Delft University of Technology

Civil Engineering and Geosciences

Transport & Planning

Master of Science Thesis

The effects of pricing and service configurations on a ride-pooling service with pick-up and drop-off points

Author:
Marko Maričić (4526406)

Thesis Committee:
Dr. O. Cats
Dr. R. Kucharski
Dr. A. Bombelli
Dr. S. R. Danda

November 6, 2021

Preface

One of my biggest fears that I had when starting my Master's studies was started and even completing the so-called Master of Science thesis. I could not fathom the concept of choosing, creating, and completing your own project that lasts around nine months. But here I am, two years and a few months later with a complete document on my Master thesis. I can truly say that it was not even that bad. As a matter of fact, it was probably the most I have ever grown as a researcher, engineer, analyst, and a person in a given time-frame. I can honestly tell anyone that has the similar fear that I had that this will be the best period in your Master study, or even your entire study career thus far. However, working on an entire project your own is quite challenging and tiresome, this made me even more grateful of the chances I had to work with my buddy Henk on the countless projects and assignments of during in our Master.

Undoubtedly, the past couple of years have been a roller coaster ride for almost everyone, myself included. I am truly grateful for the people close to me that were always there to listen to my rants, complaints, ideas, and plans while also always trying to guide me to a positive direction and mindset. A large part of the positivity in my life is due to my parents, Maja and Joca, without them none of this would be possible. They were always able to persevere through the most difficult of situations while raising my sister and I. Of course, my sister Jovana was always there for me, travelling thousands of kilometres to surprise me on special occasions while in the meantime teaching me patience, kindness, and even Italian.

Of course, the past five years in the Netherlands have allowed me to become friends, and practically family, with some of the most amazing people I have met so far. Everyone within my circle was always ready to take on new challenges, activities, and adventures that could not have went any better. Even so, I was able to create a new home here in Zuid-Holland due to the amount of great people around me. A very special person to me by the name of Roos was one of the main people that helped me keep my sanity, happiness, and motivation during this surreal pandemic situation whilst my actually family was separated by thousands of kilometres. Thank you very much for being you and I can not wait for our next set of adventures together.

Furthermore, I give an enormous amount of appreciation to the thesis committee that was there to readily provide feedback and support through various stages of my thesis. Thank you Dr. Oded Cats for taking me up on the Additional Research Project exactly a year ago, I am extremely grateful for the opportunity to be able to work with you, have you as the Chair of the committee, the regular feedback you brought, and for the effort you made with coffee meetings which allowed me to follow the progress of countless students while also keeping my presentation skills up to date. I would also like to thank Dr. Rafał Kucharski for also taking me up on the Additional Research Project a year ago, I have learned a tremendous amount in coding and data handling from our time working together and I really appreciate your patience when I was still extremely novice installing ExMAS. As well, a lot of appreciation goes to Dr. Alessandro Bombelli and Dr. Santosh Danda for taking their valuable time to be part of my committee, your extensive and valuable feedback was more than helpful and definitely helped generate new perspectives.

Оваа теза сакам да ја посветам на Стефан, Павле и Никола. Секогаш ќе бидеш во моето срце. It is dedicated to you too Ryan, I cherish the times we had together and I hope you are all at peace now.

Marko Maricic
November 6, 2021

Nomenclature

Trip characteristics

Q_i Single trip request

o_i Traveller origin

d_i Traveller destination

t_i^p Requested pick-up time

$\hat{t_i^p}$ Actual time traveller is picked up

$\tilde{t_i^p}$ Actual time traveller is picked up from pick-up node

l_i Trip distance

t_i Trip duration

$\hat{t}_i$ Trip duration in a pooled ride

$\tilde{t}_i$ Trip duration in a pooled ride with PUDO

t_{w_i} Total walking time to and from PUDO nodes

Number sets

$\mathbf{Q}$ Set of trip requests for a given simulation time

$\tilde{\mathbf{Q}}$ Set of trip requests with possible PUDO nodes

$\tilde{\mathbf{Q}}_r$ Set of trip requests within a pooled ride with PUDO points

$\tilde{\mathbf{o}}_\mathbf{i}$ Set of possible pick-up nodes that within a radius from o_i

$\tilde{\mathbf{d}}_\mathbf{i}$ Set of possible drop-off nodes that within a radius from d_i

$\mathbf{O}_r$ Ordered sequence of served trip's origins

$\mathbf{D}_r$ Ordered sequence of served trip's destinations

Pricing variables

λ^{ns} Service fare of a private ride (€ /km)

λ^{s} Service fare of a pooled ride (€ /km)

$\lambda^{\tilde{s}}$ Service fare of a pooled ride with PUDO points (€ /km)

λ Relative difference in service fare between a pooled and a private ride, also referred to as shared discount

$\tilde{\lambda}$ Relative difference in service fare between a pooled ride with PUDO points and a private ride, also referred to as PUDO discount

Utility parameters

β^c Cost-sensitivity parameter

β^d Delay-sensitivity

β^s Sharing discomfort

β^t Value-of-time

β_v^t Value-of-time of vehicle

β^w Walking-sensitivity

Utilities

U_i^{ns} Traveller Utility of a private ride

U_i^s Traveller Utility of a pooled ride

$U_i^{\tilde{s}}$ Traveller Utility of a pooled ride with PUDO points

$U_{r_v}^{\tilde{s}}$ Utility of the vehicle of a pooled ride with PUDO points

Performance indicators

$\Delta u_{i,r}^s$ Traveller-level relative difference in utility between the ride-pooling with PUDO and the door-to-door pooling option

$\Delta t_{i,r}^s$ Traveller-level relative difference in in-vehicle travel time

ΔU_r^s Ride-level relative difference in total travellers' utility of a given pooled ride

ΔT_r^s Ride-level relative difference in vehicle travel time of a given pooled ride

U System-wide total traveller utility

T_v System-wide total vehicle-hours travelled

T_q System-wide total passenger hours travelled

T_{q_r} System-wide total in-vehicle passenger hours travelled

I System-wide total revenue generated

O System-wide average occupancy of a vehicle

ΔU^{ns} System-wide relative difference in total traveller utility comparing modal split of private, door-to-door, and PUDO rides to the case where only private rides are selected

ΔU^s System-wide relative difference in total traveller utility comparing modal split of private, door-to-door, and PUDO rides to the case when only private and door-to-door rides are selected

Algorithm parameters

α Service setting that has a range $0 < \alpha < 1$ and is used to guide the algorithm to favour the vehicle ($\alpha \approx 1$) or the travellers within a ride ($\alpha \approx 0$) or both

Executive Summary

Introduction

It is a well known fact that more and more humans are moving to cities, consequently increasing congestion as people are progressively moving to cities. Traffic congestion forced Americans to travel 8.7 billion hours more with a total estimated cost of $190 billion (Lasley, 2021) and is undeniably detrimental to the environment (Pant & Harrison, 2013). Indeed, public transport modes such as bus or metro help alleviate congestion by grouping numerous travellers into a single vehicle, yet it is not enough.

A key concept for future human mobility is ride-pooling[1] where multiple travellers' trips are matched to a shared ride, with respect to their routes and schedules. Accordingly, individuals can travel at a reduced fare (compared to ride-hailing) and drivers can save on their operating costs. This could result in the reduction of operator fleet size, alleviating traffic congestion, and reducing travel times for multiple types of road users. However, the extent of these benefits are still unclear and often disputed (Buhaug & Urdal, 2013; Ke, Yang, & Zheng, 2020; W. Li, Pu, Li, & Ban, 2019; Schaller, 2021).

The route and schedule flexibility of drivers is a constraint in ride-pooling as the matching of drivers and riders could be difficult at low demand levels. The ratio of matched participants depends on the distribution density of trips in space and time. Scenarios with low density (e.g. off peak hours) could suffer of a state where the demand for trips attracts an insufficient supply and vice versa (Furuhata et al., 2013). A high matching rate is critical for the success of ride-pooling services while also minimising the effort and inconvenience of riders(Stiglic, Agatz, Savelsbergh, & Gradisar, 2015). Higher matching rates and lower operating costs are the intended results of introducing pick-up and drop-off (PUDO) points in a ride-pooling service.

Meeting points are utilised in services such as *Uber Express Pool*,[2]. Here, multiple travellers in close proximity and with similar request times are matched to the same pooled ride and are then guided to a meeting point where they are picked-up from. The travellers are then dropped off at a certain distance from their respective destination, to which they would have to walk.

State of the art of ride-pooling

Looking into the existing reach of ride-pooling with PUDO points, Stiglic et al. (2015) propose an exact algorithm in which drivers and riders of a ride-pooling service are matched on a large scale. Space time windows can be used to formulate flexible and pickup delivery locations which is solved using Lagrangian relaxation (Zhao, Yin, An, Wang, & Feng, 2018) or through dynamic programming where a query of optimal meeting points are described as NP-hard (R.-H. Li, Qin, Yu, & Mao, 2015).

Fielbaum, Bai, and Alonso-Mora (2021) analyse ride-pooling with optimised pick-up and drop-off walking locations with heuristics. Fielbaum (2021) optimises the vehicle route in a ride-pooling system with an exact and heuristic algorithm where a third of vehicle costs could be saved when ride-pooling. Goel, Kulik, and Ramamohanarao (2016) propose a ride-pooling scheme that chooses optimal pick-up locations from a set of fixed locations through maximising car occupancy and preserving user privacy. Gurumurthy and Kockelman (2020) uses an agent based simulator to explore the advantages of using PUDO locations for dynamic ride-pooling.

Flex-route transit systems share similarities with ride-pooling systems as users are able to request a PUDO point that is not located along the original bus route. Zheng, Li, Qiu, and Wei (2019) introduces meeting points that users would walk to, employing mixed integer programming to serve as many trip requests as possible while minimising the total trip time of the accepted travellers. Wang, Zeng, Ma, and Guo (2021) utilises discrete choice modelling and a vehicle routing model to generate routes and schedules of customised buses to potential travellers by proposing a PUDO point.

Whether a (flexible) bus system or a ride-pooling system, minimising the detour of a vehicle is a common objective; research has shown that the introduction of meeting points increases the number of matched requests (thus increased the number of potential users) while decreasing the costs of the system

[1]Also known as ride-sharing.

[2]URL: `https://www.uber.com/newsroom/expresspool/`

(reduction in vehicle distance travelled). Similar results were also seen in systems optimising PUDO points (Fielbaum, 2021; Fielbaum et al., 2021; Stiglic et al., 2015; Zheng et al., 2019). However, it seems that extensive research exists on the ability to tackle the ride-pooling with PUDO problem efficiently, where the current state of the art aims to perform a fast and accurate optimisation. Moreover, the benefits of incorporating PUDO points to a ride-pooling service were also discussed although the extent of these benefits with respect to the service's configuration are still missing. For instance, the effects that system or behavioural parameters and demand levels have on operator performance and level of service are yet to be explored, to the extent of author's research. In essence, the selection of PUDO points for a ride would depend on the operator and the travellers within that ride. For this, the operator would look to minimise its operating costs, and thus searching for the shortest route, while travellers aim at maximising their utility, which consists of a number of different factors such as: discount price, waiting time, or walking time to/from PUDO location. Accordingly, this thesis aims at quantifying the effects such parameters would have on to the service performance and determine the best settings for certain scenarios.

Study objectives and research questions

This thesis aims to analyse the effect that a ride-pooling with PUDO points has on the operator performance and level of service of travellers. The proposed algorithm is intended to be an extension of an existing utility-based algorithm named ExMAS (Kucharski & Cats, 2020). Thus, the objective of this thesis is to examine the effects pricing, demand levels, and service settings have on the performance of ride-pooling with PUDO.

The main research question of this thesis can be stated as:

1. *How does the pricing of ride-pooling with PUDO locations affect its level of service and operating costs when comparing to its competition such as private ride-hailing and door-to-door ride-pooling?*

The research can be furthered with the following research sub-question:

2. *How do the demand level and the service settings affect these performance indicators?*

Methodology

Exact matching of attractive shared rides

The base matching algorithm used in this thesis is named ExMAS, a demand-driven algorithm that exactly matches trips into shared rides (Kucharski & Cats, 2020). With the use of efficient graph searches, the algorithm adequately narrows the search space to determine the exact solution of attractive shared rides. With the use of utility based formulation, the current modal split within ExMAS compares the attractiveness of a trip within a pooled ride to its non-pooled counterpart.

ExMAS calculates the traveller utility of a hailed, private trip (U_i^{ns}) of service-fare λ^{ns} (€ /km) and compares it to the utility of a trip in a shared ride (U_i^s, r) of service-fare λ^s (€ /km). A traveller finds a shared ride attractive only if $U_i^s, r > U_i^{ns}$ which is achieved through a reduced service fare ($\lambda^s < \lambda^{ns}$), or summarised as relative discount $\lambda = -(\lambda^s - \lambda^{ns})/\lambda^{ns}$.

Ride-pooling with PUDO

Problem formulation

The problem of ride-pooling with PUDO can be interpreted as an optimisation of two variables, the utility of the vehicle and the utility of all travellers within a ride. The utility of a traveller in a pooled ride with PUDO points is $U_{i,r}^{\tilde{s}}$ where the service fare is now $\lambda^{\tilde{s}}$ and a walking time t_{w_i} is incorporated. These longer pooled rides where travellers must walk must be offered a lower service fare ($\lambda^{\tilde{s}} < \lambda^s$) in order to make the utility of the trip attractive ($U_{i,r}^{\tilde{s}} > U_{i,r}^s$).
The total utility of travellers within a shared PUDO ride is $U_r^{\tilde{s}} = \sum_{i \in \tilde{\mathbf{Q}}_r} U_{i,r}^{\tilde{s}}$, or simply the sum of the utilities of all travellers within the ride. For the vehicle, or the operator, its utility for a ride is

$U_{r_v}^{\tilde{s}} = \beta_v^t \tilde{t}_{r_v}$, derived from the vehicle travel time ($\tilde{t_{r_v}}$) which is the length of time from picking up the first traveller and dropping off the last traveller. The parameter β_v^t (€ /s) represents the value-of-time of the vehicle.
The utilities of travellers can be then plugged into the maximisation objective function shown below. This is done on a trip level where the utilities of all trips within a ride are maximised. To account for both operator and user, a weight (α) is set to $U_{r_v}^{\tilde{s}}$ while (1-α) is set to $U_r^{\tilde{s}}$ where $0 < \alpha < 1$. With this, we can configure for either the traveller ($\alpha \approx 0$), the vehicle ($\alpha \approx 1$), or treating the travellers and vehicle equally ($\alpha = 0.5$).

$$\text{argmax}\left(\alpha U_{r_v}^{\tilde{s}} + (1-\alpha) U_r^{\tilde{s}}\right)$$

Algorithm architecture

The algorithm seeks new routes for already matched pooled rides, the output of ExMAS. An exhaustive search with heuristics is performed by the algorithm with the flowchart visualised in Figure 1 and consists of three segments:

- Configuring ride data
  - Input demand data is generated using Albatross (Arentze, 2005) while network graph is constructed using OSMnx (Boeing, 2017). ExMAS generated trip requests are then refined by allocating each traveller a set of PUDO nodes within a predefined walking radius.
  - Each (pooled) ride is iterated through where a skim matrix of travellers' PUDO nodes is created. The distances of skim pairs are ordered in distance ascending order to improve search performance.
- Obtaining and assessing route distances
  - For a given ride, each feasible route with PUDO points is iterated through with current route distance constantly compared to the candidate route.
  - Shortest route approach is conducted until the last trip pair where the traveller utility of the given route is computed.
- Computing and assess travel and vehicle utility
  - The total utility (weighted sum of vehicle and traveller) is computed and compared to the candidate route.
  - Once search splace explored, or iteration limit reached, the candidate route with relevant information is saved.

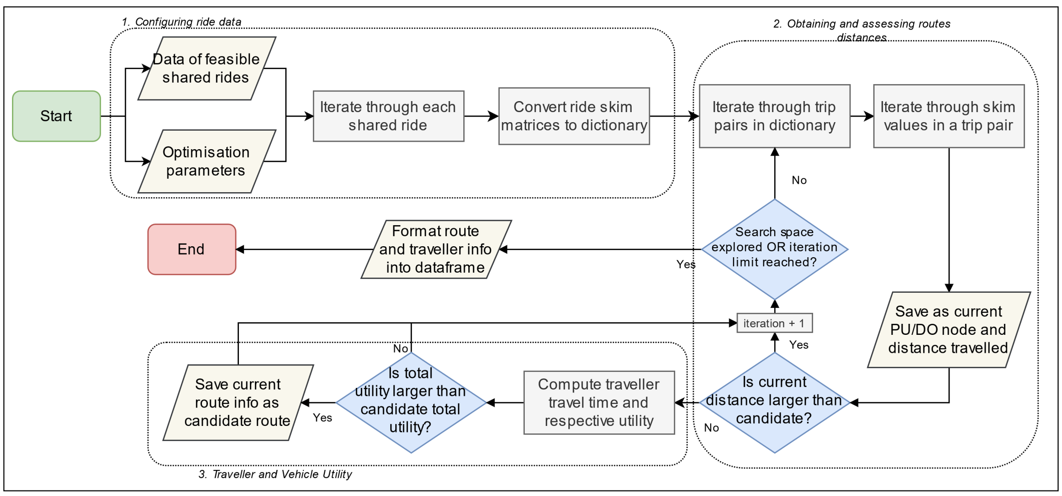

Figure 1: Flowchart visualising the sequences of the algorithm.

Application

Case study

The algorithm is applied to an Amsterdam network graph, which is generated with OSMnx (Boeing, 2017) with activity based synthetic demand created using Albatross (Arentze, 2005) that uses 2019 PC4 location data. ExMAS (Kucharski & Cats, 2020) is used to assign a series to trip requests to a set of attractive shared rides which would then be used for PUDO route search.

Experimental set-up

The experiment set-up is summarised in Table 1. The door-to-door ride-sharing λ and PUDO ride-sharing $\tilde{\lambda}$ discounts are intended to answer the main research question while the variables α and Q are intended to answer the sub research question and can be regarded as a sort of sensitivity analysis.

Table 1: Independent variables used for experimentation along with their respective ranges

Parameter	**Symbol**	**Range of values**	**Unit**
Door-to-door sharing discount	λ	{0.15, 0.2, 0.25, 0.3, 0.35}	-
PUDO sharing discount	$\tilde{\lambda}$	{0.2, 0.3, 0.4, 0.5}	-
Service setting	α	{0.5, 0.75, 0.95}	-
Demand level (No. of trips)	Q	{1000, 2000, 3000, 4000}	[trips/hr]

The results of the experiments will be quantified using performance indicators such as comparing the relative differences in utility and travel times. The relative differences in utility below are for trip and ride level, respectively. A traveller would deem a PUDO ride attractive if $\Delta u^s_{i,r} > 0$. For that ride to be selected, the sum of the travellers' must be larger than the door-to-door rival ($\Delta U^s_r > 0$).

$$\Delta u^s_{i,r} = (U^{\tilde{s}}_{i,r} - U^s_{i,r})/U^s_{i,r}$$

$$\Delta U^s_r = \textstyle\sum_{i \in \mathbf{Q_r}} \Delta u^s_{i,r}$$

Travel times can be also compared to on a traveller level and ride level. Travel times can be compared to that of the door-to-door and non-shared service (T^s_r and T^{ns}_r, respectively). With this we can summarise:

$$\Delta t^s_{i,r} = (\tilde{t}_{i,r} - \hat{t}_{i,r})/\hat{t}_{i,r}$$

$$\Delta T^s_r = (\tilde{t}_r - \hat{t}_r)/\hat{t}_r$$

System-wide indicators can also be used to quantify the performance of selected private and pooled rides (Kucharski & Cats, 2020). These are as follows:

- Total passenger utility, $U = \sum_{i \in \mathbf{Q}} U_{i,r}$
- Total vehicle-hours travelled, $T_v = \sum_{i \in \mathbf{R}} t_r$
- Total passenger-hours, $T_q = \sum_{i \in \mathbf{Q}} t_{i,r}$
- Total revenue, $I = \sum_{i \in Q} \lambda_i \cdot l_i$
- System-wide indicators of the three modes are compared using relative differences (e.g. $\Delta T^s_q = (T^{\tilde{s}}_q - T^s_q)/T^s_q$)

Results

Figure 2 visualises various system-wide indicators when the system is set to various λ and $\tilde{\lambda}$. It is seen in Figure 17a found that total passenger utility improves for increasing λ and $\tilde{\lambda}$. For larger λ utility improvement relative to a system of non-shared rides (ΔU^{ns}) while lowest relative to a system with shared rides (ΔU^{s}). This indicates that ΔU^{ns} improvement at high λ is largely due to the door-to-door pooled rides rather than pooled rides with PUDO. The ΔU^{s} is largest (around 2%) when the difference between λ and $\tilde{\lambda}$ is largest and at these points, ΔU^{s} and ΔU^{ns} are quite similar to each other.

The trends of seen with the increase in passenger hours travelled (i.e. ΔT_q^s) are almost identical to those seen with the utility improvement, seen Figure 2b. For all discounts, passenger hours increase (regardless if compared to non-shared or shared rides) as walking times are incorporated in the passenger travel time.

Figure 2c shows that the total vehicle hours reduced relative to shared rides (ΔT_v^s) has almost no correlation (or a really small correlation) with the discount prices offered. Although ΔT_v^s can be up to 2.2% when differences between λ and $\tilde{\lambda}$ are large. An almost linear negative trend is seen with ΔT_v^s and λ where largest ΔT_v^s occur when $\lambda = 0.35$. When comparing to non-shared rides, largest reductions in vehicle hours are due to the door-to-door pooled rides.

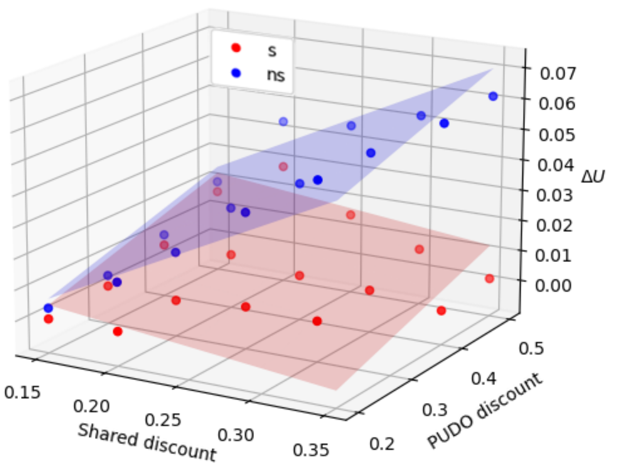

(a) Total passenger utility

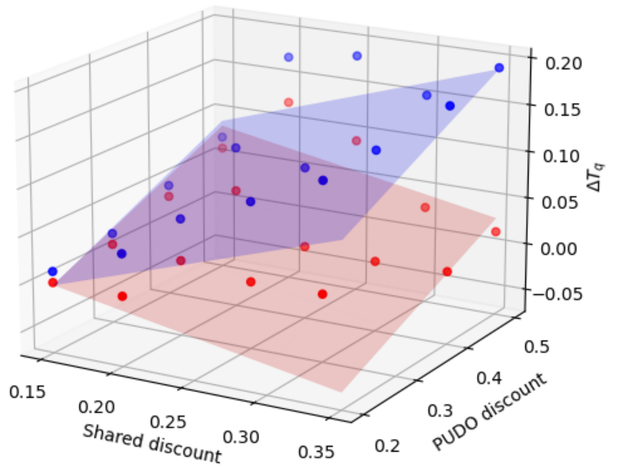

(b) Total passenger hours

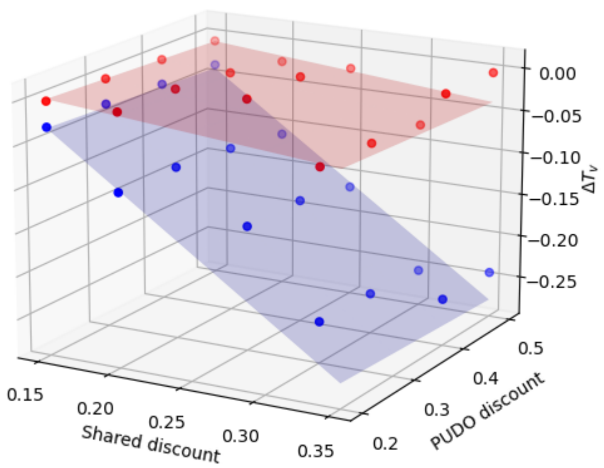

(c) Total vehicle hours

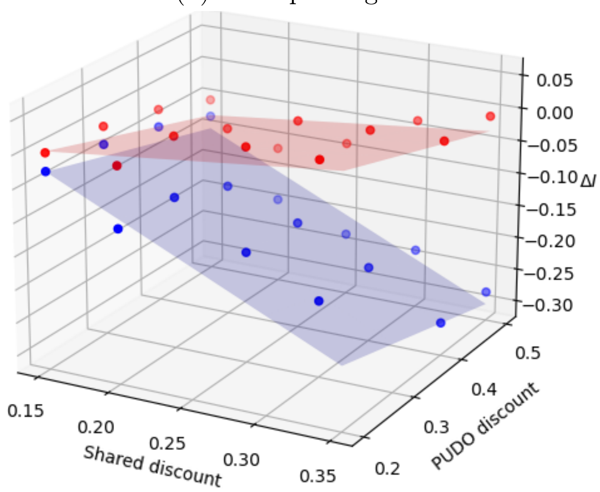

(d) Total revenue

Figure 2: Correlations of relative differences of various system-wide performance indicators with λ and $\tilde{\lambda}$. In this case, PUDO ride-pooling is compared to the door-to-door and private alternatives, visualised in red and blue respectively. $Q = 1000$ trips and $\alpha = 0.75$ for all experiments here.

Any pricing configuration produces a decrease in revenue due to the discounts offered. Nevertheless, increasing both the shared and PUDO discounts negatively affects the relative change in revenue with the ΔI^{ns} up to -28.1% when λ and $\tilde{\lambda}$ are largest. When comparing to shared rides, ΔI^{s} is up to -0.8% when $\lambda = 0.35$. When the difference between discounts is largest (i.e. $\lambda = 0.15$,$\tilde{\lambda} = 0.5$) ΔI^{s} is -5.5%. Although, when $\lambda = 0.15$ the difference between revenue generated for non-shared and shared is fairly small. This is understandable as at these low discounts, a small number of travellers actually opt for ride-pooling which lower the amount of revenue lost.

To shed light on the effect α has on the performance of ride-pooling with PUDO, when the difference between λ and $\tilde{\lambda}$ is low, minimal difference in ΔU^{s} exists between the different α values. As the difference between λ and $\tilde{\lambda}$ increases, the largest improvements in utility occur when $\alpha = 0.5$ while $\alpha = 0.75,\ 0.95$ are similar in their utility improvements. Similar patterns are seen when examining ΔT_q^s for various α values.

For ΔT_v^s, setting α to 0.5 will reduce the total vehicle hours the least. Again, $\alpha = 0.75,\ 0.95$ result in quite similar benefits where $\alpha = 0.95$ is able to reduce the total vehicle hours by a small margin. Again, The largest vehicle hour reduction, around 2.5% is seen when $\lambda = 0.2$, $\tilde{\lambda} = 0.5$, $\alpha = 0.95$. The correlation between ΔT_v^s and discounts offered range from 0-2.5%.

When looking at revenue, small differences between the λ and $\tilde{\lambda}$ result in similar revenues for all service settings. Once the difference between these two discounts becomes larger and especially when $\tilde{\lambda} = 0.5$, the resulting ΔI^{s} becomes more significant between $\alpha = 0.5$ and $\alpha = 0.75, 0.95$. The former produces the largest loss in revenue at these discount prices.

Examples of routes generated through different service settings are visualised in Figure 3. Here, a pooled ride with three travellers is further configured with the use of PUDO nodes. When $\alpha = 0.50$, the travellers are subject to lower walking times than when $\alpha = 0.95$.

(a) $\alpha = 0.50$ (b) $\alpha = 0.75$ (c) $\alpha = 0.95$

Figure 3: The route of a door-to-door pooled ride with three travellers, visualised in blue, and the new route with PUDO points in green. Each map represents a route obtained under different α values.

We also find that traveller utility increases linearly with $\tilde{\lambda}$ offered and decreases with increasing α. Larger demand levels facilitate the matching of higher degree pooled rides which help with further reducing total vehicle hours. Incorporating PUDO locations help further reducing the travel times of vehicles and even more so with increasing Q. Increasing Q correlates negatively with ΔU_r^s, likely due to the larger number of higher degree rides. The α setting could be used as a mediator for improving either the ΔU_r^s or ΔT_r^s depending on the requirements for the time of operation.

Conclusion, limitations, and recommendations

This thesis aims to build upon existing research on ride-pooling with PUDO points by examining the effect pricing configurations, service settings, and demand levels have on an example service. This was done by first building upon an existing and scalable algorithm that matches trip requests to attractive shared rides (Kucharski & Cats, 2020). The existing utility formulation was first extended and incorporated into a route search algorithm that is able to assess the potential of a ride's PUDO point configuration with respect to the vehicle and travellers' utility. Essentially, the algorithm was able to determine new routes for already matched pooled rides by introducing PUDO locations that travellers walk to and from.

To answer the first research question:

How does the pricing of ride-pooling with PUDO locations affect its level of service and operating costs when comparing to its competition such as private ride-hailing and door-to-door ride-pooling?

A series of system-wide, ride-level, and traveller-level performance indicators were used to describe the benefits and drawbacks of ride-pooling with PUDO. The system-wide indicators showed that total passenger hours increased while the total vehicle hours reduced due to travellers selected ride-pooling with PUDO, which is expected as travellers increase their travel time by walking while the vehicle minimises the detour due to sharing. The total utility of travellers generally improved with the discount offered while the revenue decreased as pooling offers a cheaper service for the same distance.

Through testing various discount configurations, ride-pooling with PUDO points performs best when the difference between the two offered discounts is large. Essentially, shared discount should be around 0.2 while PUDO discount around 0.4 as this creates the largest differences in performance between the two services. However, it was seen that at these discount prices, ride-pooling was able to reduce a further 1.0% of vehicle hours while revenue decreased a further 3.7%. For some operators, this could be a risky operation due the insufficient reduction in operating costs.

Looking at ride-level performance indicators, larger PUDO discounts increase the number of selected PUDO rides which results in a larger average ride utility. It was seen that ride utility increases linearly with the PUDO discount offered, largely due to the linear properties of the utility formula. On the other hand, higher discounts do not improve the reduction in ride vehicle travel time for both selected and non-selected rides. This is likely due to walking time only partially increasing with PUDO discount. It seems that, offering larger discounts to the service fare does not greatly increase the incentives to walk. Ultimately, this also depends on the fleet size and the capacity of a vehicle, which could be examined in future work.

To answer the research sub-question:

How do the demand level and the service settings affect these performance indicators?

For system-wide indicators, insignificant reductions or improvements were seen between the different service settings when the difference between $\tilde{\lambda}$ and λ were small. Larger differences in discount showed that setting the service setting to middle ($\alpha = 0.5$) created the largest improvements in utility and increased passenger hours the most. However, the same service setting creates the largest reductions in revenue and the smallest reductions in vehicle hours. Generally, larger α values ($\alpha \geq 0.75$) would create insignificant differences between one another in improvement or reductions in system-wide indicators.

On the ride-level performance indicators, $\alpha = 0.5$ generally results in the largest selections of pooled rides with PUDO. This is due to the generally shorter walking times imposed on traveller.

When varying the demand level, no significant differences were seen in the ride utility improvement. The utility of a ride does not seem to depend on the amount of travellers in the network. The average vehicle travel time savings seemed to slightly improve with demand but as well insignificantly. The vehicle travel time likely reduced to the larger number of higher degree rides at higher demand levels.

Limitations and recommendations for future work

The algorithm used in this report determines an optimal PUDO configuration through a shortest route search while the travellers' utility is only considered in the last step of the route search. The differences in search space between the algorithm and an exact exhaustive search that all optimum paths are not considered when service setting is more traveller focused. Furthermore, replications were not incorporated within this study as the running all the experiments once required relatively long run-times (i.e. at least one week to run all the experiments once). Optimising the run-time of the algorithm should facilitate the ability to run more replications in the same amount of time. The effect of behavioural parameters on the service was not analysed which means that the study is only applicable to a certain group of people.

Of course, there a number of further limitations which are simply beyond the scope of this report. For instance, this research takes on a planning perspective rather than a real-time optimisation which means that no traffic congestion is considered. This would undeniably alter the reported performance indicators and the actual PUDO stop locations. With traffic, such locations could not simply be set on intersections, rather they should be carefully chosen so that traffic would be minimally disrupted while a vehicle is picking up a traveller. As well, a vehicle is assumed to appear at the first traveller's pick-up point and empty vehicle travel is not considered. These are some of the points that would be extremely relevant when attempting to operate a real-time algorithm.

Contents

1 Introduction

It is a well known fact that the human population has been constantly growing in the past couple centuries and consequently increasing congestion as people are progressively moving to cities. As an example, traffic congestion forced Americans to travel 8.7 billion hours more with a total estimated cost of $190 billion (Lasley, 2021). Such additions to travel times can be detrimental to the environment as road traffic is one of the main sources of particulate matter emissions, especially in urban areas (Pant & Harrison, 2013). Indeed, public transport modes such as bus or metro help alleviate congestion by grouping numerous travellers into a single vehicle, yet it is not enough.

A key concept for future human mobility is where multiple travellers' trips are matched to a shared ride with respect to their routes and schedules, this is known as ride-pooling[3]. Accordingly, individuals can travel at a reduced fare (compared to ride-hailing) and drivers can save on their operating costs as the usage of the respective vehicle is increased. This could result in the reduction of operator fleet size, alleviating traffic congestion, and reducing travel times for multiple types of road users. However, the extent of these benefits are still unclear and often disputed (Buhaug & Urdal, 2013; Ke et al., 2020; W. Li et al., 2019; Schaller, 2021).

The route and schedule flexibility of vehicles is a constraint in ride-pooling as the matching of rides and riders could be difficult at low demand levels. The ratio of matched participants depends on the distribution density of trips in space and time. Scenarios with low density (e.g. off peak hours) could suffer of a state where the demand for trips attracts an insufficient supply and vice versa (Furuhata et al., 2013). A high matching rate is critical for the success of ride-pooling services while also minimising the effort and inconvenience of travellers (Stiglic et al., 2015). Higher matching rates and lower operating costs are the intended results of introducing pick-up and drop-off (PUDO) points in a ride-pooling service.

Meeting points are utilised in services such as *Uber Express Pool*,[4] seen in Figure 4a. Here, multiple travellers in close proximity and with similar request time are matched to the same pooled ride and are then guided to a meeting point where they are picked-up from. The travellers are then dropped off a certain distance from their respective destination. On the other hand, *ViaVan*[5] allows passengers to walk to a certain pick-up location that minimises vehicle detour while also maintaining a minimal access time, seen in Figure 4b. In this case, travellers are not grouped to a common meeting point but are rather matched with respect to their distance to the route of the vehicle. Both types of ride-pooling come with their own set of benefits and drawbacks.

This research devises an algorithm that optimises pooled rides by allowing respective travellers to walk/from potential PUDO points. The goal is to apply and test the algorithm under various system settings and examine the respective service performance. The service with PUDO points is intended to be comparable to the door-to-door and private ride competition.

The rest of this report is structured as follows, Sections 1.1 and 1.2 present the current state of ride-pooling with meeting/PUDO locations and the research questions of this thesis, respectively. Section 2 presents the mathematical formulation and algorithm architecture of ride-pooling with PUDO. Section 4 presents the results obtained through experimentation, with finally the conclusion, limitations, and recommendations for future work discussed in Section 5.

1.1 State of the art

Meeting points in a ride-pooling system can improve the percentage of matched riders, percentage of matched participants, and mileage. Ergo, travellers would have to walk a short distance and plan their departure more carefully in order to arrive at the meeting point on time. Stiglic et al. (2015) propose an exact algorithm that optimally matches in which drivers and riders are matched on a large scale. Space time windows can be used to formulate flexible and pickup delivery locations which is solved using Lagrangian relaxation (Zhao et al., 2018). Two novel solutions based on dynamic programming

[3] Also known as ride-sharing.
[4] URL: `https://www.uber.com/newsroom/expresspool/`
[5] URL: `https://viavan.amsterdam/`

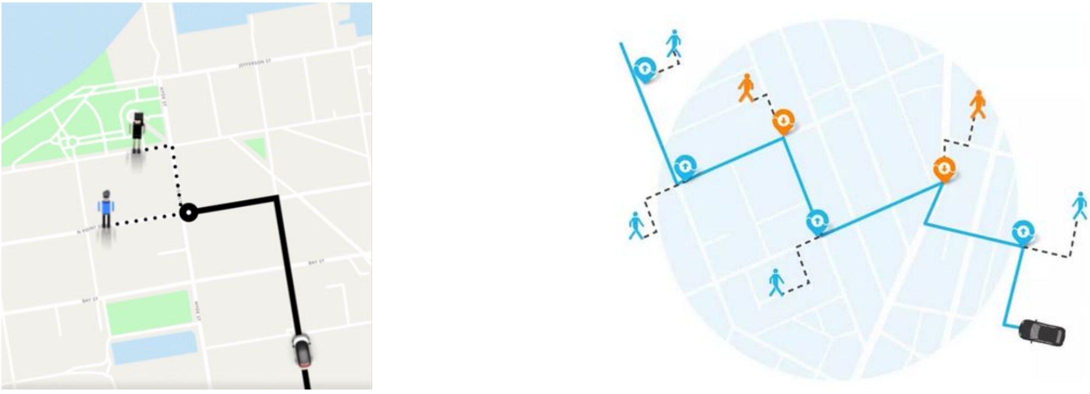

(a) After making a request, travellers would wait for an algorithm to match trips within close proximity to a single common meeting point, this particular concept is utilised by *Uber Express Pool*.

(b) Travellers assigned a PUDO location that minimises detour of a vehicle, trip requests that are in proximity of the vehicle's route are grouped together, this concept is utilised by *ViaVan*.

Figure 4

are proposed by R.-H. Li et al. (2015) to solve a query of optimal mutli-meeting-point route search in presented in where the problem is described as NP hard.

Research also exists on algorithms solving the ride-pooling problem with PUDO points. Fielbaum et al. (2021) analyse ride-pooling with optimised pick-up and drop-off walking locations where heuristics were utilised to reduce the computational time. The algorithm was first tested on a toy network and then applied to a Manhattan dataset. Fielbaum (2021) optimises the vehicle route in a ride-pooling system with an exact and heuristic algorithm. The results of these two are then compared to door-to-door counterpart where it was found that the detours reduced through implementing PUDO points to travellers can save a third of the vehicle costs when ride-pooling. Goel et al. (2016) propose a ride-pooling scheme that chooses optimal pick-up locations from a set of fixed locations which aims to maximise car occupancy and preserving user privacy. Gurumurthy and Kockelman (2020) uses an agent based simulator to explore the advantages of using PUDO locations for dynamic ride-pooling as average vehicle occupancy increases and empty vehicle miles travelled decreased with optimal PUDO locations. Jin and Xia (2021) combines transport theory with simulated annealing to create an integrated framework for solving the stop locations and passenger to stop allocation. An assessment scheme for meeting point locations is proposed in Czioska, Mattfeld, and Sester (2017) where a survey sought the stated preference of facilities at a meeting point, such as seating or lighting.

Flex-route transit systems share similarities with ride-pooling systems as users are able to request a pick-up or drop-off that is not located on the original bus route. Users are able to schedule such route deviations in certain time period before alighting. Zheng et al. (2019) introduces meeting points that user would walk to, employing mixed integer programming to serve as many trip requests as possible while minimising the total trip time of the accepted travellers. A fare reduction is given for passengers that walk to a meeting point with results demonstrating a significant reduction in rejection rate. Wang et al. (2021) utilises discrete choice modelling and a vehicle routing model to generate routes and schedules of customised buses to potential travellers by proposing a PUDO point. The disaggregated trip choice model showcases the variation of the choice probabilities of trips by modifying service fare, walking time, and travel time. The study found that the service was profitable to the provider (when compared to a service without PUDO points) as a larger number of trip requests were served.

Taking a step back, it is regarded that the general intentions of ride-pooling is to reduce the total vehicle hours and the required fleet size of the operator by matching multiple trips to a pooled ride, in theory this should alleviate traffic congestion in a city (Ke et al., 2020; W. Li et al., 2019; Z. Li, Hong, & Zhang, 2016). However, some research shows that ride-pooling systems could draw public transportation patrons while private car users (usually with higher income) would tend to switch to non-shared services due the reliability, comfort, privacy they provide (Schaller, 2021). Public transport is most frequently used by people with lower income for which ride-pooling services would be more attractive (Aberle, 2020). Travellers from different income levels and geographical backgrounds can also exhibit significant differences in choosing between a private or a pooled ride when provided with time and price trade-offs. Policies that combine promotional, curb management, and pricing strategies can be developed to even the attraction of such services amongst a large group of people of varying demographic (Lazarus, Caicedo, Bayen, & Shaheen, 2021).

Spatial disparities are also known to exist between different modalities. For instance, when looking at public transport (PT) and private cars, the former tends to have a shorter travel time when traversing shorter distances (i.e. < 3 km) and during peak hours. Furthermore, the use of cars seems to dominate during the night and in areas where PT schedule is limited (Liao, Gil, Pereira, Yeh, & Verendel, 2020). On the other hand, ride-sourcing is likely to have a significant increase in usage during peak hours where operational characteristics such as fleet size and utilisation rate are higher (Oh et al., 2020). Spatial disparities in modal choice also exist, for instance in Philadelphia, ride-pooling tends to be unbalanced in different neighbourhoods as service usage is limited to both the demographic and accessibility to infrastructure such as PT and amenities (Shokoohyar, Sobhani, & Ramezanpour Nargesi, 2020). Spatial pricing of ride-pooling and sourcing systems could be used to control the supply state of the provider in certain areas (Zha, Yin, & Xu, 2018).

Whether a (flexible) bus system or a ride-pooling system, minimising the detour of a (shared) ride is a common objective; research has shown that the introduction of meeting points increases the number of matched requests (thus increased the number of potential users) while decreasing the costs of the system (reduction in vehicle distance travelled). Similar results were also seen in systems optimising PUDO points (Fielbaum, 2021; Fielbaum et al., 2021; Stiglic et al., 2015; Zheng et al., 2019). Nonetheless, it seems that extensive research exists on the ability to tackle the ride-pooling with PUDO problem efficiently, where the current state of the art aims to create a fast and accurate optimisation. Moreover, the benefits of incorporating PUDO points to a ride-pooling service were also discussed although the extent of these benefits with respect to the service's configuration are still missing. For instance, to the extent of author's research, the effects that certain parameters and demand levels have on operator performance and level of service are yet to be explored. In essence, the selection of PUDO points for a ride would depend on the operator and the travellers within that ride. For this, the operator would look to minimise its operating costs, thus searching for the shortest route, while travellers aim to maximise their utility, which consists of a number of different factors such as: discount price, waiting time, and walking time to/from PUDO location. Accordingly, this thesis aims to quantify the effects such parameters would have on to the service performance and determine the optimal settings for certain scenarios.

1.2 Study objectives and research questions

This thesis aims to analyse the effect that a ride-pooling with PUDO configuration has on the operator performance and level of service of travellers. The proposed algorithm is intended to be an extension of an existing algorithm named ExMAS, which is an exact, replicable demand driven algorithm that bundles trips into shared rides (Kucharski & Cats, 2020). The algorithm is based on utility-based formulation where the search is bound to attractive shared rides and graph searches with a sequence of predetermined nodes which then derives an exact solution to the search space. Thus, the objective of this thesis is to examine the effects pricing, demand levels, and optimisation parameters have on the performance of ride-pooling with PUDO points.

The main research question of this thesis can be stated as:

1. *How does the pricing of ride-pooling with PUDO locations affect its level of service and operating costs when comparing to its competition such as private ride-hailing and door-to-door ride-pooling?*

The research can be furthered with the following research sub-question:

2. *How do the demand level and service settings affect these performance indicators?*

2 Methodology

This section first summarises the base algorithm, ExMAS (Kucharski & Cats, 2020), that the thesis builds upon. The formulation of the ride-pooling with PUDO problem follows and the explanation of the representative algorithm in the end.

2.1 Exact matching of attractive shared rides ExMAS

The base matching algorithm used in this thesis is named ExMAS, a demand-driven algorithm that exactly matches trips into shared rides (Kucharski & Cats, 2020). With the use of efficient graph searches, the algorithm adequately narrows the search space to determine the exact solution of attractive shared rides.

With the use of utility based formulation, the current modal split within ExMAS compares the attractiveness of a trip within a pooled ride to its non-pooled counterpart. A single trip (i) is part of a subset of n trip requests defined as $\mathbf{Q} = \{Q_1, Q_2, ..., Q_n\}$. Each trip request contains its own origin (o_i), destination (d_i), and pick-up time (t_i^p), thus $Q_i = (o_i, d_i, t_i^p)$.

For a private ride, the utility is formulated in Equation 2.1, note that a private ride caters to only one trip request.

$$U_i^{ns} = \beta^c \lambda^{ns} l_i + \beta^t t_i \tag{2.1}$$

Here, λ^{ns}, the service fare (€ /km), couples with trip distance (l_i) to compute the fee of a trip. The trip travel time (t_i) is computed as $t_i = t_{v_i} + t_{b_i} + t_{a_i}$, where in-vehicle travel time (t_{v_i}) is the duration from o_i to d_i while (t_{b_i}, t_{a_i}) represent the boarding and alighting time respectively. Multiplying the fee with the cost-sensitivity parameter (β^c) and travel time with the value-of-time parameter (β^t) allow for the non-shared trip to be expressed as a utility.

For a traveller within a pooled ride r, a few more attributes are considered for the utility computation, seen in Equation 2.2.

$$U_{i,r}^s = \beta^c \lambda^s l_i + \beta^t \beta^s \left(\hat{t}_i + \beta^d \left(\hat{t}_i^p - t_i^p \right) \right) \tag{2.2}$$

Here, a shared service fare (λ^s) is applied instead. The traveller is picked up at $\hat{t}_i^p$ while the trip travel time is now $\hat{t}_i$, which consists of the same quantities as the non-shared travel time. Realistically, the travel time in a shared ride is longer than that of a private ride ($\hat{t}_i > t_i$) while pick-up time can be either before or after the requested time ($\hat{t}_i^p \neq t_i^p$). The delay is the difference between the pick-up time and requested time $\left(\hat{t}_i^p - t_i^p\right)$. Travellers in a shared ride are further penalised with the parameter for the discomfort of sharing (β^s) while the pick-up delay is weighted with delay-sensitivity (β^d). Note that multiple trip requests (or trips) make up a single pooled ride where the number of trips served (or number of travellers in a ride) make up the degree of the pooled ride.

Longer, less comfortable shared trips tend to require some sort of compensation to increase their attractiveness. This could be done by offering a lower λ^s such that $\lambda^s < \lambda^{ns}$. By offering a lower service fare for ride-pooling, a traveller would be rewarded for the inconvenience of pooling a ride. Undoubtedly, this reduction in service fare should match the operator's savings on the operating costs of the service. Offering larger discounts could increase the attractiveness of ride-sharing but it also runs the risk of being detrimental to the provider due to generating less revenue for the same trip requests.

To determine whether a shared ride is attractive or not, the difference between the utilities of a shared ride and non-shared ride should be taken, as seen in Equation 2.3.

$$U_{i,r} = U_{i,r}^s - U_i^{ns} = \beta^c \lambda l_i + \beta^t \left(t_i + \beta^s \left(\hat{t}_i + \beta^d \left(\hat{t}_i^p - t_i^p \right) \right) \right) + \epsilon \tag{2.3}$$

Here, $\lambda = \lambda^s - \lambda^{ns}$ or the discount offered for ride-sharing. It is important to note that β^c and β^t consist of a negative value as travel time and costs are perceived negatively. Furthermore, as relative utility is utilised to determine attractiveness, only a shared ride experiences a pick-up delay and thus set to null for a private ride. Essentially, it is assumed that private ride travellers are picked up at their requested time. Finally, The ϵ represents some random distribution, allowing for the representation both probabilistic and deterministic systems. However, for this assignment, the model is assumed to be deterministic (ϵ=0) for simplicity.

For a shared ride r, a traveller i would consider the ride attractive if and only if $U_{i,r} > 0$. To assist with readability, a relative discount is used where $\lambda = -(\lambda^s - \lambda^{ns})/\lambda^{ns}$. For instance, if $\lambda = 0.2$, it implies that the fare of a shared ride is 20% discounted that the private counterpart.
With the formulation above, it is evident that such utility functions could be easily extended for other modes such as public transport or ride-sharing with PUDO. In this thesis, ride-sharing with PUDO will be considered as its own separate mode (rather than nested within ride-sharing) where its utility will be compared to that of a (door-to-door) shared and a private ride. This is further described in the following section.

2.2 Ride-pooling with PUDO points

This section presents the mathematical formulation for ride-pooling with PUDO. Figure 5 visualises ride-pooling with potential PUDO points. Here, two different pooled rides of degrees two and three (i.e. two pooled rides of two and three travellers) are visualised with the blue and pink line respectively. The origin and destination of the respective trip requests are indicated labelled. The crosses seen around each trip destination and origin indicate the feasible PUDO points a traveller could walk to or from.

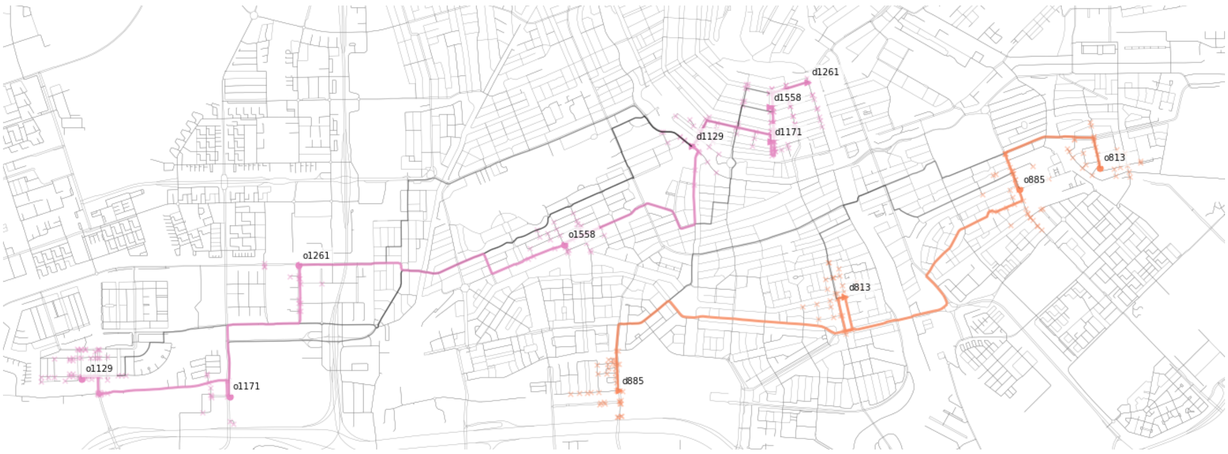

Figure 5: Pooled rides with two and four travellers, visualised by the orange and pink lines respectively. The origins and destinations are labelled for each traveller with the surrounding crosses indicating all feasible PUDO locations within a 5 minute walking radius.

2.2.1 Problem formulation

The problem of ride-sharing with PUDO can be interpreted as an optimisation of two variables, the utility of the vehicle and the utility of all travellers within a vehicle. It is important to now extend the contents of a trip request when assessing a given route PUDO locations. The extended form of trip requests with PUDO points can be described as $\tilde{\mathbf{Q}} = \{\tilde{Q}_1, \tilde{Q}_2, ..., \tilde{Q}_n\}$ where a single trip is $\tilde{Q}_i = (o_i, d_i, \tilde{\mathbf{o}}_i, \tilde{\mathbf{d}}_i, t_i^p)$. Essentially a copy of Q_i, $\tilde{Q}_i$ includes PUDO points $(\tilde{\mathbf{o}}_i, \tilde{\mathbf{d}}_i)$ which are sets of nodes within a predefined radius centred around travellers' (o_i, d_i). For a given road network graph $G(N, A)$, where N is the set of nodes and A is the set of route-able arcs, a traveller's origin and destination are obtained from $o_i, d_i \in N$ while the PUDO points are $\tilde{\mathbf{o}}_i, \tilde{\mathbf{d}}_i \subset N$.

From there a ride r that serves a number of trip requests, $\tilde{\mathbf{Q}}_r \subseteq \tilde{\mathbf{Q}}$, consists of the information seen in Equation 2.4.

$$r = (\tilde{\mathbf{Q}}_r, \mathbf{O_r}, \mathbf{D_r}, t_r^p) \tag{2.4}$$

Here, $\mathbf{O}_r$ and $\mathbf{D}_r$ are the ordered sequences of the served trips' origins and destinations, respectively. The departure time from the first origin is denoted as t_r^p. It is important to note that within ExMAS,

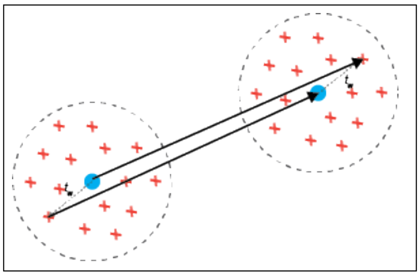

Figure 6: Origins of two travellers with two different routes of equal length. For each route, one traveller would have a zero walking time while the other would have a walking time of some seconds. Setting an exponent to t_{w_i} equalises walking times for all travellers in a ride.

all origins are served before destinations, or all travellers within a shared ride are picked up before being dropped off.

The utility for a traveller i in a shared ride with PUDO r is formulated in Equation 2.5 and is similar to the shared ride utility seen in Equation 2.2. However, now there is an additional of walking time variable (t_{w_i}) and a walking discomfort parameter (β_w). Travel time and pick-up time are now $\tilde{t}_i$ and $\tilde{t}_i^p$, respectively. The service fare is now $\lambda^{\tilde{s}}$.

$$U_{i,r}^{\tilde{s}} = \beta^c \lambda^{\tilde{s}} l_i + \beta^t \left(\beta^s \left(\tilde{t}_i + \beta^d \left(\tilde{t}_i^p - t_i^p \right) \right) + \beta_w t_{w_i}^{1.1} \right) \tag{2.5}$$

Walking time (t_{w_i}) is computed as the sum of walking times to pick-up point ($\tilde{o}_i$) from origin (o_i) and from drop-off point ($\tilde{d}_i$) to destination (d_i). Within Equation 2.5, an exponent of 1.1 is set to t_{w_i}, this allows for averaging of walking times for all travellers within a ride. Figure 6 shows an example of two traveller's origins with two different routes of identical length. If no exponent is set to t_{w_i}, then the two identical routes would have one traveller benefiting while the other disadvantaged. Adding the exponent to t_{w_i} tends to equalise walking times, generating ride-pooling routes where all travellers benefit. An example of a set of walking routes can be seen in Appendix A.

The travel time $\tilde{t}_i$ is computed in the same manner as indicated in Section 2.1. The in-vehicle time is the time from pick-up point to drop-off point. However, boarding and alight time is computed differently with PUDO ride-sharing, travellers picked up (or dropped off) from common PUDO locations will be rewarded with a lower boarding (or alighting) time. This allows the algorithm to favour meeting points through the higher utility obtained by reducing $\tilde{t}_i$ and potentially reducing passenger delay $\left(\tilde{t}_i^p - t_i^p\right)$.

It is important to note that the total travel time of utilising PUDO ride-pooling will generally be larger than door-to-door ride-pooling, namely due to the walking times of travellers. These longer journeys with walking involved will generally be less attractive than a door-to-door service. For this reason, the price of a pooled ride with PUDO points should be lower than the door-to-door alternative ($\lambda^{\tilde{s}} < \lambda^s$) which in turn provides incentives for travellers to opt for PUDO ride-sharing. Undoubtedly, the value $\lambda^{\tilde{s}}$ should be chosen so that the service or the system remains profitable which could be assessed by the magnitude of vehicle hours reduced through PUDO ride-sharing. For the sake of readability, $\lambda^{\tilde{s}}$ would rewritten as a discount relative to λ^{ns} where $\tilde{\lambda} = -(\lambda^{\tilde{s}} - \lambda^{ns})/\lambda^{ns}$. Thus Equation 2.5 can be rewritten as:

$$U_{i,r}^{\tilde{s}} = \beta^c \lambda^{ns} \left(1 - \tilde{\lambda}\right) l_i + \beta^t \left(\beta^s \left(\tilde{t}_i + \beta^d \left(\tilde{t}_i^p - t_i^p \right) \right) + \beta_w t_{w_i}^{1.1} \right) \tag{2.6}$$

The total utility of travellers within a shared PUDO ride is formulated in Equation 2.7 and is simply the sum of the utilities of all travellers within a ride. For the operator, its utility for a ride is described in Equation 2.8.

$$U_r^{\tilde{s}} = \sum_{i \in \tilde{\mathbf{Q}}_r} U_{i,r}^{\tilde{s}} \tag{2.7}$$

$$U_{r_v}^{\tilde{s}} = \beta_v^t \tilde{t}_{r_v} \tag{2.8}$$

Here, $U_{r_v}^{\tilde{s}}$ is the utility of a vehicle for a given ride and derived from the vehicle travel time, $\tilde{t}_{r_v}$, which is the length of time from picking up the first traveller and dropping off the last traveller. The parameter β_v^t (€/s) represents the value-of-time of the vehicle, or simply the operating costs. This parameter is used to convert the vehicle time to a utility, allowing it to be compared to traveller utility. It is important to note that both $U_v^{\tilde{s}}$ and $U_{r_v}^{\tilde{s}}$ are perceived negatively due to the negative signs of β^c, β^t, and β_v^t.

The utilities of travellers can be then plugged into the maximisation objective function shown in Equation 2.9. The route configuration takes place on a trip level where the utilities of all trips within a ride are maximised. To account for both operator and user, a weight, or service setting, (α) is set to $U_{r_v}^{\tilde{s}}$ while (1-α) is set to $U_r^{\tilde{s}}$ where $0 < \alpha < 1$. With this, we can seek the best route for either the traveller, the vehicle, both the traveller and the vehicle. Larger α ($\alpha \approx 1$) values will configure the route for the vehicle where larger walking times are likely to be induced onto the travellers. Smaller α ($\alpha \approx 0$) will aim to configure the route for the traveller where walking times are low, such routes could be similar to that of the door-to-door service.

$$\text{argmax}\left(\alpha U_{r_v}^{\tilde{s}} + (1-\alpha)U_r^{\tilde{s}}\right) \tag{2.9}$$

Figure 7 visualises the differences of routes when the objective function is set to a specific α value. Here, origins of travellers are visualised in points while destinations in triangles. The circle surrounding each origin and destination is the maximum walking radius of the traveller while the red crosses exhibit the available PUDO nodes. When $\alpha \approx 0$, the route resembles that of a door-to-door service as the route is configured for the traveller only. When $\alpha \approx 1$, the route usually touches the extremity of the circle as the objective function seeks the global shortest path which usually invokes the largest walking times. The travellers and vehicle are treated equally if $\alpha \approx 0.5$, the route is usually not the global shortest path but the walking times can be considerably lower; the computed route is likely to be shorter than the door-to-door service.

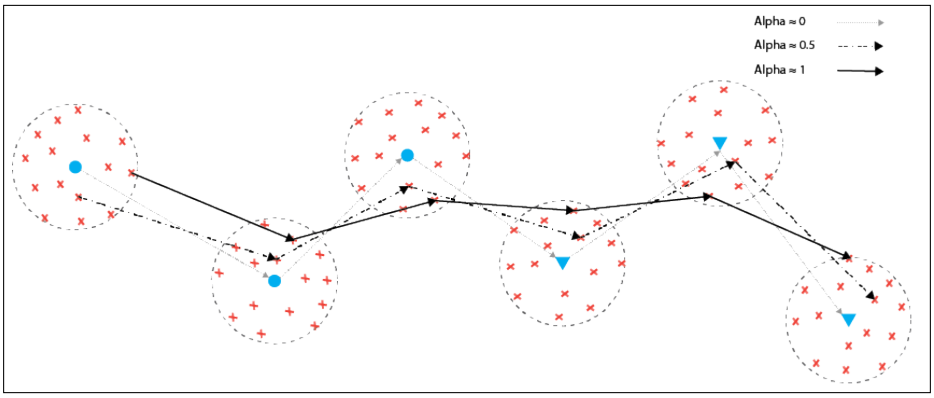

Figure 7: Visualisation of the ride-sharing with PUDO locations problem. Traveller origins and destinations are visualised with points and triangles respectively. PUDO locations are visualised with crosses which are limited to the dotted circle bounded by a walking distance radius. The different routes are obtained by setting a certain α to the objective function.

2.2.2 Algorithm Architecture

The algorithm that follows the formulation presented in Section 2.2.1. The flowchart of the algorithm in its entirety is visualised in Figure 8 with relevant notebooks (in Python) uploaded to a *GitHub* repository[6]. The entire algorithm consists of three different segments and are elaborated in this section, namely:

- Configuring ride data
- Obtaining and assessing route distances

[6]Url: `https://github.com/MarkoM-5/ExMAS/tree/PUDO-dev/ExMAS/notebooks/Afstudeer`

- Computing and assessing travel and vehicle utility

It is important to note that this algorithm does not match trips to shared rides but rather configures already matched shared rides. Essentially, trips that are matched as feasible shared rides within ExMAS are further configured by allocating PUDO points to each traveller. Therefore, the start phase in Figure 8 is only once ExMAS has has matched all possible trips into shared rides. This algorithm performs an exhaustive route search of all possible PUDO combinations with heuristics applied that improve search performance. An iteration limit is also set to prevent the route search from going to hundreds of thousands of (redundant) iterations. Setting an iteration limit could prevent the algorithm for detecting the global optimal route. It is therefore crucial to set up the ride data in a manner that allows for the algorithm to assess the best candidate routes before the iteration limit is reached.

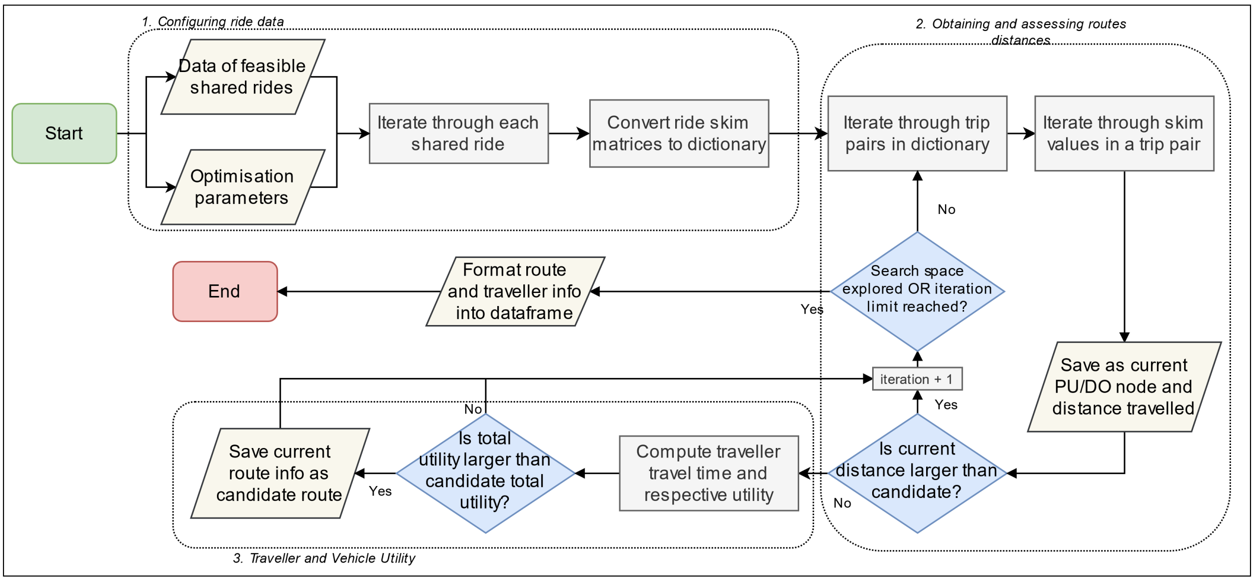

Figure 8: Flowchart visualising the sequences of the algorithm.

Configuring ride data

One of the first inputs of the algorithm, labelled as 'data of feasible shared rides' in Figure 8, is the expanded form of the trip requests generated by ExMAS where each trip is allocated a set of possible PUDO nodes ($Q_i \rightarrow \tilde{Q}_i$). The respective sequences of creating this data is visualised in Figure 9. The input demand data is generated using Albatross (Arentze, 2005) while the network is constructed using OSMnx (OpenStreetMap NetworkX, Boeing (2017)).

Figure 9 shows that the network graph and data frame of the generated trips requests are set as the input along with a certain value that constitutes as the walking radius for each traveller (i.e. 5 mins walking time). With these inputs known, each trip origin (o_i) and destination (d_i) is extracted. The network skim matrix generated by OSMnx is then reduced with respect to a trip's origin and destination. In the reduced skim matrix, we are interested in all the nodes that are within the distance set by the walking radius. These PUDO nodes are extracted for each trip requests where a new trip request data frame is created that contains the $\tilde{\mathbf{o}}_i, \tilde{\mathbf{d}}_i$ for each traveller.

With a data frame consisting of the PUDO locations for all trips, we can now go back to Figure 8. The data frame consisting of feasible shared rides is used as input, this data frame contains the information of which trips are matched to which and what type of ride. Basically, this allows for $\mathbf{O}_r, \mathbf{D}_r$ to be known. Optimisation parameters such as α, $\tilde{\lambda}^s$, and the iteration limit are also set.

We iterate through every ride and obtain the skim matrix for every trip pair ($\mathbf{O}_r + \mathbf{D}_r$). A skim matrix is constructed for each trip pair, if a ride consists of n trips then there are $n-1$ skim matrices. Each skim matrix consists of the distances from all the PUDO points of a trip request to the following trip request (i.e. $\tilde{\mathbf{o}}_1 \rightarrow \tilde{\mathbf{o}}_2 ... \tilde{\mathbf{o}}_n \rightarrow \tilde{\mathbf{d}}_1 \rightarrow \tilde{\mathbf{d}}_2 ... \tilde{\mathbf{d}}_n$). To improve search performance, the rows in a skim matrix are stacked and ordered with distances in ascending order. As well, the skim matrices of a ride are saved as a Python dictionary in order to improve look-up performance. Each trip pair is assigned its own index, i.e. first trip pair is given index 0, second trip pair is given index 1, and so on. Note that no distance computations are conducted in this algorithm rather, matrix searches and reductions from original network skim matrix are used to determine the best route.

Finally, before the route search can begin, the algorithm must be initialised. Empty lists are constructed to save distance values of each PUDO combination while predecessor values are assigned large values. This is further elaborated in the following segment.

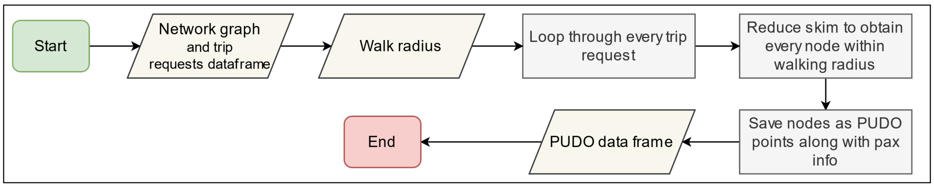

Figure 9: Flowchart visualising the procedure for assigning PUDO nodes to all travellers.

Obtaining and assessing route distances

In Figure 8, the second segment indicates its first process as iterating through every trip pair in the dictionary of skim matrices. Let us assume that there are M trip pairs ($M = n - 1$ with n being total trips in a ride) in this dictionary, $m = 0$ being the first trip pair. For every trip pair, we iterate through every skim pair, the destination node is then used to obtain the skim value for the following trip pair. In essence, there is a loop within a loop.

As the algorithm iterates through every trip pair, the distance obtained from the respective skim value is saved. Up until (and not including) the m^{th} trip pair, the current distance is compared to the 'candidate' distance. The candidate distance constitutes as the shortest distance thus far in the route search. If the current distance is higher than the candidate distance, the loop breaks and the next row in the skim matrix from the previous trip pair is used. An iteration is counted if a loop must break. Basically, the sum of distances obtained for and up to until the m^{th} trip pair is constantly compared to the sum of the distance from the candidate route. If the sum of distances turns out to be larger then the loop breaks and creates a new route from the previous trip pair.

It is important to note that in the zero-th iteration, the candidate distance is set to a very large number. The same applies for the utility of travellers and vehicle as this allows for the zero-th iteration to be saved as the candidate solution.

Once the m^{th} trip pair is reached, it the algorithm then moves onto segment three where various computations are made to evaluate the utility of the route for both travellers and vehicle.

Computing and assessing travel and vehicle utility

The third segment of the algorithm is when the m^{th} trip pair is reached. In Figure 8, the process of computing traveller travel time and respective utility is further elaborated in Figure 10. With the current route nodes as input, the walking times, boarding/alighting, and delay times for each traveller are computed. As stated in Section 2.2.1, common meeting or drop-off points are allocated with shorter boarding and alighting times.

Walking time constitutes as out of vehicle travel time while the boarding/alighting times constitute as the in-vehicle time. The in-vehicle travel time also constitutes the time it takes to travel the entire route with a given vehicle speed. With the trip order, in-vehicle, and out-of-vehicle travel time, the delay of each traveller can be computed. The delay is assumed to be the length of time that a traveller (thus vehicle is late) or vehicle (thus vehicle is early) must wait at a traveller's pick-up point. The first traveller in a pooled PUDO ride is assumed to have no delay, thus the vehicle picks up the first travellers immediately after reaching the pick-up point.[7] The following travellers in shared ride have their requested pick-up (t_i^p) time be at the time they reach their pick-up point while the actual pick-up time ($\tilde{t_i^p}$) is when the vehicle reaches that pick-up point. The vehicle can be early or late to this pick-up point, however this does not necessarily matter as the absolute value is taken.

[7] The first traveller is picked up at time $t_i^p + t_{w_i}$, therefore for this traveller $\tilde{t_i^p} - t_i^p = 0$

With all relevant times for a traveller known, the utility of each traveller can be obtained using Equation 2.5. While the utility of travellers can be obtained using Equation 2.7. The output from this is the total traveller utility.

Going back to Figure 8, the total traveller utility and the utility of the vehicle (which is obtained using Equation 2.8) is compared to the candidate utilities with Equation 2.9. If the sum of the current utilities is larger than the candidate utilities then the current route information gets saved as the candidate route information. Route information consists of PUDO nodes, travel times, and utilities for every traveller along with the vehicle travel time. The loop then breaks and goes to a previous trip pair to examine another route. If the the utilities are smaller than the candidate utility, nothing gets saved and the loop breaks and a different route is investigated.

The final decision checks whether the entire search space has been explored or whether the iteration limit has been reached. If either or a true, the algorithm stops and the candidate route information is exported as a data frame.

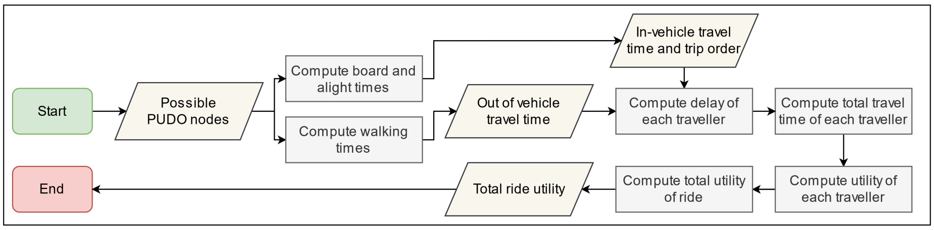

Figure 10: Flowchart visualising the procedure for computing the utility of travellers within a ride.

3 Application

To understand the benefits and consequences of using ride-pooling with PUDO, the algorithm is applied to a realistic urban context, further described in the following. This section also characterises the set-up of the independent variables used for the experimentation of the algorithm along with the relevant performance indicators used to assess the performance of the system and service.

3.1 Case study

As stated previously, the algorithm utilises a synthetic demand along with a network graph from a real city. In this case, the city in question is Amsterdam. Within the Netherlands, *Uber*, has grown as a major platform for mobility within Amsterdam and other larger Dutch cities. Out of a list of 50 countries that *Uber* operates in and arranged in terms of ride cost, the Netherlands was the second highest on the list where a 5 km ride costs \$10.07.[8] A traveller in the Netherlands could find such prices unattractive due to the country's already extensive (and cheaper) public transport network and infrastructure. Consequently, an *UberPool* service has not yet implemented due to insufficient demand which is based on the currently available *Uber* services in Amsterdam. This leads to a paradox, a lack of demand due to the high-cost of *Uber* impeding the roll-out of a cheaper ride-pooling service which could amplify the demand due to competitive prices. Ultimately, this prevents the launch of *UberPool* in Amsterdam as the company could deem it unprofitable to operate such a service with lower demand levels. Lower demand levels could result in sub-optimal trip matching, resulting in rides with large detours, unattractive for travellers, or unprofitable for the operator. Incorporating meeting or PUDO points in a ride-pooling scheme is a method of reducing detours a vehicle would make when ride-pooling therefore reducing the total travel times of both driver and user (Stiglic et al., 2015).

The travel behaviour data was obtained from Albatross (Arentze, 2005). This allows for the generation of a synthetic demand that spans over the course of 24 hours where each traveller is assigned its own mode. With a demand level set, ExMAS retrieves a random set of trip requests and matches them to attractive shared rides (Kucharski & Cats, 2020). The network graph of Amsterdam is generated using OSMnx (Boeing, 2017). The city's complex street network can then be decomposed into nodes and thus generating a skim matrix. For this paper, 2019 PC4 of the Netherlands is used to assign locations of the activities from the Albatross data which is then converted to the nodal locations using OSMnx.

It is important to note that even though Albatross assigns a mode to a traveller, some agents are given very long bike rides (i.e. longer than 10 km) or very short car car rides (i.e. less than a kilometre). For such reasons, the dataset was altered to only consider trips longer than 1 km. The mode assigned to an agent is disregarded and all agents are assumed as potential ride pooling or hailing users. The resulting dataset consists of over 170000 potential trips. On the other hand, the network graph consists of 26000 edges and 11000 nodes.

3.2 Experimental set-up

The experiments are based on the independent variables shown in Table 2. The combination of variables λ and $\tilde{\lambda}$ are intended to answer the main research question while the variables α and Q are intended to answer the sub research question and can be regarded as a sort of sensitivity analysis. The α variables were selected on the basis of the plots seen in Appendix B. Here, the search space of the algorithm (with heuristics) is compared to its exact counterpart in order to examine the consideration all suitable PUDO combinations and the effect α has on the heuristics. Since the algorithm contains heuristics that are related to the route length, the search space will find the route that provides the highest total utility (Equation 2.7) when $\alpha \approx 1$). However, Appendix B shows that lower α values return non-global optimal routes as the exact algorithm finds the true optimum at a much later iteration. For this reason, values of $\alpha < 0.5$ were not chosen for experimentation as the error could too large for a reliable conclusion.

[8] URL: `https://tinyurl.com/y5p35k28`

The parameters kept fixed throughout experimentation can be seen in **??**. The service-fare λ^{ns} was chosen to represent the price of an *Uber* ride in Amsterdam; a traveller would be indicated the price of a hailed ride or a discounted pooled ride before selection. As well, the β^t is based on the Dutch context (Kouwenhoven et al., 2014) while β_v^t was computed so that $\alpha U_{r_v}^{\tilde{s}} \approx (1-\alpha)U_r^{\tilde{s}}$ when $\alpha = 0.5$. The remaining utility parameters are based on estimates.

Table 2: Independent variables used for experimentation along with their respective ranges

Parameter	**Symbol**	**Range of values**	**Unit**
Door-to-door sharing discount	λ	{0.15, 0.2, 0.25, 0.3, 0.35}	-
PUDO sharing discount	$\tilde{\lambda}$	{0.2, 0.3, 0.4, 0.5}	-
Service setting	α	{0.5, 0.75, 0.95}	-
Demand level (No. of trips)	Q	{1000, 2000, 3000, 4000}	[trips/hr]

Table 3: Parameters that are kept fixed throughout the experimentation.

Parameter	**Symbol**	**Value**	**Unit**
Service fare	λ^{ns}	1.5	€ /km
Cost sensitivity	β^c	-1	-
Value of time	β^t	-0.0035	-
Value of time veh	β_v^t	-0.027	-
Sharing discomfort	β^s	-1.3	-
Walking discomfort	β^w	-1.1	-
Delay discomfort	β^d	-1	-
Max Vehicle Capacity	-	4	travellers
Simulation Time	-	1	hour

The results of the experiments will be quantified using performance indicators with some (such as ride utility and travel times) already mentioned in Sections 2.1 and 2.2. However, these are insufficient in demonstrating the performance of PUDO ride-sharing with respect to door-to-door ride-sharing and ride-hailing. Since we are looking to configure already matched rides, we only need to examine the differences in traveller utility between PUDO and door-to-door ride-pooling, described in Equation 3.1. Here, the relative differences in utility are for trip and ride level, respectively. A traveller would deem a PUDO ride attractive if $\Delta u_{i,r}^s > 0$. For that ride to be selected, the sum of the travellers' must be larger than the door-to-door rival ($\Delta U_r^s > 0$). For readability purposes, relative difference in traveller utility is denoted in lowercase while the utility of a ride in uppercase.

$$\begin{aligned} \Delta u_{i,r}^s &= (U_{i,r}^{\tilde{s}} - U_{i,r}^s)/U_{i,r}^s \\ \Delta U_r^s &= \textstyle\sum_{i \in \mathbf{Q_r}} \Delta u_{i,r}^s \end{aligned} \tag{3.1}$$

Travel times can be also compared to on a traveller level and ride level. Where the former is used to compute total passenger-hours (or minutes/seconds) while the latter calculates the vehicles hours, seen in Equation 3.2. To further quantify the performance of ride-sharing with PUDO, travel times can be compared to that of the door-to-door and non-shared service (T_r^s and T_r^{ns}, respectively). With this we can summarise:

$$\begin{aligned} \Delta t_{i,r}^s &= (\tilde{t}_{i,r} - \hat{t}_{i,r})/\hat{t}_{i,r} \\ \Delta T_r^s &= (\tilde{t}_r - \hat{t}_r)/\hat{t}_r \end{aligned} \tag{3.2}$$

System-wide indicators can also be used to quantify the performance of selected private and pooled rides (Kucharski & Cats, 2020). These are as follows:

- Total passenger-hours, $T_q = \sum_{i \in \mathbf{Q}} t_{i,r}$
- Total passenger in-vehicle-hours, $T_{q_r} = \sum_{i \in \mathbf{R}} t_r$

- Total vehicle-hours travelled, $T_v = \sum_{i\in\mathbf{R}} t_r$
- Total passenger utility, $U = \sum_{i\in\mathbf{Q}} U_{i,r}$
- Total revenue, $I = \sum_{i\in Q} \lambda_i \cdot l_i$
- Average Occupancy, $O = T_{q_r}/T_r$
- System-wide indicators of the three modes are compared using relative differences (e.g. $\Delta T_q^s = (T_q^{\tilde{s}} - T_q^s)/T_q^s)$[9]

On a final note regarding the experimental set-up, stochasticity for every experiment will exist as the Albatross dataset can be up 43 to 170 times larger than the retrieved number of trips for experimentation. Undoubtedly, this can result an extremely large number of trip combinations that could provide varying results. Figure 11 shows the ΔU_r^s and ΔT_r^s for every pooled PUDO ride when $Q = 1000$ trips with six replications. It can be seen that the distributions of each performance indicator for each replication are bounded within a similar shape where the differences between the two performance indicators arise when looking at the widths of the distributions. In essence, the distributions show that replications only create minor differences in the means.

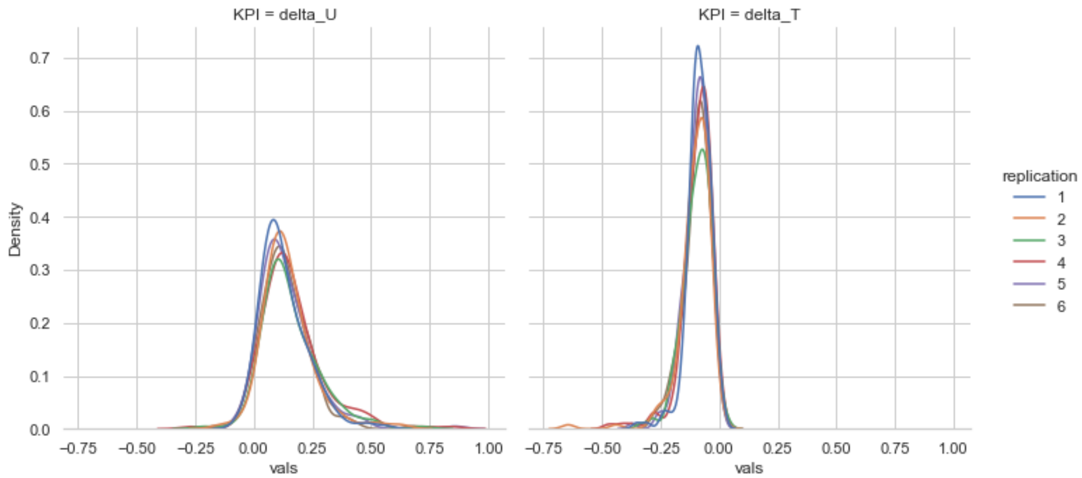

Figure 11: A set of replications conducted when $Q = 1000$ trips, $\tilde{\lambda} = 0.3$, and $\alpha = 0.95$.

It is recommended that three to five replications should be conducted when performing simulations (Law & McComas, 1991), however it is more precise to calculate the required number of replications with respect to a certain percentage deviation. Equation 3.3 shows that the number of replications n can be computed by rearranging the confidence interval formula (Robinson, 2004). Here, $\bar{X}$ is the mean of the output of the replications, S is the standard deviation of the outputs, $t_{n-1,\alpha/2}$ is retrieved from Student's t-distribution with $n - 1$ degrees of freedom and a significance of α (not be confused with service setting α), and d being the desired percentage deviation. With the outputs used to construct Figure 11, a significance level level (α) of 0.05 and percentage deviation (d) of 5%, a total of 25 and 10 replications for ΔU_r^s and ΔU_r^s would be required to achieve the desired accuracy.

$$n = \left(\frac{100 S t_{n-1,\alpha/2}}{d\bar{X}}\right)^2 \tag{3.3}$$

Now, the aforementioned experiments require considerably long run-times which would be made even longer by running 5 to 25 replications. As such, the results presented in this report are with respect to the experiments being run only once. Of course, this creates error due to stochasticity and if time allows, a handful of replications should be done.

[9]System-wide comparison are with respect to all selected rides from an optimisation. ΔT_q^{ns} compares system with PUDO pooled rides to non-shared alternative. ΔT_q^s compares system with PUDO pooled rides to EXMAS output (which involves non-shared and matched shared rides).

4 Results

With respect to the main research question posed in Subsection 1.2, the main objective of this research is to examine and quantify the outcomes that specific pricing incentives for ride-pooling (with PUDO) have on the operator and traveller. In this case, the system-wide, ride-level, and traveller-level performance indicators of ride-pooling with PUDO are compared to the door-to-door pooled and private alternative. Moreover, various pricing schemes were tested under various service settings (α) and demand levels (Q); this allows for answering the research sub-question.

Do note that the devised algorithm for this research is not intended to determine the global optimised path in the most efficient nor accurate manner possible. When setting α to lower values (i.e. $\alpha < 0.5$), the algorithm's search space is ineffective in obtaining the true global optimum path, this was discussed in Subsection 3.2. This thesis, instead, explores the application of PUDO points to an existing pooled ride with its performance analysed according to certain pricing and service setting scenarios. Bringing the results of all experiments for discussion would simply be unfeasible due to the amount of content. As a result, the following sub-sections discuss a selection of variables that were presented in Table 2.

As an example, Figure 12 visualises the resulting route optimisation for a two and four degree PUDO pooled ride when $\alpha = 0.95$ and $\tilde{\lambda} = 0.4$. It should be clear here that the pooled rides are able to further minimise their detours by appointing PUDO locations to a travellers. These reductions in detours cut the in-vehicle travel time of all the travellers within the pooled ride. As well, an instance of common meeting/drop-off points are also visible as the last two travellers in the four degree pooled ride are dropped off at a common PUDO node.

This section presents the results and analysis of the experimentation set-up presented in Section 3.2. First in Section 4.1, the system-wide performance indicators of ride-pooling with PUDO are compared to the shared (ExMAS output) and non-shared alternatives for various shared (λ) and PUDO ($\tilde{\lambda}$) discounts. Then, the effect of service setting α and pricing discounts are analysed in Section 4.2. The effect of $\tilde{\lambda}$ and α on the distribution of ride-level indicators follows, in Section 4.3. Then the effects of $\tilde{\lambda}$ and α on walking times are presented in Section 4.4. Finally, the effects of Q, $\tilde{\lambda}$, and α are summarised in Section 4.5.

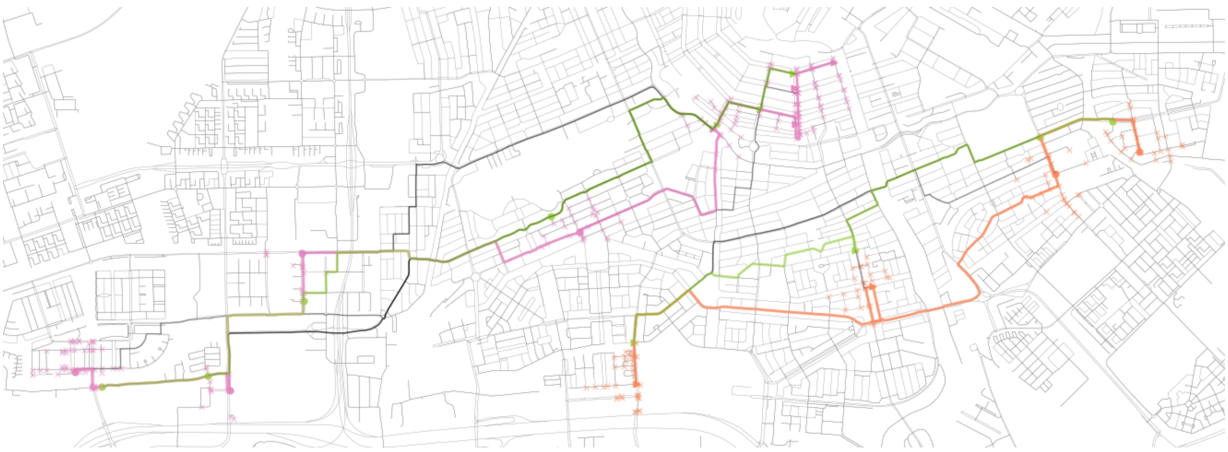

Figure 12: Pooled rides with two and four travellers, visualised by the orange and pink lines respectively where the new route with PUDO points for both is highlighted in green. Dots represent origins while triangles represent destinations.

4.1 Ride-pooling with PUDO under various pricing strategies

With the demand level Q held at a 1000 trips and α set to 0.75, the results in this section showcase the relative differences of total passenger hours, vehicle hours, passenger utility, and revenue differ for various discount configurations. First, Section 4.1.1 looks at how pricing affects a system with private and door-to-door pooled rides. Then, ride-pooling with PUDO is incorporated into the system where analysis of various pricing configurations take place, found in Section 4.1.2

4.1.1 Effects of pricing on a system with private and door-to-door pooled rides

Before quantifying and analysing the system-wide performance indicators, it is important to first examine how the modal-split between private and pooled rides vary with the λ offered. Figure 13 visualises the mode selection done by ExMAS when a certain λ is offered to travellers. Here, the larger discounts increase the attractiveness of ride-sharing which results in a larger number of travellers opting for (door-to-door) pooled rides. It can be immediately be seen that the λ offered plays a large role in whether travellers decide to share or not.

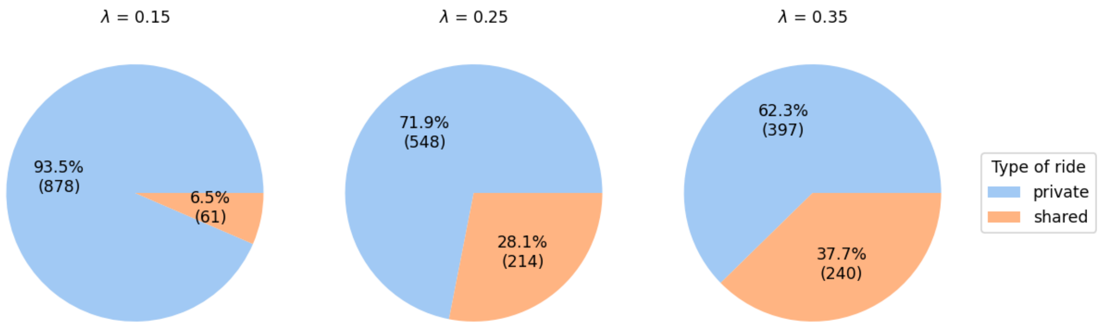

Figure 13: Modal split between private and door-to-door pooled rides for a given λ.

To further inspect the effects that reduced shared service-fares have on the system performance, Table 4 provides the relative differences of the system-wide indicators between the ExMAS optimisation (i.e. selected private and pooled rides) and the case where only private rides are selected. It is evident that higher sharing discounts tend to increase the total passenger hours travelled[10] as well as the total passenger utility and average occupancy of the vehicle. This matches the increase in the number of selected shared rides, seen in Figure 13, where higher discounts increase the attractiveness of sharing a ride. A larger number of shared rides results longer passenger travel times. As more travellers are able to fit in a smaller number of rides (higher occupancy), higher discounts provide larger savings in total vehicle hours. However, revenue decreases with these larger discounts where the reduction in revenue is around the same percentage as the reduction in vehicle hours when $\lambda = 0.2$. When $\lambda = 0.35$ the revenue is 4.6% lower than the vehicle hours saved; this is an indication that these higher discount prices may be quite unprofitable as the loss in revenue surpasses the vehicle-hour savings.

Table 4: System-wide indicators from ExMAS with $Q = 1000$ trips, the relative differences compare system of private and door-to-door shared rides to a system of only private rides for certain λ.

λ	**Pass. Hours** ΔT_q [%]	**Veh. Hours** ΔT_v [%]	**Pass. Utility** ΔU [%]	**Revenue** ΔI [%]	**Occupancy** ΔO [%]
0.20	4.4	-9.1	1.6	-9.2	14.9
0.35	17.3	-23.9	6.0	-27.5	54.2

4.1.2 Effects of pricing configurations on pooled rides PUDO points

With the effects of λ on system performance briefly discussed, we can now incorporate a new mode, namely ride-pooling with PUDO points. The incentives to opt for this type of pooled-ride now also depend on the $\tilde{\lambda}$ offered. The influence of $\tilde{\lambda}$ on ride-pooling mode choice can be seen in Figure 14. Here, λ is held at 0.2 which means that the number of private rides does not change, this is due to the fact the algorithm optimises already matched rides from ExMAS. However it should be clear from Figure 14 that increasing $\tilde{\lambda}$ increases the attractiveness of ride-pooling with PUDO as more travellers are seen to opt for that option. It is also visible that only a couple of pooled rides with PUDO are selected when $\tilde{\lambda} \approx \lambda$ while significantly more PUDO pooled rides are chosen when $\tilde{\lambda} \geq 2 \cdot \lambda$. This

[10]Note that for private and door-to-door shared rides, $T_q = T_{q_r}$ due to the entirety of travellers' journey being in-vehicle.

confirms the statement Section 2.2.1 that $\lambda^{\tilde{s}} < \lambda^{s}$ (or $\tilde{\lambda} > \lambda$) for travellers to deem ride-pooling with PUDO attractive.

Figure 14: Modal split between private, door-to-door pooled, PUDO pooled rides for a given $\tilde{\lambda}$ when $\lambda = 0.2$.

To further interpret the causes, benefits, and drawbacks of the modal splits seen in Figure 14, Tables 5 and 6 summarise the relative differences of various system-wide performance indicators with the non-shared[11] and shared[12] case, respectively.

Investigating Table 5 first, most system-wide indicators seem to follow the trends seen in Table 4 where increasing λ generally increases the magnitude of all the performance indicators. Likewise could be seen when increasing $\tilde{\lambda}$, total passenger-hours and utility improvement can increase up to four times. This is a result of accounting walking time for total passenger hours, moreover higher discounts should make walking more attractive for certain pooled rides (as total utility improves) due to more travellers opt for ride-pooling with PUDO. When looking at total traveller in-vehicle time $\Delta T^{ns}_{q_r}$, it is evident that larger $\tilde{\lambda}$ decrease the total amount of time travellers spend in a vehicle, or more specifically in a pooled ride. Intriguingly, $\tilde{\lambda} = 0.5$ is able to increase total time spent in-vehicle by only 0.6%, this indicates that travellers are able to replicate the in-vehicle travel time of a private ride if they opt for ride-pooling with PUDO.

Shifting the attention to Table 6, it should be immediately clear that the largest differences between the system types arise when difference between the two discounts is the largest. The performance of the systems are practically identical when $\tilde{\lambda} \approx \lambda$ (or $\lambda^{s} \approx \lambda^{ns}$), which is logical as virtually no pooled rides with PUDO are selected in such cases. Nevertheless, as the difference between the discounts increases, and more travellers opting for ride-pooling with PUDO, the total traveller in-vehicle travel time and total vehicle-hours shorten as detours of pooled rides are minimised. Although it seems that travellers are able to conserve their in-vehicle hours more than the vehicle itself, this results in lower average occupancy. Even though a common goal of ride-pooling is to increase vehicle occupancy, this lower average occupancy due to PUDO point incorporation is not necessarily bad-looking as travellers are able to minimise the amount of time spent within a vehicle more than the total vehicle hours reduced. Within a pandemic-like situation, this could be seen as beneficial for travellers that would still like to pool.

[11] Comparing to a system where only private rides operate.

[12] Comparing to a system where both private and door-to-door pooled rides operate.

Now, focusing on the case where $\lambda = 0.20$ and $\tilde{\lambda} = 0.4,\ 0.5$, bringing the optimal pricing to discussion. Increasing the $\tilde{\lambda}$ from 0.4 to 0.5 (i.e. 25% increase in discount offered) further decreases revenue by 7.7% while the total vehicle hours are only reduced by a further 1.2%. Indeed, travellers find such higher discounts more attractive which is seen in the 2.8% improvement in utility. However, it is quite likely that such a large decrease in revenue with a much smaller reduction in total vehicle hours could be quite unprofitable for the provider. Accordingly, to increase the likelihood of profitability, $\tilde{\lambda}$ should not be any higher than 0.4.

Table 5: Results obtained from the PUDO optimisation algorithm when $Q = 1000$ trips and $\alpha = 0.75$ for various λ and $\tilde{\lambda}$ configurations. The relative differences here compare the total system to a system of only private rides.

λ	$\tilde{\lambda}$	$\boldsymbol{\Delta T_q^{ns}}$[%]	$\boldsymbol{\Delta T_{q_r}^{ns}}$[%]	$\boldsymbol{\Delta T_v^{ns}}$[%]	$\boldsymbol{\Delta U^{ns}}$[%]	$\boldsymbol{\Delta I^{ns}}$[%]	$\boldsymbol{\Delta O^{ns}}$[%]
0.2	*0.2*	4.4	4.3	-9.2	1.6	-9.2	14.9
	0.3	5.0	4.1	-9.3	1.6	-9.5	14.8
	0.4	9.8	2.7	-10.0	2.2	-12.6	14.0
	0.5	17.1	0.6	-11.1	4.3	-19.6	13.1
0.35	*0.4*	17.7	17.2	-24.0	6.0	-27.5	54.3
	0.5	19.1	16.9	-24.2	6.1	-28.1	54.1

Table 6: Results obtained from the PUDO optimisation algorithm when $Q = 1000$ trips and $\alpha = 0.75$ for various λ and $\tilde{\lambda}$ configurations. The relative differences here compare the total system to a system of only private and door-to-door pooled rides.

λ	$\tilde{\lambda}$	$\boldsymbol{\Delta T_q^{s}}$[%]	$\boldsymbol{\Delta T_{q_r}^{s}}$[%]	$\boldsymbol{\Delta T_v^{s}}$[%]	$\boldsymbol{\Delta U^{s}}$[%]	$\boldsymbol{\Delta I^{s}}$[%]	$\boldsymbol{\Delta O^{s}}$[%]
0.2	*0.2*	-0.1	-0.1	-0.2	0.0	0.0	0.0
	0.3	0.6	-0.3	-0.2	0.1	-0.3	0.0
	0.4	5.2	-1.7	-1.0	0.7	-3.7	-0.8
	0.5	12.2	-3.7	-2.2	2.8	-11.4	-1.7
0.35	*0.4*	0.3	-0.1	-0.1	0.0	0.0	0.1
	0.5	1.6	-0.4	-0.3	0.1	-0.8	-0.1

With some pricing configurations quantitatively analysed, we can also examine the trends of certain performance indicators for all pricing configurations. Figure 15 plots the correlation of total relative difference of system-wide indicators with all λ and $\tilde{\lambda}$ offered. The relative differences shown in Tables 5 and 6 are a selection of the data plotted in Figure 15. In these figures, relative differences to private-only system and pooled/private system are highlighted with blue and red respectively.

Starting from the top-left, Figure 15a visualises the total traveller utility improvement for a certain pricing configuration. From here, it is clear that total traveller utility improves with increasing λ and $\tilde{\lambda}$, this was also seen in Table 6. For larger λ, ΔU^{ns} is the largest while ΔU^{s} is lowest. This indicates that for larger λ, the improvement in ΔU^{ns} is a result of the passenger utility of door-to-door pooled rides rather than pooled rides with PUDO points. The ΔU^{s} improvement is largest when the difference between λ and $\tilde{\lambda}$ is largest, however at this pricing configuration there is a minor difference between ΔU^{ns} and ΔU^{s}. This is due to low number of selected door-to-door pooled rides which was seen in Figure 13.

Figure 15b visualises the correlation of the mentioned discounts with the relative difference of total passenger hours travelled where the trends are similar to those seen in Figure 15a. It is deducible that that utility improvements correlate with passenger hours travelled as travellers that opt for PUDO ride-pooling will increase their total travel time due to walking to and from their designated PUDO points.

Now, the improvements in total vehicle hours for each pricing configuration are visualised in Figure 15c. The trend of ΔT_v^s should instantly stand out, or rather the lack thereof. Varying pricing configurations does not seem to create much difference in ΔT_v^s (when comparing to ΔT_v^{ns}), although vehicle reductions do exist up to %2 (as seen in Table 6). On the other hand, an almost (negative) linear trend is seen between shared discount and $\Delta T_v{}^{ns}$ with largest reduction in vehicle hours when $\lambda = 0.35$. This shows that when comparing to a system of only private rides, like ΔU^{ns} and T_q^{ns}, the reduction in total vehicle hours is mostly due to the matched shared rides from ExMAS.

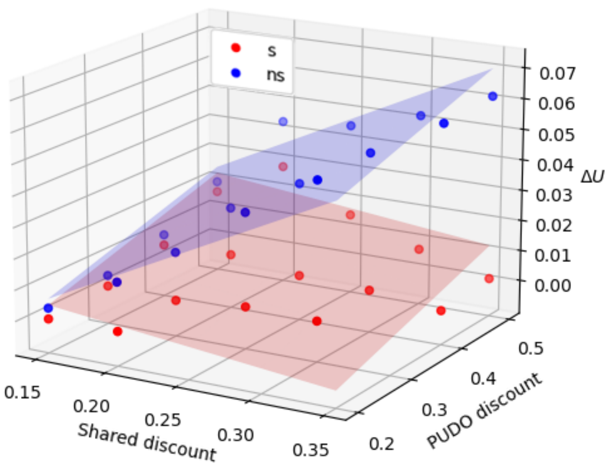

(a) Total passenger utility

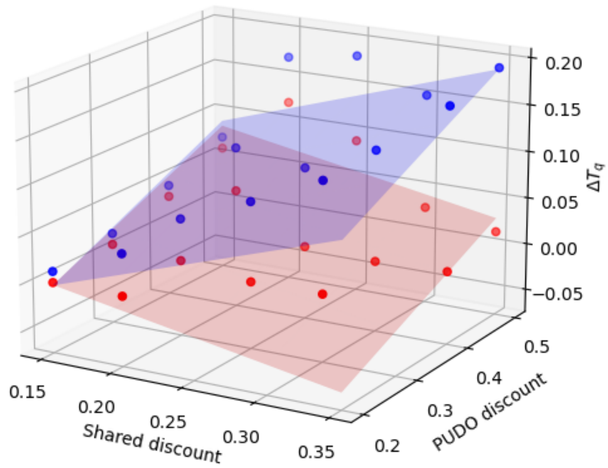

(b) Total passenger hours

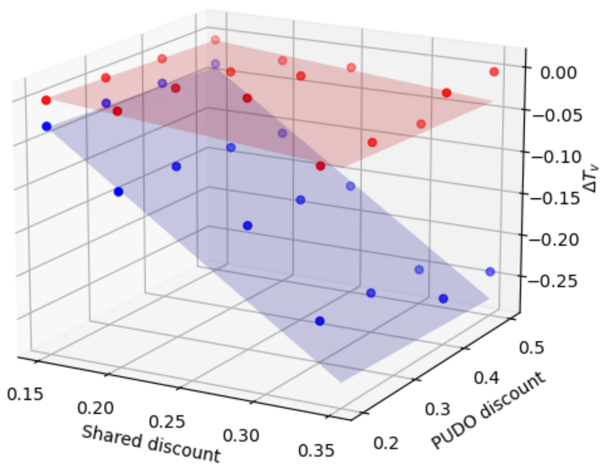

(c) Total vehicle hours

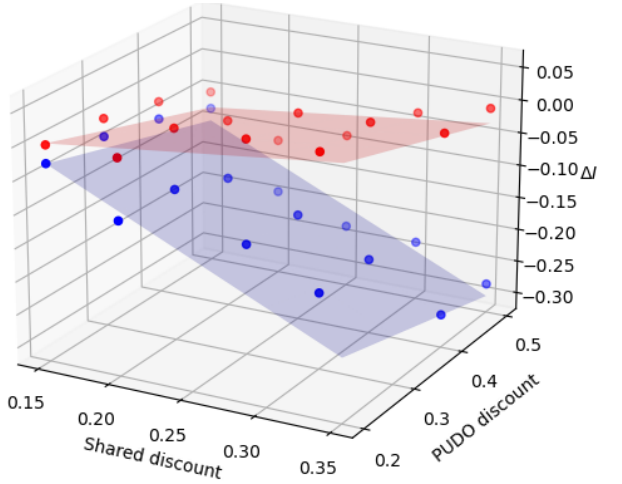

(d) Total revenue

Figure 15: Correlations of relative differences of various system-wide performance indicators with λ and $\tilde{\lambda}$ discounts. In this case, PUDO ride-pooling is compared to the door-to-door and private alternatives, visualised in red and blue respectively. $Q = 1000$ trips and $\alpha = 0.75$ for all experiments here.

Finally, Figure 15d plots the variation of revenue with pricing configurations. It should be immediately clear that any pricing configuration produces a decrease in revenue as travellers are offered a reduced service fare for sharing. It is logical to see that a system of only private rides produces the largest amount of revenue as this system operates at full service fare. Nevertheless, increasing both the λ and $\tilde{\lambda}$ negatively affect ΔI^{ns}. The largest loss in revenue occurs when λ and $\tilde{\lambda}$ offered are highest, this occurs due to the large number of selected PUDO pooled rides at this pricing configuration. Else-ways, ΔI^s is least negative when $\lambda = 0.35$, this was also seen in Table 6 where minor differences where seen between the two systems when $\lambda \approx \tilde{\lambda}$. When the difference between discounts is largest (i.e. $\lambda = 0.15$ and $\tilde{\lambda} = 0.5$) ΔI^s is greatest, however when $\lambda = 0.15$ the difference between ΔI^s and ΔI^{ns} is fairly small. This is understandable as at these low discounts, a small number of travellers actually opt for ride-pooling, the magnitude of lost revenue is dictated by the number of PUDO pooled rides.

4.2 Ride-pooling with PUDO under various service settings

This section characterises the modal split and performance of ride-pooling with PUDO when varying both the service setting α and PUDO discount $\tilde{\lambda}$. The value of α guides the algorithm to favour either the vehicle, the traveller, or both.

Figure 16 showcases the modal split between private, door-to-door pooled, and PUDO pooled rides for increasing α when the service is held to a certain pricing configuration. It is visible here that α does not change the amount of selected private rides, which is expected as α is used to optimise door-to-door pooled rides. Still, Figure 16 shows that almost double the number of pooled rides with PUDO are selected when $\alpha = 0.50$ than when $\alpha \geq 0.75$, which is expected as at $\alpha = 0.50$ the algorithm considers the vehicle and the travellers within equally. Interestingly, more pooled rides with PUDO are selected when $\alpha = 0.95$ than when $\alpha = 0.75$; the former considers mostly only the vehicle during the route optimisation (i.e. seeks for global shortest route for a given set of PUDO points).

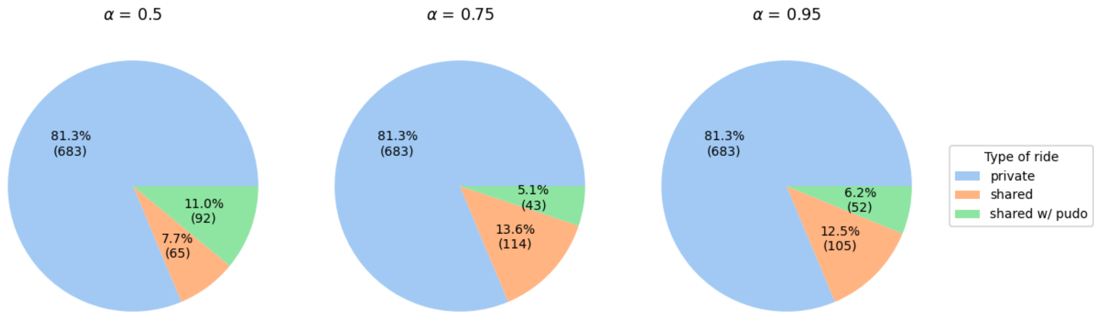

Figure 16: Modal split between private, door-to-door pooled, PUDO pooled rides for various service settings α when $\tilde{\lambda} = 0.4$ and when $\lambda = 0.2$.

To shed light on the effect α has on the performance of ride-pooling with PUDO, Figure 17 visualises the relative differences of system-wide indicators with respect to a system with private and door-to-door pooled rides. Here, various pricing configurations are also incorporated. Figure 17a plots the utility improvement for various service and pricing settings. When the difference between λ and $\tilde{\lambda}$ is small (i.e. $\lambda - \tilde{\lambda} \leq 0.15$), insignificant difference in ΔU^s exists between the different α values. As the difference between λ and $\tilde{\lambda}$ increases, it becomes clear that the largest improvements in utility occur when $\alpha = 0.50$ while $\alpha = 0.75, 0.95$ are similar in their ΔU^s. Moreover, the largest improvement in utility is visible when $\lambda = 0.2$, $\tilde{\lambda} = 0.5$, and $\alpha = 0.50$, although the utility improvement is considerably lower when $\lambda = 0.15$. This likely due to the low number of selected pooled rides at such low λ thus impeding any significant utility improvements.

The the increase in total passenger hours for service and pricing configuration is plotted in Figure 17b. Similar to Figures Figure 15a and 15b, it is clear that the correlations of ΔT_q^s are akin to those seen in Figure 17a. Varying α when $\lambda - \tilde{\lambda} \leq 0.15$ creates minor differences in ΔT_q^s. Larger differences between discounts offered result in larger ΔT_q^s, with largest occurring when $\alpha = 0.50$. This is not necessarily due to travellers walking more but rather more travellers opting for PUDO ride-pooling, as seen in Figure 16. It is also visible here that $\alpha = 0.75$ produces the lowest amount of passenger hours on some occasions. However, the resulting ΔT_q^s is only 0.5-1% larger than when $\alpha = 0.9$.

As the generated route and the number of selected PUDO pooled rides change with α and discount offered, it is important to examine how the total vehicle hour reduction vary with each setting. Figure 17c plots ΔT_v^s for various λ, $\tilde{\lambda}$, and α. It is immediately visible that setting α to 0.5 will reduce the total vehicle hours the least, which is sensible as $\alpha = 0.50$ should generate routes with shorter walking times and thus preventing large detour reductions. Again, $\alpha = 0.75,\ 0.95$ result in quite similar benefits where $\alpha = 0.95$ is generally able to reduce total vehicle hours the most by a small margin (when compared to $\alpha = 0.75$). The largest vehicle hour reduction, around 2.5% is seen when $\lambda = 0.2$ and $\tilde{\lambda} = 0.5$, or when the difference between λ and $\tilde{\lambda}$ is large. Undoubtedly, the correlation between ΔT_v^s and discounts offered are much clearer in Figure 17c than in Figure 15c but do note that ΔT_v^s ranges from 0-2.5%.

Finally, Figure 17d shows how the generated revenue for a certain α varies with discounts offered.

It is also visible here that when the difference between the shared and PUDO discounts small, the select α result in ΔI^s. Once $\lambda - \tilde{\lambda}$ becomes larger and especially when $\tilde{\lambda} = 0.5$, the resulting ΔI^s becomes much more significant between $\alpha = 0.50$ and $\alpha = 0.75,\ 0.95$. The former produces the largest loss in revenue at these discount prices. This is due to more travellers opting for ride-pooling with PUDO (seen in Figure 17a) when $\alpha = 0.50$ which results in more travellers paying less for a ride.

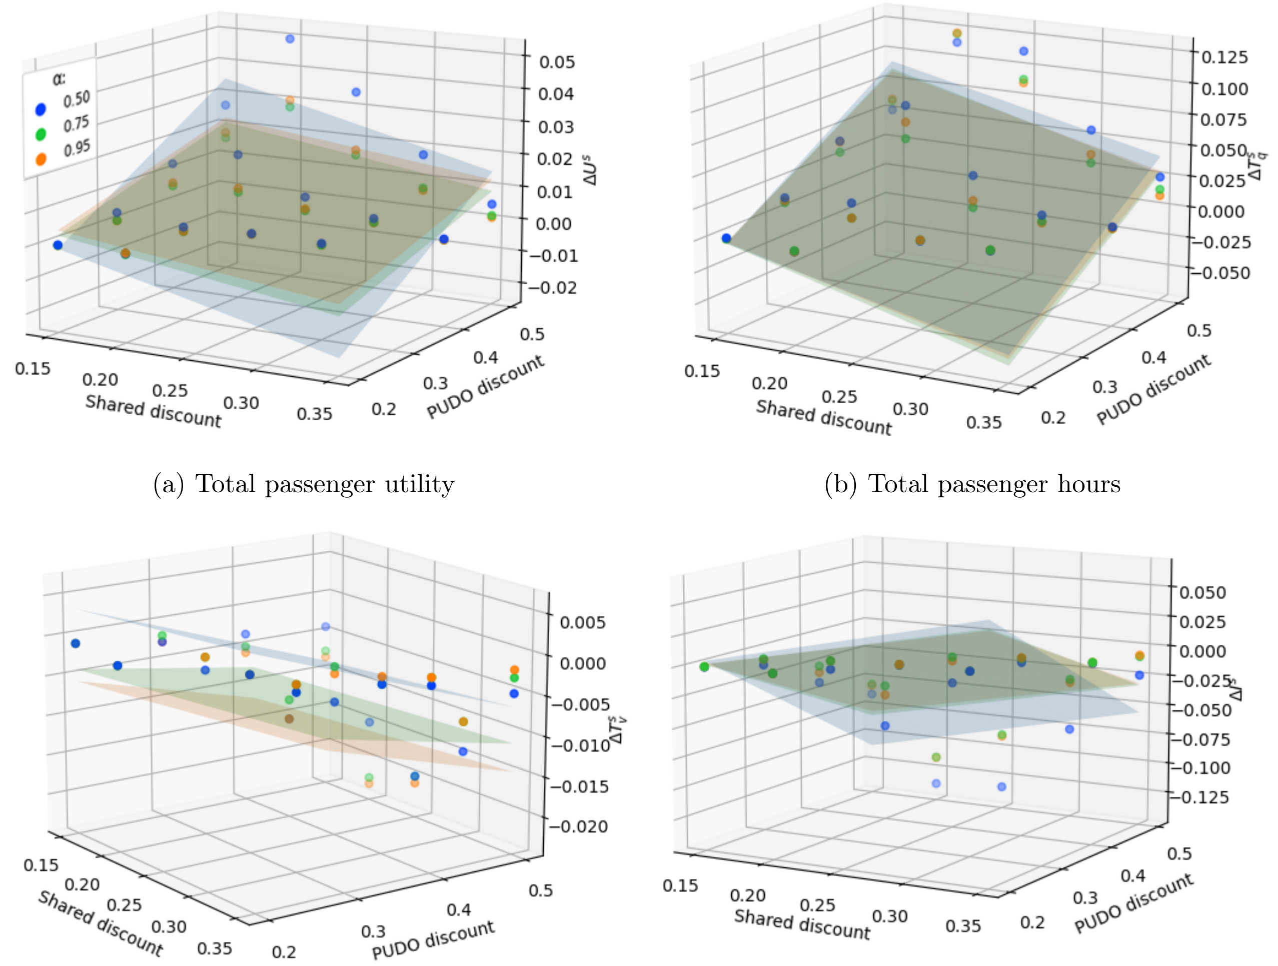

(a) Total passenger utility

(b) Total passenger hours

(c) Total vehicle hours

(d) Total revenue

Figure 17: Correlations of relative differences of various system-wide performance indicators with λ and $\tilde{\lambda}$ discounts. In this case, PUDO pooled rides are only compared the door-to-door alternative while the system wide performances of different optimisation parameters are also visualised. For $\alpha = 0.5, 0.75, 0.95$ the respective colours are red, green, and orange.

To aid with understanding of the different possibilities of a generated route of a pooled ride, Figure 18 visualises the routes of a three degree ride under various α. Here, the original pooled ride is visualised in blue while the respective generated route with PUDO is highlighted in green. Note that when $\alpha = 0.5$ the algorithm treats travellers in a pooled ride and the vehicle equally while when $\alpha \approx 1$ the algorithm favours the vehicle, as seen in Equation 2.9. The route seen in Figure 18a is significantly different to those seen in Figures 18b and 18c. This is mostly due to the ride dropping off the second last passenger closer to their actual destination. As well, it is visible that the last traveller is dropped at their true destination when $\alpha = 0.50$ while for higher α values, the last traveller is dropped off at some distance from the true destination. Also, the routes generated when $\alpha = 0.75$ and $\alpha = 0.95$ exhibit extremely minor differences.

(a) $\alpha = 0.50$ (b) $\alpha = 0.75$ (c) $\alpha = 0.95$

Figure 18: The route of a door-to-door pooled ride with three travellers, visualised in blue, and the optimised route with PUDO points in green. Each map represents a route obtained under different α values.

4.3 Distribution of selected pooled rides with PUDO

In this and the following sub-section, the performance of ride-pooling with PUDO is examined at a disaggregate level, as ride and traveller-level performance indicators are discussed, respectively. Disaggregate effects of pricing configurations and service settings on the PUDO pooled ride system can further enlighten the results presented in Sections 4.1 and 4.2.

On that note, Figure 19 visualises the distributions the relative difference of ride-utility[13] and the ride travel time[14] between a PUDO pooled ride and the door-to-door alternative. Figure 19 visualises the distributions of all optimised pooled-rides when $Q = 2000$ trips, and $\lambda = 0.2$. The distribution of each $\tilde{\lambda}$ and α is displayed, also each distribution is categorised as selected[15] or non-selected.

Figure 19a summarises the distributions of ride-utility ΔU_r^s. It is easy to see that low PUDO discounts ($\tilde{\lambda} \leq 0.3$) are quite unattractive for traveller as practically no travellers opt for PUDO pooled rides. With the exception when $\alpha = 0.50$ where a few PUDO pooled rides can be seen to be selected. With larger discounts, the distributions of selected rides become larger where mean ΔU_r^s is also seen to improve. The improvement of ΔU_r^s with discount holds for both selected and non-selected PUDO pooled rides. When $\alpha = 0.50$, the mean ΔU_r^s is generally larger than when $\alpha \geq 0.75$ while also creating a more attractive service as more rides are selected at $\alpha = 0.50$. It can be seen that the least amount of non-selected rides occur when at $\tilde{\lambda} = 0.5$ this alpha value.

The distributions show that both $\tilde{\lambda}$ and α have a considerable effect on the attractiveness of pooled rides with PUDO. Since low α values favour travellers by minimising walking-time where possible, it is sensible to see a greater number of selected PUDO pooled rides. As well, there is another indication that insignificant difference exists between $\alpha = 0.75$ and $\alpha = 0.95$.

Turning the attention to the difference in ride travel time, Figure 19b summarises the distributions of $\Delta T_{v_r}^s$. With this performance indicator, increase in $\tilde{\lambda}$ slightly reduces the mean travel time of non-selected rides, this is clearer when looking at $\alpha > 0.5$. Intriguingly, the selected PUDO rides actually have a smaller mean reduction of vehicle travel time that non-selected rides. When $\alpha = 0.50$, $\tilde{\lambda}$ seems to play a small role in the distribution of $\Delta T_{v_r}^s$ where selected rides only have a slightly worse $\Delta T_{v_r}^s$ (around 1%). For larger α, a larger difference is seen in $\Delta T_{v_r}^s$ between the selected and non-selected rides.

It seems that travellers are not concerned with the magnitude of reduced vehicle travel time as it is clear that they do not opt for the rides with the lowest amount of detour. It is interesting to see that for larger α, non-selected PUDO optimised pooled rides could reduce up to 10 of the vehicle

[13] The sum of utilities of all travellers within a ride, Equation 2.7.

[14] Time between first traveller picked-up and last traveller dropped-off $\tilde{t}_{r_v}$

[15] A PUDO ride is selected if the total traveller utility is greater than the door-to-door alternative, $\Delta U_r^s > 0$

travel time. This is likely due to the longer walking times involved in PUDO rides that are able to substantially reduce the vehicle travel time, this is also visible in the larger attractiveness for PUDO ride-pooling when $\alpha = 0.50$. The walking time to and from PUDO points appear to play a large role in determining the attractiveness of an optimised ride, this further examined in the following sub-section.

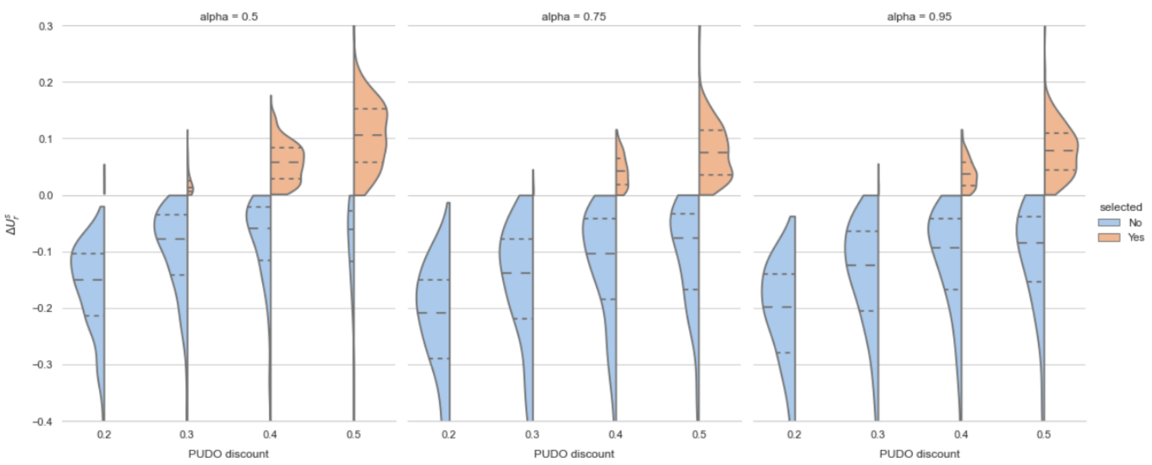

(a) Ride level utility improvement ΔU_r^s

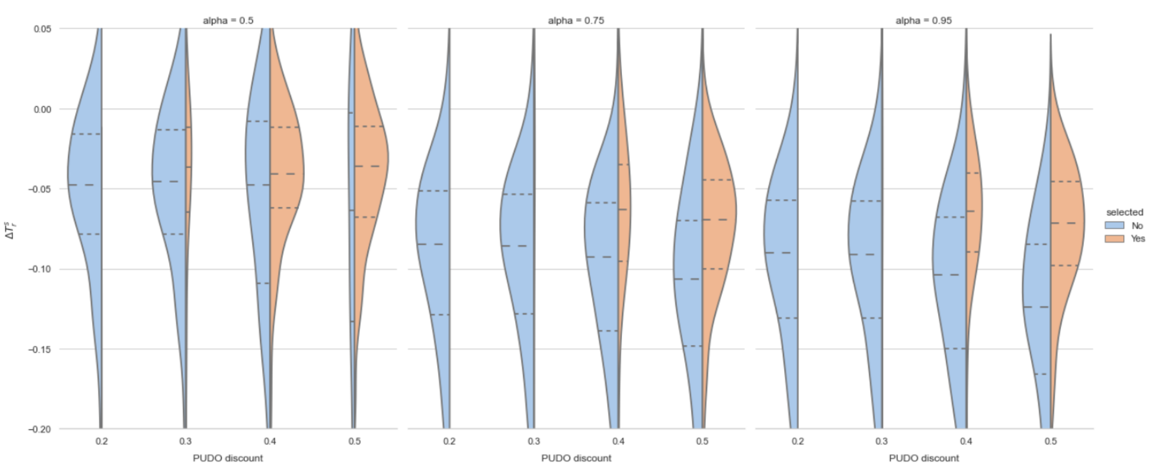

(b) Ride Vehicle hour reduction $\Delta T_{v_r}^s$

Figure 19: With $Q = 2000$ trips and $\lambda = 0.2$, distributions of ride-level (with respect to shared rides) indicators for PUDO optimised pooled rides for various $\tilde{\lambda}$ and α. Distributions are split into selected and non-selected PUDO pooled rides categories and shaded with orange and blue, respectively.

4.4 Walking times for travellers

As there was indication that traveller's prefer rides that do not provide the largest reductions in vehicle travel time, or the largest cost-saving, it is then necessary to zoom into the behavioural characteristics of the travellers themselves. Figure 20 displays the walking time distribution of all travellers for a given $\tilde{\lambda}$ and α. Walking time is generally the shortest when $\alpha = 0.50$, at this service setting selected rides resort to t_{w_i} that are on average up to 50 seconds shorter than the non-selected rides. For all α values presented, t_{w_i} increases with $\tilde{\lambda}$, which indicates that PUDO locations (and therefore t_{w_i}) are influenced by the discount offered. Longer t_{w_i} are generally seen for larger α; average t_{w_i} is around 100 seconds longer when $\alpha \geq 0.75$ than when $\alpha = 0.50$. When $\alpha \geq 0.75$, the selected rides invoke, on average, 350 seconds of walking time (and are considerably shorter than the non-selected rides) although the variance of walking time is much larger than the selected rides when $\alpha = 0.50$. Essentially, Figure 20 provides evidence that favouring the vehicle (i.e. $\alpha \approx 1$) leads the algorithm to finding the shortest route possible for the given set of PUDO points of the travellers within a ride. It is clear that t_{w_i} plays a significant role in determining the attractiveness of a ride with PUDO points. Routes generated when $\alpha = 0.5$ produce the largest traveller utility improvements (and largest number of selected PUDO pooled rides) due to the significantly shorter walking times invoked.

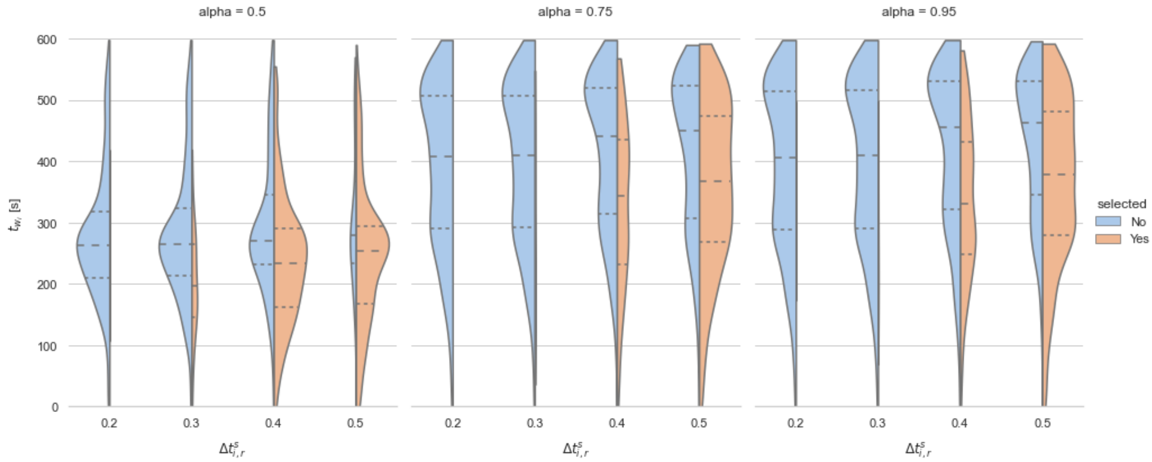

Figure 20: With demand $Q = 2000$ trips and $\lambda = 0.2$, distributions of walking times for travellers are split into selected and non-selected PUDO pooled rides, coloured with orange and blue respectively.

Matching travellers to attractive pooled rides depends of on various factors and results in travellers matched to rides of various degree. Walking times of travellers in varying ride degrees could differ due to the origins, destinations, and area accessibility. Undoubtedly, rides of higher degree also result in a larger search space due to a larger set of PUDO nodes. Figure 21 plots the normalised distributions of traveller walking times for various ride degrees and for the position of a traveller within a ride. These figures show t_{w_i} for all PUDO pooled rides (both selected and non-selected) when $Q = 4000$ trips, $\lambda = 0.2$, and $\tilde{\lambda} = 0.4$. This analysis does not look into the effects of varying $\tilde{\lambda}$ as this variable produces minimal effect on the distribution of t_{w_i} when splitting per ride degree or ride position. Furthermore, on-normalised distributions that depict the amount ride of degree two, three, or four can be found in Appendix C, here it can be seen that two degree rides are most common while four degree rides are least common. This is largely due to the matching done by ExMAS where higher degree rides are more common with higher demands (i.e. $Q > 4000$ trips).

Figure 21a shows that rides of degree two and three follow practically identical t_{w_i} distributions when $\alpha = 0.5$, At this service setting, travellers in rides of degree four have considerably longer walking times. The t_{w_i} distributions of all ride degrees when $\alpha \geq 0.75$ are fairly similar to one another. The exceptions here are that two degree rides have a double peak and that rides of degree four are almost symmetrically distributed. On a general note, the average t_{w_i} of each ride degree are longer for $\alpha \geq 0.75$, which was also seen in Figure 20.

Figure 21b shows that when $\alpha = 0.50$, the distributions of t_{w_i} are almost identical for all positions within a ride. This shows that when treating the vehicle and traveller equally, the algorithm aims to equalise the t_{w_i} for all travellers within a ride. When $\alpha = 0.75$, the range of t_{w_i} of travellers increases along with their mean t_{w_i}, however at this service setting no traveller position exhibits a particularly unique distribution. The only exception is for the first travellers in a ride due to the double peak t_{w_i} distribution, this sheds light on the source of two degree ride distribution seen in Figure 21a. With $\alpha = 0.95$ it becomes clearer that the first traveller in a ride tends to have the t_{w_i}. This occurs due to algorithm aiming to obtain the global shortest path at this service setting, which results in the ride picking up the first and dropping off the last passenger on the extremity of their walking radius. An example of this could be seen in Figure 18c where the last traveller is dropped-off at a further distance from their true destination.

To further investigate the effects of α on t_{w_i}, Figure 22 visualises the correlation between traveller walking and in-vehicle time. The demand level $Q = 2000$ trips , $\lambda = 0.2$, and $\tilde{\lambda} = 0.4$. In this case, when $\alpha = 0.5$, a neutral correlation exists between t_{w_i} and $\tilde{t}_i$. For $\alpha \geq 0.95$, travellers become more sensitive to their travel time as now a positive correlation exists between t_{w_i} and $\tilde{t}_i$. Selected rides that consist of a long $\tilde{t}_i$ tend to have longer t_{w_i} while shorter $\tilde{t}_i$ have shorter t_{w_i}. This can also be seen with the non-selected rides as here shorter $\tilde{t}_i$ with longer t_{w_i} are deemed unattractive.

Figure 22 shows that when $\alpha \approx 1$, travellers place more value on their t_{w_i}. Essentially, longer t_{w_i} are only compensated for longer $\tilde{t}_i$. This is comparable to evaluating your airport transfer time to your total flight time. For instance, travellers are likely to find a 4 hour transfer time for a 2 hour flight less attractive than a 4 hour transfer time for a 10 hour flight.

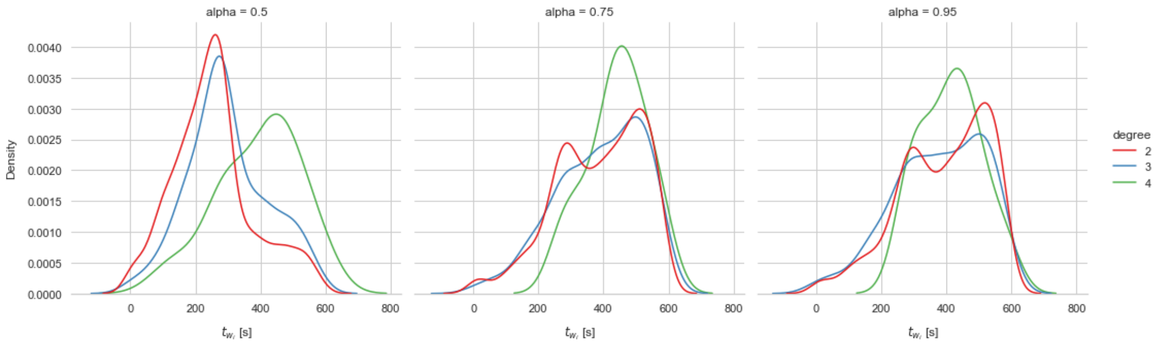

(a) Ride degrees ranging from two to four which constitutes as rides with two to four travellers.

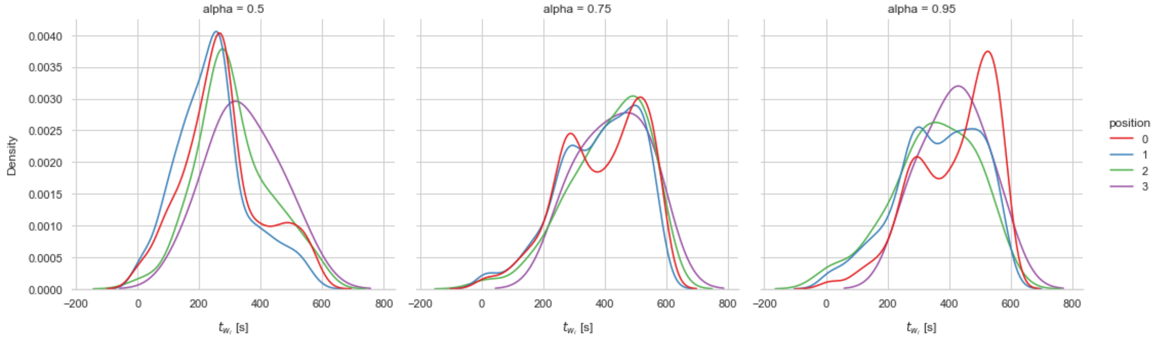

(b) Position within a pooled ride, position 0 represents the first traveller to enter a pooled ride and position 3 represents the fourth traveller to enter a pooled ride.

Figure 21: Normalised distributions of walking times for all PUDO pooled rides rides. Number of trips $Q = 4000$ trips, shared discount $\lambda = 0.2$, $\tilde{\lambda} = 0.4$.

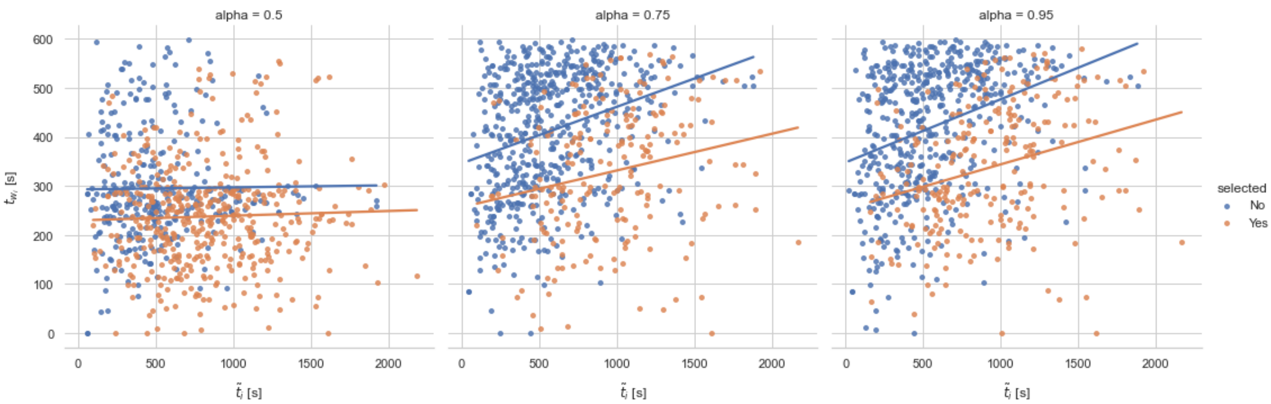

Figure 22: For a demand level $Q = 2000$ trips, $\lambda = 0.2$ of $\tilde{\lambda} = 0.4$, scatter plots visualise the correlation of traveller walking time with in-vehicle time for certain α and categorised with selected PUDO pooled ride.

4.5 Examining the effect of demand on ride-level indicators

The demand level, or the number of trip requests in an hour of simulation, plays a vital role in the scalability of a ride-pooling network. A pooling service must be able to match trips accordingly such that the operator is able to save on operating costs while users still deem the service as attractive. In general, higher demand levels allow for a larger number of higher degree door-to-door pooled rides due

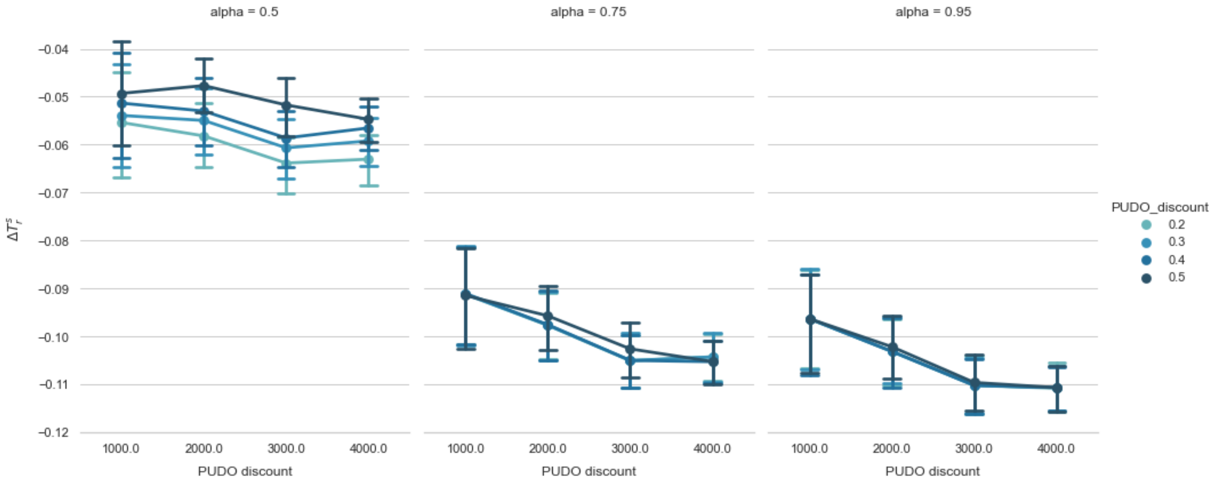

(a) Ride-level vehicle travel time improvement ΔT_r^s

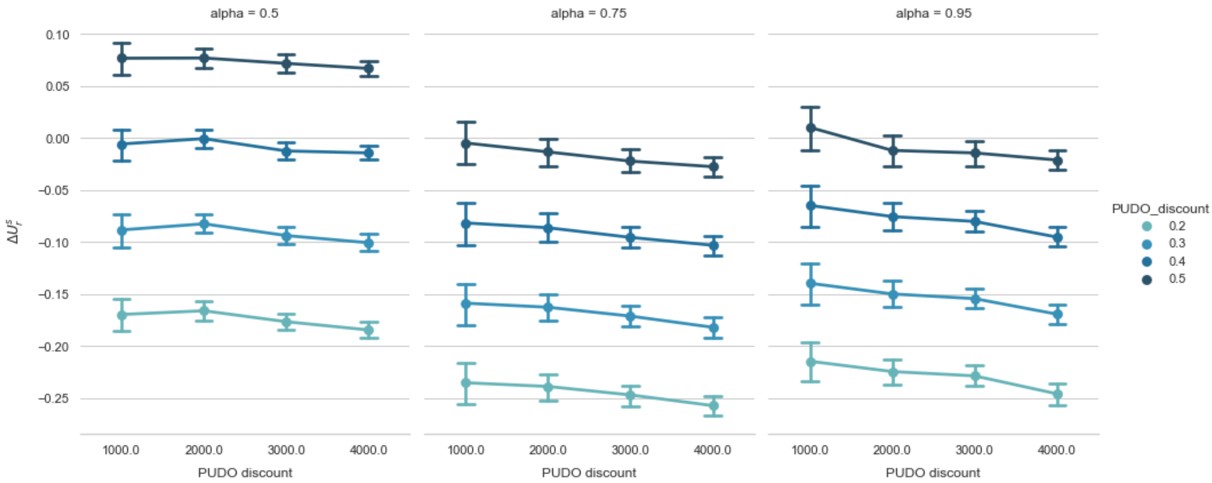

(b) Ride-level traveller utility improvement ΔU_r^s

Figure 23: Box and whisker plots visualising mean ride-level performance indicators and their variance for various Q, $\tilde{\lambda}$, and α. The $\lambda = 0.2$ for all experiments visualised.

.

to the higher trip matching possibilities, in turn allowing the operator to further reduce vehicle hours (Kucharski & Cats, 2020). Now, for the case of ride-pooling with PUDO, Figure 23 plots the average ride-level vehicle travel time and utility improvements for all selected and non-selected PUDO pooled rides when the service is set to various Q, $\tilde{\lambda}$, and α. Figure 23 also visualises the variance of each experiment, resulting in a so-called box and whisker plot.

Looking first at the average vehicle-hour savings, Figure 23a plots correlations of ΔT_r^s with Q for the set of $\tilde{\lambda}$ and α. It is immediately clear that ΔT_r^s improves with α (around 4-6% further travel time reductions) and that increasing $\tilde{\lambda}$ can diminish ΔT_r^s (mostly notable when $\alpha = 0.50$), which was also seen in Figure 19b. Increasing Q is improves ΔT_r^s as a larger number of higher degree rides, although ΔT_r^s improves mostly from $1000 \leq Q \leq 3000$ where insignificant difference is seen between $Q = 3000$ and $Q = 4000$. When $\alpha = 0.50$, larger demand levels could improve[16] ΔT_r^s by another 1% while when $\alpha = 0.75$ the further improvements could be up to 1.5%. It is also important to note that the variance of ΔT_r^s decreases with larger demand levels since sample sizes are larger.

The correlations of average ΔU_r^s with Q are showcased in Figure 23b. Here it is evident that ΔU_r^s decreases linearly with Q where trends are more prominent when $\alpha \geq 0.75$. Again, the decrease in attractiveness is likely due to the larger number of higher degree rides which usually invoke larger detours. Furthermore, as seen in Figure 19a, larger α and lower $\tilde{\lambda}$[17] significantly reduce the average

[16] When comparing average values.

[17] ΔU_r^s decreases linearly with $\tilde{\lambda}$ due its linear application in Equation 2.6

ΔU_r^s. It seems that $\alpha = 0.50$ creates a weak linear trend between ΔU_r^s and Q due to the minimised and equalised walking times (see Figure 21) at this service setting. Setting $\alpha \geq 0.75$ creates the most attractive rides when $\lambda \tilde{=} 0.5$ and $Q = 1000$, which is when most selected pooled rides consist of two travellers. The larger walking times induced by $\alpha \geq 0.75$ would make higher degree rides less attractive because of the higher likelihood of passenger delay (due to the larger number of travellers in a ride) which results in longer total travel time.

In all, it is important to remember that Q is an exogenous variable used to dictate the required number of trip requests that are extracted from the Albatross dataset, resulting in $\mathbf{Q}$. Essentially, Q provides the number of trip requests that are available for ExMAS to match to attractive shared rides. Larger demand levels facilitate the matching of higher degree pooled rides which help with further reducing total vehicle hours. Figure 23 shows that incorporating PUDO locations helps further reducing the travel times of vehicles and even more so with increasing Q, although ΔT_r^s seems to converge to a maximum reduction when $Q > 3000$. Albeit, increasing Q correlates negatively with ΔU_r^s, likely due to the larger number of higher degree rides. The service setting α could be used as a mediator for improving either the ΔU_r^s ($\alpha = 0.50$) or ΔT_r^s ($\alpha = 0.95$) depending on the requirements for the time of operation.

5 Conclusions, implications, limitations, and recommendations

This thesis aims to build upon existing research on ride-pooling with PUDO points by examining the effects pricing configurations, service settings, and demand levels have on an example service. This was done by first building upon an existing and scalable algorithm that matches trip requests to attractive shared rides (Kucharski & Cats, 2020). The existing utility formulation was first extended and incorporated into a route search algorithm that is able to assess the potential of a ride's PUDO point configuration with respect to the vehicle and travellers' utility. Essentially, the algorithm was able to further optimise already matched shared rides by introducing a walking factor for travellers. The method and code of the algorithm can be found on the respective *GitHub*[18] repository.

The algorithm was applied to a synthetic data-set which represented the activity based demand of around 100,000 agents in the city of Amsterdam. The objectives of this thesis were carried out with respect the case of Amsterdam. The results obtained in this thesis therefore apply mostly to the Dutch context, discrepancies should be expected when examining the performance of such a system in different cities or countries.

The main research question sought to examine the outcomes pricing configurations had on a ride-pooling service with PUDO nodes. To do this, the algorithm was applied to a total of 20 different door-to-door sharing discounts (λ) and PUDO sharing discounts ($\tilde{\lambda}$) combinations. The discount is the reduction of service fare, relative to the base fare of a non-shared ride (λ^{ns}). The door-to-door discount is utilised to match the trip requests into attractive shared rides while PUDO discount gave was intended to 'reward' travellers walking to and from a PUDO point. In order to make ride-pooling with PUDO an attractive alternative, the PUDO discount must be larger than the door-to-door discount ($\tilde{\lambda} > \lambda$).

A series of system-wide, ride-level, and traveller-level performance indicators were used to describe the benefits and drawbacks of ride-pooling with PUDO. The system-wide indicators showed that total passenger hours increased while the total vehicle hours reduced due to travellers selecting ride-pooling with PUDO, which is expected as travellers increase their travel time by walking while the vehicle minimises the detour due to sharing. The total utility of travellers generally improved with the discount offered while the revenue undoubtedly decreased as pooling offers a cheaper service for the same distance.

The magnitude of the improvement or reduction of system-wide indicators greatly depended of the discounts offered and were vastly different between ExMAS optimised and non-shared rides. The experiments showed that total utility improvement and increase in passenger hours followed similar patterns. The total utility (with respect to ExMAS optimised rides) improved as the difference between shared and PUDO discount increased. With respect to non-shared rides, utility improves mostly on the basis of the shared discount offered, however large PUDO discounts also improved the utility (and increase passenger hours) the most. Regarding the operator, minor differences are seen between the discount configurations when looking at the vehicle hours (range from 0.5-2.2%) reduced (with respect to ExMAS optimised rides). When comparing to non-shared rides, vehicle hours are reduced linearly with increasing shared discount. The difference in revenue (with respect to ExMAS optimised rides) was largest when the difference between the two discounts was largest, however at this point the difference in revenue with respect to private rides was the smallest. When comparing to private rides, the loss in revenue was largest when sharing and PUDO discounts were both high.

Through testing various discount configurations, ride-pooling with PUDO points performs best when the difference between the two offered discounts is large. Setting shared discount to 0.2 and PUDO discount to 0.4 creates some of the largest differences in performance between the two services. It was seen that at these discount prices, ride-pooling was able to reduce a further 1.0% of vehicle hours while revenue decreased a further 3.7%. For some operators, this could be a risky operation due the insufficient reduction in operating costs.

Looking at ride-level performance indicators, larger PUDO discounts increase the amount of selected PUDO rides which results in a larger average ride utility. It was seen that ride utility increases linearly with the PUDO discount offered, largely due to the linear properties of the utility formula. On the other

[18]URL: `https://github.com/MarkoM-5/ExMAS/tree/PUDO-dev/ExMAS/notebooks/Afstudeer`

hand, higher discounts do not improve the reduction in ride vehicle travel time for both selected and non-selected rides. This is likely due to walking time only partially increasing with PUDO discount. It seems that, offering larger discounts to the service fare does not greatly increase the incentives to walk. Ultimately, this also depends on the fleet size and the capacity of a vehicle, which could be examined in future work.

With the effects of pricing configurations characterised, attention can be shifted to the sub-research question, which sought to characterise the outcomes of larger demand levels and various service settings on a ride-pooling service with PUDO points. The service setting instructed the algorithm who to optimise a ride for the travellers, the vehcile, or both. This is known as the parameter α and essentially adds adds a weight to the utility of the operator and the utility of the travellers within a ride. When the service setting is vehicle orientated, the algorithm searches for the shortest route possible which induces longer walk times for travellers. When the service setting is traveller oriented, the algorithm searches for a route where where walking times are minimised. In this thesis, the service setting ranges from middle (vehicle and travellers treated equally) to high ($0.50 \leq \alpha \leq 0.95$).

For system-wide indicators, insignificant reductions or improvements were seen between the different service settings when the difference between PUDO and shared discount was small. Larger differences in discount offered showed that setting the service setting to middle ($\alpha = 0.5$) created the largest improvements in utility and increased passenger hours the most. However, the same service setting would create the largest reductions in revenue and the smallest reductions in vehicle hours. Generally, larger α values ($\alpha \geq 0.75$) would create insignificant differences between one another in improvements or reductions in system-wide indicators.

On the ride-level performance indicators, $\alpha = 0.5$ generally results in the largest selections of pooled rides with PUDO. This is due to the generally shorter walking times imposed on travellers with walking times around 100 seconds shorter when treating the operator and traveller equally. Furthermore, larger α introduced inequalities in walking times where the first traveller of a ride usually walk longer than the subsequent travellers. When $\alpha = 0.5$, all travellers within a ride were more likely to have equal walking times.

It was also interesting to see how travellers' perception of travel time changed with different service settings. When $\alpha = 0.5$, a neutral correlation exists between the time spent in a selected pooled PUDO ride and the total walking time. While when $\alpha \geq 0.75$, travellers preferred opt out of PUDO pooled rides that were of short trip duration and of long walking time. At this high service setting, travellers were willing to walk longer only if the trip duration was relatively longer. This is similar to real-life scenarios where travellers with long flights (i.e. 10+ hours) would be fine with long layovers (i.e. 4+ hours) while flights of shorter duration with long layovers would be undesirable.

The demand level set for an experiment dictates the available trips that can be matched to attractive shared rides. Larger demand levels increases the number of trips that can be matched which in turn facilitates the increased likelihood of matching more travellers within a single pooled ride (i.e. pooled rides with more than two travellers). Generally, higher degree rides allow for further reductions in vehicle travel time and incorporating PUDO nodes into such rides helps with further reduction of vehicle travel time. Increasing demand levels reduce the average vehicle travel to a certain point. However, increasing demand level correlates negatively with traveller utility improvement. This is likely due to the larger number of higher degree rides that result in a longer passenger travel time due to longer delay times as more travellers are now in a single pooled ride. The service setting could be used as a mediator for either improving the traveller utility or the vehicle traveller time, depending on the requirements for the time of operation.

5.1 Implications

So what do the effects of pricing, service setting, and demand level mean for the city of Amsterdam? The results in this thesis showed that offering competitive discounts for ride-pooling with PUDO[19] can further reduce the total vehicle hours of the system, in turn helping alleviate the congestion and vehicle pollution in the city. Although, the maximum system-wide vehicle hour reductions were around 2%,

[19] I.e. larger difference between PUDO and door-to-door discount

allowing travellers to walk to PUDO location would help lower the amount of vehicles in the narrow streets of Amsterdam's inner ring. It was also seen that at some service settings, travellers would walk further to/from a PUDO node for longer trip durations, this could provide higher accessibility to areas of a city with poor existing PT infrastructure. As well, ride-pooling with PUDO could be a good alternative for travellers that are deterred from overcrowded PT lines during pandemic scenarios or when a PT line is under sudden closure.

Now, what does this all mean for an operator such as *Uber* that wants to implement ride-pooling with PUDO in Amsterdam? Ignoring the fact that, at the time of this thesis being written, *Uber* currently does not have a door-to-door ride-pooling service (or *UberPool*) in Amsterdam, there are still various conclusions from this study that could be implemented in a future ride-pooling with PUDO service. It is clear that further operating costs can be saved by introducing PUDO points however, this only occurs when the difference between PUDO and door-to-door discount is large. But, If PUDO discount is too large (i.e. $\tilde{\lambda} = 0.5$), *Uber* runs the risk of losing more revenue than saving on operating costs. The benefits and drawbacks of ride-pooling with PUDO seen in this thesis are supplements to already matched and attractive door-to-door pooled rides, further utility improvements or vehicle hour reductions could have been seen if trips were matched to pooled rides with PUDO or if PUDO points were inserted to non-selected door-to-door rides.

Moreover, this thesis shows that a company like *Uber* can control its supply with the use of a service setting. The use of the service setting showed that ride-pooling with PUDO can be made much more attractive to travellers by setting fair PUDO points to walk to by sacrificing vehicle travel time savings. Such a traveller orientated service setting would be useful when there supply exceeds demand. When demand exceeds supply, a vehicle orientated service setting could be used maximise vehicle travel time reductions by inducing longer walking times to travellers which also limits the attractiveness of the service or the amount of travellers opting for that service. This is similar to the concept of surge pricing used by *Uber* when demand exceeds supply at certain locations, such as airports.

Finally, it was also seen the larger demand levels are able to provide greater reductions vehicle travel times. This reinforces the idea that *Uber* must be sure that the demand for pooling exists before rolling out such a service. Of course, demand could be generated by attracting users with competitive prices however, it was also seen the general attractiveness of a ride decreased with increasing demand. Now this is likely due to making the duration of higher degree rides even longer due to PUDO incorporation but it is still something that must carefully analysed and discussed prior to deciding to roll-out a pooling service.

5.2 Limitations and recommendations for future work

Throughout the report a number of limitations were indicated and discussed. To emphasise, this thesis looked into the various effects that PUDO points had on already matched and selected door-to-door pooled rides. Examining the possibilities of making non-selected, or unattractive, pooled rides attractive with the use of PUDO points would be worthwhile. That way we could treat ride-pooling with a PUDO as a mode that is operates parallel with door-to-door ride-pooling rather than as competition.

The results obtained in this research are from a single algorithm and could vary for different algorithms used, such as the ones presented by Fielbaum (2021) and Fielbaum et al. (2021). The algorithm used in this report determines a candidate PUDO configuration through a shortest route search while the travellers' utility is only considered in the last step, or last trip pair. The differences in search space between the algorithm and an exact exhaustive search that all optimum paths are not considered when service setting is more traveller focused.

Furthermore, replications were not incorporated within this study as the running all the experiments once required relatively long run-times (i.e. at least one week to run all the experiments once). Since it was calculated that at least 10 replications would be necessary to increase the reliability of the experiments, future studies using this algorithm should aim to incorporate replications. Optimising the run-time of the algorithm (such as maximising the amount of experiments that are able to run in parallel) should facilitate the ability to run more replications.

This research only examined the effect system settings and demand levels have on the performance of ride-pooling with PUDO. The effect of behavioural parameters on the service was not analysed. The behavioural characteristics of humans varies with geography, demographic, income, etc. Future work with this algorithm would be able to incorporate behavioural studies for ride-pooling with PUDO. Using Albatross (Arentze, 2005), the demand for virtually any city could be generated where the algorithm could be easily applied to any network graph. This could easily allow for the comparison of performance between different cities.

Moreover, spatial analyses could also be conducted with this or similar algorithms. Different cities, and different areas within a city have varying infrastructure and PUDO stop possibilities. This could mean that some parts of a city or even some cities in general are better equipped for the ride-pooling with PUDO which would definitely help an operator with rolling out such a service.

Temporal studies would also be useful as timing between vehicle and traveller is important with ride-pooling with PUDO points. This algorithm does compute the amount of time a vehicle or a traveller would have to wait for one another while the first traveller is also given zero delay. In a real world scenario, travellers and vehicles should be penalised for long waiting times where departure time should be carefully selected.

Undoubtedly, there a number of further limitations which are simply beyond the scope of this report. For instance, this research takes on a planning perspective rather than a real-time optimisation which means that no traffic congestion is considered. This would undeniably alter the reported performance indicators and the actual PUDO stop locations. With traffic, such locations could not simply be set on intersections, rather they should be carefully chosen so that traffic would be minimally disrupted while a vehicle is picking up a traveller. As well, a vehicle is assumed to appear at the first traveller's pick-up point and empty vehicle travel is not considered. These are some of the points that would be extremely relevant when attempting to operate a real-time algorithm.

Bibliography

Aberle, C. (2020). Who benefits from mobility as a service? a gis-based investigation of the population served by four ride-pooling schemes in hamburg, germany. *KN-Journal of cartography and geographic information*, *70*(1), 25–33.

Arentze, T. (2005). Albatross: A learning-based transportation oriented simulation system. *Via-Via*, *36*(3), 49–51.

Boeing, G. (2017). Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, *65*, 126–139.

Buhaug, H., & Urdal, H. (2013). An urbanization bomb? population growth and social disorder in cities. *Global environmental change*, *23*(1), 1–10.

Czioska, P., Mattfeld, D. C., & Sester, M. (2017). Gis-based identification and assessment of suitable meeting point locations for ride-sharing. *Transportation research procedia*, *22*, 314–324.

Fielbaum, A. (2021). Optimizing a vehicle's route in an on-demand ridesharing system in which users might walk. *Journal of Intelligent Transportation Systems*, 1–20.

Fielbaum, A., Bai, X., & Alonso-Mora, J. (2021). On-demand ridesharing with optimized pick-up and drop-off walking locations. *Transportation research part C: emerging technologies*, *126*, 103061.

Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M.-E., Wang, X., & Koenig, S. (2013). Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, *57*, 28–46.

Goel, P., Kulik, L., & Ramamohanarao, K. (2016). Optimal pick up point selection for effective ride sharing. *IEEE Transactions on Big Data*, *3*(2), 154–168.

Gurumurthy, K. M., & Kockelman, K. (2020). How much does greater trip demand and aggregation at stops improve dynamic ride-sharing in shared autonomous vehicle systems? *Preprint at https://www. researchgate. net/publication/343451580*.

Jin, K., & Xia, C. (2021). Data-driven design of microtransit services via optimal transport and simulated annealing. In *Proceedings of the 22nd international workshop on mobile computing systems and applications* (pp. 179–181).

Ke, J., Yang, H., & Zheng, Z. (2020). On ride-pooling and traffic congestion. *Transportation Research Part B: Methodological*, *142*, 213–231.

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.

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.

Lasley, P. (2021). 2021 urban mobility report.

Law, A. M., & McComas, M. G. (1991). *Secrets of successful simulation studies* (Tech. Rep.). Institute of Electrical and Electronics Engineers (IEEE).

Lazarus, J. R., Caicedo, J. D., Bayen, A. M., & Shaheen, S. A. (2021). To pool or not to pool? understanding opportunities, challenges, and equity considerations to expanding the market for pooling. *Transportation Research Part A: Policy and Practice*, *148*, 199–222.

Li, R.-H., Qin, L., Yu, J. X., & Mao, R. (2015). Optimal multi-meeting-point route search. *IEEE Transactions on Knowledge and Data Engineering*, *28*(3), 770–784.

Li, W., Pu, Z., Li, Y., & Ban, X. J. (2019). Characterization of ridesplitting based on observed data: A case study of chengdu, china. *Transportation Research Part C: Emerging Technologies*, *100*, 330–353.

Li, Z., Hong, Y., & Zhang, Z. (2016). Do on-demand ride-sharing services affect traffic congestion? evidence from uber entry. *Evidence from Uber Entry (August 30, 2016)*.

Liao, Y., Gil, J., Pereira, R., Yeh, S., & Verendel, V. (2020). Disparities in travel times between car and transit: Spatiotemporal patterns in cities. *Scientific reports*, *10*(1), 1–12.

Oh, S., Kondor, D., Seshadri, R., Zhou, M., Le, D.-T., & Ben-Akiva, M. (2020). Spatiotemporal characteristics of ride-sourcing operation in urban area. *arXiv preprint arXiv:2011.07673*.

Pant, P., & Harrison, R. M. (2013). Estimation of the contribution of road traffic emissions to particulate matter concentrations from field measurements: A review. *Atmospheric environment*, *77*, 78–97.

Robinson, S. (2004). *Simulation: the practice of model development and use* (Vol. 50). Wiley Chichester.

Schaller, B. (2021). Can sharing a ride make for less traffic? evidence from uber and lyft and implications for cities. *Transport policy*, *102*, 1–10.

Shokoohyar, S., Sobhani, A., & Ramezanpour Nargesi, S. R. (2020). On the determinants of uber accessibility and its spatial distribution: Evidence from uber in philadelphia. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *10*(4), e1362.

Stiglic, M., Agatz, N., Savelsbergh, M., & Gradisar, M. (2015). The benefits of meeting points in ride-sharing systems. *Transportation Research Part B: Methodological*, *82*, 36–53.

Wang, L., Zeng, L., Ma, W., & Guo, Y. (2021). Integrating passenger incentives to optimize routing for demand-responsive customized bus systems. *IEEE Access*, *9*, 21507–21521.

Zha, L., Yin, Y., & Xu, Z. (2018). Geometric matching and spatial pricing in ride-sourcing markets. *Transportation Research Part C: Emerging Technologies*, *92*, 58–75.

Zhao, M., Yin, J., An, S., Wang, J., & Feng, D. (2018). Ridesharing problem with flexible pickup and delivery locations for app-based transportation service: mathematical modeling and decomposition methods. *Journal of Advanced Transportation*, *2018*.

Zheng, Y., Li, W., Qiu, F., & Wei, H. (2019). The benefits of introducing meeting points into flex-route transit services. *Transportation Research Part C: Emerging Technologies*, *106*, 98–112.

Appendices

A Walking routes

Table 7: Tables showing how walking times of two travellers with two different routes with either no exponent or an exponent of 1.1 set to the walking time. Setting an exponent allows the algorithm choose routes where walking times are (more) equal between the travellers. This is seen in 7b as the sum of walking times for route is smallest when $t_{w_1} = t_{w_2}$.

(a) Without setting an exponent.

	t_{w_1}	t_{w_2}	**sum**
route 1	150	150	300
route 2	300	0	300

(b) Exponent set to 1.1.

	t_{w_1}	t_{w_2}	**sum**
route 1	248	248	495
route 2	531	0	531

B Comparison of exact and heuristic

The results shown here visualise the performance of the algorithm when searching for an optimum route. The algorithm described in Section 8 is compared to an exact, exhaustive search with no break points.

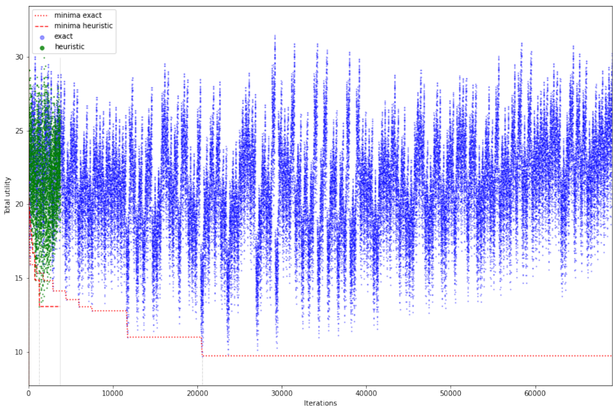

(a) $\alpha = 0.1$.

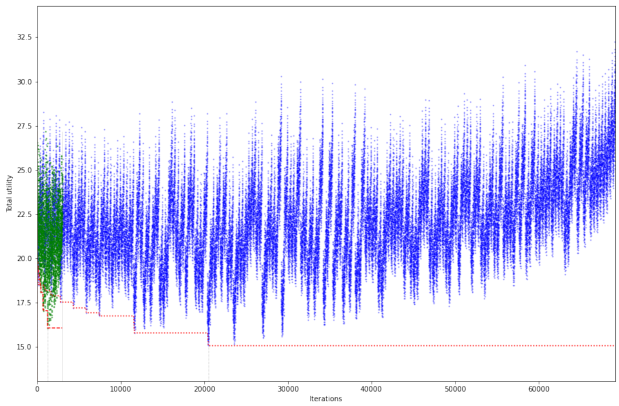

(b) $\alpha = 0.5$.

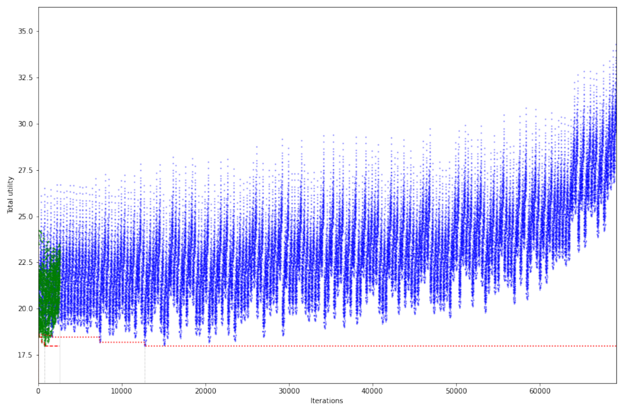

(c) $\alpha = 0.9$.

Figure 24: Disutilities of travellers and vehicle of a 2 degree ride for every possible route configuration with the heuristic in green and exact in blue.

Figure 24 visualises the results of the optimisation of a two degree ride with $\alpha = 0.9, 0.5, 0.1$ where total utility is plotted against iterations. In this case the total utility is sum of Equation 2.9 for each PUDO configuration. Moreover, the Figures indicate the disutility of a PUDO configuration thus the lowest resulting utility is preferred. The heuristic results are indicated with a green shade while the exhaustive search in blue. The dashed and dotted lines indicate the current minimum of the heuristic and exhaustive search, respectively. A ride of degree two is selected for optimisation as higher degree rides will result in 10^7 iterations as the search space is the product of the size all the PUDO nodes in a given ride. For this two degree ride, around 70000 possible routes were considered for the exact exhaustive search and up to 5000 for the heuristic.

Figure 24c shows the computed disutilities when $\alpha = 0.9$, thus the vehicle utility has the highest weight. A positive trend is seen for the exact exhaustive search due the trip-pair skim dictionary being sorted in ascending order. The absolute minimum utility is obtained at around 12000 iterations. On the other hand, the heuristic is able to reach this absolute minimum in less than a 1000 iterations. With $\alpha \approx 1$ it is evident that the algorithm is able to find the true optimum route.

Figure 24b shows the result when the α is set to 0.5. In this case, the travellers within the ride and the vehicle are treated equally, which means that their utilities are equally considered. The algorithm now aims to determine the shortest route whilst also slecting the most attractive PUDO nodes to walk to and from. A positive trend is still seen for the exact exhaustive search albeit with a smaller gradient than in Figure 24c. Furthermore, the absolute minimum disutility is now found at a later iteration count, namely at around 20000 iterations where this disutility is considerably lower than in Figure 24b. This shows that a certain PUDO configuration greatly influences the utility of travellers. The heuristic is also seen to have a lower disutility however, it never finds the actual minimum. This is expected as break points in the algorithm are mostly based on the route distance while only the final trip pair computes and compares the utilities of travellers. Still, the heuristic is able to go through the search space efficiently and produce a result with a relatively small error.

With $\alpha = 0.1$, Figure 24a shows that the exhaustive search has a neutral trend when optimising for travellers. The minimum disutility is still found at the same iteration count as in Figure 24a, however being much lower. This is sensible as the optimum PUDO configurations for the travellers could already be found when $\alpha = 0.5$ where the lower disutility could be as result of the vehicle utility being weighted much less. The heuristic is seen to have a much larger error now as the algorithm now optimises practically only for travellers whilst many of the break points are distance related.

C Non-normalised walking time distributions

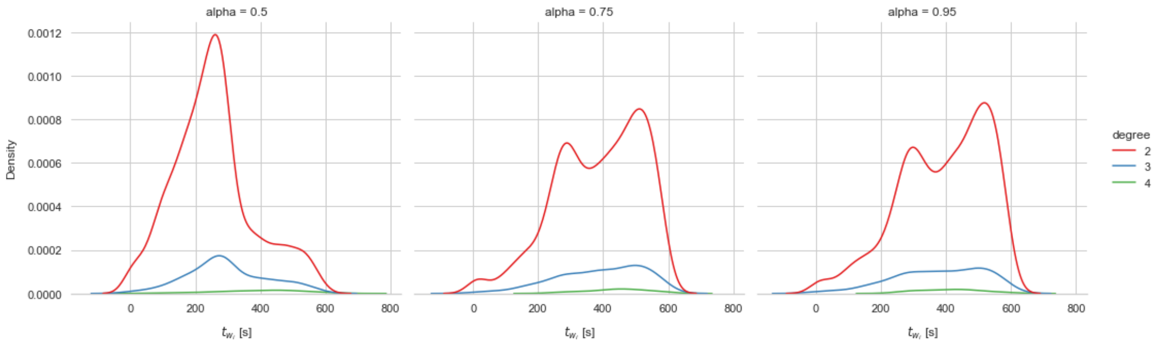

(a) Ride degrees ranging from two to four which constitutes as rides with two to four travellers.

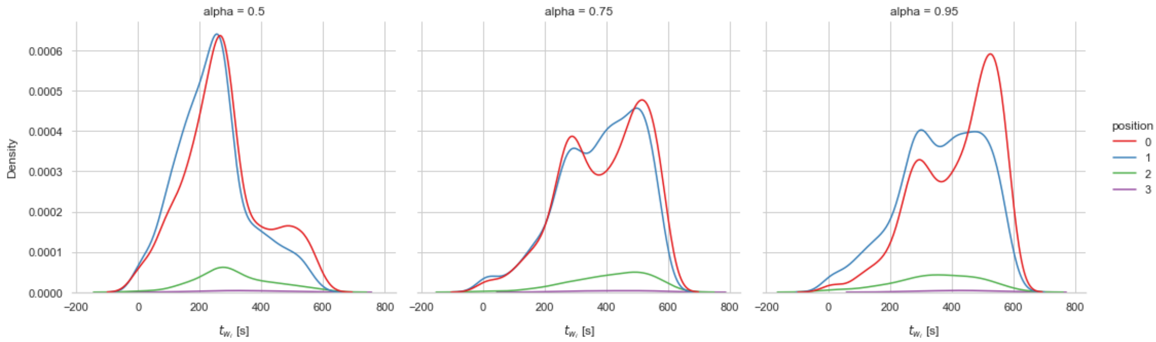

(b) Position within a pooled ride, position 0 represents the first traveller to enter a pooled ride and position 3 represents the 4th traveller to enter a pooled ride.

Figure 25: Non-normalised distributions of walking times for all pudo optimised rides. Number of trips $Q = 4000$, shared discount λ set to 0.2, PUDO discount $\tilde{\lambda}$ set to 0.4.