

On the quality requirements of demand prediction for dynamic public transport

Peled, Inon ; Lee, Kelvin; Jiang, Yu ; Dauwels, Justin ; Pereira, Francisco C.

DOI

[10.1016/j.commtr.2021.100008](https://doi.org/10.1016/j.commtr.2021.100008)

Publication date

2021

Document Version

Final published version

Published in

Communications in Transportation Research

Citation (APA)

Peled, I., Lee, K., Jiang, Y., Dauwels, J., & Pereira, F. C. (2021). On the quality requirements of demand prediction for dynamic public transport. *Communications in Transportation Research*, 1, 1-11. Article 100008. <https://doi.org/10.1016/j.commtr.2021.100008>

Important note

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

Copyright

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

Takedown policy

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



Contents lists available at ScienceDirect

Communications in Transportation Research

journal homepage: www.journals.elsevier.com/communications-in-transportation-research

On the quality requirements of demand prediction for dynamic public transport

Inon Peled^{a,*}, Kelvin Lee^b, Yu Jiang^{a,**}, Justin Dauwels^c, Francisco C. Pereira^a^a Danmarks Tekniske Universitet (DTU), Technology, Management and Economics Dept., Kgs. Lyngby, 2800, Denmark^b Nanyang Technological University (NTU), Graduate College, 50 Nanyang Avenue, Singapore, 637553^c Delft University of Technology (TU Delft), Microelectronics Dept., 2600, Netherlands

ARTICLE INFO

Keywords:

Dynamic public transport
Demand forecasting
Non-Gaussian noise
Predictive optimization

ABSTRACT

As Public Transport (PT) becomes more dynamic and demand-responsive, it increasingly depends on predictions of transport demand. But how accurate need such predictions be for effective PT operation? We address this question through an experimental case study of PT trips in Metropolitan Copenhagen, Denmark, which we conduct independently of any specific prediction models. First, we simulate errors in demand prediction through unbiased noise distributions that vary considerably in shape. Using the noisy predictions, we then simulate and optimize demand-responsive PT fleets via a linear programming formulation and measure their performance. Our results suggest that the optimized performance is mainly affected by the skew of the noise distribution and the presence of infrequently large prediction errors. In particular, the optimized performance can improve under non-Gaussian vs. Gaussian noise. We also find that dynamic routing could reduce trip time by at least 23% vs. static routing. This reduction is estimated at 809,000 €/year in terms of Value of Travel Time Savings for the case study.

1. Introduction

Public Transport (PT) has traditionally used static itineraries that remain unchanged for months (Ceder, 2016). However, as autonomous mobility advances and the vision of Smart Cities takes shape, the day approaches when PT becomes dynamic, so that some itineraries are adapted to real-time transport demand (i.e., are demand-responsive) (Bösch et al., 2018; Horáždovský et al., 2018). Meanwhile also, predictive models of transport demand are increasingly used for both long-term and short-term traffic management (Hashemi and Abdelghany, 2015). Future Public Transport should thus naturally employ predictive models for timely adaptation of service per expected transport demand.

The effective operation of demand-responsive PT thus requires that transport demand be accurately estimated ahead of time. For example, more accurate demand predictions can yield better utilization of PT resources, e.g., so that fewer vehicles are used while travel times are also cut shorter. Conversely, errors in predicted demand might lead to sub-optimal routing, thus resulting in the waste of energy, longer waiting times and profit loss.

The main goal of this work is to study the impact of demand prediction accuracy on subsequent performance of demand-responsive PT.

To this end, we conduct a hypothetical case study of a demand-responsive PT “pilot” experiment in Metropolitan Copenhagen, Denmark. We do so by first simulating prediction errors through various stochastic perturbations of travel demand, as observed through real-world PT trips. Using the noisy predictions, we then simulate and optimize demand-responsive fleets via a linear programming formulation, for various bus capacities and portions of dynamically routed buses. Finally, we analyze the results and draw conclusions on the impact of prediction accuracy on optimization quality.

Our case study is thus a form of Sensitivity Analysis, as we review in the next Section. However, contrary to most other Sensitivity Analysis studies on demand-responsive PT optimization, we account for both sides of demand prediction and dynamic routing and do so separately. Also in contrast to previous works, we consider predictive noise independently of any specific demand estimation models, to study the impact of prediction errors once any such model has already been fitted. We use different types of independent noise, both Gaussian and non-Gaussian, with a wide variety of statistical parameter settings. This allows us to investigate various correlation structures in the origin-destination pairs, e.g., as associated with non-recurrent traffic disruptions.

The motivation for this work is thus to offer several novelties over

* Corresponding author.

** Corresponding author.

E-mail addresses: inonp@dtu.dk, inon.peled@gmail.com (I. Peled), yujiang@dtu.dk (Y. Jiang).

<https://doi.org/10.1016/j.commtr.2021.100008>

Received 25 August 2021; Received in revised form 16 October 2021; Accepted 16 October 2021

Available online xxx

2772-4247/© 2021 The Authors. Published by Elsevier Ltd on behalf of Tsinghua University Press. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

previous research into emergent, dynamically routed PT services. On the methodological side, we take a high-level view of the impact of predictive accuracy on such PT, rather than use specific prediction models as in other works. This allows us to study such impacts without binding to any particular family of models and its specific properties as in other works. For example, models that rely on lagged information (autoregressive models, Recurrent Neural Networks, etc.) generally share some similarities in the effect of such parameters on predictive accuracy, and so our goal is to study the error distributions themselves, rather than particular parameters, learning methods, hyper-parameter tuning, etc. In turn, this also facilitates an examination of the relationship between optimal service quality and deviations from Gaussian predictive noise, as commonly assumed. On the practical side, this methodology allows us to explicitly quantify gains and losses in optimizing such services under a wide variety of possible noise distributions, without committing to a particular source of noise. That is, while predictive noise can arise from multiple sources – e.g., modeling choices, non-recurrent traffic disruptions and changes in mode preferences – we quantify its practical impact on the optimized service regardless of its sources.

The remainder of this work is organized as follows. Section 2 reviews gaps in current studies on demand-responsive PT optimization and uncertainty analysis. Section 3 details our experiments with simulated noise distributions and prediction-based optimization. Section 4 then provides the experimental results and their analysis. Finally, Section 5 recaps the work by discussing the goal, methodology and main results, and Section 6 concludes with a list of key findings and future steps.

2. Literature review

2.1. Public transport optimization

Works on PT planning and design rely traditionally on point estimates of future travel demand, obtained through manual data collection (e.g., via transport surveys). These estimates are thus realizations from a latent (i.e., unknown) transport demand distribution, and so are subject to uncertainty and errors. While such errors are widely acknowledged, their subsequent effects on PT performance are less quantified and discussed.

Modern works on PT optimization take advantage of advancements in big data and machine learning for demand estimation (Krishnakumari et al., 2020; Toole et al., 2015) and real-time fleet management (Hadji-dimitriou et al., 2020).¹ As these advancements gain traction in the PT optimization field (Iliopoulou and Kepaptsoglou, 2019; Wang and Qing-dao-er ji, 2019), it becomes increasingly important to study the effects of demand prediction accuracy on optimization quality. Further, these effects should also be evaluated in the context of future PT, which will employ dynamically routed vehicles for better demand-responsiveness (Bösch et al., 2018).

Motivated by these needs, we experiment with a case study of a PT fleet, whose itineraries are dynamically changed per predicted travel demand, in hourly intervals. We optimize service times under constraints of passenger preferences and service quality via Mixed-Integer Linear Programming (MILP), as investigated in many PT optimization studies (An and Lo, 2015; Luathep et al., 2011; Szeto and Jiang, 2014; Tong et al., 2015; Wang and Lo, 2010).

The fleet in this work thus resembles Mobility-on-Demand services and similarly relies on predicted transport demand, albeit with less flexible routes and stop locations. However, whereas many such services operate on the basis of pre-booked rides (Alonso-Mora et al., 2017; Hyland and Mahmassani, 2018), we do not assume any particular source of demand observations. Accordingly, PT demands in this case study are observed only upon their realization, i.e., only when passengers board or alight.

¹ Ref. (Iliopoulou and Kepaptsoglou, 2019) for a comprehensive review of big data applications in PT.

Contrary to previous works on PT optimization under demand uncertainty (Ukkusuri et al., 2007), we study a wide range of possible distribution properties (e.g., skew and kurtosis). Some works further attempt to mitigate errors via robust optimization techniques (An and Lo, 2015) and chance constraints (Wang et al., 2015). As error mitigation is not the focus of our case study, we consider this and other extensions for future work (Section 6).

2.2. Sensitivity analysis for demand-responsive public transport

By quantifying the impact of changes in demand prediction accuracy on subsequent PT performance, our case study is a form of Sensitivity Analysis (SA) (Saltelli et al., 2004). Nevertheless, this work differs in several respects from other SA studies in the field of demand-responsive PT optimization, as follows.

First, most SA studies on demand-responsive PT either treat demand as given or change it per predefined levels. For example, Huang et al. (2020) conduct an SA study for operating a demand-responsive transit service under deterministic changes in total demand, while Winter et al. (2018) do so for an automated demand-responsive transport system. Manasra and Toledo (2019) incorporate demand predictions within an optimization formulation and measure its sensitivity to predefined changes in demand. They note that for practical usefulness, the robustness of their formulation needs to be further studied under non-recurrent disruptions and demand surges.

Second, SA studies on demand-responsive transit often do not explicitly deal with stochastic noise in demand estimation. For example, Nickkar et al. (2020) measure the effects of varying demand on performance of autonomous fleets, but do not deal with errors in demand estimation. Similarly, Lee (2006) optimizes a transit network under various pre-defined scenarios, without checking possible deviations from optimality due to demand estimation errors. Ibeas et al. (2014) optimize bus capacity and headway per elastic demand, as collected from historical records, and without simulating stochastic noise.

Contrary to the above, we conduct SA under stochastic errors in demand estimation for demand-responsive PT. Moreover, the above works focus on developing applicable formulations for PT optimization, which motivates their choice of conducting SA under predefined levels of changes. Conversely, we do not aim to propose a directly applicable PT optimization method, but rather use the simulated PT system to study predictive noise effects from a more general perspective, as we next explain.

2.3. Uncertainty in modeling and the normality assumption

Predictive noise follows from the existence of uncertainty in modeling (Beaudrie et al., 2016) and yields residuals, i.e., differences between modeled predictions and actually observed values.² Consequently, when fitting predictive models, a normality assumption is commonly employed, whereby residuals are expected to be identically and Normally distributed (Seber and Lee, 2012). In some modeling contexts, however, residuals can be non-Gaussian (Mak, 2000). It may also be impractical to fit a model with Gaussian residuals, depending on data size and quality (Jackson and White, 2018).

The probability density of residuals can thus vary considerably in standard deviation (dispersion around the mean), skew (asymmetry around the mean) and kurtosis (weight of tails). Multiple methods have been devised for detecting deviations from normality, including plots, comparison of moments, and statistical tests (Thode, 2002). However, there are far fewer works on the impact of such deviations (Ivković et al., 2020), most of which concentrate on errors in model parameter estimates.

² Definitions of residuals differ by context and specificity (Gourieroux et al., 1987); we use a general definition that befits the context in this work.

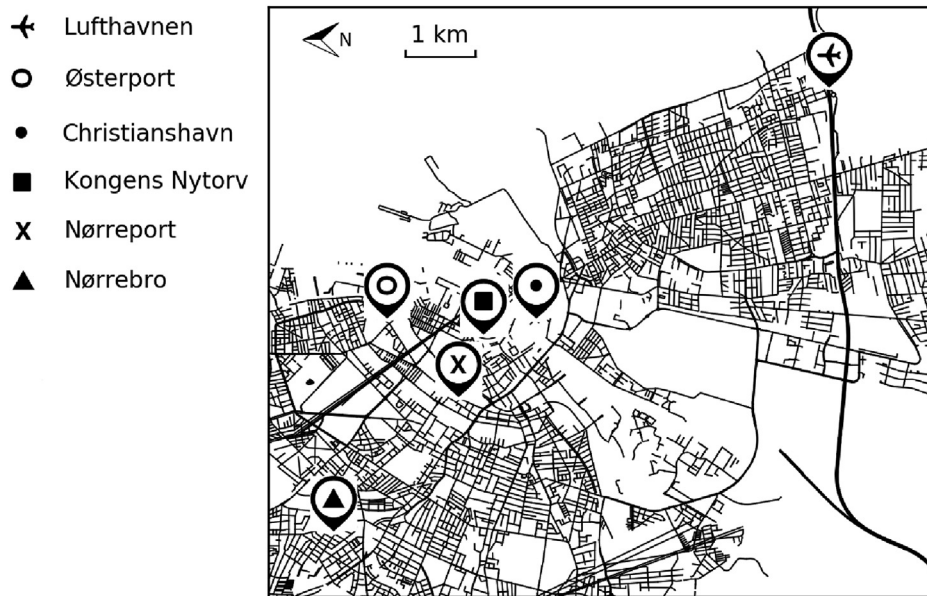


Fig. 1. The 6 most active PT stations in Metropolitan Copenhagen, Denmark.

Nelson and Granger (1979) discovered that for economic time series, Autoregressive Moving Average (ARMA) models often yield residuals with markedly non-Gaussian skew and kurtosis. Davies et al. (1980) expanded on this and concluded that the use of non-Gaussian residuals actually allows for a larger selection of models that better represent time series. In the context of chemical analysis, Wolters and Kateman (1989) used Monte Carlo simulations to quantify errors in Least Squares parameters under small deviations from normality.

More recently, He and Raghunathan (2009) use simulated data to examine sequential regression imputation methods under shifted and scaled non-Gaussian distributions: Uniform, Lognormal and t -distribution. They find that mean performance remains quite robust, despite noticeable instability in regression coefficients. To predict debt and bankruptcy of Serbian companies, Ivković et al. (2020) simulate Exponential and Weibull distributed residuals, study the resulting errors in LR parameter estimates, and devise transformations to reduce these errors. Pernot et al. (2020) study non-Gaussian errors in Quantum Machine Learning models and find that mean measures of prediction error depend significantly on the shape of the error distribution. They also note the need for more research into the impact of error distributions on model reliability in general.

Similarly to the above studies, this work uses simulated perturbations of real data to study the impact of prediction errors under different noise distributions. However, none of the previous studies apply directly to the transport domain, and most of them deal with deviations from normality in linear modeling. In contrast, we neither presume any particular modeling form nor try to mitigate uncertainty, but rather examine its effect on demand-responsive PT. Per common modeling practices, we assume only that the predictive model has low bias and is evaluated through mean error measures, as detailed in the following Sections.

3. Experiments

3.1. Data

The data for the case study consists of PT trips from 1-Jan-2017 to 21-Dec-2018 in Metropolitan Copenhagen, Denmark. These trips are conducted with electronic travel cards (“Rejsekort”) and so account for approximately 1/3 of all bus and train trips. As dynamically routed Public Transport (PT) is still emerging, we consider a hypothetical “pilot” experiment, where dynamically routed buses serve the 6 most active PT

stations, depicted in Fig. 1. For these $6^2 - 6 = 30$ Origin-Destination (OD) pairs, the data contains 2.15 million trips, based on which we simulate various closed PT systems with urban-scale travel distances.

We aggregate the data by counting hourly trip starts for each OD pair. Then, we draw at random $N = 100$ h from 2018-Dec-1 00:00, ..., 2018-Dec-21 23:00. The Ground Truth observations are thus the hourly trip counts for each OD pair in these N hours. Based on the Ground Truth observations, we next simulate noisy predictions of PT demand.

3.2. Generation of noisy demand predictions

Similarly to the bulk of existing literature on non-Gaussian residuals (Section 2), we assume that the noise is generated from continuous distributions. As such, negative PT demand values are possible and could indicate preference to use other transport modes (e.g., bike, car, walking, etc.) over PT. Further, we use independent noise distributions, to deal with the challenging case of unpredictable changes in correlation structure. For instance, given two OD’s that are typically positively correlated, a non-recurrent disruption (e.g., a road block) might cause a surge in the usage of one of them along with a steep decline in the other. The noise distributions we experiment with are all homoscedastic with standard deviation (SD) σ , for $\sigma = 0.5, 1.0, 2.0, 3.0$,³ as follows.

1. Gaussian, i.e., Normal: $\mathcal{N}(0, \sigma^2)$.
2. Uniform: $\mathcal{U}[0, \sqrt{12} \sigma]$.
3. Exponential (\mathcal{E}) with scale σ .
4. Negated Exponential, namely, $-\mathcal{E}$.
5. Weibull (\mathcal{W}) with scale 1 and shape that corresponds to SD σ ; reduces to \mathcal{E} for $\sigma = 1$.
6. Negated Weibull, namely, $-\mathcal{W}$.

We shift back each distribution by its expected value to obtain zero mean, as illustrated in Fig. 2. Table 1 compares the noise distributions with $\mathcal{N}(0, 1)$, the standard Gaussian, in terms of their 3rd and 4th standardized moments, namely, skew and kurtosis. This Table shows that our experiments cover a range of distribution properties: platykurtic, leptokurtic and mesokurtic – i.e., having kurtosis below, above, or equal to

³ We have also experimented with a finer grained range of σ and obtained consistent results, which we thus omit for brevity.

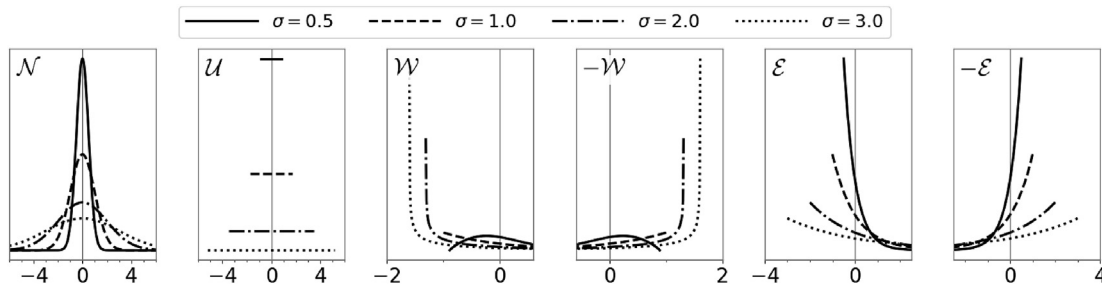


Fig. 2. Probability Density Functions (PDFs) of the noise distributions with zero mean. Note the different horizontal scale and support of each distribution. Vertical scales differ too and are omitted, because the density value is inconsequential in this work.

Table 1

Skew and kurtosis of noise distributions, marked in comparison to $\mathcal{N}(0, 1)$ as: **below**, *above*, equal.

σ	Property	\mathcal{N}	\mathcal{U}	\mathcal{E}	$-\mathcal{E}$	\mathcal{W}	$-\mathcal{W}$
0.5	Skew	0.00	1.30	0.75	-0.75	1.46	-1.46
	Kurtosis	0.19	1.80	1.50	1.50	2.35	2.35
1.0	Skew	0.00	10.39	6.00	-6.00	6.00	-6.00
	Kurtosis	3.00	28.80	24.00	24.00	24.00	24.00
2.0	Skew	0.00	83.14	48.00	-48.00	47.75	-47.75
	Kurtosis	48.00	460.80	384.00	384.00	624.50	624.50
3.0	Skew	0.00	280.59	162.00	-162.00	183.31	-183.31
	Kurtosis	243.00	2,332.80	1,944.00	1,944.00	4,975.26	4,975.26

$\mathcal{N}(0, 1)$ – as well as positive, negative and zero skew. The distributions also vary in PDF support, which is either finite (\mathcal{U}), semi-infinite (\mathcal{E} , $-\mathcal{E}$, \mathcal{W} , $-\mathcal{W}$) or infinite (\mathcal{N}).

For each OD pair (o, d) , shifted noise distribution \mathcal{D} and standard deviation σ , we independently draw N samples, $\delta_{1od}^{\mathcal{D},\sigma}, \dots, \delta_{Nod}^{\mathcal{D},\sigma} \sim \mathcal{D}$. Then for all $i = 1 \dots N$, we let g_{iod} denote the corresponding Ground Truth observation, and generate noisy predictions as follows:

$$\rho_{iod}^{\mathcal{D},\sigma} := g_{iod} + \sigma_g \cdot \delta_{iod}^{\mathcal{D},\sigma}, \tag{1}$$

where σ_g is the sample SD of all ground truth observations. The stochastic noise is thus measured in units of the SD of the observations themselves. Furthermore, every $\rho_{iod}^{\mathcal{D},\sigma}$ is unbiased, as all noise distributions have zero mean.

3.3. Fleet optimization

Having generated noisy predictions of PT demand, we now simulate the demand-responsive PT services with varying no. of buses, bus capacity, and percentage of dynamically routed buses. For capacity, our reference value is $110/3 \approx 33$, as the average bus capacity in Metropolitan Copenhagen is approx. 110, and Rejsekort accounts for approx. 1/3 of all PT trips. We vary capacity around this value as $\gamma = 10, 20, 30, 40$; the results later show that this provides sufficient insight into the effect of capacity variability. For percentage of dynamically routed buses, we use $\alpha = 0\%, 10\%, \dots, 40\%$. In particular, $\alpha = 0\%$ corresponds to a completely statically routed fleet, whose performance is thus independent of predictions and noise. Table 4 later shows that the mean and standard deviation of trip time per passenger remain stable as α increases from 30% to 40%, hence we do not increase α further. For this case study then, conversion to dynamic routing beyond $\alpha = 40\%$ does not yield significantly better routes.

Finally, we solve a mixed-integer linear programming (MILP) formulation to obtain optimal routes and fleet size. This formulation

receives as input a directed graph, where nodes correspond to PT stations, and an arc between two nodes exists if a bus can travel between the two stations. Each edge is weighted by the corresponding in-vehicle travel time. The input also includes a set of bi-directional routes R , where for every combination of 2 or more nodes, R contains the shortest acyclic path connecting these nodes. Given a total of π buses, the objective is to minimize both the number of buses actually deployed and passengers' total trip time (i.e., waiting time and in-vehicle travel time) in a given time horizon, by assigning routes to $\pi \leq \pi$ of the buses. In particular, we constrain π to fulfill two separate requirements: 1) π buses can satisfy the predicted demand in the given time horizon, 2) π buses can satisfy demand even in the 100 historically worst-case scenarios – i.e., the 100 busiest hours, by number of trips, before 1-Dec-2018 – if they are routed to serve all stops. In this manner, the optimized fleet is likely to be feasible for the actual, ground truth demand in the given time horizon.

We assume that passengers choose the shortest route from origin to destination, and waiting times and in-vehicle travel times are pre-specified. The optimizer first truncates any negative predictions to zero, as negative demand implies no passengers. Note that the noise PDFs themselves are not truncated, but rather the predictions derived from them.

Formulation 1 Fleet Optimization

$$\begin{aligned} & \text{Minimize} \\ & C\pi + \sum_{(o,d) \in Q, r \in R} c_{st}^{(1)} x_{rst}^{od} + \sum_{(o,d) \in Q, r \in R, s \in V_r, k=1.. \pi} c_{rk}^{(2)} b_{rks}^{od} \tag{2} \\ & \text{subject to} \\ & \sum_{k=1.. \pi} b_{rks}^{od} + \sum_{t:(s,t) \in E} x_{rst}^{od} - \sum_{t:(t,s) \in E} x_{rts}^{od} - a_{rs}^{od} = 0 \quad \forall (o,d) \in Q, r \in R, s \in V_r \tag{3} \end{aligned}$$

$$\sum_{\substack{r \in R, \\ k=1.. \pi}} b_{rks}^{od} - \sum_{r \in R} a_{rs}^{od} = \begin{cases} \rho_{iod}^{D,\sigma} & s=o- \\ \rho_{iod}^{D,\sigma} & s=d \\ 0 & s \notin \{o,d\} \end{cases} \quad \forall (o,d) \in Q, s \in V_r \quad (4)$$

$$\sum_{s \in V_r, (o,d) \in Q} b_{rks}^{od} - M y_{rk} \leq 0 \quad \forall r \in R, k = 1.. \pi \quad (5)$$

$$\sum_{(o,d) \in Q} x_{rst}^{od} - \gamma k y_{rk} \leq 0 \quad \forall r \in R, (s,t) \in E \quad (6)$$

$$\sum_{k=1.. \pi} y_{rk} \leq 1 \quad \forall r \in R \quad (7)$$

$$\sum_{r \in R, k=1.. \pi} k y_{rk} - \pi \leq 0 \quad (8)$$

$$\sum_{k=1.. \pi} k y_{rk} - (1 - \alpha) \pi \geq 0 \quad \forall k = 1.. \pi \quad (9)$$

$$\pi \leq \pi \quad (10)$$

$$b_{jr\pi s}^{od} + \sum_{t:(s,t) \in E} x_{jr\pi st}^{od} - \sum_{t:(t,s) \in E} x_{jr\pi ts}^{od} - a_{jr\pi s}^{od} = \begin{cases} \rho_{jod} & s=o - \rho_{jod} \\ 0 & s=d \\ 0 & s \notin \{o,d\} \end{cases} \quad \forall (o,d) \in Q, s \in V_r, j=1 \dots 100 \quad (11)$$

$$\sum_{(o,d) \in Q} x_{jr\pi st}^{od} - \gamma \pi \leq 0 \quad \forall (s,t) \in E, j = 1 \dots 100 \quad (12)$$

$$y_{rk} \in \{0,1\}, x_{rst}^{od}, b_{rks}^{od}, a_{rs}^{od}, \pi \geq 0, \pi \text{ is integral } \forall r \in R, s,t \in V_r, (o,d) \in Q, k=1.. \pi \quad (13)$$

For any hour $i = 1 \dots N$, Formulation 1 defines the linear program per the notation in Table 2. The objective and constraints are similar to those commonly used for fleet optimization, as follows. In objective (2), the first term is the operational cost of deploying a fleet of π buses. We use a large cost coefficient C per bus, to obtain the smallest fleet that still satisfies the demands. The second term is the total passengers' in-vehicle travel time, and the third term is total passengers' waiting time. Constraint (3) imposes flow conservation, such that the flow into any node of any route (from any previous node in the route) either proceeds into the next node in the route or flows out of the route through a_{rs}^{od} . This ensures that the itinerary of each passenger is complete from their origin to their destination. Constraint (4) sets the flow per boarding and alighting at a stop s . This ensures that the total flow for OD pair (o, d) into node o across all routes sums to the demand, and the total flow for OD pair (o, d) out of node d across all routes sum to the same demand. Constraint (5) ensures that when a passenger boards any route, the route and the corresponding no. of buses are selected, i.e., $y_{rk} = 1$. For example, if the flow b_{rks}^{od} is non-zero, then y_{rk} must be equal to one, indicating that exactly k buses are allocated to route r . Constraint (6) defines the capacity along all segments of all routes. Constraint (7) ensures that only one y_{rk} is set to 1 for each route, such that the number of buses allocated to each route is deterministic. Constraint (8) then guarantees that the total number of buses allocated to all routes does not exceed the deployed fleet size π , while constraint (9) dictates that at least $(1 - \alpha)\pi$ buses serve route

Table 2
Notation for formulation 1.

Sets	V	nodes
	E	arcs
	Q	OD pairs
	R	routes
	V_r	nodes in route r
Indices	s, t	nodes
	(o, d)	OD pair
	r	route
	\underline{r}	base route that serves all nodes
	k	no. buses allocated to a route
Parameters	α	portion of buses that can be dynamically routed
	γ	capacity of each bus
	$c_{st}^{(1)}$	in-vehicle travel time from s to t
	$c_{rk}^{(2)}$	average waiting time for route r , when r is allocated k buses
	M	large number
	C	operational cost of each bus
	π	maximum fleet size
Decision Variables	π	fleet size
	x_{rst}^{od}	flow of passengers traveling from s to t in route r , as part of a trip from o to d
	b_{rks}^{od}	flow of passengers who board route r at s , as part of a trip from o to d , when r is allocated k buses
	a_{rs}^{od}	flow of passengers who alight from route r at s , as part of a trip from o to d
	y_{rk}	binary indicator of allocation of k buses to route r

\underline{r} , which is route that serves all stops. Constraint (10) imposes the maximum fleet size. Constraints (11) and (12) ensure that the fleet of π buses, when routed to serve all stops, can satisfy the demand in each of the 100 busiest hours before 1-Dec-2018 (denoted by subscript j). Finally, constraint (13) defines the possible set of values for each decision variable. We use $M = 1\ 00\ 000$, $C = 1\ 00\ 000$ and $\pi = 40$ in this case study.

4. Results

The optimized fleet sizes $\pi = 38, 19, 13, 10$ are consistent across all generated noisy demands for $\gamma = 10, 20, 30, 40$, respectively. Table 3 further specifies the no. dynamically and statically routed buses for each γ and α . This percentage is seen to correspond to the PDFs in Fig. 2; e.g., as σ increases, $-\mathcal{W}$ yields the fewest negative predictions as its PDF shifts towards the positives.

We evaluate prediction quality through two commonly used, unitless measures of mean error: Mean Absolute Percentage Error (MAPE) and Rooted Mean Squared Normalized Error (RMSNE). These measures are defined as follows:

$$\begin{aligned} \text{MAPE}(\mathcal{D}, \sigma) &:= \frac{1}{N} \sum_{i=1}^N \frac{1}{|Q|} \sum_{(o,d) \in Q} \frac{|\rho_{iod}^{D,\sigma} - g_{iod}|}{g} \\ &= \frac{\sigma_g}{g|Q|N} \sum_{i=1}^N \sum_{(o,d) \in Q} |\delta_{iod}^{D,\sigma}|, \end{aligned} \quad (14)$$

Table 3
No. dynamically (D) and statically routed (S) buses.

γ	$\alpha = 0\%$		$\alpha = 10\%$		$\alpha = 20\%$		$\alpha = 30\%$		$\alpha = 40\%$	
	D	S	D	S	D	S	D	S	D	S
10	0	38	4	34	8	30	12	26	16	22
20	0	19	2	17	4	15	6	13	8	11
30	0	13	2	11	3	10	4	9	6	7
40	0	10	1	9	2	8	3	7	4	6

$$\text{RMSNE}(\mathcal{D}, \sigma) := \sqrt{\frac{1}{N} \sum_{i=1}^N \frac{1}{|Q|} \sum_{(o,d) \in Q} \left(\frac{\rho_{iod}^{\mathcal{D}, \sigma} - g_{iod}}{g} \right)^2} = \frac{\sigma_g}{g \sqrt{|Q|N}} \sqrt{\sum_{i=1}^N \sum_{(o,d) \in Q} (\delta_{iod}^{\mathcal{D}, \sigma})^2}, \quad (15)$$

where g is the mean of all ground truth observations. Each right-hand side follows from (1) and shows the dependency on sampled noise (where the fractional coefficient is noise-independent).

Fig. 3 illustrates the MAPE and RMSNE after the optimizer truncates negative predictions. As expected, both MAPE and RMSNE increase as σ increases. When ranking \mathcal{D} from best to worst, we obtain $-\mathcal{W} \leq \mathcal{W} \leq -\mathcal{E} \leq \mathcal{E} \leq \mathcal{N} \leq \mathcal{U}$ for MAPE vs. $-\mathcal{W} \leq -\mathcal{E} \leq \mathcal{N} \leq \mathcal{U} \leq \mathcal{E} \leq \mathcal{W}$ for RMSNE. Let us next analyze how closely the optimization performance follows any of these rankings.

We evaluate optimization performance through several measures, all of which are based on the objective value (i.e., total trip time). For generalizability, our results are mostly given in relative terms. However, to give a basic sense of scale for this case study, we begin with an absolute measure of trip time (min) per passenger. For this, let

$$T_i(\mathcal{D}, \sigma, \alpha, \gamma) := \frac{F_i(\mathcal{D}, \sigma, \gamma, \alpha)}{P_i}, \quad (16)$$

where for all $i = 1 \dots N$, P_i is total no. passengers observed in the i 'th hour, and $F_i(\mathcal{D}, \sigma, \gamma, \alpha)$ is their total trip time using the optimized fleet.

Table 4 summarizes $T_i(\mathcal{D}, \sigma, \alpha, \gamma)$ through its mean and SD over all i, \mathcal{D}, σ . The mean and SD are both seen to decrease as γ decreases and/or α increases. Hence as expected, trip times improve if more buses with lower capacity are used and/or more of the buses are dynamically routed. We also note that the mean and SD do not vary much when further separating by \mathcal{D} and σ , for any fixed α, γ . This is a possible consequence of using the same marginal noise distribution for all OD pairs, so that their predicted demands shift similarly, thereby balancing each other out during fleet optimization.

We now proceed to measure how much time per passenger is theoretically lost when optimizing with noisy vs. perfect predictions, as:

Table 4
Mean (\pm SD) of $T_i(\mathcal{D}, \sigma, \alpha, \gamma)$, in minutes. Lower is better.

α	$\gamma = 10$	$\gamma = 20$	$\gamma = 30$	$\gamma = 40$
0%	15.0 (± 6.6)	16.8 (± 7.1)	18.4 (± 7.6)	19.9 (± 8.1)
10%	13.2 (± 5.9)	15.5 (± 6.7)	16.8 (± 7.2)	19.4 (± 8.0)
20%	11.8 (± 5.5)	14.2 (± 6.3)	15.9 (± 7.1)	18.1 (± 7.8)
30%	11.2 (± 5.4)	13.5 (± 6.3)	15.5 (± 7.1)	17.3 (± 7.8)
40%	11.1 (± 5.4)	13.4 (± 6.2)	15.4 (± 7.0)	17.2 (± 7.8)

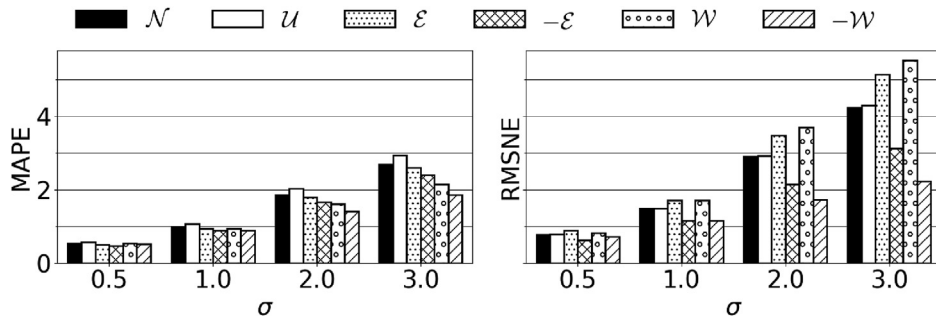


Fig. 3. MAPE and RMSNE after the optimizer truncates negative predictions. Lower is better.

$$T_{\text{loss}}(\mathcal{D}, \sigma, \gamma, \alpha) := \frac{1}{N} \sum_{i=1}^N \frac{F_i(\mathcal{D}, \sigma, \gamma, \alpha) - F_i(\mathcal{D}, 0, \gamma, \alpha)}{P_i}. \quad (17)$$

Fig. 4 illustrates $T_{\text{loss}}(\mathcal{D}, \sigma, \gamma, \alpha)$ using a grid of plots, where rows are ordered by percentage of dynamic buses (α) and columns are ordered by bus capacity (γ). We see that $T_{\text{loss}}(\mathcal{D}, \sigma, \gamma, \alpha)$ increases when either σ or γ increases, as expected. We also see that for any fixed $\sigma > 0$, a partial order on \mathcal{D} emerges as γ and α increase, so that $-\mathcal{W}$ and $-\mathcal{E}$ are significantly better than all other \mathcal{D} , and $-\mathcal{W} < -\mathcal{E} < \mathcal{N}$. These properties hold similarly for RMSNE but not MAPE, as seen earlier in Fig. 3. By properties of RMSNE, this suggests that as more dynamically routed buses are involved, the quality of routing becomes more sensitive to the presence of even a few large prediction errors – i.e., consistent predictive accuracy is indeed essential.

Interestingly, Fig. 4 also shows that $T_{\text{loss}}(\mathcal{D}, \sigma, \gamma, \alpha)$ generally increases when α increases, i.e., when more dynamic buses are used for the same fleet size. This holds also when normalizing T_{loss} by the theoretical trip time with perfect predictions, namely:

$$T_{\text{loss}}^{\text{rel}}(\mathcal{D}, \sigma, \gamma, \alpha) := \frac{1}{N} \sum_{i=1}^N \frac{F_i(\mathcal{D}, \sigma, \gamma, \alpha) - F_i(\mathcal{D}, 0, \gamma, \alpha)}{F_i(\mathcal{D}, 0, \gamma, \alpha)}, \quad (18)$$

as detailed in Table 7. Still, T_{loss} and $T_{\text{loss}}^{\text{rel}}$ are purely theoretical measurements, because observations cannot realistically be used before they manifest, and the predictive mean rarely captures them perfectly.

Next, we measure the time saved per passenger when using dynamic buses ($\alpha > 0$) vs. a completely statically routed fleet ($\alpha = 0$), namely:

$$T_{\text{gain}}(\mathcal{D}, \sigma, \gamma, \alpha) := \frac{1}{N} \sum_{i=1}^N \frac{F_i(\mathcal{D}, \sigma, \gamma, 0\%) - F_i(\mathcal{D}, \sigma, \gamma, \alpha)}{P_i}, \quad (19)$$

as illustrated in Fig. 5. We see that T_{gain} generally improves as capacity and noise SD decrease and the percent of dynamic buses increases. The only exception is $T_{\text{gain}}(\mathcal{D}, \sigma, 20, 10\%) > T_{\text{gain}}(\mathcal{D}, \sigma, 30, 20\%)$, for any \mathcal{D} and σ . This is explained by Table 3: for $\alpha = 10\%$, both $\gamma = 20, 30$ have the same no. dynamic buses, yet there are more statically routed buses for $\gamma = 20$ than for $\gamma = 30$.

We also see in Fig. 5 that as γ and α increase for any σ , a partial order on \mathcal{D} again emerges, with the same properties detected above for T_{loss} and RMSNE. This further supports our conclusion, that the consistency of predictive accuracy becomes more essential as the fleet becomes increasingly dynamic. The best is $T_{\text{gain}}(\mathcal{D}, \sigma, 10, 40\%) = 4$ min, for $\sigma \leq 0.5$ and any \mathcal{D} . The lowest is $T_{\text{gain}}(\mathcal{D}, 3.0, 40, 10\%) = 0.1$ min, for $\mathcal{D} \in \{\mathcal{N}, \mathcal{U}, \mathcal{E}\}$.

We further convert T_{gain} to an economic measure per the Danish

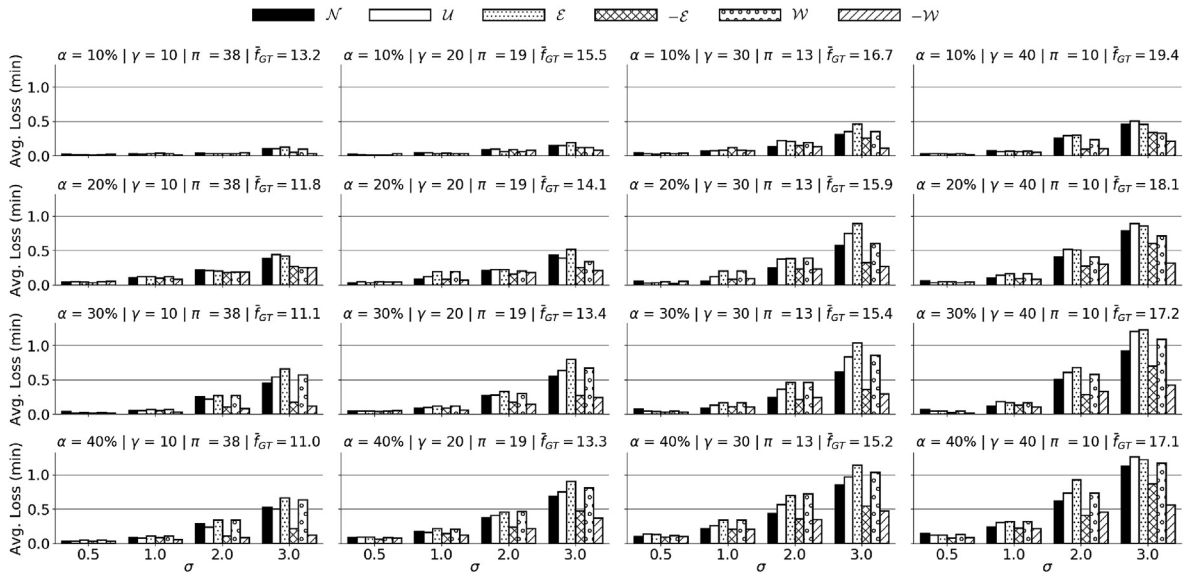


Fig. 4. Average minutes lost with noisy vs. perfect predictions (f_{GT} in plot titles). Lower is better. Numeric results are in Table 6.

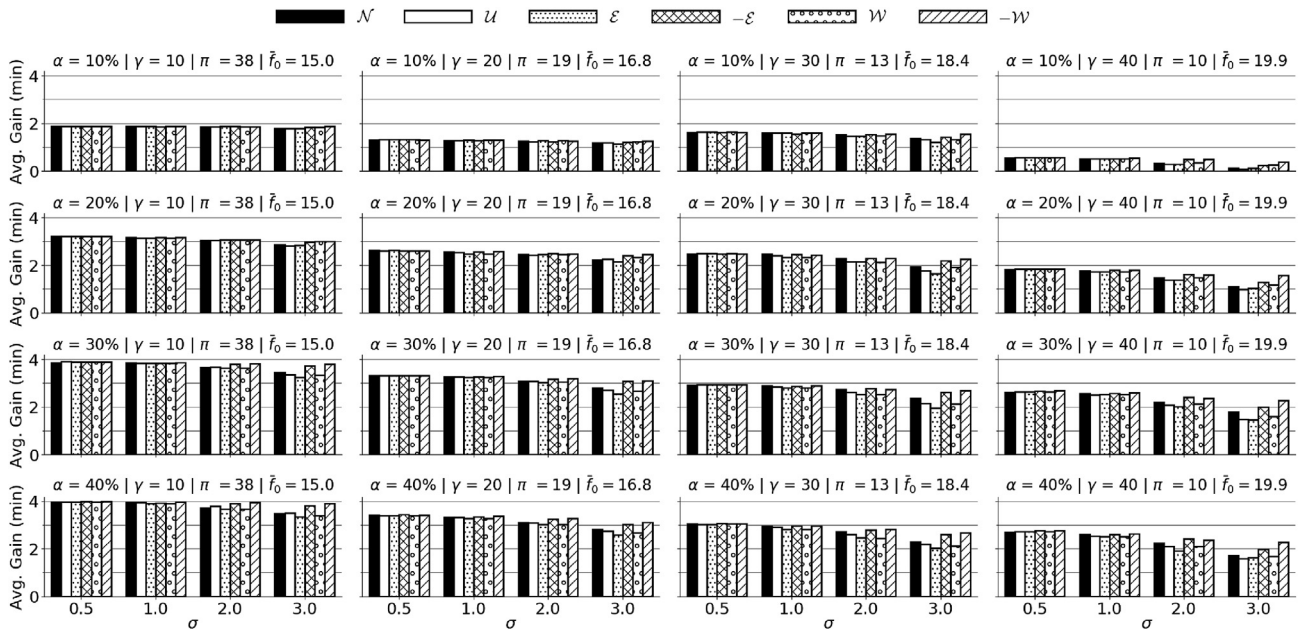


Fig. 5. Average minutes gained with dynamic vs. completely static routing (f_0 in plot titles). Higher is better. Numeric results are in Table 8.

Table 5
Minimum yearly economic gains (€). Higher is better.

α	$\gamma = 10$	$\gamma = 20$	$\gamma = 30$	$\gamma = 40$
10%	429,000	274,000	293,000	17,000
20%	677,000	515,000	392,000	238,000
30%	783,000	615,000	469,000	354,000
40%	809,000	623,000	488,000	382,000

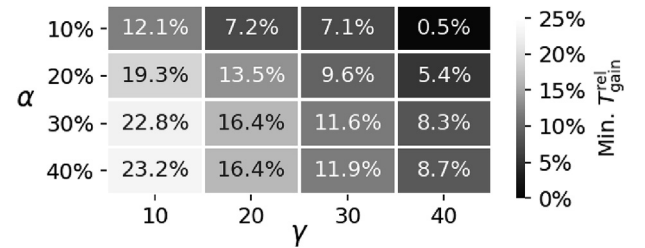


Fig. 6. Minimum $T_{gain}^{rel}(D, \sigma, \gamma, \alpha)$ over all D, σ . Higher is better.

Value of Travel Time Savings (VTTs), which has recently been estimated at $\nu = 13.43\text{€}/\text{h}$ (Rich and Vandet, 2019). For each α and γ , we thus take the smallest $T_{gain}(D, \sigma, \gamma, \alpha)$ over all D, σ and multiply it by both ν (in $\text{€}/\text{min}$) and average no. trips per year in the studied OD pairs. Table 5 provides the results, rounded to 1000€/year, where the best yearly gains are at least 809,000€ (for $\alpha = 40\%$, $\gamma = 10$).

Lastly, we measure the average relative gain when using a dynamic vs. completely static fleet, as:

$$T_{\text{gain}}^{\text{rel}}(\mathcal{D}, \sigma, \gamma, \alpha) := \frac{1}{N} \sum_{i=1}^N \frac{F_i(\mathcal{D}, \sigma, \gamma, 0\%) - F_i(\mathcal{D}, \sigma, \gamma, \alpha)}{F_i(\mathcal{D}, \sigma, \gamma, 0\%)}. \quad (20)$$

This is detailed in Table 9, where in particular, the best values of $T_{\text{gain}}^{\text{rel}}$ are close to 27.5% while the worst are close to 5%. As a concise reference for choosing operational parameters, Fig. 6 summarizes $T_{\text{gain}}^{\text{rel}}$ in terms of its minimum over all noise distributions. For instance, if operational parameters are chosen conservatively, so that each bus is large ($\gamma = 40$) and only $\alpha = 20\%$ of buses are dynamic, then at least 5% relative gain is achieved. However, if the fleet consists of vehicles with low capacity ($\gamma = 10$), many of which are dynamic ($\alpha = 40\%$), then at least 23% relative gain is achieved.

5. Discussion

The main goal of this work is to quantify the effects of demand prediction accuracy on the performance of demand-responsive Public Transport (PT). For this, we have used hourly observations of PT trips as a proxy for transport demand, and conducted simulation experiments in two steps. First, we have simulated the output of demand prediction models by perturbing the observations per various distributions, which cover a wide range of statistic properties. Based on these noisy predictions, we have then used a mixed-integer linear program (MILP) with commonly used objective and constraints, to simulate and optimize PT fleets with varying no. statically and dynamically routed buses.

We have obtained that the differences in noise distributions do not account for much variability in trip time per passenger. However, the noise distributions differ noticeably in two other measures of time per passenger: 1) time theoretically lost with noisy vs. perfect predictions, 2) time gained with dynamic vs. completely static routing. The worst loss per passenger is 1.3 min, while the best time gain per passenger is 4 min, which is more than 27% in relative terms. In economic terms of Value of Time Savings (VTTS), the best gains in this case study are at least 809,000€/year.

Also in terms of time gains and losses, we have obtained that the noise distributions rank more similarly to Rooted Squared Mean Normalized Error (RMSNE) than Mean Absolute Percentage Error (MAPE) of predictions. As seen in (14) and (15), RMSNE is dominated by exceptionally large prediction errors due to squaring, unlike MAPE. It thus appears that exceptionally large prediction errors, even if few, can strongly influence the performance of dynamic PT optimization.

Finally, we have seen that when the common normality assumption is violated, optimization performance can not only worsen but also improve. E.g., compared to the average gains and losses of the Gaussian distribution (\mathcal{N}), the Uniform (\mathcal{U}) and Weibull (\mathcal{W}) are mostly worse, whereas the Negated Weibull ($-\mathcal{W}$) and Negated Exponential ($-\mathcal{E}$) are mostly better. In conjunction with Table 1, we find that this corresponds well to skew (rather than kurtosis), as \mathcal{U} and \mathcal{W} have positive skew, while $-\mathcal{W}$ and $-\mathcal{E}$ have negative skew.

Appendix. numeric results

Table 6
 $T_{\text{loss}}(\mathcal{D}, \sigma, \gamma, \alpha)$ in min, lower is better.

α	σ	$\gamma = 10$						$\gamma = 20$						$\gamma = 30$						$\gamma = 40$							
		\mathcal{N}	\mathcal{U}	\mathcal{E}	$-\mathcal{E}$	\mathcal{W}	$-\mathcal{W}$	\mathcal{N}	\mathcal{U}	\mathcal{E}	$-\mathcal{E}$	\mathcal{W}	$-\mathcal{W}$	\mathcal{N}	\mathcal{U}	\mathcal{E}	$-\mathcal{E}$	\mathcal{W}	$-\mathcal{W}$	\mathcal{N}	\mathcal{U}	\mathcal{E}	$-\mathcal{E}$	\mathcal{W}	$-\mathcal{W}$		
10%	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
	2.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.2	0.1	0.2	0.1	0.3	0.3	0.3	0.1	0.2	0.1			
	3.0	0.1	0.1	0.1	0.1	0.1	0.0	0.1	0.1	0.2	0.1	0.1	0.1	0.3	0.4	0.5	0.3	0.3	0.1	0.5	0.5	0.5	0.3	0.3	0.2		
20%	0.5	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.1	0.0	0.1	0.1	0.0	0.0	0.0	0.0	0.0			

(continued on next column)

6. Conclusion

In conclusion, the key findings of this work are as follows.

1. As a PT fleet becomes increasingly dynamic, the consistency of predictive accuracy becomes more essential for an effective service, i.e., to decrease temporal and economical losses and increase corresponding gains.
2. However, these gains still generally increase when using more dynamically routed buses (i.e., higher α) with less capacity (i.e., lower γ), regardless of noise distribution.
3. The minimum relative gain is 5% when conservatively choosing $\alpha = 20\%$ and $\gamma = 40$ vs. 23% when more liberally choosing $\alpha = 10\%$ and $\gamma = 10$. Fig. 6 gives a fuller reference for choosing these operational parameters.
4. Violations of the common normality assumption can in fact result in more reliable predictions, e.g., as in the case of error distributions with a negative skew. This encourages further research to identify conditions where predictive errors can actually improve predictive quality.

For future work, we plan to enhance the current study in several respects, as follows. We plan to also study cases where the noise distribution incorporates correlations between the OD's, by constructing their joint distribution (Peled et al., 2019). We also plan to extend to the more general case of a full-scale PT fleet that serves more stations, and study how the predictive errors can be mitigated via robust optimization with chance constraints (Wang et al., 2015). We can also incorporate travelers' route choice behaviour by extending to a bi-modal optimization formulation (Jiang, 2021; Jiang and Ceder, 2021). We further plan to incorporate truncated noise distributions (Greene, 2006), to account for the inherently limited observability of PT usage (Gammelli et al., 2020), as well as investigate conditions where predictive errors improve predictive quality. Given more computational resources, we also plan to measure the effects of higher percentages of dynamically routed buses on trip time minimization. We can also consider a stochastic optimization formulation with recourse actions, to explicitly account for consequences of decisions made before the true travel demand is realized.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Data for this research was obtained by kind permission of the Danish Transport, Construction and Housing Authority, Movia Transit Agency, the Danish National Rail Company and the Danish Metro Company.

Table 9
 $T_{gain}^{rel}(D, \sigma, \gamma, \alpha)$ in %, higher is better.

α	$\gamma = 10$				$\gamma = 20$				$\gamma = 30$				$\gamma = 40$			
	σ	\mathcal{N}	\mathcal{U}	$-\mathcal{W}$	\mathcal{N}	\mathcal{U}	$-\mathcal{W}$	$-\mathcal{E}$	\mathcal{N}	\mathcal{U}	$-\mathcal{W}$	$-\mathcal{E}$	\mathcal{N}	\mathcal{U}	$-\mathcal{W}$	$-\mathcal{E}$
10%	0.0	12.7	12.7	12.7	8.2	8.2	8.2	8.2	9.5	9.5	9.5	9.5	3.1	3.1	3.1	3.1
	0.5	12.6	12.6	12.6	8.0	8.1	8.1	8.1	9.3	9.3	9.3	9.3	3.0	2.9	3.0	3.0
	1.0	12.6	12.5	12.6	7.9	7.9	8.0	7.9	9.1	9.0	9.0	8.9	2.8	2.8	2.8	2.8
	2.0	12.5	12.5	12.5	7.7	7.5	7.8	7.7	8.8	8.2	8.3	8.7	1.7	1.6	1.6	1.9
20%	3.0	12.1	12.1	12.3	7.4	7.5	7.6	7.3	7.8	7.8	7.1	8.0	0.7	0.5	0.8	1.2
	0.0	22.2	22.2	22.2	16.3	16.3	16.3	16.3	14.3	14.3	14.3	14.3	10.0	10.0	10.0	10.0
	0.5	22.0	21.9	22.0	16.1	16.1	16.1	16.1	14.1	14.2	14.2	14.1	9.8	9.9	9.8	9.8
	1.0	21.5	21.3	21.4	15.8	15.6	15.4	15.9	14.1	13.7	13.5	14.0	9.6	9.3	9.3	9.7
30%	2.0	20.7	20.6	20.7	15.2	14.8	14.9	15.3	13.1	12.1	12.3	13.1	8.0	7.3	7.5	8.8
	3.0	19.7	19.3	19.4	13.9	14.1	13.5	14.7	11.3	10.4	9.6	12.6	5.9	5.5	5.4	7.0
	0.0	26.8	26.8	26.8	20.8	20.8	20.8	20.8	17.1	17.1	17.1	17.1	14.4	14.4	14.4	14.4
	0.5	26.7	26.7	26.7	20.6	20.6	20.5	20.6	16.8	16.9	16.9	16.9	14.2	14.2	14.2	14.2
40%	1.0	26.5	26.5	26.4	20.3	20.2	20.3	20.2	16.6	16.4	16.3	16.6	13.9	13.6	13.7	14.0
	2.0	25.2	25.2	25.0	19.3	18.8	18.8	19.7	15.8	14.8	14.6	16.0	11.8	11.3	11.1	13.0
	3.0	23.7	23.3	22.8	17.8	17.4	16.4	19.0	13.7	12.9	11.6	15.1	9.5	8.4	8.3	10.9
	0.0	27.5	27.5	27.5	21.5	21.5	21.5	21.5	17.9	17.9	17.9	17.9	15.1	15.1	15.1	15.1
50%	0.5	27.3	27.3	27.3	21.1	21.1	21.0	21.2	17.4	17.3	17.3	17.4	14.5	14.6	14.6	14.7
	1.0	27.1	27.1	26.9	20.6	20.7	20.4	20.9	16.8	16.7	16.4	17.0	14.0	13.7	13.6	14.2
	2.0	25.7	26.0	25.3	19.5	19.0	18.8	20.2	15.7	14.7	14.3	16.2	12.0	11.4	10.6	13.1
	3.0	24.0	24.4	23.2	17.7	17.3	16.4	18.7	13.3	12.9	11.9	15.0	9.2	8.7	9.0	10.8

References

Alonso-Mora, J., Wallar, A., Rus, D., 2017. Predictive routing for autonomous mobility-on-demand systems with ride-sharing. In: IEEE International Conference on Intelligent Robots and Systems. IROS), pp. 3583–3590.

An, K., Lo, H.K., 2015. Robust transit network design with stochastic demand considering development density. *Transp. Res. Part B Methodol.* 81, 737–754.

Beaudrie, C.E., Kandlikar, M., Ramachandran, G., 2016. Using expert judgment for risk assessment. In: *Assessing Nanoparticle Risks to Human Health*, second ed. Elsevier Inc., Cambridge MA, pp. 91–119.

Bösch, P.M., Becker, F., Becker, H., Axhausen, K.W., 2018. Cost-based analysis of autonomous mobility services. *Transport Pol.* 64, 76–91.

Ceder, A., 2016. *Public Transit Planning and Operation: Modeling, Practice and Behavior*. CRC press, Boca Raton FL.

Davies, N., Spedding, T., Watson, W., 1980. Autoregressive moving average processes with non-normal residuals. *J. Time Anal.* 1 (2), 103–109.

Gammelli, D., Peled, I., Rodrigues, F., Pacino, D., Kurtaran, H.A., Pereira, F.C., 2020. Estimating Latent Demand of Shared Mobility through Censored Gaussian Processes, vol. 120. *Transportation Research Part C: Emerging Technologies*.

Gourieroux, C., Monfort, A., Renault, E., Trognon, A., 1987. Generalised residuals. *J. Econom.* 34 (1–2), 5–32.

Greene, W., 2006. Censored data and truncated distributions. *Econometric Theory*. In: Mills, T.C., Patterson, K. (Eds.), *Palgrave Handbook of Econometrics*, first ed., vol. 1. Palgrave Macmillan, Basingstoke (chapter 20).

Hadjimitsriou, N.S., Lippi, M., Mamei, M., 2020. A data driven approach to match demand and supply for public transport planning. *IEEE Trans. Intell. Transport. Syst.* 1–11.

Hashemi, H., Abdelghany, K., 2015. Real-time traffic network state prediction for proactive traffic management: simulation experiments and sensitivity analysis. *Transport. Res. Rec.* 2491 (1), 22–31.

He, Y., Raghunathan, T.E., 2009. On the performance of sequential regression multiple imputation methods with non normal error distributions. *Commun. Stat. Simulat. Comput.* 38 (4), 856–883.

Horázdovský, P., Novotný, V., Světek, M., 2018. Data-driven management of dynamic public transport. In: *IEEE 2018 Smart City Symposium Prague. SCSP*, pp. 1–5.

Huang, A., Dou, Z., Qi, L., Wang, L., 2020. Flexible route optimization for demand-responsive public transit service. *J. Transport. Eng., Part A: Systems* 146 (12).

Hyland, M., Mahmassani, H.S., 2018. Dynamic autonomous vehicle fleet operations: optimization-based strategies to assign avts to immediate traveler demand requests. *Transport. Res. C Emerg. Technol.* 92, 278–297.

Ibeas, A., Alonso, B., dell’Olio, L., Moura, J.L., 2014. Bus size and headways optimization model considering elastic demand. *J. Transport. Eng.* 140 (4).

Iliopoulou, C., Kepaptsoglou, K., 2019. Combining ITS and optimization in public transportation planning: state of the art and future research paths. *European Transport Research Review* 11 (1).

Ivković, I., Rajič, V., Stanojević, J., 2020. Coverage probabilities of confidence intervals for the slope parameter of linear regression model when the error term is not normally distributed. *Commun. Stat. Simulat. Comput.* 49 (1), 147–158.

Jackson, D., White, I.R., 2018. When should meta-analysis avoid making hidden normality assumptions? *Biom. J.* 60 (6), 1040–1058.

Jiang, Y., 2021. Reliability-based equitable transit frequency design. *Transportmetrica: Transport. Sci.* 1–31.

Jiang, Y., Ceder, A.A., 2021. Incorporating personalization and bounded rationality into stochastic transit assignment model. *Transport. Res. C Emerg. Technol.* 127, 103–127.

Krishnakumari, P., van Lint, H., Djukic, T., Cats, O., 2020. A data driven method for od matrix estimation. *Transport. Res. C Emerg. Technol.* 113, 38–56.

Lee, Y.J., 2006. Transit network sensitivity analysis. *Journal of Public Transportation* 9 (1), 21–52.

Luathep, P., Sumalee, A., Lam, W.H., Li, Z.C., Lo, H.K., 2011. Global optimization method for mixed transportation network design problem: a mixed-integer linear programming approach. *Transp. Res. Part B Methodol.* 45 (5), 808–827.

Mak, T.K., 2000. Heteroscedastic regression models with non-normally distributed errors. *J. Stat. Comput. Simulat.* 67 (1), 21–36.

Manasra, H., Toledo, T., 2019. Optimization-based operations control for public transportation service with transfers. *Transport. Res. C Emerg. Technol.* 105, 456–467.

Nelson, H., Granger, C., 1979. Experience with using the box-cox transformation when forecasting economic time series. *J. Econom.* 10 (1), 57–69.

Nickkar, A., Lee, Y.J., Meskar, M., 2020. Sensitivity analysis for the optimal automated demand responsive feeder transit system. In: *Automated People Movers and Automated Transit Systems 2020*. American Society of Civil Engineers (ASCE), Reston VA, pp. 85–96.

Peled, I., Lee, K., Jiang, Y., Dauwels, J., Pereira, F.C., 2019. Online Predictive Optimization Framework for Stochastic Demand-Responsive Transit Services, pp. 1–34 arXiv preprint arXiv:1902.09745.

Pernot, P., Huang, B., Savin, A., 2020. Impact of non-normal error distributions on the benchmarking and ranking of quantum machine learning models. *Mach. Learn.: Sci. Technol.* 1 (3).

Rich, J., Vandet, C.A., 2019. Is the value of travel time savings increasing? analysis throughout a financial crisis. *Transport. Res. Pol. Pract.* 124, 145–168.

Saltelli, A., Tarantola, S., Campolongo, F., Ratto, M., 2004. *Sensitivity Analysis in Practice: a Guide to Assessing Scientific Models*, first ed. Wiley Online Library.

Seber, G.A., Lee, A.J., 2012. *Linear Regression Analysis*, second ed. John Wiley & Sons, Hoboken NJ.

Szeto, W., Jiang, Y., 2014. Transit route and frequency design: Bi-level modeling and hybrid artificial bee colony algorithm approach. *Transp. Res. Part B Methodol.* 67, 235–263.

Thode, H.C., 2002. *Testing for Normality*, first ed. CRC press, Boca Raton FL.

Tong, L., Zhou, X., Miller, H.J., 2015. Transportation network design for maximizing space–time accessibility. *Transp. Res. Part B Methodol.* 81, 555–576.

Toole, J.L., Colak, S., Sturt, B., Alexander, L.P., Evsukoff, A., González, M.C., 2015. The path most traveled: travel demand estimation using big data resources. *Transport. Res. C Emerg. Technol.* 58, 162–177.

Ukkusuri, S.V., Mathew, T.V., Waller, S.T., 2007. Robust transportation network design under demand uncertainty. *Comput. Aided Civ. Infrastruct. Eng.* 22 (1), 6–18.

Wang, D.Z., Lo, H.K., 2010. Global optimum of the linearized network design problem with equilibrium flows. *Transp. Res. Part B Methodol.* 44 (4), 482–492.

Wang, H., Lam, W.H.K., Zhang, X., Shao, H., 2015. Sustainable transportation network design with stochastic demands and chance constraints. *International Journal of Sustainable Transportation* 9 (2), 126–144.

Wang, X., Qing-dao-er ji, R., 2019. Application of optimized genetic algorithm based on big data in bus dynamic scheduling. *Cluster Comput.* 22 (6), 15439–15446.

Winter, K., Cats, O., Correia, G., van Arem, B., 2018. Performance analysis and fleet requirements of automated demand-responsive transport systems as an urban public transport service. *International Journal of Transportation Science and Technology* 7 (2), 151–167.

Wolters, R., Kateman, G., 1989. The performance of least squares and robust regression in the calibration of analytical methods under non-normal noise distributions. *J. Chemometr.* 3 (2), 329–342.



Yu Jiang received his PhD in civil engineering from the University of Hong Kong (HKU) in 2014. He is currently an Assoc. Professor in the Technical University of Denmark (DTU), specializing in modeling and optimizing public transport system.



Justin Dauwels received his PhD in electrical engineering from ETH, Zurich, in 2005. He is an Assoc. Professor at the TU Delft. His main research interests lie in probabilistic modeling, intelligent transportation systems and human behavior analysis. He is also co-founder of several AI related spin-off companies.



Inon Peled is a Postdoctoral Researcher in the Machine Learning for Smart Mobility research group in the Technical University of Denmark (DTU), where he received his PhD in 2021. His main research interests lie in predictive modeling for abnormal conditions in the transport domain.



Francisco Câmara Pereira received his PhD in computer science from the University of Coimbra, Portugal, in 2005. He is a Full Professor and leader of the Machine Learning for Smart Mobility group in DTU. His main research interests lie in applying ML for predicting and improving transport systems.



Kelvin Lee is a PhD candidate with the Graduate College at Nanyang Technological University (NTU), Singapore. He has received his BEng degree in electrical and electronic engineering from NTU. His main research interests lie in predictive optimization methods with applications in transportation.