy* 48, 169–182. doi:[10.1016/j.tranpol.2016.02.013](#).
- Alonso-Gonzalez, M.J., Cats, O., Niels, V.O., Sascha, H.-L., Serge, H., 2020. On the influence of cost and time on the willingness to share a ride: a scenario analysis. *Transp. Res. Rec.* in press.
- Alonso-Gonzalez, M.J., Hoogendoorn-Lanser, S., van Oort, N., Cats, O., Hoogendoorn, S., 2020. Drivers and barriers in adopting mobility as a service (maas) a latent class cluster analysis of attitudes. *Transp. Res. Part A* 132, 378–401. doi:[10.1016/j.tra.2019.11.022](#).
- Alonso-Gonzalez, M.J., van Oort, N., Cats, O., Hoogendoorn-Lanser, S., Hoogendoorn, S., 2020. Value of time and reliability for urban pooled on-demand services. *Transp. Res. Part C* 115, 102621. doi:[10.1016/j.trc.2020.102621](#).
- Alonso-Mora, J., Samaranyake, S., Wallar, A., Frazzoli, E., Rus, D., 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceed. Natl. Acad. Sci.* 114 (3), 462–467. doi:[10.1073/pnas.1611675114](#).
- Arentze, T.A., Timmermans, H.J., 2004. A learning-based transportation oriented simulation system. *Transp. Res. Part B* 38 (7), 613–633. doi:[10.1016/j.trb.2002.10.001](#).
- Bicocchi, N., Mamei, M., 2014. Investigating ride sharing opportunities through mobility data analysis. *Pervasive and Mobile Computing* 14, 83–94. doi:[10.1016/j.pmcj.2014.05.010](#). Special Issue on Pervasive Education Special Issue on The Social Car: Socially-inspired Mechanisms for Future Mobility Services.
- Big Data Innovation Team, 2019. The transportation impacts of vehicle-for-hire in the city of Toronto. In partnership with: University of Toronto transportation research institute (UTTRI).
- Cáp, M., Alonso Mora, J., 2018. Multi-objective analysis of ridesharing in automated mobility-on-demand. *Proceedings of RSS 2018: Robotics-Science and Systems XIV*.
- Conway, M., Salon, D., King, D., 2018. Trends in taxi use and the advent of ridehailing, 1995–2017: evidence from the us national household travel survey. *Urban Sci.* 2 (3), 79. doi:[10.3390/urbansci2030079](#).
- Cordeau, J.-F., 2006. A branch-and-cut algorithm for the dial-a-ride problem. *Oper. Res.* 54 (3), 573–586. doi:[10.1287/opre.1060.0283](#).
- Frei, C., Hyland, M., Mahmassani, H.S., 2017. Flexing service schedules: assessing the potential for demand-adaptive hybrid transit via a stated preference approach. *Transp. Res. Part C* 76, 71–89. doi:[10.1016/j.trc.2016.12.017](#).
- Friedrich, M., Hartl, M., Magg, C., 2018. A modeling approach for matching ridesharing trips within macroscopic travel demand models. *Transportation* 45 (6), 1639–1653.
- Furuhata, M., Dessouky, M., Ordez, F., Brunet, M.-E., Wang, X., Koenig, S., 2013. Ridesharing: the state-of-the-art and future directions. *Transp. Res. Part B* 57, 28–46. doi:[10.1016/j.trb.2013.08.012](#).
- Hosni, H., Naoum-Sawaya, J., Artail, H., 2014. The shared-taxi problem: formulation and solution methods. *Transp. Res. Part B* 70, 303–318. doi:[10.1016/j.trb.2014.09.011](#).
- Kalczynski, P., Miklas-Kalczynska, M., 2019. A decentralized solution to the car pooling problem. *Int. J. Sustain. Transp.* 13 (2), 81–92. doi:[10.1080/15568318.2018.1440674](#).
- Kouwenhoven, M., de Jong, G.C., Koster, P., van den Berg, V.A., Verhoef, E.T., Bates, J., Warffemius, P.M., 2014. New values of time and reliability in passenger transport in the netherlands. *Research in Transportation Economics* 47, 37–49. doi:[10.1016/j.retrec.2014.09.017](#). Appraisal in Transport.
- Lavieri, P.S., Bhat, C.R., 2019. Investigating objective and subjective factors influencing the adoption, frequency, and characteristics of ride-hailing trips. *Transp. Res. Part C* 105, 100–125. doi:[10.1016/j.trc.2019.05.037](#).
- Li, S., Tavaafoghi, H., Poolla, K., Varaiya, P., 2019. Regulating tncs: should uber and lyft set their own rules? *Transp. Res. Part B* 129, 193–225. doi:[10.1016/j.trb.2019.09.008](#).
- Li, W., Pu, Z., Li, Y., Ban, X.J., 2019. Characterization of ridesplitting based on observed data: a case study of chengdu, china. *Transp. Res. Part C* 100, 330–353. doi:[10.1016/j.trc.2019.01.030](#).
- LLC, T. S., 2019. Nyc taxi & ride-hailing stats dashboard.
- Mourad, A., Puchinger, J., Chu, C., 2019. A survey of models and algorithms for optimizing shared mobility. *Transp. Res. Part B* 123, 323–346. doi:[10.1016/j.trb.2019.02.003](#).
- Narayan, J., Cats, O., van Oort, N., Hoogendoorn, S., 2019. Does ride-sourcing absorb the demand for car and public transport in amsterdam? In: *2019 6th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pp. 1–7.
- Pandey, V., Monteil, J., Gambella, C., Simonetto, A., 2019. On the needs for maas platforms to handle competition in ridesharing mobility. *Transp. Res. Part C* 108, 269–288. doi:[10.1016/j.trc.2019.09.021](#).
- Pelzer, D., Xiao, J., Zehe, D., Lees, M.H., Knoll, A.C., Aydt, H., 2015. A partition-based match making algorithm for dynamic ridesharing. *IEEE Trans. Intell. Transp. Syst.* 16 (5), 2587–2598.
- Qian, X., Zhang, W., Ukkusuri, S.V., Yang, C., 2017. Optimal assignment and incentive design in the taxi group ride problem. *Transportation Research Part B: Methodological* 103, 208–226. doi:[10.1016/j.trb.2017.03.001](#). Green Urban Transportation.
- de Ruijter, A., Oded, C., Javier, A.-M., Serge, H., 2020. Ride-sharing efficiency and level of service under alternative demand, behavioral and pricing settings. *Transp. Res. Rec.* in press.
- Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S.H., Ratti, C., 2014. Quantifying the benefits of vehicle pooling with shareability networks. *Proceed. Natl. Acad. Sci.* 111 (37), 13290–13294. doi:[10.1073/pnas.1403657111](#).
- Santos, D.O., Xavier, E.C., 2013. Dynamic taxi and ridesharing: A framework and heuristics for the optimization problem. In: *Twenty-Third International Joint Conference on Artificial Intelligence*.
- Simonetto, A., Monteil, J., Gambella, C., 2019. Real-time city-scale ridesharing via linear assignment problems. *Transp. Res. Part C* 101, 208–232. doi:[10.1016/j.trc.2019.01.019](#).
- Tachet, R., Sagarra, O., Santi, P., Resta, G., Szell, M., Strogatz, S.H., Ratti, C., 2017. Scaling law of urban ride sharing. *Sci. Rep.* 7 (1), 42868. doi:[10.1038/srep42868](#).
- Vazifeh, M.M., Santi, P., Resta, G., Strogatz, S.H., Ratti, C., 2018. Addressing the minimum fleet problem in on-demand urban mobility. *Nature* 557 (7706), 534–538. doi:[10.1038/s41586-018-0095-1](#).
- Wang, H., Yang, H., 2019. Ridesourcing systems: a framework and review. *Transp. Res. Part B* 129, 122–155. doi:[10.1016/j.trb.2019.07.009](#).