

Dette resumé er publiceret i det elektroniske tidsskrift

Artikler fra Trafikdage på Aalborg Universitet

(Proceedings from the Annual Transport Conference
at Aalborg University)

ISSN 1603-9696

www.trafikdage.dk/artikelarkiv



Deep Survival Modelling for Shared Mobility

Bojan Kostic, boko@dtu.dk, Transport Division, DTU Management, DTU

Mathilde Pryds Loft, maplo@adm.dtu.dk, Transport Division, DTU Management, DTU

Filipe Rodrigues, rodr@dtu.dk, Transport Division, DTU Management, DTU

Stanislav S. Borysov, stabo@dtu.dk, Transport Division, DTU Management, DTU

Francisco C. Pereira, Transport Division, DTU Management, DTU

Abstrakt

With an increased focus on minimizing traffic externalities in metropolitan areas, a growing interest in environmentally friendly and shared mobility systems has emerged, such as electric car-sharing systems. However, increasing demand and larger area coverage often make it difficult to keep cars available where and when customers need them. This problem can be alleviated by predicting for how long cars stay vacant at given pick-up/drop-off locations. To maximize their usage, it can be more beneficial to relocate the cars at certain periods to more desired locations. In this paper, we tackle the problem of predicting time-to-pickup for shared cars in a probabilistic way as a function of time by applying time-to-event modelling through survival analysis. Both statistical and deep neural network approaches to survival regression were investigated. The Cox proportional hazards model (CPH) is compared to the deep neural network model DeepSurv. To predict survival times, a two-step approach was formulated, where in the upper level a classification is used to classify cars into two groups based on idle time duration, whereas in the lower level for each given group time-to-event modelling is applied. DeepSurv method demonstrated a stronger fit compared to CPH. The two-step approach resulted in over 15% improvement in performance, comparing to the one-step approach, where no classification is used.

Introduction

Shared mobility and car-sharing services have become increasingly popular over the last 10 years. Shared mobility is a rapidly increasing form of transportation worldwide, in order to reduce traffic congestion and provide more climate-friendly modes of transportation. However, with increasing demand, companies providing these types of transportation are facing challenges, such as keeping the vehicles accessible to customers at their desired locations. This issue is amplified when the system operates in a free-flowing regime, i.e. when customers are free to drop off cars at any place in a limited area where is most convenient for them. One solution to tackle this issue is to relocate cars manually to more desired locations, especially in the busy areas. Being able to predict which cars will be vacant the longest, will decrease vacancy time, thus increasing revenue and user satisfaction.

The main goal of this paper is to explore methods for predicting shared cars' vacancy time. The methodology used is general and can be applied to any type of shared mobility system, such as bike-

sharing. The core concept of the analysis is to model the idle time as *survival* time. The time interval a car is vacant will be modeled as how long it *survives*. Therefore, one objective of this research is to investigate how well survival analysis can be applied in this context, and to provide comparison between several survival models.

Background

There are various types of shared mobility concepts, such as car-sharing (station-based and free-floating), bike-sharing or on-demand services. For a broad overview of shared mobility, we recommend Machado et al. 2018 [1]. These shared mobility systems have attracted attention of the research community, having increasingly more research being done in the recent years. An increasing presents of companies entering the market, due to promising business opportunities, mainly drives this trend. For example, the market for car-sharing has been examined, in terms how likely is that individuals will join the car-sharing system in Switzerland, to identify new market opportunities, done by Juschten et al. 2019 [2]. Another reason is the need for sustainable transport system and cities in general. Different stakeholders may have different objectives: operators want to maximize their profit, authorities want to provide regulation, transport researchers to evaluate its impact and contribute in directing future developments. The shared mobility systems have been analysed from both a simulation and data-driven perspective, as more data are being available from private fleet operator companies.

Methods

As mentioned before, we want to model idle time of cars using survival analysis. Survival analysis [3] provides time-to-event (TTE) analysis with censored observations. TTE is expressed through two events: the *birth* event (which is our case is the drop-off) and the *death* event (which is the subsequent pick-up), hence the term “survival”. Time between the two events is survival time, denoted by T , which is treated as a non-negative random variable. Given the *birth* event, it represents the probability distribution of the *death* event happening in time. In addition, observations of true survival times may not always be available, thus creating distinction between uncensored (true survival times available) and censored observations (true survival times unknown, but only the censorship time). This information is expressed through an event indicator E , with binary outcome. The censorship time is assumed to be non-informative, and T and E are independent. The survival function $S(t)$ can be defined from the cumulative distribution function $F(t)$, as $S(t) = 1 - F(t)$, which in terms of probabilities can be formulated as $S(t) = P(T > t)$. The survival function can be interpreted in terms of the gradient (slope) of a line, with steeper line segments meaning an event is more likely to happen in the corresponding interval.

Most commonly used statistical methods for survival analysis is the Kaplan-Meier method [4], which is univariate counting-based method for general insights. To estimate the effect of features, the Cox proportional hazards model (CPH) introduced by Cox 1972 [5] is used, which is a survival regression model. It provides interpretability but lacks accuracy with larger datasets. To improve on accuracy recently deep neural network have been used, thus creating deep survival models. One of them is DeepSurv, by Katzman et al. 2018 [6], which extends CPH to account for non-linear relationships among features.

The concordance index (C-index) will be the main measure of quality and performance. The C-index can be seen as an equivalent to the R^2 score in regression. The C-index measures the correctness of the ordering of predicted survival times against the truly observed times. A value of 1 represents a perfect score, random guessing corresponds to 0.5. In general, one can expect to achieve a value between 0.6 and 0.7 for an average fit, above 0.7 is generally considered a strong fit.

Data

Data used for this research were kindly provided by DriveNow Copenhagen, a car-sharing company. The data comes from a fleet of electric cars operating as a free-floating car-sharing system. The dataset consists of several hundred thousand observations. Each observation is one realized trip. A trip is described by a set of features, such as vehicle ID and pick-up and drop-off data (location, time and battery level). The dataset

includes hundreds of cars in the period of around a year, collected in 2017 and 2018. The allowed area for car locations covers most of the city, and is divided into zones of various granularity. The zone can be inferred from the location coordinates. The zone ID will be given attention later in the training process, as it will turn out to be a decent predictor variable, both individually as well as jointly with other features.

As an objective with shared mobility systems is to maximize fleet usage, vehicles can be relocated from their current drop-off locations to more attractive locations if their estimated vacant time in the current location is high. This can be easily detected from the data, as pick-up location of one trip is different from the drop-off location of the previous trip the car had. Since from the data it is unknown at what time the car has been relocated, it is assumed that the car is picked up immediately after being moved. These observations will be treated as censored, as it is unknown how long the cars survived. We know that the car was parked in a specific zone and was vacant until the point it was relocated. It is, however, unknown for how long it would have stayed parked in the zone, had it not been relocated.

Exploratory analysis and general insights

We investigate here the significance individual features have on survival times. Their average effect is typically estimated using the Kaplan-Meier (KM) model, which is used to compute average survival functions for each features values, which then allows for easy comparison. The survival times used in these analyses include values of less than or equal to 50 hours, which comprise more than 99% of the data, hence all the survival functions converge almost to zero.

Fig. 1 shows the results of investigating the effect the battery level and distance to the city centre have on the survival function. We can see that if the battery is charged more than 50% when the car was dropped off, it has a higher probability of being picked up comparing to the cars with lower battery levels, as the blue line decreases faster than the orange line. Similar effect can be observed in the other case, where cars closer to the city centre (≤ 5 km) will be picked up faster than the other cars.




Figure 1. The effect of the battery level and distance to the city centre on survival times

Results

Here we test and compare the regression models, CPH and DeepSurv, in predicting survival times. The results of the survival time predictions are shown in Fig. 2. Two approaches were investigated: the one-step approach applies survival analysis on the whole dataset, whereas the two-step approach first conducts the classification of survival times based on the certain threshold and then applies survival analysis on the two classified groups of data samples. The one-step approach results start with “1” in the legend and are depicted using dashed lines, and the two-step approach starts with “2” and is depicted using solid lines. For both of them, three variants are investigated: “Cox baseline” (in red) is the CPH model using only Zone ID variable as a single feature for prediction, “Cox” (in blue) is the CPH model using all the features and “DeepSurv” (in black) is a deep neural network model using all the features. The two-step approach contains classification error, therefore we also evaluate the survival predictions on the true split (i.e.

perfect classification), where we manually split survival times of the cars on the specific threshold and then run survival analysis on those splits (these are denoted with “TRUE FIT”). In this way, we can see the effect of the classification error on the final prediction accuracy. We tested various split thresholds (treated in hours), shown on the x-axis, whereas on the y-axis is the C-index, which is our accuracy measure. We can see that for shorter thresholds we can obtain much higher accuracy with the two-step model, and as the threshold increases, i.e. more data belong to one class, the two-step approach converges to the one-step approach. For the split at 1 hour, i.e. where we classify the idle times of cars to be less than or greater than 1 hour, and then run DeepSurv, we can get the most accurate prediction of the survival times of the cars.




Figure 2. The comparison of the one-step approach and two-step approach performance

Conclusion

We successfully applied survival analysis, and especially deep survival analysis, in the context of shared mobility systems. We proved that these methods could be very informative and provide predictions that can be used for optimizing shared vehicle fleets. In addition, applying the two-step approach improved the performance of all the three different models. DeepSurv was overall a stronger predictor than CPH. In general, the two-step approach is more demanding to fit, as it requires more hyper-parameter tuning. Next steps would be to investigate joint prediction for neighboring vehicles, to account for the effect that when one vehicle is picked up, it then affects the probability of its neighboring vehicles being picked up.

References

- [1] Machado, C. A. S, Hue, N. P. M. S., Berssaneti, F. T., Quintanilha, J. A. An overview of shared mobility. *Sustainability*, 10(4342), 2018.
- [2] Juschten, M., Ohnmacht, T., Thao, V. T., Gerike, R., Hössinger, R. Carsharing in Switzerland: identifying new markets by predicting membership based on data on supply and demand. *Transportation*, 46(4):1171–1194, 2019.
- [3] Harrell, F. *Regression Modeling Strategies*. 2nd ed. Springer Series in Statistics, Switzerland: Springer International Publishing, 2015.
- [4] Kaplan, E. L. and Meier, P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481, 1958.
- [5] Cox, D. R. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220, 1972.
- [6] Katzman, J. L., Shaham, U., Cloninger, A., Bates, J., Jiang, T., Kluger, Y. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Medical Research Methodology*, 18(24), 2018.