

Development of Machine Learning Models for Predicting Trip Cancellation in Ride-Hailing Platforms.

Annunciata Nduku Kyalo*, Dr. Mvurya Mgala, Dr. Kennedy Hadullo

Institute of Computing and Informatics, Technical University of Mombasa, Mombasa, Kenya

Corresponding Author: Annunciata Nduku Kyalo

DOI: <https://doi.org/10.51584/IJRIAS.2026.11080062>

Received: 19 August 2026; Accepted: 24 August 2026; Published: 05 September 2026

ABSTRACT

Trip cancellations remain a significant operational challenge in ride-hailing platforms, affecting service reliability, driver utilization, customer experience, and platform efficiency. The objective of this study was to develop multiple machine learning models for predicting trip cancellation outcomes in ride-hailing platforms. The study adopted the Obtain, Scrub, Explore, Model, and iNterpret (OSEMN) framework and used 103,024 historical Ola and Uber booking records. Data preparation involved cleaning, handling missing values, feature engineering, encoding, and feature selection. Nine predictor variables were retained for model development: Vehicle Turn Around Time, Customer Turn Around Time, Customer Cancellation Reason, Driver Cancellation Reason, Ride Status, Payment Method, Ride Distance, Driver Ratings, and Customer Rating. The dataset was partitioned using an 80:20 stratified split, resulting in 82,419 training and 20,605 testing records. Class weighting was applied during model training to reduce bias towards the majority class. Four supervised machine learning models; Logistic Regression, Decision Tree, Gradient Boosting, and Extreme Gradient Boosting (XGBoost), were developed and evaluated using accuracy, precision, recall, F1-score, and confusion matrix. The models achieved training accuracies of 85.20%, 87.72%, 90.03%, and 92.52%, respectively. On the testing dataset, the corresponding accuracies were 84.00%, 86.02%, 89.06%, and 91.62%. XGBoost achieved the strongest testing performance, with 91.62% accuracy, 87.12% precision, 91.41% recall, and 89.12% F1-score. The relatively small differences between training and testing performance indicate that the developed models generalized well to unseen data. The study concludes that supervised machine learning can effectively predict trip cancellation outcomes, with XGBoost demonstrating the strongest predictive capability among the four developed models.

Keywords: Machine Learning, Trip Cancellation, Ride-Hailing, XGBoost, Predictive Model.

INTRODUCTION

Ride-hailing platforms have transformed urban transportation by providing technology-enabled services that connect riders with available drivers. The platform coordinates ride requests, driver-rider matching, trip execution, and vehicle repositioning within a continuous service process (Wen et al., 2024). Despite the convenience offered by these platforms, trip cancellations remain a significant operational challenge. Cancellations disrupt the matching process between drivers and customers, reduce service reliability, and negatively affect customer satisfaction, driver earnings, platform efficiency and revenue (Liu et al., 2024; Sun et al., 2023; Wang et al., 2019).

Trip cancellations can occur at different stages of the booking process and may be initiated by either customers or drivers. Previous research has associated cancellation with factors such as travel time, distance between vehicles and passengers, road congestion, availability of alternative transport, and spatial and temporal mismatches between supply and demand (Li et al., 2012; Nuzzolo et al., 2018; Sun et al., 2023; Zhang & Fan, 2020). Frequent cancellations may increase passenger waiting time, reduce platform trust, and create inefficiencies in driver utilization. Predicting cancellation can therefore provide useful information for proactive interventions such as improved driver-passenger matching, vehicle repositioning, rerouting, and incentive adjustments (Liu et al., 2024; Ngo & Mishra, 2023; Xu et al., 2023).

Machine learning provides an opportunity to analyze historical ride-hailing data and identify patterns associated with trip cancellation. Existing studies have applied machine learning, deep learning, reinforcement learning, and optimization approaches to ride-hailing prediction and operational problems. Sun et al. (2022) and Sun et al. (2023) applied deep learning approaches to estimate cancellation probability, while Gupta et al. (2024) investigated predictive modelling for Uber ride cancellation and price estimation. Other studies have addressed cancellations indirectly through vehicle repositioning, driver-rider matching, demand forecasting, and reinforcement learning approaches (Ngo & Mishra, 2023; Wen et al., 2024; Xu et al., 2023).

Despite these developments, existing research has largely examined individual global platforms such as Uber, Lyft, and Didi, while limited attention has been given to predictive modelling across multiple ride-hailing platforms. Previous studies have also predominantly focused on passenger-initiated cancellations, with limited consideration of driver-initiated cancellations. Furthermore, there is a need to comparatively develop multiple machine learning classification algorithms using a common modelling process to determine their applicability to trip cancellation prediction (Gupta et al., 2024; Sun et al., 2022).

Accordingly, this study focused on developing multiple machine learning models for predicting trip cancellation using historical booking data from Ola and Uber. Four supervised machine learning algorithms: Logistic Regression, Decision Tree, Gradient Boosting, and Extreme Gradient Boosting (XGBoost), were developed to classify ride-hailing bookings as either completed or cancelled. The study therefore addresses the objective of developing multiple machine learning models for predicting trip cancellation while providing a multi-platform perspective that incorporates both customer-initiated and driver-initiated cancellations.

METHODOLOGY

Research Design

The study adopted a quantitative experimental research design followed the Obtain, Scrub, Explore, Model, and iNterpret (OSEMN) framework. OSEMN provides a structured approach for carrying out data science and machine learning studies through five stages: obtaining the data, preparing it, exploring relevant patterns, developing models, and interpreting the results (Kumari et al., 2020). In this study, the framework guided the preparation of ride-hailing data and the development of machine learning models for predicting trip cancellation. Data preprocessing, analysis, feature selection, model development, and performance evaluation were carried out using Python in Google Colab.

Obtain

The study used the publicly available Ola and Uber Ride Booking and Cancellation dataset obtained from Kaggle and provided by Mengar (2025). The dataset is available under the Creative Commons CC0 1.0 Universal Public Domain Dedication license. It contained 103,024 historical ride-hailing booking records and 20 original variables covering booking information, vehicle and customer information, trip characteristics, cancellation information, payment details, ratings, and booking status. Each record represented an individual ride booking together with the associated trip information. Booking_Status was used as the target variable for trip cancellation prediction.

Table 1: Description of Dataset Variables.

Variable	Description
Date	Date on which the booking was made.
Time	Time of the booking.
Booking_ID	Unique booking identifier.
Booking_Status	Indicates whether the booking was completed or cancelled.
Customer_ID	Unique customer identifier.
Vehicle_Type	Type of vehicle requested.
Pickup_Location	Customer pickup location.
Drop_Location	Destination location.

Vehicle Turn Around Time	Driver arrival time to pick up location.
Customer Turn Around Time	Time taken by the customer to arrive for the trip.
Customer Cancellation Reason	Reason provided by the customer for cancelling a ride.
Driver Cancellation Reason	Reason provided by the driver for cancelling a ride.
Ride Status	Indicates whether previous rides were incomplete.
Ride Status Reason	Reason for incomplete rides.
Booking Value	Total booking fare.
Payment Method	Payment option used for the trip.
Ride Distance	Distance of the trip (km).
Driver Ratings	Driver rating score.
Customer Rating	Customer rating score.
Vehicle Images	Images associated with the vehicle type.

Scrub

During the Scrub stage, the dataset was cleaned and prepared for analysis and model development. The data were first examined for missing values, duplicate records, inconsistencies, and inappropriate data types. Missing values were handled according to the characteristics of the affected variables. Booking_ID and Customer_ID were removed because they served as unique identifiers rather than useful predictors, while Vehicle Images was removed as a non-predictive variable.

Feature engineering was then carried out to simplify selected variables and make them more meaningful for modelling. Temporal variables were grouped into broader time categories, vehicle types and locations were consolidated into relevant groups, and numerical variables such as turnaround time, booking value, ride distance, and ratings were transformed into meaningful categories. Payment methods were also grouped into cash and digital payment categories. These transformations reduced the complexity of the original data while retaining information relevant to trip cancellation prediction.

Feature encoding was done to convert categorical variables into numerical values to make them suitable for the machine learning algorithms. Booking_Status was converted into a binary target variable, where completed trips were represented by 0 and unsuccessful trips by 1. Unsuccessful trips included customer cancellations, driver cancellations, and cases where a driver was not found.

The prepared dataset contained 63,967 completed trips (62.09%) and 39,057 cancelled trips (37.91%). The original class distribution was retained, while class weighting was applied during model training to reduce bias towards the majority class. As He and Garcia, (He & Garcia, 2009) described in their article, *“The fundamental issue with the imbalanced learning problem is the ability of imbalanced data to significantly compromise the performance of most standard learning algorithms.”*




Figure 1: Distribution of Booking Status on the Dataset.

The dataset was then divided into training and testing sets using an 80:20 stratified split. Stratification maintained the proportion of completed and cancelled trips in both subsets. This resulted in 82,419 records for training and 20,605 records for testing. The training data were used to develop the models, while the testing data were reserved for assessing their performance on unseen records.

Table 2: Summary of Dataset Partitioning.

Dataset	Number of records	Percentage (%)
Training dataset	82,419	80
Testing dataset	20,605	20
Total	103,024	100

Explore

During the Explore stage, the distribution of the variables and their relationships with trip cancellation were examined. Feature selection was carried out to identify the variables that provided the most useful information for model development while reducing variables with limited predictive contribution.

Three complementary techniques were used: correlation analysis, Mutual Information, and Random Forest feature importance. Correlation analysis examined the strength of the relationship between individual predictors and Booking_Status. Mutual Information was used to identify both linear and nonlinear dependencies, while Random Forest feature importance ranked the predictors according to their contribution to prediction (Breiman, 2001). The rankings from the three methods were combined, and Recursive Feature Elimination with Cross-Validation (RFECV) was used as an additional check of the retained variables.

Nine predictor variables were retained for model development: Vehicle Turn Around Time, Customer Turn Around Time, Customer Cancellation Reason, Driver Cancellation Reason, Ride Status, Payment Method, Ride Distance, Driver Ratings, and Customer Rating.




Figure 2: Comparison of Feature Relevance Across Feature Selection Methods.




Figure 3: Comparison of Feature Selection Methods with RFECV Selection.

Model

Four supervised machine learning algorithms were developed: Logistic Regression, Decision Tree, Gradient Boosting, and Extreme Gradient Boosting (XGBoost). The same selected predictor variables and training data were used to develop the four models so that their performance could be compared under a common modelling process.

Logistic Regression was used as the baseline classification model because of its suitability for binary classification (Hosmer Jr et al., 2013). Decision Tree was used to capture decision rules and nonlinear relationships within the data. Gradient Boosting builds a sequence of weak learners, with each learner attempting to improve errors made by the preceding learners (Friedman, 2001). XGBoost extends the gradient boosting approach through regularization and computational optimization, making it suitable for complex classification problems (Chen & Guestrin, 2016).

Feature scaling using StandardScaler was applied only to Logistic Regression because of its sensitivity to differences in variable magnitudes. The scaler was fitted using the training data and then applied to the testing data to avoid information leakage. Decision Tree, Gradient Boosting, and XGBoost were trained using the original encoded values.

iNterpret

The final OSEMN stage involved evaluating and interpreting the performance of the developed models. Each model was assessed using the training and testing datasets. Accuracy, precision, recall, F1-score, and confusion matrix were used to measure predictive performance (Sokolova & Lapalme, 2009).

Accuracy measured the proportion of all bookings that were correctly classified as completed or cancelled:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Precision measured how many of the trips predicted as cancelled were actually cancelled:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall measured how many of the actual cancelled trips were correctly identified by the model:

$$\text{Recall} = \frac{TP}{TP + FN}$$

The F1-score provided a balance between precision and recall and was particularly useful for considering both false-positive and false-negative predictions:

$$\text{F-score} = \frac{(1 + \beta^2) TP}{(1 + \beta^2) TP + \beta^2 FN + FP}$$

For the commonly used **F1-score** (where $\beta=1$ \betaeta = 1β=1)

The confusion matrix provided a more detailed view of the predictions by showing True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN). A True Positive represented a cancelled trip correctly predicted as cancelled, while a True Negative represented a completed trip correctly predicted as completed. A False Positive occurred when a completed trip was predicted as cancelled, whereas a False Negative occurred when an actual cancellation was predicted as a completed trip (Sokolova & Lapalme, 2009).

Table 3: Confusion Matrix for Classification Model Evaluation.

Actual Class	Predicted Cancelled (1)	Predicted Success (0)
Cancelled (1)	True Positive (TP)	False Negative (FN)
Success (0)	False Positive (FP)	True Negative (TN)

RESULTS

Four supervised machine learning models, Logistic Regression, Decision Tree, Gradient Boosting, and XGBoost, were developed for predicting trip cancellation. The models were trained using 82,419 booking records and evaluated on a testing dataset of 20,605 records. Training accuracy was first examined to determine how well each model learned from the training data, while accuracy, precision, recall, and F1-score were used to assess performance on the testing data.

Logistic Regression

Logistic Regression achieved a training accuracy of 85.20%. On the testing dataset, the model recorded an accuracy of 84.00%, precision of 76.96%, recall of 82.49%, and F1-score of 79.63%. The small difference between the training and testing accuracy indicated that the model maintained relatively consistent performance when applied to unseen data.

Decision Tree

The Decision Tree classifier achieved a training accuracy of 87.72%. On the testing dataset, it recorded an accuracy of 86.02%, precision of 79.59%, recall of 84.88%, and F1-score of 82.15%. The testing accuracy remained close to the training accuracy, showing that the model retained good predictive performance on unseen booking records.

Gradient Boosting

Gradient Boosting achieved a training accuracy of 90.03%. On the testing dataset, the model recorded an accuracy of 89.06%, precision of 84.01%, recall of 87.85%, and F1-score of 85.89%. The testing accuracy was only slightly lower than the training accuracy, indicating consistent performance across the two datasets.

XGradient Boosting

XGBoost achieved the highest training accuracy at 92.52%. It also recorded the strongest performance on the testing dataset, with an accuracy of 91.62%, precision of 87.12%, recall of 91.41%, and F1-score of 89.12%. The testing accuracy remained close to the training accuracy, showing that the model maintained strong predictive performance on unseen data.




Figure 4: XGBoost Testing Confusion matrix.

Comparison of Model Performance

The training accuracies for Logistic Regression, Decision Tree, Gradient Boosting, and XGBoost were 85.20%, 87.72%, 90.03%, and 92.52%, respectively. On the testing dataset, the models achieved accuracies of 84.00%, 86.02%, 89.06%, and 91.62%, respectively. The results showed a steady improvement from Logistic Regression to the boosting models, with XGBoost recording the highest testing accuracy, precision, recall, and F1-score.

Table 4: Training and Testing Performance (%) of the Machine Learning Models.

Model	Training Accuracy	Testing Accuracy	Testing Precision	Testing Recall	Testing F1-Score
Logistic Regression	85.20	84.00	76.96	82.49	79.63
Decision Tree	87.72	86.02	79.59	84.88	82.15
Gradient Boosting	90.03	89.06	84.01	87.85	85.89
XGBoost	92.52	91.62	87.12	91.41	89.12




Figure 5: Comparison of Testing Performance of Machine Learning Models.

Overall, all four models demonstrated the ability to predict trip cancellation, with testing accuracies ranging from 84.00% to 91.62%. XGBoost produced the strongest overall testing performance among the four developed models.

DISCUSSION

The findings show that all four supervised machine learning models developed in this study were able to predict trip cancellation, although their performance differed across the algorithms. Testing accuracy ranged from 84.00% for Logistic Regression to 91.62% for XGBoost. The relatively small differences between training and testing accuracy across the models also indicate that the models maintained their predictive performance when applied to unseen testing data.

Logistic Regression provided a useful baseline, achieving a training accuracy of 85.20% and testing accuracy of 84.00%. Its testing recall of 82.49% was higher than its precision of 76.96%, indicating that the model was reasonably effective at identifying actual trip cancellations but produced more false cancellation predictions than the other models. Although its performance was lower than that of the tree-based models, the results demonstrate that a conventional binary classification approach can still capture useful patterns associated with trip cancellation.

Decision Tree improved on Logistic Regression, recording a training accuracy of 87.72% and testing accuracy of 86.02%. This improvement suggests that allowing nonlinear relationships and decision rules among the

predictor variables provided additional predictive value. Decision Tree approaches have also been applied to trip cancellation classification in previous research (Krishnaveni & JayaTharunigha, 2022).

The boosting models produced stronger results. Gradient Boosting achieved 90.03% training accuracy and 89.06% testing accuracy, while XGBoost achieved 92.52% and 91.62%, respectively. XGBoost also recorded a testing precision of 87.12%, recall of 91.41%, and F1-score of 89.12%. The stronger performance of the boosting models indicates that combining learners sequentially was better suited to capturing the complex relationships within the ride-hailing data than the individual Logistic Regression and Decision Tree classifiers.

These findings support earlier research showing the value of data-driven approaches in predicting trip cancellation. Sun et al. (2022) used a Seq2Seq deep learning model with multi-source data to predict cancellation probability and reported improved performance compared with models based on a single data source. Sun et al. (2023) similarly applied a deep residual network to predict online ride-hailing order cancellation probability and obtained higher accuracy than baseline deep learning models. Gupta et al. (2024) also demonstrated the usefulness of predictive modelling for Uber ride cancellation using factors such as request time, pickup location, trip distance, and price estimation.

The present study extends this body of work by developing and comparing four supervised machine learning models using a common modelling process and data covering both Ola and Uber. Previous studies reviewed in this research were largely based on individual platforms and mainly considered passenger-initiated cancellations (Gupta et al., 2024; Sun et al., 2022, 2023). In contrast, the present study incorporated unsuccessful trips arising from customer cancellation, driver cancellation, and driver-not-found cases. This provides a broader basis for developing machine learning models for trip cancellation prediction.

Overall, the results demonstrate that historical ride-hailing data can be used effectively to develop machine learning models for predicting trip cancellation. While all four algorithms produced useful predictive performance, the improvement observed from Logistic Regression and Decision Tree to Gradient Boosting and XGBoost shows the advantage of ensemble boosting approaches for this prediction task. These models can provide a basis for data-driven approaches aimed at identifying potential trip cancellations and supporting operational decisions in ride-hailing platforms.

CONCLUSION

This study developed four supervised machine learning models, Logistic Regression, Decision Tree, Gradient Boosting, and XGBoost, for predicting trip cancellation using historical booking data from Ola and Uber. The models were developed through a structured process involving data preparation, feature selection, training, and testing using the same set of selected predictor variables.

The findings showed that all four models were able to predict trip cancellation, although their predictive performance differed. Logistic Regression provided a useful baseline, while Decision Tree showed improved performance. The boosting models performed better overall, with Gradient Boosting and XGBoost showing stronger predictive capability. XGBoost produced the best overall performance among the models developed. The small differences observed between training and testing performance also showed that the models maintained consistent performance when applied to unseen testing data. These findings demonstrate that machine learning can provide an effective approach for predicting trip cancellation from historical ride-hailing data.

Based on these findings, ride-hailing platforms can consider incorporating machine learning models into their operational systems to support early identification of trips that are likely to be cancelled. The stronger performance of XGBoost suggests that it can provide a useful basis for trip cancellation prediction where similar booking data are available. Predictions from such models could support operational decisions such as improved driver-passenger matching, trip reallocation, and targeted interventions aimed at reducing cancellations. The

models should also be periodically retrained as new booking data become available so that they remain responsive to changes in travel patterns and user behavior.

Future research should validate the developed model using additional datasets from different ride-hailing platforms and operational environments to further assess its generalizability and robustness. Additional contextual factors, including weather conditions, traffic congestion, road incidents, special events, and real-time driver availability, should also be investigated as these factors may influence trip cancellation. Future studies could further evaluate advanced approaches such as deep learning and hybrid machine learning models to determine whether they offer improvements over the models developed in this study. Finally, the integration of trip cancellation prediction models into real-time ride-hailing systems should be explored to support proactive decision-making, operational efficiency, and service reliability.

Ethical Considerations

The study used secondary data that did not contain personally identifiable information, and no primary data were collected directly from human participants.

Ethical Approval: Ethical approval for the study was obtained from the Technical University of Mombasa Scientific and Ethics Review Committee (TUMSERC). Research authorization was also obtained from the National Commission for Science, Technology and Innovation (NACOSTI).

Conflict of Interest: The authors declare no conflict of interest.

Data Availability

The dataset used in this study is publicly available on Kaggle. The Ola & Uber: Ride Booking & Cancellation Data dataset was provided by Mengar (2025) and contains historical ride-hailing booking records used for data preprocessing and machine learning model development. The dataset is distributed under the Creative Commons CC0 1.0 Universal Public Domain Dedication license, which permits unrestricted use, modification, and distribution. It can be accessed through the Kaggle repository and is also cited in the reference list.

Link: <https://www.kaggle.com/datasets/hetmengar/ola-and-uber-ride-booking-and-cancellation-data>

ACKNOWLEDGEMENTS

I acknowledge the Technical University of Mombasa for providing me with the opportunity to conduct this research. I also extend my appreciation to my colleagues and peers for their valuable discussions and support, and to my family and friends for their continued encouragement throughout the study.

REFERENCES

1. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
2. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. 785–794.
3. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
4. Gupta, A., Pipaliya, A., Jacob, L., & Balagangadhar, K. (2024). Predictive modeling for Uber ride cancellation and price estimation: An integrated approach. 1–8.
5. He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263–1284.
6. Hosmer Jr, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression*. John Wiley & Sons.

7. Krishnaveni, S., & JayaTharunigha, M. (2022). Predict the cancellation of trips using classifiers and time series modeling. *JETIR*, 9(6).
8. Kumari, K., Bhardwaj, M., & Sharma, S. (2020). OSEMN approach for real time data analysis. *International Journal of Engineering and Management Research*, 10(2), 107–110.
9. Li, X., Pan, G., Wu, Z., Qi, G., Li, S., & Zhang, D. (2012). Prediction of urban human mobility using large-scale taxi traces and its applications. *Frontiers of Computer Science*, 6(1), 111–121.
10. Liu, Y., Yang, H., & Qin, X. (2024). An incentive strategy for the retention of impatient passengers in ride-sourcing markets. *Transportation Research Part E: Logistics and Transportation Review*, 189, 103645.
11. Mengar, H. (2025). Ola & Uber: Ride booking & cancellation data [Dataset]. Kaggle. <https://www.kaggle.com/datasets/hetmengar/ola-and-uber-ride-booking-and-cancellation-data>
12. Ngo, H. H., & Mishra, S. (2023). Dynamic carpooling and repositioning for ridehailing services with mixed strategies.
13. Nuzzolo, A., Comi, A., Papa, E., & Polimeni, A. (2018). Understanding taxi travel demand patterns through floating car data. 445–452.
14. Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management*, 45(4), 427–437.
15. Sun, H., Lv, Z., Li, J., Xu, Z., & Sheng, Z. (2022). Prediction of cancellation probability of online car-hailing orders based on multi-source heterogeneous data fusion. 168–180.
16. Sun, H., Lv, Z., Li, J., Xu, Z., & Sheng, Z. (2023). Will the order be canceled? Order cancellation probability prediction based on deep residual model. *Transportation Research Record*, 2677(6), 142–160.
17. Wang, X., Liu, W., Yang, H., Wang, D., & Ye, J. (2019). Customer behavioural modelling of order cancellation in coupled ride-sourcing and taxi markets. *Transportation Research Procedia*, 853–873.
18. Wen, D., Li, Y., & Lau, F. C. (2024). A survey of machine learning-based ride-hailing planning. *IEEE Transactions on Intelligent Transportation Systems*.
19. Xu, M., Yue, P., Yu, F., Yang, C., Zhang, M., Li, S., & Li, H. (2023). Multi-agent reinforcement learning to unify order-matching and vehicle-repositioning in ride-hailing services. *International Journal of Geographical Information Science*, 37(2), 380–402.
20. Zhang, W., & Fan, Y. (2020). Spatiotemporal characteristics and self-organization of urban taxi dispatch. *Complexity*, 1–11.