

## Highlights

### **Predicting the probability distribution of bus travel time to move towards reliable planning of public transport services**

Léa Ricard, Guy Desaulniers, Andrea Lodi, Louis-Martin Rousseau

- We propose probabilistic models for travel time density prediction
- We compare two types of probabilistic models and an ensemble model
- We select an optimal temporal aggregation level
- Probabilistic models prevail in predicting the density of the travel time

# Predicting the probability distribution of bus travel time to move towards reliable planning of public transport services

Léa Ricard<sup>a,\*</sup>, Guy Desaulniers<sup>b</sup>, Andrea Lodi<sup>b</sup> and Louis-Martin Rousseau<sup>b</sup>

<sup>a</sup>Department of Computer Science, Université de Montréal, Montréal, Canada

<sup>b</sup>Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal, Canada

---

## ARTICLE INFO

### Keywords:

Long-term prediction  
Travel time  
Probabilistic model  
Public transport  
Reliability

## ABSTRACT

An important aspect of the quality of a public transport service is its reliability, which is defined as the invariability of the service attributes. Preventive measures taken during planning can reduce risks of unreliability throughout operations. In order to tackle reliability during the service planning phase, a key piece of information is the long-term prediction of the density of the travel time, which conveys the uncertainty of travel times. We introduce a reliable approach to one of the problems of service planning in public transport, namely the Multiple Depot Vehicle Scheduling Problem (MDVSP), which takes as input a set of trips and the probability density function (p.d.f.) of the travel time of each trip in order to output delay-tolerant vehicle schedules. This work empirically compares probabilistic models for the prediction of the conditional p.d.f. of the travel time, as a first step towards reliable MDVSP solutions. Two types of probabilistic models, namely similarity-based density estimation models and a smoothed Logistic Regression for probabilistic classification model, are compared on a dataset of more than 41,000 trips and 50 bus routes of the city of Montréal. The result of a vast majority of probabilistic models outperforms that of a Random Forests model, which is not inherently probabilistic, thus highlighting the added value of modeling the conditional p.d.f. of the travel time with probabilistic models. A similarity-based density estimation model using a  $k$  Nearest Neighbors method and a Kernel Density Estimation predicted the best estimate of the true conditional p.d.f. on this dataset.

---

## 1. Introduction

In order to increase the ridership and attract new users, public transport agencies put increasing emphasis on improving the quality of the service they provide and particularly its regularity, also referred to as reliability (Ma, Ferreira and Mesbah, 2014). Studies show that a majority of passengers put more value on a reduction of the travel time (TT) variability than on a reduction of TT itself (Bates, Polak, Jones and Cook, 2001). However, day-to-day disruptions and delays are considered unavoidable on the day of operation (Amberg, Amberg and Kliewer, 2019; Kramkowski, Kliewer and Meier, 2009) due to the randomness of incidents and to the variation in demand and capacity factors. Bus sharing the road with other road-based vehicles (e.g., cars, bikes and trucks) are likely to have even higher degrees of variability, because they are subject to the same - morning and evening peaks - traffic patterns (Comi, Nuzzolo, Brinchi and Verghini, 2017).

TT variability is explained by yearly, monthly, day-to-day and hourly variability as well as vehicle-to-vehicle variability (B. Anil, Vanajakshi and Subramanian, 2014; Büchel and Corman, 2018; Kieu, Bhaskar and Chung, 2015). Daily variability is the fluctuation between the TT of a trip at a specific time on different days whereas hourly variation is the variation of a route's TT during the course of the day. A bus route is defined by an ordered sequence of road segments and bus stops, where the first and the last stops are called terminals. A bus line usually has two associated routes, each one going in opposite directions (e.g., North-South or East-West axis). A trip is defined as a unique travel which follows a specific route according to a predefined schedule. TTs are usually longer during peak periods because of denser traffic and/or passenger flows. Indeed, dwell time increases with the number of passengers, because the loading and unloading are more laborious (Büchel and Corman, 2018). Bus driver behavior and vehicle type (e.g., standard or articulated) are examples of underlying causes of vehicle-to-vehicle variability.

---

\*Corresponding author

Email addresses: lea.ricard@umontreal.ca (L. Ricard); guy.desaulniers@gerad.ca (G. Desaulniers); andrea.lodi@polymtl.ca (A. Lodi); louis-martin.rousseau@polymtl.ca (L. Rousseau)  
ORCID(s): 0000-0002-2061-1170 (L. Ricard)

Reliability can be addressed prior to the operational phase. At the strategic planning level, adding reserved lanes for buses can increase the reliability of the service, while during operations, bus holding is a popular solution to alleviate risks of bus bunching. The latter consists in holding a bus at key locations along the route if it is ahead of time. However, service reliability is rarely taken into account at the tactical planning level step, when the detailed planning of the service is computed (van Oort, 2011). Service planning in public transport is usually divided in five problems that can be addressed in a sequential or (partially) integrated order: network design, headway setting between two consecutive trips along a same route or timetabling setting, vehicle scheduling, duty scheduling and crew rostering. The second problem specifies for a set of trips their departure times from terminals and all subsequent stops. The vehicle scheduling problem assigns vehicles to cover the latter set of trips, in such a way that every trip is covered exactly once and at minimal costs. When the operator's fleet is spread in two or more depots, it is referred to as the Multiple Depot Vehicle Scheduling Problem (MDVSP). Duty scheduling consists in building anonymous duties, which are sequences of trip segments and breaks forming a driver workday and respecting various working rules, to cover all trips to operate on a given day. Finally, crew rostering assigns those duties to the drivers over a one-week or one-month horizon.

One of the key information for all five problems or one in particular (vehicle scheduling) is the TT prediction. There are two types of TT prediction: short-term and long-term. Both types can predict either a segment or a complete trip TT. The former is usually performed less than one hour before a trip and uses online information as well as external factors (e.g., weather). This type of prediction can be integrated to the operator's operations control system and provides online information to the users about the estimated arrival time of a bus. On the other hand, the long-term TT prediction can be performed a few days before the trip and helps for transit planning.

In this work, we explore methods for the long-term prediction of public bus TT in order to address reliability when solving the MDVSP. We define the reliable MDVSP with stochastic TTs, which takes as input a set of timetabled trips and their conditional TT probability density function (p.d.f.) in order to output cost-efficient and delay tolerant vehicle schedules. We frame the long-term prediction of the density of the TT (PDTT) as a supervised learning problem which aims at predicting, for each trip in a set of timetabled trips, the complete conditional p.d.f. of its TT where the condition corresponds to the trips' characteristics. The objective of this problem is to predict the complete distribution of the TT instead of the expected TT, in order to capture the inherent uncertainty of the duration of bus trips. Models that fit the conditional p.d.f. of the TT are compared in order to select the one that outputs the best estimate of the true conditional p.d.f. of the TT given a trip feature vector. Probabilistic models are compared to a Random Forest model, which provided the most promising results among the three regression models studied in the work of Moreira, Jorge, Sousa and Soares (2012). We introduce two types of probabilistic models, namely the similarity-based density estimation model and the smoothed logistic regression model for probabilistic classification, and present experimental results on a large-scale dataset of more than 41,000 trips and 50 bus routes. Our contribution is two-fold:

- The state-of-the-art for long-term prediction of public bus TT is almost nonexistent (Moreira-Matias, Mendes-Moreira, de Sousa and Gama, 2015). This work tries to fill this gap and, in addition, it is to our knowledge the first work to propose probabilistic models for the long-term prediction of public bus TT.
- To the best of our knowledge, it is the first study in the field of public transport that empirically studies such a large number of bus route's TTs simultaneously. We hope this improves the transferability of our results to other bus networks.

The remainder of this paper is organized as follows. Section 2 defines the reliable MDVSP with stochastic TTs and sets out its subproblem, the PDTT, which consists in predicting the long-term conditional p.d.f. of the TT. In Section 3, we review the literature on TT analysis. The dataset used and its features are discussed in Section 4. Section 5 describes the methodology that can be applied for the PDTT. In Section 6, data preparation as well as features and parameters selection are presented, before comparing the performance of all models for the prediction of the p.d.f. of the TT. Section 7 summarizes our findings.

## 2. The reliable MDVSP with stochastic TTs

The MDVSP outputs a set of vehicle schedules, which are sequences of trips and waiting times starting and ending at the same depot, such that each travel is either a deadhead trip (e.g., between a depot and a terminal or between two terminals) or a timetabled trip. A deadhead trip between two terminals enables the connection of two timetabled trips

ending and starting at different terminals. The reliability of a vehicle schedule is often assessed by its delay tolerance, which is defined as the aptitude to prevent delays to propagate to timetabled trips downstream. This type of delay is called secondary delay, whereas deviation from the planned duration of a timetabled trip caused by a disruption or variability during operation is called primary delay.

The reliable MDVSP with stochastic TTs takes the set  $\mathcal{V}$  of  $n$  timetabled trips and the long-term prediction of the TT of each of these trips in input. Let  $D$  be the set of depots,  $S$  the set of all feasible vehicle schedules, and  $S^d$  the subset of schedules starting and ending at depot  $d$ . The problem is to find a subset of vehicle schedules in  $S$  that covers exactly once each timetabled trip while respecting the number of available buses  $b_d$  at each depot  $d \in D$  and minimizing a weighted sum of the total planned vehicle operating cost and the total expected secondary delay. To formulate this problem, we define for each trip  $i \in \mathcal{V}$  and schedule  $s \in S$  a binary parameter  $a_{is}$  which is equal to 1 if schedule  $s$  covers trip  $i$  and 0 otherwise, and denote by  $c_s$  the cost of schedule  $s$  (including expected delay penalties). Furthermore, we introduce for each schedule  $s \in S$ , a binary variable  $y_s$  that takes value 1 if schedule  $s$  is selected in the solution and 0 otherwise. The reliable MDVSP with stochastic TTs can then be expressed as the following integer linear program:

$$\min \sum_{s \in S} c_s y_s \quad (1)$$

$$\text{s.t.} \quad \sum_{s \in S} a_{is} y_s = 1, \forall i \in \mathcal{V} \quad (2)$$

$$\sum_{s \in S^d} y_s \leq b_d, \forall d \in D \quad (3)$$

$$y_s \in \{0, 1\}, \forall s \in S. \quad (4)$$

Constraints (2) ensure that each timetabled trip is covered by a selected vehicle schedule, whereas constraints (3) impose vehicle availability at each depot.

The objective function (1) minimizes the total cost of the selected schedules which combines planned operational costs and expected secondary delay penalties. The planned costs usually include a fix cost per vehicle used and a variable cost that depends on the traveled distance and the waiting time attended by a bus driver. Consider a vehicle schedule  $s = \{v_1, v_2, \dots, v_{m_s}\} \in S$  of  $m_s$  timetabled trips. For notational conciseness, we denote trip  $v_i$  by  $i$  directly in the following. The total cost  $c_s$  of the vehicle schedule  $s$  is a weighed sum of the planned costs  $q_s$  and the sum of the expected secondary delay  $\mathbb{E}(R_i)$  of all trips  $i$  covered by the schedule  $s$ , weighted by a factor  $\beta$ :

$$c_s = q_s + (\beta \sum_{i=1}^{m_s} \mathbb{E}(R_i)). \quad (5)$$

The secondary delay  $R_i$  of a trip  $i$  is the difference between its actual departure time  $D_i$  and the planned departure time  $d_i$ , computed as

$$R_i = D_i - d_i. \quad (6)$$

The random variable  $D_i$  is a convolution of the previous trip's actual departure time ( $D_{i-1}$ ), actual TT ( $T_{i-1}$ ) and the minimum in-between time between trips  $i$  and  $i-1$  ( $l_{i-1,i}$ ):

$$D_i = \max\{D_{i-1} + T_{i-1} + l_{i-1,i}, d_i\} \quad (7)$$

$$D_1 = d_1. \quad (8)$$

Precisely,  $l_{i-1,i}$  accounts for the deadhead travel between terminals, if the trip  $i-1$  ends at a different terminal than the departure terminal of trip  $i$ , and the minimum break time for drivers.

The expected value of  $R_i$  is computed by integrating over the p.d.f. of the secondary delay of trip  $i$ . Indeed, the expected value of a continuous random variable  $X$  with a p.d.f.  $p(\cdot)$  is  $\mathbb{E}(X) = \int xp(x)dx$ . We have that the p.d.f. of

$R_i$  is recursively computed using the p.d.f. of  $T_1$  to  $T_{i-1}$ . Indeed, the expected values of  $T_1$  to  $T_{i-1}$  are not sufficient to compute the expected value of  $R_i$ . Thus, instead of being interested in the long-term prediction of the expected TT of a set of trips, we will rather focus on the long-term prediction of the p.d.f. of the TT of a set of trips. We call the latter approach the PDTT.

We want to learn a predictive model for the long-term p.d.f. of the TT in order to minimize the total expected secondary delays when solving the reliable MDVSP with stochastic TTs. The PDTT computes for each trip  $i$  in a set of unseen timetabled trips (test set), an estimate of its conditional p.d.f.,  $\hat{p}(T_i | \mathbf{x}_i)$ , given  $\mathbf{x}_i$  the set of characteristics of trip  $i$ . It is assumed that the random variables  $\{T_1, \dots, T_n\}$  are independent and identically distributed. The estimator  $\hat{p}(\cdot)$  is learned from observed trips in a training set, such that  $\hat{p}(\cdot)$  is fitted in order to minimize the negative log-likelihood (NLL) of the training set (containing  $n_{train}$  points), computed as

$$NLL_{train} = - \sum_{i=1}^{n_{train}} \log(\hat{p}(t_i | \mathbf{x}_i)), \quad (9)$$

with  $t_i$  the true TT of trip  $i$ . If possible, the performance of the PDTT should be evaluated by the test mean square error (MSE) of the expected secondary delay

$$MSE = \frac{1}{n_{test}} \sum_{i=1}^{n_{test}} (r_i - \mathbb{E}(\hat{R}_i))^2, \quad (10)$$

with  $n_{test}$  is the number of test points,  $\mathbb{E}(\hat{R}_i)$  the estimated value of the expected secondary delay computed using  $\hat{p}(T_i | \mathbf{x}_i)$  and  $r_i$  the true secondary delay of trip  $i$ . Indeed, the purpose of the PDTT is to improve the reliability of the MDVSP solutions and this starts by accurate expected secondary delays. We will fit PDTT models on a sample of the observed timetabled trips, namely the ones of the routes with the largest number of recorded trips (see Section 4). Since this dataset was collected from real buses circulating in the city of Montréal, some actual vehicle schedules contain timetabled trips from routes not in the sample studied. For such vehicle schedules, it is not possible to compute  $D_i$  and  $R_i$  for most of the trips  $i$  in the schedule. It is thus impossible to evaluate the PDTT models in terms of the expected secondary delay error of all trips in the test set in this work. For that reason, the performance of the PDTT will be evaluated in terms of the NLL score as well, but of the test set. Generally, if a PDTT model achieves a better test NLL, it also achieves a lower expected secondary delay error. The NLL of the test set is analogue to the one of the training set and is computed as

$$NLL_{test} = - \sum_{i=1}^{n_{test}} \log(\hat{p}(t_i | \mathbf{x}_i)). \quad (11)$$

To summarize, a PDTT model outputs for every trip  $i$  in a set of trips an estimate of the conditional p.d.f. of its TT,  $\hat{p}(T_i | \mathbf{x}_i)$ . This differs from *point prediction* models (Dutordoir, Salimbeni, Deisenroth and Hensman, 2018), as here we are interested in predicting the complete probability distribution instead of only the expected value of the TT. This will enable us to compute and minimize the expected secondary delays of a set of timetabled trips when solving the reliable MDVSP with stochastic TTs.

### 3. Related works and contributions

The introduction of automatic vehicle location (AVL) data has given rise to a flourishing number of studies in the field of public transport on speed, arrival time and TT analysis. Because TT and arrival time measures are closely related, studies on both measures are treated without distinction. Indeed, the arrival time  $A_i$  of a trip  $i$  is given by

$$A_i = D_i + T_i. \quad (12)$$

In this section, three topics are covered: long-term TT prediction, TT variability analysis and TT distribution modeling. The latter fits the TT distribution of trips that occurred during a given period in order to analyze the shape

and nature of the p.d.f. of the TT, without trying to predict future events, as in the PDTT. We extend the field of the first topic to all road-based transport, but the subsequent topics are restricted to the public transport field. Approaches proposed for the PDTT are inspired by lessons learned through the review of the literature on these topics.

Compared with the literature on short-term TT prediction, studies on long-term TT prediction are rare and to the best of our knowledge, only the works of Chen, Liang and Chu (2020), Moreira et al. (2012) and Klunder, Baas and Op de Beek (2007) proposed or reviewed long-term TT prediction methods. In a survey on improving the planning of public transit using AVL data, Moreira-Matias et al. (2015) suggested that the long-term TT prediction should be valid for an horizon of at least the entire forecasting period. They divided models found in the literature for short-term TT prediction in four categories: (i) machine learning and regression, (ii) state-based and time-series, (iii) traffic theory-based and (iv) historical databased, and suggested that some regression algorithms applied for the short-term TT prediction could be adapted to the long-term TT prediction. Gradient boosting is an example of a model that was successfully applied to the short-term TT prediction and then adapted by Chen et al. (2020) to the long-term TT prediction of trips on a freeway segment in Taiwan. The features were ranked in order of relative importance: time of the day, day of the week, national holiday, day of long / consecutive holiday, big event / activity, electronic tool collection fee promotion and narrowing of roadway. Three regression models, namely a Projection Pursuit Regression, a Support Vector Machine and a Random Forest, were compared by Moreira et al. (2012) for the long-term TT prediction of trips of one public bus route in Porto. With a basic pre-processing work, the Random Forest had better results, but was slightly outperformed by a Projection Pursuit Regression when the authors added an instance selection step. Klunder et al. (2007) trained a  $k$  Nearest Neighbours algorithm ( $k$ NN) with only time-based variables for the long-term TT prediction on a motorway network in the Netherlands.

The reasons for TT variability can be external or internal (Yetiskul and Senbil, 2012) and related to demand or capacity (Mazloumi, Currie and Rose, 2010) (see Figure 1). In an early study, Abkowitz and Engelstein (1983) suggested that shorter routes may have reduced TT variability. Also, they reported that a running time deviation at the beginning of a route tends to propagate downstream. Hence, control actions to correct early deviations on a route could reduce TT variability. Strathman and Hopper (1993) reported that the afternoon peak period has higher TT variability, in particular because of the higher passenger demand. In a study on TT variability in the city of Ankara, Yetiskul and Senbil (2012) found major differences in regional TT variability and suggested that bus-stop spacing should depend on the neighborhood density. Comi et al. (2017) performed a time series decomposition of the TT and compared it to the temporal traffic patterns. The two are reported to have similarities. Also, the seasonality of the time series decomposition was most significant for the hour of the day.

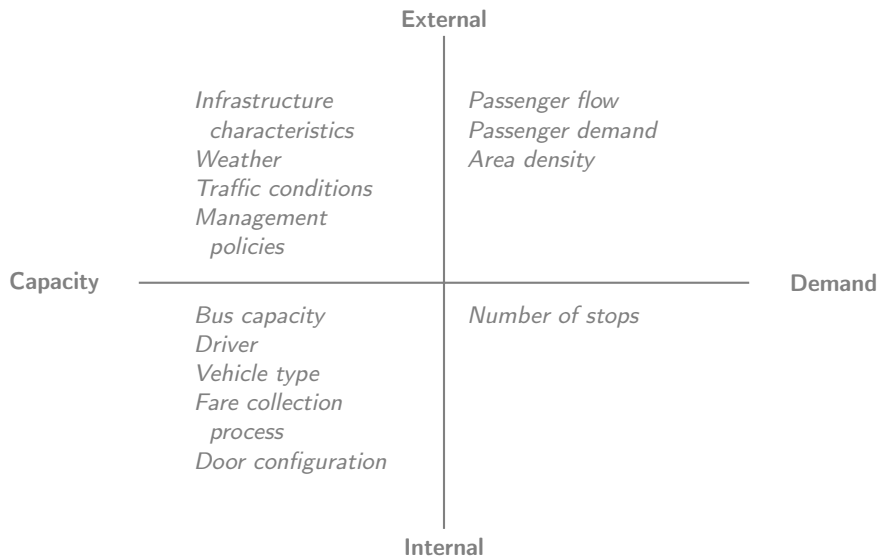


Figure 1: Reasons of TT variability

TT distribution modeling has been studied mostly with the objective of quantifying the reliability of a transit

service. Most works on TT distribution modeling in public transit occurred after the introduction of AVL systems. We focus our review on the work of Mazloumi et al. (2010), Ma, Ferreira, Mesbah and Zhu (2016) and Büchel and Corman (2018), as they are, in our view, the most comprehensive studies, from which useful lessons can be learned for the PDTT. Mazloumi et al. (2010) assessed the shape and nature of the TT distribution over the course of the day and for different levels of temporal aggregation on segments of a bus route. The authors measured the level of temporal aggregation by the length of the departure time windows (DTW), which are time slots for which trips departing during the slot are aggregated for subsequent analyzes (e.g., 15 minutes, 30 minutes or 1 hour). The study concluded that for shorter DTWs, the TT distribution follows a Normal distribution. For longer DTWs, this result holds for peak periods, but not for off-peak periods. For the latter, the Log-Normal distribution fits better. Also, the contribution of a set of features to the TT was assessed through a linear regression analysis. The land use (industrial vs. residential) and the length of the segment were those affecting the most the TT variance. Ma et al. (2016) studied intensively the influence of temporal and spatial aggregation on the TT distribution, with the objective of providing common grounds for modeling and performance evaluation metrics. To this end, several settings of temporal and spatial aggregations were assessed and an evaluation approach based on a statistical hypothesis test was proposed. The TT distribution was broken down in a running time distribution and a dwelling time distribution. The TT of a trip starting at stop  $i$  and ending at stop  $j$  was decomposed in dwelling times  $DT_k$  at each stop  $k$  of the trip and running times  $RT_{(k,k+1)}$  between each pair of stops

$$T_{(i,j)} = \sum_{k=i}^{j-1} DT_k + RT_{(k,k+1)}. \quad (13)$$

Results concerning the normality of the TT distribution were in line with the ones of Mazloumi et al. (2010). Also, the analysis suggested that spatial aggregation tends to decrease the multimodality of TT distribution. A multimodal distribution is defined as a probability distribution with several modes. A Gaussian mixture model (GMM) was proposed to address the multimodality of the link level TT distribution. Büchel and Corman (2018) found that the Log-Normal distribution was, out of four unimodal statistical distributions, the best fit for the TT distribution modeling.


## 4. Data

The dataset used for this study was collected during a 2-month period from 08/28/2017 to 10/29/2017 by in-car Advanced Public Transport Systems (APTS) installed in buses running in the city of Montréal, Canada. Those systems collect automatically at every stop of a trip the corresponding trip identifier, route identifier, direction identifier, stop identifier, date, scheduled departure time, scheduled arrival time, actual departure time, actual arrival time and number of passengers loading or unloading, among other things. The scheduled departure, scheduled arrival, actual departure and actual arrival times are stored in milliseconds. The actual TT of a trip is the difference between the actual arrival time and the actual departure time at the terminals. Hence, for every route, only the first and last stops (i.e., terminals) data is kept. Since the APTS were embedded in approximately 20% to 30% of the vehicles at that time, weekends and holidays had an insufficient number of trips recorded. Indeed, during weekends and holidays, the service is reduced and thus the number of trips recorded during those days is too small to conduct relevant data-driven analysis. For that reason, weekends and holidays are not studied and are removed from the dataset. After removing weekends and holidays, the dataset has more than 116,000 trips. Of the 408 routes in the dataset, only the 50 most frequent are kept for the remainder of the study, resulting in a dataset of over 41,000 trips. Figure 2 shows that the 50 selected routes run between 4:00AM to 1:59AM (+1 day) during weekdays. To facilitate the notation, we add two extra hours to the usual 24-hour daily period. Thus, we say that the selected routes run between 4:00AM to 25:59PM.

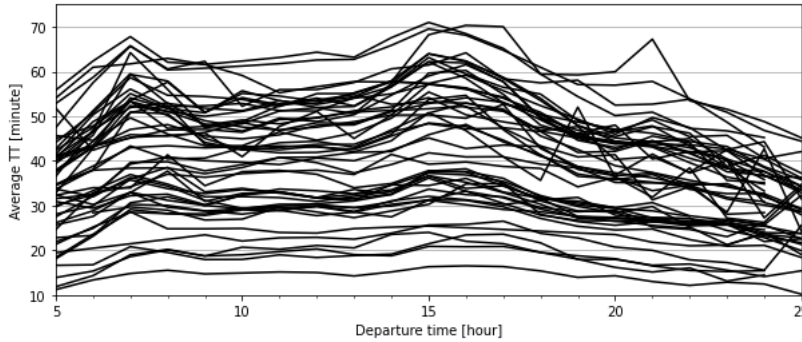
### 4.1. Route's characteristics

The distribution of the average TT per hour is presented in Figure 3, where each broken line represents the evolution of the average TT of a route. It is possible to distinguish the morning peak approximately from 6:00AM to 8:59AM, and the afternoon peak, approximately from 14:00PM to 17:59PM for all routes. As for the average TT, the number of recorded trips is also higher during peak periods (see Figure 2). The afternoon peak usually has an average TT higher than the morning peak and the average TT as well as the number of recorded trips per hour are generally decreasing after 17:00PM.

### Predicting the probability distribution of bus travel time



**Figure 2:** Total number of recorded trips per hour in the dataset



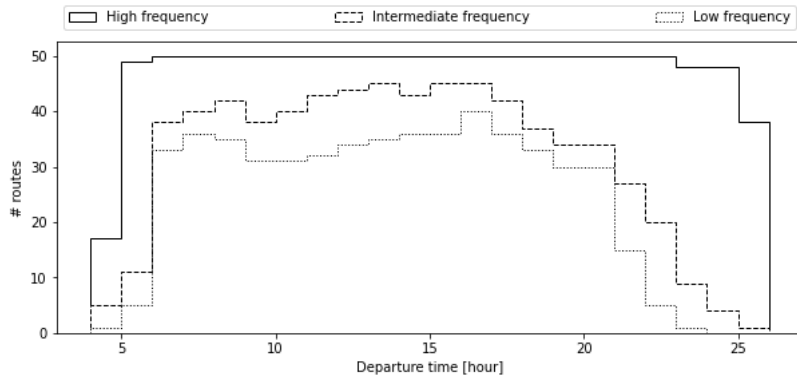
**Figure 3:** Average TT per hour

The main characteristics of the 50 selected routes, namely the number of stops, the distance traveled and the type of operational region, are presented in Table 1. Each route has a unique combination of line identifier and direction, such that *A-East* and *A-West* are two different routes of the same bus line, but in opposite directions. We categorized the type of operational region in 6 categories: residential areas, crossing city center (CC), from city center (to a residential area), to city center (from a residential area), from an industrial (indust.) area (to a residential area), to an industrial area (from a residential area). Bus line B is the only one crossing the city center; it starts in a residential neighborhood, crosses the city center and ends in another residential neighborhood. Industrial areas are characterized by a high density of factories. The city center and industrial areas are usually regions where a large number of people commute to work every day. The majority of bus routes operate in residential areas (32 out of 50). Those bus routes may, for example, connect two residential areas or a residential area to a subway or train station. The number of stops per bus route ranges from 17 to 74 stops, while the distance traveled ranges from 3.0 km to 15.3 km. In general, as the number of stops goes up, the distance traveled goes up as well. Line P is a counterexample, because it has a large number of stops close to each other. Note that the number of stops and the distance traveled of two routes of the same line are generally not equal, as the path in one direction is usually not symmetric to the path in the other direction (e.g., because some streets are one-way).

The scheduled departure time of a trip is fixed when solving the timetabling or headway setting problem. If the trip is headway-based, the duration between the departure time of two consecutive trips is constant for a given time segment. Otherwise, if the trip is timetabled, the scheduled departure time is established in order to fulfill a given minimum level of service. In general, a stable frequency helps passenger transparency regarding the passing time of buses. In the case of the 50 selected routes, a mix of headway-based and timetabled trips were recorded. Figure 4 presents the average route frequency breakdown. The frequency is divided in 3 categories: high frequency (6 trips per hour and more), intermediate frequency (between 3 trips and 5 trips per hour) and low frequency (less than 3 trips per hour). For every route, the average frequency per hour is identified and classified in one category. Then, we count the number of routes per category and per hour. In Figure 4, we can see that a majority of the selected routes have a high frequency during daytime (6:00AM to 20:59PM). In contrast, the majority of the routes have a lower frequency during the early morning (4:00AM to 5:59AM) and during the end of the day (21:00PM to 25:59PM). The lines O and P are the only ones with a low frequency all day long.

**Table 1**  
Characteristics of bus routes studied

Line	Dir.	#stops	Dist. (km)	Region	Line	Dir.	#stops	Dist. (km)	Type of Region
A	East	52	13.9	Residential	M	West	18	4.0	Residential
A	West	49	14.5	Residential	N	East	28	5.3	Residential
B	East	46	13.2	Cross CC	N	West	29	5.3	Residential
B	West	46	12.0	Cross CC	O	North	35	3.0	To indust.
C	North	40	12.1	Residential	O	South	40	3.4	From indust.
C	South	45	12.3	Residential	P	East	74	9.0	Residential
D	North	33	10.3	From indust.	P	West	67	8.5	Residential
D	South	36	10.4	To indust.	Q	East	40	7.0	Residential
E	East	50	14.2	Residential	Q	West	38	7.0	Residential
E	West	52	13.3	Residential	R	East	37	5.3	Residential
F	East	34	7.8	Residential	R	West	35	5.3	Residential
F	West	36	7.7	Residential	S	East	47	11.8	From indust.
G	North	17	4.6	Residential	S	West	51	11.6	To indust.
G	South	19	4.3	Residential	T	North	34	8.5	To indust.
H	North	37	9.3	Residential	T	South	30	8.5	From indust.
H	South	40	10.8	Residential	U	North	46	11.1	Residential
I	East	71	15.3	Residential	U	South	42	10.7	Residential
I	West	68	15.3	Residential	V	East	46	9.5	Residential
J	North	28	7.1	From CC.	V	West	49	8.5	Residential
J	South	30	7.1	To CC	W	East	43	10.3	Residential
K	East	53	11.1	To CC	W	West	47	11.6	Residential
K	West	51	11.4	From CC	X	North	30	6.6	From CC
L	East	35	5.9	Residential	X	South	34	7.5	To CC
L	West	36	6.0	Residential	Y	North	30	8.0	To CC
M	East	18	4.4	Residential	Y	South	28	8.0	From CC



**Figure 4:** Average route frequency, with frequency defined by high frequency: 6 trips per hour and more; intermediate frequency: between 3 trips and 5 trips per hour; low frequency: less than 3 trips per hour

#### 4.2. Features analysis


The PDTT has to be based upon features (i.e., explanatory variables) that are available a few days or weeks in advance. For example, meteorological conditions are likely to influence the TT duration. However, since it is an information that is not available when solving service planning problems, it is not considered. Likewise, the TT of the previous trip is not considered. The list of possible features includes the day of the week, type of region, route identifier, distance, number of stops, scheduled departure time, week number and year. Possible values and types of

**Table 2**  
Long-term features

Feature	Type	Possible values
Day of the week	Categorical	{Monday, Tuesday, ..., Friday}
Region	Categorical	{residential, crossing CC, ..., to indust.}
Route identifier	Categorical	{A East, A West, ..., Y South}
Distance (km)	Non-categorical	[3, 15.3]
Number of stops	Non-categorical	{17, 18, ..., 74}
Scheduled departure time	Non-categorical	[4:00AM, 25:59PM]
Week number	Non-categorical	{35, 36, ..., 44}
Year	Non-categorical	{2017}

features (categorical or non-categorical) are listed in Table 2.

The feature year is discarded because our dataset is spread over 2017 only. The statistical significance of the features scheduled departure time, day of the week and week number can be analyzed visually by looking at Figures 3, 5 and 6, which present the average TT per route depending on each of the feature respectively. Figure 3 suggests that the scheduled departure time has a high importance. The relationship between the TT and the scheduled departure time is not linear. Generally, the average TT of a bus route increases during peak hours and is steady between the morning and the afternoon peaks. Second, Figure 5 suggests that the relationship between the TT and the day of the week is less important. This is hardly surprising given that Saturdays and Sundays are not considered. Interestingly, there is no common pattern between the routes; for example some routes have a slightly higher average TT on Tuesdays than on Mondays and Wednesdays, while some others have an inverse pattern (i.e., the average TT on Tuesdays is slightly lower than on Mondays and Wednesdays). Third, Figure 6 suggests that the relationship between the TT and the week number seems significant only for a handful of bus routes. In Figures 3, 5 and 6, we can observe that the average TT differs greatly from one bus route to another, unsurprisingly as each route has its own characteristics (as discussed earlier).



**Figure 5:** Average TT per day of the week

Furthermore, the statistical significance of all the features can be assessed by the permutation feature importance (Breiman, 2001), which measures the increase of the MSE of a Random Forest estimator when the values of a feature are permuted. The importance of a feature  $\ell$  is the difference between the MSE of a Random Forest estimator on the original dataset and the average (over 10 shuffles) MSE of the same estimator on a corrupted dataset (with the values of feature  $\ell$  permuted). Results of this analysis are presented in Table 3. It indicates that the two most important features are the number of stops and the distance traveled, which are two main characteristics of a bus route. When the number of stops and the distance traveled is known, the route identifier is less statistically significant. It also confirms that the scheduled departure time is an important feature to predict the TT. The week number is also significant, but impacts the prediction error to a smaller extent. The day of the week and the type of operational region have a limited impact.

## Predicting the probability distribution of bus travel time




Figure 6: Average TT per week number

Table 3

Permutation feature importance

Feature	Importance	Relative importance (%)
Number of stops	185.76	51.45
Distance	62.37	17.28
Scheduled departure time	56.95	15.77
Route identifier	33.73	9.34
Week number	8.51	2.36
Day of the week	7.74	2.14
Region	5.96	1.65

## 5. Methodology

Several approaches have been proposed to compute the conditional p.d.f. in a probabilistic fashion. Gaussian process based models are among the most popular. The work of Dutordoir et al. (2018) proposed a Gaussian process-based model for estimating the conditional p.d.f. using latent variables in order to model non-Gaussian probability distributions. Also, Bishop (1994) developed Mixture Density Networks, which is a type of artificial neural network predicting multimodal conditional density distributions. The main drawback of Mixture Density Networks is that they perform poorly when the size of the dataset is not large enough. In the work of Yeo, Melnyk, Nguyen and Lee (2018), the continuous p.d.f. prediction is converted into a classification task by using a discretization technique. This simplifies the learning task and traditional probabilistic classifiers can be used to predict the probability mass function, which can be smoothed later on into a p.d.f.. The focus of this paper is on frequentist models (Koller and Friedman, 2009). In the remainder of this section two approaches for the PDTT are presented. The first one estimates the p.d.f. of the TT of a set of similar trips using parametric, semi-parametric or non-parametric density estimation models. The second approach, namely the smoothed Logistic Regression for probabilistic classification, is similar to that of Yeo et al. (2018), but fits a Logistic Regression instead of a Recurrent Neural Network estimator.

### 5.1. Similarity-based density estimation

Similarity-based density estimation models are a two-step process: for each trip, (1) find the set of similar trips and (2) estimate the density of this particular set, by fitting a parametric, semi-parametric or non-parametric model. Next, we will define two similarity-based methods and introduce some density estimation models.

#### 5.1.1. Similarity-based methods

Consider a trip  $i$  with a feature vector  $(\mathbf{x}_1^{(i)}, \dots, \mathbf{x}_d^{(i)})$ . We want to select, based upon one of the following similarity-based methods, the set of trips in the (reduced) training set that have similar attributes:

- Equivalent DTW (eDTW) : select all trips from the same route that have a scheduled departure time in the same

departure time window (DTW) (Mazloumi et al., 2010; Büchel and Corman, 2018; Ma et al., 2016) than trip  $i$ .

- $k$  Nearest Neighbors ( $k$ NN): select the  $k$  nearest neighbors of the trip  $i$ . The distance between trip  $i$  and trip  $j$  is the Euclidean distance between their feature vectors and is computed as

$$\text{dist}(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{\sum_{\ell=1}^d (x_{\ell}^{(i)} - x_{\ell}^{(j)})^2}. \quad (14)$$

### 5.1.2. Density estimation models

The conditional probability  $p(T_i = t_i | \mathbf{x}_i)$  is estimated by fitting a given density estimation model on points close to  $i$  in the (reduced) training set, which are either trips in the same eDTW or close neighbors.

#### Parametric models

Parametric density estimation considers a restricted set of common probability distributions. Each of these distributions has a small number of parameters that have to be estimated from the data. We consider the Normal, Log-Normal, Logistic, Log-Logistic, Gamma, and Cauchy probability distributions. The Gamma distribution is a family of probability distributions containing the Exponential, Erlang and Chi-Squared distributions. In the work of Ma et al. (2016), the first four distributions were successful at modeling the TT distribution at a route level. For each trip, parameters of these statistical distributions are found using the Maximum Likelihood Estimation (MLE) algorithm.

#### Semi-parametric model: Gaussian Mixture Model

GMMs are a sub-category of mixture models composed of  $K$  normal components. GMMs are relevant when the population modeled is multimodal and has undefined subpopulations or states, such that each component represents a state. It is common in transport to use three components, one for each of the traffic states: free flow, recurrent and non-recurrent traffic (Ma et al., 2016). The p.d.f. of a  $K$ -components GMM is given by

$$\hat{p}(T_i = t_i | \mathbf{x}_i) = \sum_{k=1}^K \pi_{ik} \mathcal{N}(t_i | \mu_{ik}, \sigma_{ik}^2). \quad (15)$$

The vector of positively defined coefficients  $\pi_i = (\pi_{i1}, \dots, \pi_{iK})$ , such that  $\sum_{k=1}^K \pi_{ik} = 1$ , and the vector of the model's  $k^{\text{th}}$  component's parameters  $(\mu_{ik}, \sigma_{ik}^2)$ , are found by applying the expectation maximization (EM) algorithm.

#### Non-parametric model: Kernel Density Estimation (KDE)

A KDE model infers the p.d.f. of a random variable based on a sample of its population. To estimate the p.d.f. of a trip  $i$ , the model uses  $m$  points close to  $i$  in the training set. It is a data smoothing problem that allows to find, in a non-parametric fashion, the curve of the p.d.f. given a sample. The Gaussian kernel,  $K(\cdot)$ , is the most widely used, but any function that integrates to unity ( $\int K(t)dt = 1$ ) can replace it. The smoothness of the estimator is adjusted by the bandwidth parameter  $h$  as

$$\hat{p}(T_i = t_i | \mathbf{x}_i) = \frac{1}{m} \sum_{j=1}^m K\left(\frac{t_i - t_j}{h}\right). \quad (16)$$

## 5.2. Smoothed Logistic Regression for probabilistic classification (LR-PC)

Probabilistic classifiers are a type of machine learning model that can predict the probability that a given input belongs to a set of classes, instead of only predicting the class with the highest probability. When a numerical discretization is applied to the random variable  $T$ , such that the TT is categorized in bins of 1 minute, the PDTT task can be translated into a probabilistic classification one: estimate the probability that  $T$  takes a value that falls into class  $c \in \{0, \dots, C-1\}$ . A model is fitted per bus route, because this setting yields better experimental results than learning a unique model for all bus routes (see Section 6.2).

A question arises: how to choose the number of classes for a bus route. Our approach was to use, for a given bus route, the difference between the trip in the training set with the shortest duration,  $t_{min}$ , and the trip with the longest

duration,  $t_{max}$ , as  $C$ , the number of classes. Thus,  $P(T_i = c)$  is the probability that  $T_i$  takes a value in  $[c + t_{min}, c + 1 + t_{min}]$ . We disregarded the fact that trips in the test set can have shorter or longer duration than  $t_{min}$  and  $t_{max}$  respectively, as the smoothing discussed later implicitly solves this issue.

Multinomial Logistic Regression is naturally probabilistic and is commonly used for probabilistic classification tasks. Classes' probabilities of a multinomial Logistic Regression are defined as

$$\hat{P}([T_i] = c + t_{min} \mid \mathbf{x}_i, \mathbf{w}_c) = \frac{\exp(\mathbf{w}_c^T \mathbf{x}_i)}{\sum_{c'=1}^C \exp(\mathbf{w}_{c'}^T \mathbf{x}_i)}, \quad (17)$$

where  $\mathbf{w}_c$  is the vector of parameters of the class  $c$ , found using a stochastic average gradient descent solver.

Logistic Regression outputs a probability mass function (p.m.f.) that can be smoothed into a p.d.f. subsequently. As proposed by Yeo et al. (2018), the output of the multinomial Logistic Regression, which takes the form of a p.m.f., is passed through a one-dimensional convolution layer. This step has the effect of enforcing a spatial correlation in the output. The convolution layer is analogous to a KDE with a bandwidth  $h$  and a kernel  $K(\cdot)$ , but uses the p.m.f,  $\hat{P}([T_i] \mid \mathbf{x}_i, \mathbf{w}_c)$ , instead of a sample of trips:

$$\hat{p}(T_i = t_i \mid \mathbf{x}_i) = \sum_{c'=0}^{C-1} \left[ K \left( \frac{t_i - (c' + t_{min})}{h} \right) \times \hat{P}([T_i] = c' + t_{min} \mid \mathbf{x}_i, \mathbf{w}_{c'}) \right]. \quad (18)$$

## 6. Experimental results


Using the dataset described in Section 4, we next explain how we fit probabilistic models for the PDTT, before comparing the performance of these models. First, we describe how the data is filtered and split for the PDTT. Second, we go through features and parameters selection for each type of model and detail the selection of the temporal aggregation level. Finally, all probabilistic models are compared to a Random Forest model in terms of their performance on the test set.

### 6.1. Data preparation

The dataset is filtered in order to remove erroneous information and special situations that we do not want to cover in the PDTT. Incomplete trips, vias and trips with detours are discarded. A via is a trip that deviates from the main trip or an express trip. To remove erroneous trips which were not filtered by the previous step, the Median Absolute Deviation or MAD (Hellerstein, 2008) with a 6-delta criterion is used. A trip is discarded if it has a TT longer or shorter than the median TT value of the trips associated with the same route plus (minus) 6 times the corresponding standard deviation. This method also removes trips with extended TT due to exceptional scenarios (e.g., bus failure) that the PDTT problem should not cover, because when such exceptional scenarios occur, an additional bus is usually dispatched to recover the schedule and prevent severe delay propagation.

Figure 7 summarizes the dataset split. We split the dataset such that the training data starts on 08/28/2017 and ends on 10/08/2017. The set of test data is composed of the trips from 10/16/2017 to 10/29/2017. Hence, a complete week, from 10/09/2017 to 10/16/2017 is discarded to simulate real-life settings where the planning is done at least a few days ahead. The training set is split again in a validation set and a reduced training set by slicing the last trips recorded per route from the original training set. The reduced training set contains 80% of the trips in the original training set and the validation set contains the remaining 20%. The validation set is used for features and parameters selection, which are detailed in the next section.

Figure 7: Dataset split



**Table 4**

NLL (lower is better) of the similarity-based density estimation models using eDTW method with different levels of temporal aggregation

Model	5 periods		60 minutes		30 minutes	
	Train	Validation	Train	Validation	Train	Validation
Cauchy	3.24	3.30	2.72	2.89	2.60	2.96
Gamma	3.08	<b>3.15</b>	2.58	2.83	2.47	5.98
Normal	3.09	3.16	2.59	2.87	2.48	6.57
Log-Normal	3.09	<b>3.15</b>	2.57	2.81	2.46	5.72
Logistic	3.09	3.16	2.59	2.77	2.48	2.91
Log-Logistic	3.09	<b>3.15</b>	2.58	<b>2.75</b>	2.47	<b>2.88</b>
GMM	3.07	3.16	2.26	5.81	1.82	24.96
KDE	2.99	<b>3.15</b>	2.49	2.76	2.77	2.89

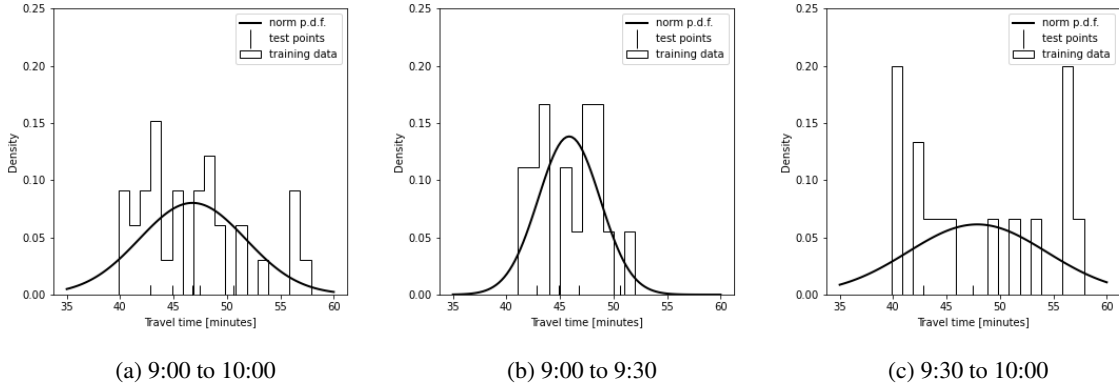
## 6.2. Features and parameters selection

Similarity-based density estimation models use one of the similarity-based methods, namely the eDTW or the  $k$ NN method. While the eDTW method does not require feature selection, as the features used are always the route identifier and the scheduled departure time, the  $k$ NN method does require feature selection. Indeed, the distance between neighbors depends on the specified feature vector. Estimating the TT density of a trip  $i$  does not require feature selection; it fits a probability density to a sample containing trips similar to trip  $i$ . Thus, here the features selection problem is reduced to finding a feature vector for the  $k$ NN method. Features considered are the ones listed in Table 3. The selection of features is carried out in parallel with the selection of parameters; we are looking for the combination of features and parameters that yields the best NLL of the validation set. The parameters of the similarity-based methods are the DTW duration for the eDTW method and the number  $k$  of neighbors for the  $k$ NN method. For the KDE, the validation set is used to select the bandwidth  $h$  and the kernel function.

We found that similarity-based density estimation models fit the data better when a  $k$ NN is fitted per bus route. By doing so, the features describing the bus route characteristics, namely the number of stops, distance traveled, route identifier and type of region, become uninformative to the model. Indeed, all trips used to train the model of a given bus route have exactly the same values for these features. The remaining features to consider are the scheduled departure time, the week number and the day of the week. For the  $k$ NN method, a feature vector composed of the two former features gives the best results. Thus, the closest neighbors of a trip  $i$  are defined as trips of the same bus route with a scheduled departure and a week number in a temporal horizon close to trip  $i$ . Temporal aggregation is a fundamental aspect of the PDTT since it has been shown to affect the shape and nature of the TT probability distribution (Mazloumi et al., 2010; Ma et al., 2016). Thus, the parameter associated to it, namely the DTW duration, is studied carefully.


We select the DTW duration by analyzing how the density estimation models perform on the validation set for different levels of temporal aggregation. DTWs considered are, going from the most aggregated to the least aggregated, 5 periods per day (before morning peak, morning peak, in-between morning and afternoon peaks, afternoon peak and after afternoon peak), 60 minutes and 30 minutes. Table 4 compares the performance of models using the eDTW method, with the values in bold indicating the best NLL of the validation set for each level of temporal aggregation. For all models except the GMM, the NLL score of the validation set is better at DTWs of 60 minutes. The Log-Logistic model obtains the best results at all aggregation levels, matched by the Gamma, Log-Normal and KDE models at an aggregation level of 5 periods per day. The GMM has a similar performance than the parametric models for the most aggregated level, but it also has a poor performance for lower levels of temporal aggregation, both 60 and 30 minutes. Since the performance on the reduced training set is good, it indicates that the GMM overfits the training data. For the DTWs considered, the conditional p.d.f. of the TT is most likely not multimodal. Thus, this model is discarded for the rest of the study. Interestingly, the training NLL decreases when the temporal aggregation level increases for all models except the KDE, while the NLL of the validation set increases from DTWs of 60 minutes to DTWs of 30 minutes for all models. This suggests that models are overfitting more at DTWs of 30 minutes than at DTWs of 60 minutes. Indeed, when the level of temporal aggregation decreases, the DTW duration and thus the number of training points per DTW decreases, which can result in a predicted conditional p.d.f. that does not fit the validation set points.

Figure 8 illustrates an overfitting case, where Figure 8a is the predicted p.d.f. of the TT of a trip when the DTW used for the trip selection is 60 minutes and Figures 8b and 8c are the corresponding p.d.f.s of the TT for DTWs of 30 minutes. Indeed, when the set of points between 9:00AM and 10:00AM is divided in two DTWs of 30 minutes, the second portion (Figure 8c) is flat with two peaks at the ends. The model estimation of the p.d.f. of the TT is poor. This is probably due to the small number of data collected between 9:30AM and 10:00AM, which results in an overrepresentation of the tails of the probability distribution.



**Figure 8:** Example of the fit of a Normal distribution with (a) 60 minutes and (b), (c) 30 minutes DTWs, route H-South.

We can conclude that, between the three levels of temporal aggregation compared, the best one is the one with DTWs of 60 minutes. We denote the eDTW method with DTWs of 60 minutes as eDTW\*. For the second similarity-based method, namely the  $k$ NN, the value of  $k$  can be chosen similarly to the duration of DTWs, by assessing the performance of the similarity-based density estimation models on the validation set for different values of  $k$ . Figure 9 shows that the NLL of the validation set decreases significantly for all models when the value of  $k$  increases up to approximately  $k = 13$ . After that, the NLL stays almost constant. Thus, we set the number  $k$  of neighbors to 13 and denote the  $k$ NN method with  $k = 13$  as  $k$ NN\*.



**Figure 9:** NLL (lower is better) of the similarity-based density estimation models using  $k$ NN method with different numbers of neighbors

The LR-PC model also yields better performance when it is fitted per bus route and when it has a feature vector containing only the scheduled departure time and some transformations thereof. The scheduled departure time is transformed in order to capture a non-linear relationship with the TT. First, it is categorized in bins of 1 hour and 30 minutes using a one-hot encoding. Second, the sine and cosine of the scheduled departure time are computed. The total dimension of the feature vector is 69 (22 for the one-hot encoding of bins of 1 hour, 44 for the one-hot encoding of bins of 30 minutes, 2 for the sine and cosine of the scheduled departure time and 1 for the original scheduled departure time). The bandwidth and the kernel function of the LR-PC model are selected based on the performance on the validation set, along with the regularization strength of the Logistic Regression.

**Table 5**

NLL (lower is better) of the training and test sets of similarity-based density estimation, LR-PC and Random Forest models

Model	similarity method	Training set	Test set
Cauchy	eDTW*	2.73	2.87
	kNN*	2.65	2.84
Gamma	eDTW*	2.59	2.79
	kNN*	2.51	2.80
Normal	eDTW*	2.61	2.82
	kNN*	2.52	2.84
Log-Normal	eDTW*	2.59	2.77
	kNN*	2.51	2.78
Logistic	eDTW*	2.61	2.74
	kNN*	2.53	2.72
Log-Logistic	eDTW*	2.59	2.72
	kNN*	2.52	2.70
KDE	eDTW*	2.61	2.73
	kNN*	2.41	<b>2.69</b>
LR-PC	-	2.48	2.70
Random Forest	-	<b>2.45</b>	<b>2.80</b>

\*DTWs = 60 minutes or  $k = 13$

The NLL of the validation set is minimized for the Random Forest model with a feature vector composed of the number of stops, distance traveled, scheduled departure time, route identifier, week number and day of the week. Thus, only the feature encoding the type of region is discarded. This is in line with the results of the permutation feature importance (see Table 3). The number of trees in the Random Forest model, the maximum number of features considered when branching, the maximum depth of each tree and the minimum number of samples required to split an internal node are selected based on the performance on the validation set.

### 6.3. Models comparison

Table 5 presents the NLL of the training set (including the validation set) and of the test set for all models. For the parametric models, the Cauchy, Logistic and Log-Logistic distributions have better results on the test set when using the  $kNN^*$  method, while the Gamma, Normal and Log-Normal have rather the opposite results. The NLL of the training set is smaller for all models when the  $kNN^*$  method is used. This indicates that the  $kNN^*$  method fits better the training points than the eDTW\* method, but may also indicate a mild overfitting (especially for Gamma, Normal and Log-Normal distributions which have a better NLL of the test set when applied with the eDTW\* method). The KDE model performs better when it uses the  $kNN^*$  method than when it uses the eDTW\* method. Overall, the KDE model using  $kNN^*$  method is the one which yields the best result on the test set, closely followed by the Log-Logistic with  $kNN^*$  and the LR-PC models.

Similarity-based density estimation models and the LR-PC model are compared with a Random Forest model. As shown in Table 5, the Random Forest is outperformed by all the models except the Gamma distribution using  $kNN^*$  method as well as the Normal and the Cauchy distributions using eDTW\* or  $kNN^*$  method, as it has a larger test NLL. A Random Forest is not inherently probabilistic; it is usually used for *point prediction* (Dutordoir et al., 2018). Indeed, the emphasis is on modeling the mapping between an input  $\mathbf{x}$  to its output  $y$  rather than on predicting the conditional p.d.f.  $p(y | \mathbf{x})$  (Dutordoir et al., 2018). When optimized with a sum-of-squares function, a Random Forest model can be interpreted as fitting a Gaussian on the conditional distribution of the TT, where the mean of the Gaussian is given by the output of the Random Forest and the variance is constant for all the inputs and equal to the MSE of the output on the training set. Thus, these experimental results show that there is an added value of modeling the conditional p.d.f. of the TT using probabilistic models. A key factor could be that in probabilistic models the variance is not constant for every trip. It is interesting to see that models using a non-Gaussian probability distribution, either a Gamma, Log-Normal, Logistic or Log-Logistic distribution outperform the one using a Normal distribution, which

**Table 6**  
MSE of the TT on the test set

Model	similarity method	MSE
Log-Logistic	$k$ NN*	16.23
KDE	$k$ NN*	16.18
LR-PC	-	16.94
Random Forest	-	15.71

\* $k = 13$

call the normality of the conditional p.d.f. of the TT into question.

The PDTT objective is to predict the complete conditional p.d.f. of the TT given a trip's feature vector. However, we can assess the performance of the models developed for the PDTT for a different type of prediction: *point prediction* (Dutordoir et al., 2018), which aims at predicting the expected value of the TT of a trip. A Random Forest optimized with a sum-of-squares function is fitted in order to minimize the MSE of the difference between the prediction (i.e., the expected value of the TT) and the true TT duration. We compare the *point prediction* performance of the Random Forest model to the models yielding the best NLL score, namely the Log-Logistic model using  $k$ NN\* method, the KDE model using  $k$ NN\* method and the LR-PC model. Table 6 shows that these all have an MSE score of the same order of magnitude as the Random Forest model. Furthermore, after the Random forest model, the KDE model using  $k$ NN\* method has the lowest MSE. Thus, the KDE model using  $k$ NN\* predicts a complete conditional p.d.f. that corresponds better to the true region in which a trip duration is likely to lie than the Random Forest model and in addition it is able to compute *point predictions* almost as accurately as the Random Forest model (less than 3% increase in MSE).

## 7. Conclusions

In public transport, reliability has become a key challenge for operators wishing to attract new users. We proposed to address the reliability of bus services at the planning stage, and more precisely when computing one of the problem of service planning, namely the vehicle scheduling problem. To that end, we defined the reliable MDVSP with stochastic TTs, a version of the vehicle scheduling problem taking in input a set of timetabled trips and the conditional p.d.f. of their TT in order to output reliable vehicle schedules (i.e., vehicle schedules with small expected secondary delays for all timetabled trips). The prediction of the p.d.f. of the TT, that we referred to as the PDTT, is done a few days ahead of operations in order to incorporate the prediction in the planning phase. We investigated if probabilistic models can predict more accurately the complete conditional p.d.f. of the TT than a Random Forest model. The latter is not inherently probabilistic, as it is typically used for *point prediction*, but under some conditions it can be interpreted as fitting a Gaussian on the conditional distribution of the TT given the characteristics of a set of trips.

We used a 2-month dataset collected by buses equipped with APTS in the city of Montréal, involving 50 bus routes and a total of over 41,000 trips. The bus routes studied have various attributes (e.g., number of stops, frequency, traveled distance, etc.) and constitute a diverse sample from which we hope to obtain results transferable to other bus networks. Based on previous works on TT variability analysis, we determined a set of features that we ranked in order of decreasing significance as follows: number of stops, distance, scheduled departure time, route identifier, week number, day of the week, type of region.

In order to compute the conditional p.d.f of the TT, we proposed two types of probabilistic models, namely similarity-based density estimation models and the LR-PC. The former is a two-step process that firstly find, for each trip, the set of similar trips and then estimate the density of this set using parametric, semi-parametric or non-parametric density estimation models. We proposed two types of similarity-based methods, namely the eDTW and the  $k$ NN, for which the temporal aggregation level and the number of neighbors had to be set, respectively. The LR-PC model applies a numerical discretization to the TT before fitting a Logistic Regression classifier per bus route. The output of the Logistic Regression is then smoothed into a p.d.f. using a convolution layer analogous to a KDE.

Previous works on TT distribution modeling indicated that the level of temporal aggregation greatly affects the shape and nature of the TT distribution. Thus, we carefully selected the DTW duration based on the performance on

the validation set. The GMM model had a poor performance for DTWs of 60 minutes and 30 minutes and thus we concluded that the conditional p.d.f. of the TT is most likely not multimodal. This result is aligned with the one of Ma et al. (2016) which observed that the multimodality of the TT decreases with spatial aggregation. Finally, models were compared in terms of their test NLL. From all the models tested, the density-based estimation model using *k*NN method and a KDE yielded the best NLL score. The Random Forest model was outperformed by most of probabilistic models and had comparable *point prediction* accuracy (measured by the MSE of the TTs) to models with the best NLL scores. This result indicated that there is an added value in modeling the conditional p.d.f. of the TT using probabilistic models. In particular, probabilistic models allow the variance of the p.d.f. of the TT to be different for every trip, as opposed to the Random Forest model as we interpreted it. Also, as similarity-based density estimation models using the Normal distribution were outperformed by several other parametric distributions, the normality of TTs was questioned, supporting the benefit of using probabilistic models for the PDTT.

*Declaration of interest:* None for all authors.

## Acknowledgements

The authors would like to thank Charles Fleurent and his team at GIRO Inc. for their valuable help and contributions and the Société de transport de Montréal (STM) for sharing the dataset. We gratefully acknowledge the financial support provided by GIRO Inc. and the Natural Sciences and Engineering Research Council of Canada under the grant CRDPJ 520349-17 as well as the financial support provided by the Institute for Data Valorization (IVADO).

## References

- Abkowitz, M.D., Engelstein, I., 1983. Factors affecting running time on transit routes. *Transportation Research Part A: General* 17A, 107 – 113. doi:[https://doi.org/10.1016/0191-2607\(83\)90064-X](https://doi.org/10.1016/0191-2607(83)90064-X).
- Amberg, B., Amberg, B., Kliewer, N., 2019. Robust efficiency in urban public transportation: Minimizing delay propagation in cost-efficient bus and driver schedules. *Transportation Science* 53, 89–112. doi:<https://doi.org/10.1287/trsc.2017.0757>.
- B. Anil, K., Vanajakshi, L., Subramanian, S.C., 2014. Pattern-based bus travel time prediction under heterogeneous traffic conditions, in: *Colloquium on Transportation Systems Engineering and Management*. doi:<https://doi.org/10.13140/RG.2.1.2338.5448>.
- Bates, J., Polak, J., Jones, P., Cook, A., 2001. The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review* 37, 191 – 229. doi:[https://doi.org/10.1016/S1366-5545\(00\)00011-9](https://doi.org/10.1016/S1366-5545(00)00011-9).
- Bishop, C.M., 1994. Mixture Density Networks. Technical Report NCRG/94/004. Department of Computer Science and Applied Mathematics, Aston University, UK.
- Breiman, L., 2001. Random forests. *Machine Learning* 45, 5–32. doi:10.1023/A:1010933404324.
- Büchel, B., Corman, F., 2018. Modelling probability distributions of public transport travel time components, in: *18th Swiss Transport Research Conference*. doi:<https://doi.org/10.3929/ETHZ-B-000263929>.
- Chen, C.M., Liang, C.C., Chu, C.P., 2020. Long-term travel time prediction using gradient boosting. *Journal of Intelligent Transportation Systems* 24, 109–124. doi:<https://doi.org/10.1080/15472450.2018.1542304>.
- Comi, A., Nuzzolo, A., Brinchi, S., Verghini, R., 2017. Bus travel time variability: some experimental evidences. *Transportation Research Procedia* 27, 101 – 108. doi:<https://doi.org/10.1016/j.trpro.2017.12.072>.
- Dutordoir, V., Salimbeni, H., Deisenroth, M.P., Hensman, J., 2018. Gaussian process conditional density estimation, in: *32nd Conference on Neural Information Processing Systems*.
- Hellerstein, J.M., 2008. Quantitative Data Cleaning for Large Databases. Technical Report. UNECE.
- Kieu, L.M., Bhaskar, A., Chung, E., 2015. Public transport travel-time variability definitions and monitoring. *Journal of Transportation Engineering* 141. doi:[https://doi.org/10.1061/\(ASCE\)TE.1943-5436.0000724](https://doi.org/10.1061/(ASCE)TE.1943-5436.0000724).
- Klunder, G., Baas, P., Op de Beek, F., 2007. A long-term travel time prediction algorithm using historical data. Technical Report. TNO.
- Koller, D., Friedman, N., 2009. *Probabilistic Graphical Models: Principles and Techniques*. The MIT press.
- Kramkowski, S., Kliewer, N., Meier, C., 2009. Heuristic methods for increasing delay-tolerance of vehicle schedules in public bus transport, in: *VIII Metaheuristics International Conference*.
- Ma, Z., Ferreira, L., Mesbah, M., 2014. Measuring service reliability using automatic vehicle location data. *Mathematical Problems in Engineering* 2014, 1–12. doi:<https://doi.org/10.1155/2014/468563>.
- Ma, Z., Ferreira, L., Mesbah, M., Zhu, S., 2016. Modeling distributions of travel time variability for bus operations. *Journal of Advanced Transportation* 50, 6–24. doi:<https://doi.org/10.1002/atr.1314>.
- Mazloumi, E., Currie, G., Rose, G., 2010. Using gps data to gain insight into public transport travel time variability. *Journal of Transportation Engineering-asce - J TRANSP ENG-ASCE* 136, 623–631. doi:[https://doi.org/10.1061/\(ASCE\)TE.1943-5436.0000126](https://doi.org/10.1061/(ASCE)TE.1943-5436.0000126).
- Moreira, J., Jorge, A., Sousa, J., Soares, C., 2012. Comparing state-of-the-art regression methods for long term travel time prediction. *Intelligent Data Analysis* 16, 427 – 449. doi:<https://doi.org/10.3233/IDA-2012-0532>.
- Moreira-Matias, L., Mendes-Moreira, J., de Sousa, J.F., Gama, J., 2015. Improving mass transit operations by using avl-based systems: A survey. *IEEE Transactions on Intelligent Transportation Systems* 16, 1636–1653. doi:<https://doi.org/10.1109/TITS.2014.2376772>.

## Predicting the probability distribution of bus travel time

- van Oort, N., 2011. Service Reliability and Urban Public Transport Design. Ph.D. thesis.
- Strathman, J.G., Hopper, J.R., 1993. Empirical analysis of bus transit on-time performance. *Transportation Research Part A: Policy and Practice* 27, 93 – 100. doi:[https://doi.org/10.1016/0965-8564\(93\)90065-S](https://doi.org/10.1016/0965-8564(93)90065-S).
- Yeo, K., Melnyk, I., Nguyen, N., Lee, E.K., 2018. De-rnn: Forecasting the probability density function of nonlinear time series, in: 2018 IEEE International Conference on Data Mining (ICDM), pp. 697–706. doi:<https://doi.org/10.1109/ICDM.2018.00085>.
- Yetiskul, E., Senbil, M., 2012. Public bus transit travel-time variability in ankara (turkey). *Transport Policy* 23, 50 – 59. doi:<https://doi.org/10.1016/j.tranpol.2012.05.008>.