



HAL
open science

Impact of Mobility Patterns on Federated Learning applied to Human Mobility Prediction

João Paulo Esper, Aline Carneiro Viana, Jussara M. Almeida

► To cite this version:

João Paulo Esper, Aline Carneiro Viana, Jussara M. Almeida. Impact of Mobility Patterns on Federated Learning applied to Human Mobility Prediction. NetMob 2023, Oct 2023, Madrid, Spain. , 2023. ⟨hal-04699120⟩

HAL Id: hal-04699120

<https://inria.hal.science/hal-04699120v1>

Submitted on 16 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Impact of Mobility Patterns on Federated Learning applied to Human Mobility Prediction

João Paulo Esper*, Aline Carneiro Viana†, Jussara M. Almeida*

*Universidade Federal de Minas Gerais, Brazil. † INRIA, France.

E-mail: {joaopauloesper, jussara}@dcc.ufmg.br*, aline.viana@inria.fr†

Abstract—The Federated Learning (FL) framework has been applied in multiple domains, offering solutions that provide both accuracy and data privacy protection. Yet, specifically for human mobility prediction, prior solutions have been analyzed on mobility datasets that are spatially and temporally sparse, AND neglected the impact of the *heterogeneity* of users’ mobility patterns. Heterogeneity on the (fine-grained) spatial and temporal mobility patterns directly impact prediction, hardening the FL performance analysis. As such, prior evaluations of FL on mobility prediction are limited and may overestimate the robustness of the proposed solutions. We here aim to fill this gap by analyzing the impact that different mobility patterns (e.g., repetitive and/or exploratory patterns) have on the performance of FL-based human mobility prediction models, in terms of both model effectiveness and efficiency.

Index Terms—Mobility prediction, federated learning, privacy.

I. INTRODUCTION

Human mobility prediction drives the solution to a diverse range of complex problems, including resource allocation, traffic management, and epidemic prevention, to name a few [1]. Human mobility is typically represented as a time series describing an ordered temporal sequence of visited locations by a user on a daily basis, i.e., a circadian trajectory. The prediction task considered in this paper is the inference of the next location to be visited by a user at the next timestamp.

Recently proposed human mobility prediction models often exploit machine learning techniques, notably deep learning models [2], [3], to deliver state-of-the-art prediction accuracy. Yet, such approaches present concerning privacy issues that are rooted in their *centralized* architecture, which requires the uploading of sensitive data (e.g., individuals’ visited locations) to a server where the prediction model training occurs.

Towards meeting users’ increasing demands for privacy guarantees in various domains (notably mobility prediction), the decentralized Federated Learning (FL) framework was proposed [4], [5]. In FL, a shared model is first *decentralized* trained on multiple devices, using only local data. That is, private location data is only stored and analyzed locally on the corresponding device and used to train a local prediction model. The local prediction models are then uploaded to a central server. Then, a global model is generated by the server, that first aggregates the received local parameters, then chooses candidate devices and finally distributes the updated global model. The previous steps are repeated until convergence, and an optimal model for mobility prediction is generated [4]. Note that only the locally trained models, i.e., model parameters (e.g., weights), naturally less sensitive than users’ raw location data, are transferred to the server. As such, FL offers the possibility of producing (accurate) predictions while still restricting access to users’ sensitive data.

FL has been applied to various problems, from image classification to next word prediction [6], [7]. For mobility prediction specifically, some prior efforts adapted solutions from other domains to mobility problems. For example, in [6], an image classification model was used for transportation mode prediction by converting coordinates into pixels. Yet, directly adapting solutions from other domains can be both challenging and inefficient due to the spatial and temporal related specificities of mobility prediction. Indeed, visits to certain point of interests are directly correlated to the users’ routines and preferences. These are hard to embed in models for *next word* prediction or image classification, which are more concerned with grammatical structures of a language and recognition of patterns on a static low-dimensional space.

There are only a few FL solutions that were designed for mobility prediction [4], [5]. Yet, prior analyses of them neglected the impact that the naturally heterogeneous human patterns may have on FL effectiveness. Also, the usage of social network datasets, that are both sparse in space and coarse in time, challenges routine and mobility patterns characterization.

In this work, we aim to fill this gap by analyzing the performance of alternative FL-based mobility predictions in scenarios with users with varying mobility patterns, identified in real and less sparse human mobility data. We aim to answer the following question: *How do existing FL-based mobility prediction models perform for users with very different mobility patterns, such as very repetitive behavior (e.g., routines) or more exploratory visiting patterns (e.g., tourists)?* Our analyses comprise both model effectiveness (accuracy) and efficiency (resource usage and execution time), and offer insights into possible improvements to current FL solutions.

II. EVALUATION

A. Methodology: Models and dataset

After searching for FL models designed for mobility prediction in the literature, we settled with two different models introduced in [5], namely *GRU-Spatial* and *Flashback*, as these were the only ones with publicly accessible code which we were able to run with no execution errors. *GRU-Spatial* uses Gated Recurrent Units (GRUs) which, given a sequence, learn which data is relevant keeping or not. GRU was proposed as an evolution from Recurrent Neural Networks (RNNs), which had their learning process impacted by very large/small gradient values, an issue tackled by the gating mechanism in GRU. *Flashback*, in turn, was first proposed as a centralized model [3] and later adapted to a federated setup [5]. It is an RNN-based prediction model proposed for sparse mobility data,

Table I: Federated results considering 5 clients and samples with 3,000 users. Resource measured in vRAM consumption.

	Acc@1 Flashback	Acc@1 GRU-Spt.	Acc@5 Flashback	Acc@5 GRU-Spt.	Resource Flashback	Resource GRU-Spt.	Time Flashback	Time GRU-Spt.
Scouters	0.111 ± 0.02	0.301 ± 0.02	0.270 ± 0.02	0.538 ± 0.02	1,662 MB	1,746 MB	26 min 17 sec	14 min 41 sec
Regulars	0.160 ± 0.01	0.395 ± 0.02	0.322 ± 0.03	0.618 ± 0.01	1,604 MB	1,686 MB	21 min 41 sec	11 min 45 sec
Routiners	0.213 ± 0.02	0.500 ± 0.02	0.384 ± 0.02	0.721 ± 0.01	1,510 MB	1,538 MB	19 min 10 sec	10 min 18 sec
Mixed	0.169 ± 0.01	0.397 ± 0.05	0.349 ± 0.02	0.635 ± 0.04	1,622 MB	1,706 MB	21 min 31 sec	11 min 33 sec

Table II: User profiles and metrics of their trajectories.

	Trajectory length	Stationarity	Diversity
Scouters	420.66 ± 4.11	0.391 ± 0.01	0.840 ± 0.00
Regulars	345.21 ± 1.72	0.566 ± 0.01	0.598 ± 0.00
Routiners	320.53 ± 2.27	0.668 ± 0.01	0.404 ± 0.01

which explores spatial-temporal contexts, searching historical hidden states with equivalent context to improve its prediction.

Our study uses a fully anonymized *Call Detail Record* (CDR) dataset gathered by a major telecom operator in a metropolitan area in China. It contains records of 58,502 users spanning over a two week period. Unlike real raw CDRs, the recorded location is not the cell tower’s coordinates the user was attached to, but rather the centroid of a 200 m^2 cell in a grid representation nearest to the tower the user stayed mostly attached for each hour. No user identity or the corresponding operator’s cell tower infrastructure is provided. To deal with missing records and homogenize the number of records per user, the data was handled as in [1], resulting in 41,460 users.

To characterize the users in our dataset, we followed the clustering mechanism proposed in [1] to group them into three mobility profiles based on their trajectories: (i) *Scouters* are users more prone to explore and discover new areas; (ii) *Regulars* are users who constantly alternate between explorations and revisits; and (iii) *Routiners* are users who rarely explore and prefer to stick to their known places. We also used two metrics, *Stationarity* and *Diversity* [8], to analyze each user trajectory. *Stationarity* measures the number of records for which the user stays continuously in the same place. *Diversity* quantifies the number of distinct segments of trajectories given a time-ordered sequence of locations visited by a user.

B. Results

We started by clustering the 41,460 users in our dataset into the three aforementioned mobility profiles and characterizing their trajectories. We found 5,743 (14%) users characterized as *Scouters*, 25,032 (60%) users as *Regulars*, and 10,685 (26%) users as *Routiners*. Table II presents a broad characterization of the trajectories of users in each cluster, including average trajectory length, *Stationarity* and *Diversity*. *Scouters* tend to have much longer and diverse trajectories, while more data may favor model training, greater *Diversity* also makes pattern learning more challenging, motivating the study of model effectiveness for different profiles.

We have carried out a set of experiments evaluating effectiveness and efficiency of both *Flashback* and *GRU-Spatial* for different subsets of users in various scenarios. For illustration purposes, Table I presents some of our results for a sample of 3,000 users from each profile as well as a sample of 3,000 users with mixed profiles, keeping the fractions of users of each profile the same as in the original dataset. The table shows results for both model effectiveness (accuracy) and efficiency

(resource usage and total execution time). The users were federated into 5 clients, each client containing data from 600 users, which is a setup comparable to the one used in [5].

III. DISCUSSIONS AND FUTURE WORK

Considering the proposed scenarios, we noticed the considerable impact that different mobility patterns can have on both the effectiveness and efficiency of the FL models. Easier to predict users (i.e., *Regulars* and *Routiners*) experience great improvements on the accuracy of both models, accelerated the learning process and reduced resource consumption. *Scouters*, in turn, really challenged both models on every aspect, specially accuracy. Indeed, even a small fraction of *Scouters* (14%) greatly impacted both models in the *Mixed* scenario, evidencing the impact that heterogeneity has on the solutions.

These initial results offer many future directions of exploration. For example, we aim to study strategies to explicitly incorporate into model training the different properties of the mobility profiles, as well as investigate the performance trade-offs of favoring one particular profile over the others. Studying how the FL solutions can adapt to fluctuations in mobility patterns (and even profiles) over time is also worth pursuing. Finally, prior work has discussed the reconstruction of trajectories in an FL environment by inferring patterns given the model weights [4]. Analyzing how such security issues may correlate with the mobility patterns and how they impact model effectiveness is also an interesting avenue to pursue.

ACKNOWLEDGEMENTS

This work was supported by *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq), *Fundação de Amparo à Pesquisa do Estado de Minas Gerais* (FAPEMIG), *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES) and the CAPES-STIC-AMSUD 22-STIC-07 LINT project.

REFERENCES

- [1] L. Amichi *et al.*, “Understanding individuals’ proclivity for novelty seeking,” in *Proc. of the 28th International Conference on Advances in Geographic Information Systems*, 2020.
- [2] J. Feng *et al.*, “Deepmove: Predicting human mobility with attentional recurrent networks,” in *Proc. of the world wide web conference*, 2018.
- [3] D. Yang *et al.*, “Location prediction over sparse user mobility traces using RNNs,” in *Proc. of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2020.
- [4] J. Feng *et al.*, “PMF: A privacy-preserving human mobility prediction framework via federated learning,” *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, no. 1, 2020.
- [5] C. E. J. Ezequiel *et al.*, “Federated Learning for Privacy-Aware Human Mobility Modeling,” *Frontiers in Artificial Intelligence*, vol. 5, 2022.
- [6] F. Yu *et al.*, “Privacy-preserving federated learning for transportation mode prediction based on personal mobility data,” *High-Confidence Computing*, vol. 2, no. 4, 2022.
- [7] T. Yang *et al.*, “Applied federated learning: Improving google keyboard query suggestions,” *arXiv preprint arXiv:1812.02903*, 2018.
- [8] D. do Couto Teixeira *et al.*, “On estimating the predictability of human mobility: the role of routine,” *EPJ Data Science*, vol. 10, no. 1, 2021.