



HAL
open science

Survey of Federated Learning Models for Spatial-Temporal Mobility Applications

Yacine Belal, Sonia Ben Mokhtar, Hamed Haddadi, Jaron Wang, Afra
Mashhadi

► **To cite this version:**

Yacine Belal, Sonia Ben Mokhtar, Hamed Haddadi, Jaron Wang, Afra Mashhadi. Survey of Federated Learning Models for Spatial-Temporal Mobility Applications. *ACM Transactions on Spatial Algorithms and Systems*, In press, 10.1145/3666089 . hal-04613220

HAL Id: hal-04613220

<https://hal.science/hal-04613220v1>

Submitted on 15 Jun 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 Attribution 4.0 International License

SURVEY OF FEDERATED LEARNING MODELS FOR SPATIAL-TEMPORAL MOBILITY APPLICATIONS

A PREPRINT

✉ **Yacine Belal**
National Institute of Applied Sciences
Lyon, France
yacine.belal@insa-lyon.fr

✉ **Sonia Ben Mokhtar**
National Institute of Applied Sciences
CNRS
Lyon, France
sonia.benmokhtar@insa-lyon.fr

✉ **Hamed Haddadi**
Imperial College London
London, UK
h.haddadi@imperial.ac.uk

Jaron Wang
University of Washington
Bothell, USA
jierui98@uw.edu

✉ **Afra Mashhadi**
University of Washington
Bothell, USA
mashhadi@uw.edu

June 15, 2024

ABSTRACT

Federated learning involves training statistical models over edge devices such as mobile phones such that the training data is kept local. Federated Learning (FL) can serve as an ideal candidate for training spatial temporal models that rely on heterogeneous and potentially massive numbers of participants while preserving the privacy of highly sensitive location data. However, there are unique challenges involved with transitioning existing spatial temporal models to federated learning. In this survey paper, we review the existing literature that has proposed FL-based models for predicting human mobility, traffic prediction, community detection, location-based recommendation systems, and other spatial-temporal tasks. We describe the metrics and datasets these works have been using and create a baseline of these approaches in comparison to the centralized settings. Finally, we discuss the challenges of applying spatial-temporal models in a federated setting and by highlighting the gaps in the literature we provide a road map and opportunities for the research community.

Keywords Spatio-Temporal Applications, Federated Learning, Privacy-Preserving, Mobility Prediction, Transportation, Location Based Social Networks

1 Introduction

Spatio-temporal (ST) mobility data collected by location-based services (LBS) [72] and other means such as Call Data Records (CDR), WiFi hotspots, smart watches, cars, etc. is very useful from a socio-economical perspective as it is at the heart of many useful applications (*e.g.*, navigation, geo-located search, geo-located games) and it allows answering numerous societal research questions [86]. For example, CDR has been successfully used to provide real-time traffic anomaly as well as event detection [155, 161], and a variety of mobility datasets have been used in shaping policies for urban communities [47] or epidemic management in the public health domain [129, 128]. From an individual-level perspective, users can benefit from personalized recommendations when they are encouraged to share their location data with third parties [37].

While there is no doubt about the usefulness of location-based applications, privacy concerns regarding the collection and sharing of individuals' mobility traces or aggregated flow of movements have prevented the data from being utilized to their full potential [151, 17, 89]. A mobility privacy study conducted by De Montjoye et al. [33] illustrates that four spatio-temporal points are enough to identify 95% of the individuals, which exacerbates the user re-identification risk and could be the origin of many unexpected privacy leakages. Indeed, various studies have shown that numerous threats

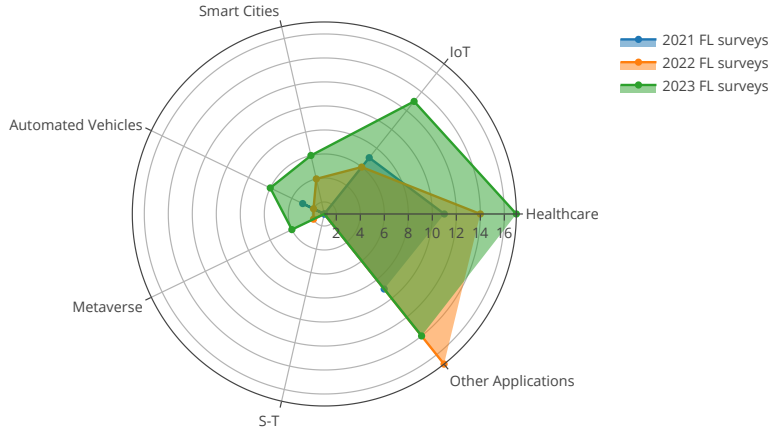


Figure 1: Number of FL Surveys and the underlying applications tackled. **Note the absence of surveys on FL-ST works**, which contrasts with the ever-growing number of FL-ST methodological papers (e.g., 26 papers in 2023).

arise if location data falls into the hands of inappropriate parties. These threats include re-identification [111], the inference of sensitive information about users [89, 165] (e.g., their home and work locations, religious beliefs, political interests or sexual orientation). In some extreme cases, sharing geo-located data may even endanger users’ physical integrity (e.g., the identification of protesters in dictatorial regimes or during wars¹) or their belongings (e.g., robbery²).

One way to consider addressing privacy challenges is to break from centralized data collection and maintain the location data on user devices. In this decentralized paradigm, a variety of solutions are plausible. A possible approach is to create accurate enough synthetic location data to train spatial-temporal models and then fine-tune them on users’ devices to perform best for the individuals’ location data [67]. Although this is a promising approach, the research has shown that mobility trajectory generation is an extremely difficult task that may not portray data heterogeneity accurately [88, 107]. Moreover, trajectories from different locations, periods, or user groups may have distinct characteristics that the pre-trained model does not handle effectively. Although this is an ongoing field of research, to date there exists little evidence as to whether models can transfer mobility knowledge across cities [67], as with a significant distribution shift between the data used for pre-training and the data for the target task, transfer learning may not generalize well.

In parallel, an alternative approach that is increasingly considered a promising approach is *Federated Learning* (FL) [117]. FL relies on clients (be it users’ smartphones or edge devices) to train a machine learning model on their local training data and share the model weights with a central server (called the FL server) that aggregates the received clients’ contributions. As such, FL can inherently alleviate some privacy concerns [149] while enabling clients to collaboratively train and utilize a global model without ever sharing their private data. This characteristic renders FL more compliant with recent privacy regulations (e.g., General Data Protection Regulation (GDPR) [158], California Consumer Privacy Act (CCPA) [63]) and an increasing user awareness [156]. Additionally, it facilitates the emergence of new use cases (e.g., digital healthcare [144, 92]). Indeed, several real-world applications leveraging FL have already emerged including drug discovery [133] and fraud detection [4].

Compared with other applications of FL, ST data and models possess certain characteristics and properties that introduce challenges. For instance, in domains involving ST data, the observations made at nearby times and locations cannot hold the independent property due to the auto-correlation (e.g., changes in traffic activity occur smoothly over time and space). As a result, classical FL algorithms that assume independence among observations are not suited for ST applications [11]. Furthermore, ST data is not identically distributed, that is, it does not meet the assumption that every instance has the same probability to belong to a set. These two properties are referred to as *iid* (independent and identically distributed). Finally, ST data is often both spatially and/or temporally sparse and often created at high frequencies, especially in real-time applications, resulting in a high velocity and volume of data generation.

Related Surveys Since FL emerged half a decade ago, many survey papers [95, 164, 34, 49, 125, 140] have reviewed relevant literature from various perspectives. In 2023 alone, 91 surveys on federated learning were published [164, 13, 109] with the majority being on applications-specific topics, predominantly focused on healthcare [142, 58, 138, 163], finance [130], blockchain [74, 200, 135, 126], and the Internet of Things (IoT) [131, 6, 29, 157]. Figure 1 presents the

¹<https://www.independent.co.uk/news/ukraine-ap-russia-gps-kyiv-b2093310.html>

²<https://pleaserobme.com>

growing number of application-specific FL surveys³, which is paralleled by an emergence of papers on FL-ST topics. In motivating the need for this survey, we discuss the perspectives that IoT-related and generic FL surveys have covered and highlight their shortcomings when applied to ST models. More specifically, FL Surveys on smart cities including those of IoT applications are closest to our work. Zheng et al. [199] reviewed applications of FL in smart cities. They provided a broad overview of various applications of smart cities including IoT, transportation, communication, medical care, and finance. They created a glossary of the papers in these domains without discussing the diversity of existing approaches or drawing comparisons between these works. Their work mainly focused on security and privacy challenges and outlined algorithmic efficiency as a future direction but did not consider the heterogeneity and sparsity often associated with the location data. The closest survey to ours is the work of Pandya et al. [131], which has a subsection discussing the FL transportation applications in smart cities. However, only one of the research works discussed in their survey is revisited in our work, as their focus predominantly revolves around a different type of use cases (*e.g.*, number plate recognition, railway ... etc.).

Other FL surveys that are *not* application-specific and review methodological approaches of FL regarding heterogeneity, model convergence, and personalization fall short in drawing a comparison between these different techniques when applied to ST data. Although some of the challenges such as non-iid are well covered in those surveys, the other characteristics of the ST data are not covered. Finally, we note that despite its unique properties and the increasing number of spatial-temporal research papers emerging in FL, no survey paper has attempted to categorize, review, and compare this emergence of FL-ST methodological approaches (Figure 1).

Our Contribution The contributions of this survey are as follows:

- In this survey, we study 38 existing spatial-temporal works that have been tailored for the FL paradigm. As we review these recent advancements, we provide brief yet detailed descriptions of each approach, supplemented by a meta-analysis comparing their efficacy in addressing the encountered challenges.
- As part of this survey we create an open-source repository⁴ where we host the existent FL-ST algorithms that are available as well as sample datasets. We also include available mobility data handling (pre-processing, analysis, and visualization) software. We hope that this open-source repository enables the research community to contribute their work and enable a greater number of baselines.
- We discuss open challenges and a roadmap that we envision for the research community to explore in the coming years. Some of the research questions that emerge from our survey are: *How can FL effectively achieve better personalization through the integration of semantic information?* (See Section 5.1) *What novel techniques and solutions can be developed to detect and mitigate Byzantine behaviors in FL-ST?* (See Section 5.2) *How can current ST mobility frameworks be optimized to enhance communication efficiency?* (See Section 5.3) *How does fairness of models behave in FL-ST and more importantly, how can it (*i.e.*, fairness) be defined in this context?* (See Section 5.4) *How can the ST mobility research community establish and promote standardized practices?* (See Section 5.5) *How can realistic cross-silo spatial-temporal datasets be built in order to address ownership challenges and support benchmarking?* (See Section 5.6) and finally, *what gaps must be addressed to facilitate a seamless transition towards real-world FL-ST applications?* (See Section 5.7).

Organization of this survey The remainder of this paper is structured as follows. First, we present a background on FL in Section 2. Then, we present a set of spatial-temporal applications under study in Section 3 and detail the used FL approaches for these applications in Section 4. Finally, we discuss open research challenges in Section 5 and conclude the paper in Section 6.

2 Federated Learning : Preliminaries

Federated Learning (FL) is a paradigm to perform distributed machine learning at the edge [20, 80, 116, 87]. In FL, a global (joint) model is trained in a decentralized fashion without the need to collect and process user data centrally. The central service provider, often called aggregator, distributes a shared ML model to multiple users for training on local data, and then aggregates the resulting models into a single, more powerful model, using an aggregation method (*e.g.*, Federated Averaging [116]).

More formally, let θ be the global model that an FL instance aims to learn and $F(\theta)$ be the cost function that evaluates θ . let $\mathcal{N} = \{1, 2, \dots, n\}$ be the set of users, with each user i possessing a local dataset \mathcal{D}_i . Note that we use ‘users’ and ‘participants’ interchangeably throughout the paper.

³Our methodology in harvesting data about all survey papers on FL is outlined in the Appendix.

⁴<https://github.com/YacineBelal/Survey-of-Federated-Learning-Models-for-Spatial-Temporal-Mobility-Applications>

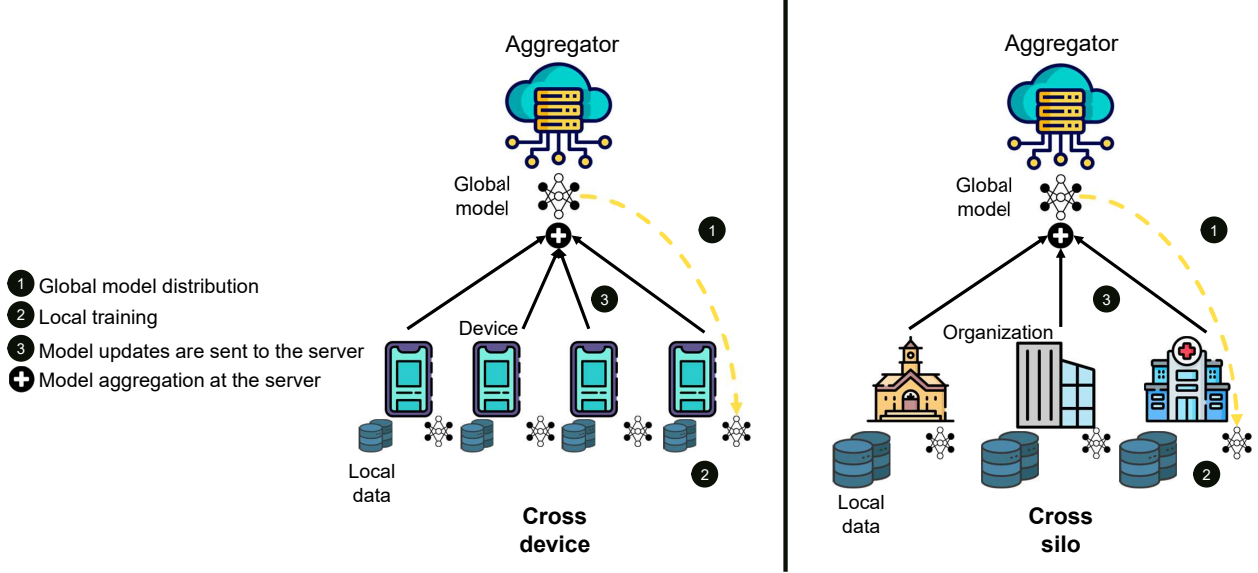


Figure 2: Federated Learning general workflow. On the left, a cross-device setting where each client is a device (e.g., smartphone) and on the right is cross-silo setting where each client represents an organization. The main difference between the two is the the magnitude and availability, both in terms of clients and local data.

The aggregator initiates by randomly sampling $\mathcal{S} = \{1, 2, \dots, m\}$ with $\mathcal{S} \subset \mathcal{N}$ being a subset of users chosen to participate to a specific round t . The aggregator transmits the global model θ^t to \mathcal{S} . Each selected participant i trains θ^t on \mathcal{D}_i to obtain an updated model, denoted θ_i^{t+1} . Then, i send the model updates $\nabla F(\theta_i^t, \mathcal{D}_i)$ (or θ_i^{t+1}) back to the aggregator. Communication overhead can be reduced by applying a random mask to the model weights [87]. The central server then aggregates the received updates [141] to create the new global model, θ^{t+1} . This round-based mechanism is repeated until some convergence criteria of θ is reached. Equation 1 presents a Federated Averaging-based aggregation step.

$$\theta^{t+1} = \theta^t - \frac{\eta}{\sum_{i=1}^m |\mathcal{D}_i|} \sum_{i=1}^m |\mathcal{D}_i| \cdot \nabla F(\theta_i^t, \mathcal{D}_i) \quad (1)$$

Where η denotes the global learning rate and $|\cdot|$ denotes the set cardinality operation.

This described algorithm represents the general FL workflow. However, it is essential to note that there exist several flavors of FL. Based on the order of magnitude of the number of participants and the assumptions regarding their computational capabilities, two FL settings are distinguished in this paper: cross-device FL (CD FL) and cross-silo FL (CS FL) [70].

CD FL typically refers to scenarios where the number of users is in the order of millions of individuals, each possessing limited and heterogeneous computational resources (e.g., personal laptops, cell phones). In contrast, cross-silo FL operates at a (much) smaller scale of the number of users, representing each a large entity (e.g., hospitals, cellular operators, and data centers), with substantial computational capabilities. Figure 2 presents an overview of both schemes of FL.

2.1 Challenges

There are various challenges in FL that the research community has been studying in the past years (see [80] for full reference). Two of these challenges impacting FL models are generally regarding **data heterogeneity** and **model personalization**.

- Data heterogeneity refers to the scenario where clients' data are not-iid which could lead to domain shift problems making learning a generalizable representation a difficult task. Consider two users i and j , with data following probability distributions P_i and P_j , respectively. These probabilities are typically rewritten as $P_i(Y|X)P_i(X)$ and $P_j(Y|X)P_j(X)$, respectively, where X and Y are the features' (e.g., visited POIs)

and labels' (e.g., next POI) probability distributions. Considering this notation, data can be non-iid in several ways: the marginal distribution of features can vary between clients (i.e., $P_i(X) \neq P_j(X)$), such is the case for clients who have significantly different preferences in a POI recommendation use case. This is known as a covariate shift. Alternatively, the marginal distribution of labels can be different, also known as the prior distribution shift (i.e., $P_i(Y) \neq P_j(Y)$). This is especially the case for clients tied to different geo-regions. Finally, there can also be a concept shift, that is, the conditional distributions $P(X|Y)$ and $P(Y|X)$ can highly vary across clients and context (e.g., weather, time... etc). As we will discuss in the next sections, the human trajectory is extremely unique [33] and strongly influenced by different meta-parameters. This leads to the coexistence of different types of non-iid property and calls for an increased need for personalization.

- **Model Personalization** on the other hand refers to the process of assigning different model parameters (weights) to different clients. The objective behind this is twofold: i) improve the global model's individual performance and ii) mitigate the impact of non-iid property. To this end, various methodologies have emerged. One approach involves incorporating context features into the model during training, known as personalization through featurization. Another technique, meta-learning, focuses on adapting the global model to the local learning task. In its simplest form, it boils down to fine-tuning the global model or a subset of it. Understanding what level of generalization can be learned globally and what layers of the models need to be personalized locally is an active area of research.

There are several other FL challenges that we consider throughout this paper. Firstly, concerns over data and model parameter privacy, which are addressed through methods such as Differential Privacy (DP) [118], Trusted Execution Environments (TEEs) [124] and encryption-based methods [152, 110, 69] (i). Secondly, Byzantine Resilience [59], defined as the capability to train an accurate statistical model despite arbitrary behaviors, presents another significant challenge (ii). Finally, Communication Efficiency stands as a crucial aspect to consider (iii).

In addition to these challenges, there are specific considerations within the realm of ST mobility applications. These include: 1) the constraints on client resources, 2) the online learning setting - specifically, whether clients train while simultaneously collecting/creating data, and 3) the comparison of solutions with at least one federated competitor, aside from standalone FL. This comprehensive evaluation framework has allowed us to conduct a meta-comparison, analyzing the strengths and limitations of the main works within each approach. Moreover, it has facilitated the identification of potential research directions (See Section 5).

2.2 Frameworks

As FL establishes itself as the standard machine learning paradigm, numerous frameworks have emerged to offer scalable and flexible solutions for privacy-preserving machine learning in distributed settings. Selecting an appropriate FL framework for designing new algorithms is a non-trivial task. Therefore, in this section, we define specific criteria to facilitate the selection process of an FL framework. The following criteria will be considered:

- **Supported Frameworks:** Most FL frameworks can be extended to integrate with various machine learning frameworks. However, this process is often complex. Hence, we will outline the machine learning frameworks inherently supported by each FL framework and we will qualify by "native" each framework that provides its own optimization library.
- **Privacy tools:** As privacy stands as a central motivation for FL, it is imperative to recognize that standalone FL lacks privacy [140, 80]. Within this context, we previously mentioned a formal model, DP, as well as hardware-based solutions like TEEs. Additionally, a range of privacy-enhancing technologies are based on cryptographic tools such as Secure Multi-Party Computation (SMPC) [152, 110], Homomorphic Encryption (HE) [69, 182] and Secret Sharing [150]. Evaluating which of these tools an FL framework supports is crucial due to the complexity of integrating them posthoc.
- **Scenario:** This criterion assesses if the framework was primarily designed for real-world deployments, addressing physical environments, networking challenges, and scalability, or if its focus was solely on simulations.
- **Benchmarking:** To enhance the efficiency of the research workflow, the majority of Federated Learning frameworks come equipped with predefined machine learning tasks, including models, datasets, and benchmarks. These tasks will be highlighted for each framework. It's important to note that while frameworks come with suggested tasks, they are not restricted solely to those tasks.

Table 1 presents a summary of the specified evaluation criteria for the main FL frameworks. Notably, most frameworks are extensively benchmarked for tasks in computer vision and Natural Language Processing (NLP), with relatively limited consideration for spatial mobility applications. *FedScale* [90] distinguishes itself by offering starting pack

Table 1: Available functionalities and bench-marking across different FL frameworks.

Name	Supported Frameworks	Privacy tools	Scenario	Benchmarking
Tensorflow Federated (TFF)	Tensorflow	DP SMPC	Simulation	Computer Vision NLP
Pysift & Pygrid	Tensorflow Pytorch	DP SMPC	Simulation Real	Computer Vision
FATE	Tensorflow Pytorch Sickit-learn	HE	Simulation Real	Computer Vision
Flower	Framework-agnostic	DP	Simulation	Computer Vision
PaddleFL	Native	DP SMPC Secret Sharing	Simulation Real	Computer Vision NLP Recommendation
FederatedScope	Tensorflow Pytorch	DP SMPC	Simulation Real	Computer Vision NLP Recommendation Graph Learning
LEAF	Tensorflow	✗	Simulation	Computer Vision NLP
FedML	Framework-agnostic	SMPC	Simulation Real	Computer Vision NLP
FedScale	Tensorflow Pytorch	DP SMPC	Simulation Real	Computer Vision NLP Recommendation Mobility Prediction ⁵ RL
FS-Real	Tensorflow Pytorch	✗	Simulation Real	Speech Recognition
MindSpore Federated	MindSpore	SMPC LDP	Simulation Real	Computer Vision NLP Recommendation

mobility datasets, including a task for taxi trajectory prediction utilizing the TaxiPorto dataset [120]. Additionally, it incorporates the Waymo Motion Dataset, comprising object trajectories and corresponding 3D maps for 103,354 scenes, utilized in autonomous driving scenarios.

A noteworthy feature is observed in *FederatedScope* [167], which introduces benchmarks for graph learning, potentially applicable to Graph Neural Network (GNN)-based approaches for flow prediction tasks.

Regarding machine learning framework support, it’s notable that only two of the existing solutions offer a framework-agnostic approach to Federated Learning—namely, *Flower* [18] and *FedML* [65].

Considering privacy enhancements, most frameworks integrate DP and SMPC functionalities. However, *LEAF* [24] does not provide support for privacy-enhancing technologies. On the other hand, *FATE* [104] stands out by offering unique Homomorphic Encryption support.

In terms of compatibility with physical networking infrastructures, device heterogeneity, and client scalability—referred to as real-world scenarios- *FedScale* [90], *FS-Real* [51], and *Mindspore Federated* [190] emerge as the only frameworks capable of accommodating such requirements. Notably, *FS-Real* builds upon *FedScale*, enhancing its scalability to accommodate up to 100k devices and improving compatibility with heterogeneous hardware and software configurations across devices. However, the integration of *FedScale*’s privacy tools into *FS-Real* remains unclear. In contrast, *Mindspore Federated* achieves comparable scalability by leveraging Elastic Load on the cloud, while incorporating both client-level DP (*i.e.*, LDP) and SMPC techniques. However, it does require the usage of *Mindspore*’s machine learning framework.

For an in-depth comparative analysis across the described FL frameworks, refer to [102]. Additionally, extensive comparisons of peer-to-peer FL frameworks have been conducted in [15].

⁵FedScale is the only framework that comes with an off-the-shelf mobility dataset benchmark.

3 Applications

Mobility Prediction Mobility prediction can be defined as algorithms and techniques to estimate the future locations of users. Predicting the next location of users can help with a range of applications including networking (e.g., handover management), pandemic management (e.g., contact tracing), etc. This type of prediction is performed on individual users’ traces where the historical trend of the user’s visited locations can help in predicting the likelihood of their next location.

Transportation With the rising availability of transportation data collected from various sensors like road cameras, GPS probes, and IoT devices, there is an enormous opportunity for city planners to leverage these types of data to facilitate various tasks such as traffic flow prediction. Different from *trajectory* data that records a sequence of locations and time in each trip, crowd flow data only have the start and end locations of a trip, and how many people flow in and out of a particular region can be counted. Indeed, traffic flow prediction using spatial-temporal data has been one of the main focuses of the research community (See comprehensive survey [77]). In the context of transportation, this problem is often considered as forecasting which is to predict traffic speed or traffic flow of regions or road segments based on historical aggregated mobility data.

Community Detection Community detection is an important aspect of urban planning as it allows researchers and planners to identify patterns and trends in human movement. By identifying groups of individuals or locations that are highly connected, researchers and planners can gain insight into how people move through a city, which can inform the design of transportation systems and urban spaces. Additionally, community detection can help identify areas of a city that are at risk of overcrowding or under-utilization, allowing for proactive measures to be taken to address these issues. The underlying enabler of identifying urban communities [54, 47] is spatial temporal data that presents the amount of time spent in different parts of the city. Researchers have shown that human mobility exhibits a strong degree of non-linearity [32] and models that rely on non-linear clustering algorithms, such as the one proposed by Ferreira et al. [46], to detect urban communities have been shown to outperform traditional approaches such as Principal Component Analysis (PCA), and Model-Based Clustering (MBC) and DB-SCAN techniques on a variety of centralized geospatial traces.

Location Based Social Networks LBSNs such as Foursquare and Flickr are social networks that use GPS features to locate the users and let the users broadcast their locations and other content from their mobile devices. LBSNs do not merely mean that the locations are added to the user-generated content (UGC) in social networks so that people can share their location information but also reshape the social structure among individual users that are connected by both their locations in the physical world and their location-tagged social media content in the virtual world. LBSNs contain a large number of user check-in data which consists of the instant locations of each user. Such social networks could also be thought of as the underlying application of Location-based recommendation systems.

4 Approaches

In this section we summarize FL-ST approaches in four main categories of i) trajectory predictive approaches which focus on the next-point prediction of user’s trajectories, ii) traffic flow prediction approaches, iii) Top-K Location-based recommendation approaches and iiiii) other approaches. For each approach, we present the evaluation metrics, the datasets used and the various FL strategies. Table 2 provides an overview of the most important works.

4.1 Trajectory Predictive Approaches

Given the trajectory of a user, these approaches aim at predicting the next position. In a centralized setting where the training data containing all users’ trajectories are available, Recurrent Neural Networks (RNN)-based approaches including Long Short Term Memory (LSTM) and Gated Recurrent Unit (GRU) can be broadly applied in dealing with trajectories and predictive tasks [99, 42, 52, 101]. As this type of data (*i.e.*, trajectory) is highly privacy sensitive, one of the challenges that the research community has been focusing on is creating models that can be tuned for privacy and utility, namely, PUTs (Privacy-Utility Trade-off). Models such as [180, 139, 37] leverage the centralized training data to enhance the privacy level of the traces by reducing user’s re-identification risk and at the same time optimizing the utility of the predictions (*i.e.*, higher accuracy of next point predictions).

The challenges of trajectory predictive approaches in FL are different. Firstly, distributed/decentralized learning creates a heterogeneity problem, as different statistical distributions between clients can lead to a phenomenon called *client drift*, where clients’ model updates/gradients have opposing directions, making their aggregate senseless [170]. As stated in the introduction, this challenge can, in principle, be mitigated to some extent [172]. However, the unique properties of

people’s mobility [33] as well as the different factors that dictate the local distributions of mobility data [192, 83, 100] (*e.g.*, preferences, temporal, spatial, urban characteristics ... etc.) render generic techniques ineffective. Second and as a result of the first challenge, creating a global model for predicting the next location of users that works equally well for all the users becomes an extremely challenging task. That is, one must decide to what extent should clients adopt the global model and when to opt-in for a purely personalized model. In this section we review the existing works in this domain, review how they account for the mentioned challenges and more generic ones (See Table 3 for the meta-analysis). Finally, we compare their performances in Table 4 against centralized predictive approaches namely ST-RNN [101], MCARNN [99], DeepMove [42], and VANext [52].

Table 2: CD denotes Cross-Device approaches and CS for cross-silo

Model	Year	Approach	Dataset	Federated Strategy
Trajectory Predictive Approaches				
Fan et al. [40]	2019	Transfer-learning	Private Mobile Phone Traces	CD
PMF [43]	2020		Foursquare [171], DenseGPS [42], Twitter [183]	CD, Attack-resilient
STSAN [94]	2020	ST Attention Layer	Foursquare[171], Twitter [184], Yelp [105]	CD Adaptive Model Fusion
Ezequiel et al.[39]	2022	GRU-Spatial and Flashback	Foursquare [171], Gowalla [27]	CD
STLPF [162]	2022	AutoEncoder with Global/Local attention	Foursquare[171]	CD
Flow Predictive Approaches				
FedGRU [103]	2020	GRU	PeMS [25]	CS, FedAvg
FedTSE [178]	2022	Reinforcement Learning	England Freeway Dataset[36]	CS, FedAvg
FedSTN [177]	2022	GNN	Taxi-NYC [1], Taxi-BJ [75]	CS, Vertical FL
CNFGNN [121]	2021	GNN	PeMS-BAY [98], METR-LA [75]	CS
CTFL [185]	2022	GNN	PeMSD4, PeMSD7	Clustered FL, CS
MVFF [38]	2022	GRU+GNN	Yelp [105], NY-Bike [2]	Vertical FL, CS
Community Detection Approaches				
F-DEC [114]	2021	Deep Embedded Clustering	GeoLife [198]	CD, FedAvg
Other				
EDEN [84]	2021		Privacy Optimization	CD
PREFER [60]	2021		Location Rec Sys	CD
PEPPER [14]	2022		Location Rec Sys	CD, Gossip Learning
MTSSFL [186]	2021		Transport Mode Inference	CD
Fed-DA [181]	2021		Network Traffic	CS
Fed-NTP [147]	2022		Network Traffic	CS

4.1.1 Metrics

Trajectory prediction models are commonly benchmarked on a handful of available datasets using $Acc@K$ metric which is computed as an average of how many times the correct location was within the top-K predicted places (sorted by the model’s predictions). For example, for an $Acc@5$ metric, the target (or actual output) is compared against a vector of the top-5 most probable locations output by the model. If the target is an element of the top-5 vector, the prediction is correct (or true positive). In practice, papers commonly report an Average Percentile Rank (APR).

4.1.2 Federated Trajectory Predictive Approaches

Table 3 offers a concise overview of the challenges addressed by the existing literature in trajectory predictive approaches. In this application domain, the non-iid nature stemming from a substantial number of users, each with unique characteristics, as well as privacy considerations, takes precedence over other challenges. Notably, while the former challenge appears to have been adequately addressed by the majority of the reviewed works, the latter has only been specifically tackled in [43]. In addition, it is noteworthy that most works have not taken into account federated

competitors, preventing a comprehensive evaluation of their impact. This concern is pivotal not only for the current task but for all ST-FL, as elaborated in Section 5.5. Similar apprehensions extend to the robustness against Byzantine behaviors. However, this criterion holds less significance in the context of CD FL compared to the CS FL setting. In the following, we delve into a more comprehensive discussion of these works.

Fan et al. [40] proposed a federated attention-based personalized human mobility prediction. They apply a *few-shot* learning human mobility predictor that makes personalized predictions based on a few records for each user using an attention-based model. Furthermore, they take advantage of pre-training strategies where the predictor is trained on another smaller mobility dataset to accelerate the FL training on devices. However, even with the pre-training and attention-based strategy, the model requires over 1000 rounds of data communication rounds and is not sufficiently robust for the irregular nature of the human movement.

STSAN [94] Li et al. in 2020 proposed a cross-silo personalized next point prediction model named STSAN (Spatial-Temporal Self-Attention Network) which integrates *AMF* (Adaptive Model Fusion Federated Learning) for offering a mixture of a local and global model. The spatial attention layer allows for capturing the user’s preference for geographic location, and temporal attention captures the user’s temporal activity preference. To overcome the non-iid challenge, the AMF function enables the algorithm to learn specific personalization at each aggregation step on the FL server. The approach is evaluated on Foursquare, Twitter, and Yelp datasets. Table 4 reports its performance on the Foursquare dataset against centralized prediction approaches and shows the superior performance of 99% APR and 6% increase in Acc@1 compared to VANext [52].

PMF [43] Feng et al. proposed PMF, a privacy-preserving mobility prediction framework that uses FL to train general mobility models in a privacy-aware manner. In PMF, every participating device trains locally a representation of the global (centralized) model by using only the locally available dataset at each device. The framework also accounts for attack cases in the mobility prediction task and uses a group optimization algorithm on mobile devices to tackle these attacks. In the group optimization procedure, the whole model is divided into the risky group trained with protected data and the secure group trained with normal data. Furthermore, an efficient aggregation strategy based on robust convergence and an effective polling schema for fair client selection in the centralized server. The results of this model on DenseGPS and Twitter dataset show similar top-1 performance to DeepMove [42] but on Foursquare Tokyo it performs poorer than centralized baselines and other FL approaches.

STLPP [162] Wang et al. proposed a spatial-temporal location prediction framework (STLPP) where the next point prediction algorithm is trained based on a self-attention layer that enables information to be learned between long sequences in both local and global models. Furthermore, as part of the framework, the authors propose an approach that enables clients to cooperatively train their models in the absence of a global model. Their evaluation shows marginal improvement in APR on the Foursquare dataset when compared to the centralized approaches, and their approach performs similarly to STSAN [94] (See Table 4).

Ezequiel et al. [39] develop two implementations of GRU-Spatial and Flashback on FL for predicting the next location in human mobility. To the best of our knowledge, they are currently the only work that has worked on baselining these different approaches in an FL framework, namely Flower [18] and measuring the computational complexity of the model. They evaluate their model on Foursquare NY and Gowalla datasets.

Building upon our exploration of trajectory predictions, it is noteworthy to acknowledge a distinct but closely related research direction—mobility mode prediction. In this realm, the focus shifts towards understanding and dissecting individuals’ personal transportation modes, a facet that holds paramount importance in the efficient management of public transportation systems. Particularly in bustling metropolis cities marked by substantial populations and

Table 3: Summary of the challenges tackled by the reviewed Trajectory Predictive works. Observe that current works do not consider byzantine resilience.

	Privacy	Byzantine Resilience	Non-IIDness	Resource constraints	Overhead assessment	Federated competitors	Online Learning
Fan et al. [40]	✗	✗	✓	✗	✓	✗	✓
STSAN [94]	✗	✗	✓	✗	✗	✗	✗
PMF [43]	✓	✗	✓	✗	✓	✗	✗
Ezequiel et al. [39]	✗	✗	✗	✗	✗	✗	✗
STLPP [162]	✗	✗	✓	✗	✗	✓	✗

Table 4: Baselines for next location prediction models

	Foursquare NY		Foursquare Tokyo	
	Acc@1	APR	Acc@1	APR
Centralized Baselines				
ST-RNN [101]	0.2633	0.9431	0.2567	0.9536
MCARNN [99]	0.3167	0.9595	0.2770	0.9532
DeepMOVE [42]	0.3010	0.9221	0.2668	0.9257
VaNext [52]	0.3627	0.9792	0.3436	0.9735
Federated Approaches				
STSAN [96]	0.4297	0.9902	0.3906	0.9847
STLPF [162]	0.4067	0.9893	0.3887	0.9856
PMF [43]	0.2407	0.6182	0.2310	0.5842
Ezequiel [39]	0.1133	NA	NA	NA

heavy traffic volumes [112]. Decoding users’ traffic modes and aggregating traffic patterns on a broader scale can profoundly influence decision-making processes. For instance, it provides valuable insights for adaptively allocating public transportation resources, and strategically assigning more support during peak times or on congested streets to alleviate traffic burdens.

Within this avenue, the work presented by [174] takes center stage. They propose a federated VGG-like model expressly designed for predicting users’ transportation modes, utilizing the GeoLife dataset. This distinctive approach contributes to the evolving landscape of mobility mode prediction, offering insights that extend beyond trajectory predictions and bear significance in refining resource allocation strategies and bolstering the overall management of public transportation systems. Mensah et al. [123] extends this setting by considering three different architectures (GRU, LSTM and 1D Convolutional NN (CNN)) trained simultaneously, obtaining three global models after convergence. These global models’ outputs are combined through an ensemble learning technique (*i.e.*, stacking) with a Multi Layer Perceptron (MLP) playing the role of the meta-learner. The objective of this MLP is to find the most confident label. (*i.e.*, majority vote). The approach is tested against a federated version of each of the three architectures, showing clear improvement.

4.1.3 Datasets

The following mobility datasets are used to evaluate the trajectory predictive approaches in the research community:

- **Foursquare [171]:** It comprises two sub-datasets, Tokyo and New York data. The Tokyo dataset contains 0.5 million check-ins in Tokyo while the New York one contains over 0.2 million, both collected over a span of about 10 months (from 12 April 2012 to 16 February 2013). Each check-in includes an anonymized user ID, timestamp, and location information, e.g., GPS coordinates and semantic meaning (represented by fine-grained venue categories).
- **Twitter [184]:** It contains around 1.1 million geo-tagged tweets from Los Angeles. These tweets are collected from 1 August 2014 to 30 November 2014. Every geo-tagged tweet consists of four parts, e.g., an anonymized user ID, location information (GPS coordinates), timestamp, and the message published by the user. Compared with the other two platforms, Twitter data is very sparse when location service is not a frequently-used function for Twitter users.
- **Gowalla [27]:** It consists of two parts: a check-in dataset and a friendship network dataset. The check-in dataset contains over 6.4 million check-ins contributed by more than 196,000 users, collected over the period of February 2009 to October 2010. Similarly to Foursquare, each check-in includes an anonymized user ID, latitude and longitude coordinates, a timestamp, and a location ID. As for the friendship network dataset, it contains information about social relationships between users, represented by over 95,000 undirected edges.
- **Brightkite [27]:** This dataset is similar to Gowalla in that it includes check-in data and a friendship network. However, this dataset is significantly sparser, as check-ins were deliberately shared by users, leading to a sparser dataset. Quantitatively, the dataset includes nearly 4.5 million check-ins and 58,228 users, collected between April 2008 and October 2010.
- **Weeplaces [188]:** This dataset has been sourced from Weeplaces, a website that provides visual representations of users’ check-in activities on location-based social networks (LBSNs). The platform has been integrated with other location-based social networking services, such as Facebook Places, Foursquare, and Gowalla, through APIs. This dataset contains 7,658,368 check-ins from 15,799 users, as well as their friends present on Weeplaces.

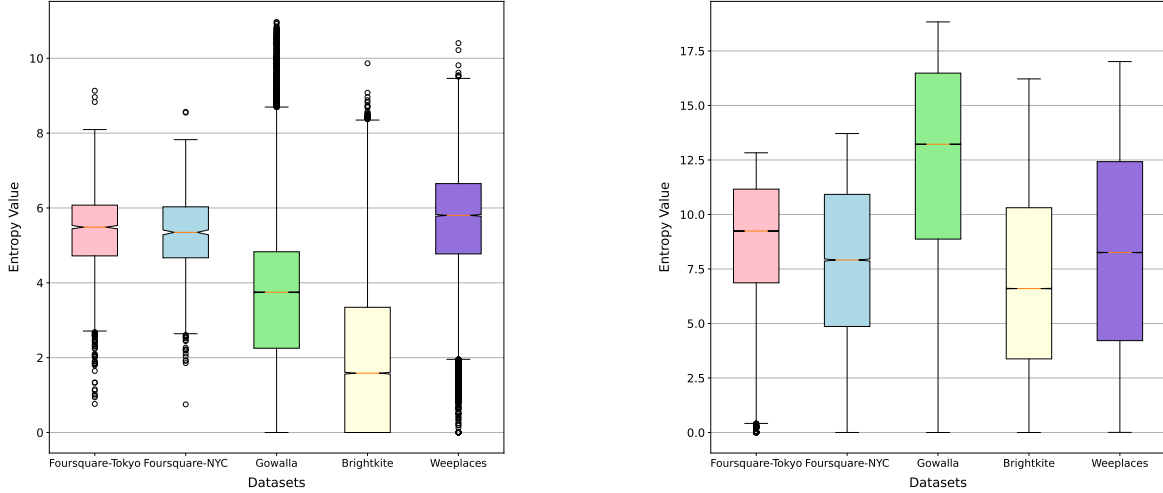
(a) Horizontal entropy H_u across different check-in datasets.(b) Vertical entropy H_l across different check-in datasets.

Figure 3: Entropy Experiments.

Table 5: Global Statistics across check-in datasets.

	N_u	N_l	N_c	\tilde{N}_u	\tilde{N}_l
Foursquare-Tokyo	2293	61858	573703	3.43	92.44
Foursquare-NYC	1083	38333	227428	2.37	84.05
Gowalla	107092	1280969	6442892	3.11	37.18
Brightkite	50686	772780	4491080	1.39	21.16
Weeplaces	15799	971308	7658368	2.72	166.64

- **GeoLife [198]:** It is a GPS trajectory dataset that was collected from 182 users over a span of five years (from April 2007 to August 2012). This dataset comprises 18,670 trajectories, each represented by a sequence of time-stamped points containing latitude, longitude, and altitude information. These trajectories capture a diverse range of users' outdoor movements, including routine activities (e.g., going home), as well as leisure activities (e.g., shopping, hiking, and cycling). Recently, 69 out of the 182 users have labeled their trajectories with transportation modes, such as driving, taking a bus, riding a bike, and walking. The labels for transportation mode are stored in a separate file for each user's folder.
- **Yelp [105]:** It is a collection of businesses, reviews, and user data extracted from the Yelp platform. This dataset is regularly updated and contains almost 7 million reviews of over 150,000 businesses located in 11 metropolitan areas across the United States and Canada. The data includes information on individual users such as their name, the number and nature of their reviews, and their list of friends. Additionally, the dataset also includes check-ins, which provide information about the frequency and duration of customer visits to businesses.
- **Priva'Mov [16]:** It comprises data collected from multifarious sensors, including WIFI, GPS, and Cellular, and contains around 286.7 million records, where each record is a timestamped trajectory point containing latitude, longitude, userID. It was collected from 100 users over a period of 15 months and primarily focuses on urban mobility around Lyon city of France.
- **DenseGPS [42]:** It is a dataset that includes private data from a major mobile application provider in China with 5000 users with one-month dense location records. This dataset is not available for the research community to use.

4.1.4 Data sparsity and heterogeneity in trajectory prediction tasks

Mobility literature defines the highest potential accuracy of predictability of any individual, termed as “maximum predictability” (Π_{max}) [108]. Maximum predictability is defined by the entropy of information of a person’s trajectory (frequency, sequence of location visits, etc.). We adopt this measure to compute the non-iid property of the mobility traces. We use Shannon’s Entropy $H(x)$ to get a sense of both the sparsity and non-iidness of several check-in datasets. In this process, we adjust the metric proposed by [145] to extend the individual entropy of users (i.e., horizontal) with an individual entropy of point-of-interests (i.e., vertical). We argue that having these two dimensions for the entropy is necessary to draw conclusions w.r.t to sparsity and heterogeneity. To quantify this relationship, we first measure the horizontal entropy in the following manner:

$$H_u(x) = - \sum_{i=1}^n P_u(x_i) \log_2[P_u(x_i)] \quad (2)$$

Where n is the number of POIs and is the size of the probability vector. $P_u(x_i)$ is the probability of an individual user u visiting location x_i considering exclusively spatial pattern.

We compute the vertical entropy as follows:

$$H_l(x) = - \sum_{u \in U} P_u(l) \log_2[P_u(l)] \quad (3)$$

Where U is the set of all users. $P(x)$ is the probability of user u visiting location l .

In Figure 3, a comparison of entropy levels is shown for different check-in datasets. Meanwhile, Table 5 provides a summary of more conventional statistics about these datasets. More specifically, N_u denotes the number of users, N_l denotes the number of POIs, N_c denotes the number of check-ins, \tilde{N}_u denotes the average number of users that visited each point of interest and \tilde{N}_l denotes the average number of POIs visited by each user. Figure 3 seems to indicate that Brightkite has the least entropy, as expected. In fact, Subfigure 3a even shows that a significant proportion of Brightkite users have zero entropy, indicating that some users have only one check-in. This correlates with the high sparsity levels of the dataset (refer to Table 5). This downward trend is also evident when viewed through the lens of vertical entropy, although it may be less pronounced. Considering these observations, we see this dataset as an appropriate option to tackle the sparsity problem in distributed learning on mobility data. Conversely, Weeplaces exhibits two important trends: Firstly, there is a considerably higher level of horizontal entropy compared to what is seen on Brightkite, without a corresponding increase in vertical entropy. This suggests that in general users are more active and less predictable, but have not explored a much wider spectrum of points-of-interest. Secondly, Figure 3a exhibits a significant number of user outliers, particularly those with low entropy, indicating the presence of a short tail. These observations, correlated with the results of the table 5, which shows a high N_l and a quite low N_u , indicate low heterogeneity levels within this dataset. Oppositely, Gowalla exhibits more horizontal entropy, which indicates that users are more active globally while also having a higher vertical entropy, indicating that users have substantially less in common than in WeePlaces. These two observations suggest a more heterogeneous, yet less sparse dataset. This makes Gowalla a perfect candidate for works tackling heterogeneity. However, the presence of a significant proportion of highly mobile outlier users should also be noted. Finally, Foursquare datasets show good entropy levels on both axes, making them an excellent choice for sanity tests.

4.2 Traffic Flow Prediction Approaches

Traffic Flow Prediction (TFP) is a problem with multifarious applications. For instance, it allows to mitigate traffic, evaluate air pollution, estimate travel time, and improve driving experience. Approaches in predicting the traffic patterns range from parametric approaches, which are based on statistical metrics (e.g., Auto-Regressive Integrated Moving Average (ARIMA) [56]), to more advanced ML models such as those based on RNNs and GNNs. For a comprehensive survey of deep learning models for traffic predictions in centralized settings refer to [77]. TFP approaches require slightly different settings in FL than trajectory prediction. Unlike individual predictive approaches where each client is assumed to correspond to an individual with their mobility traces, in flow predictive approaches, the clients are often considered to be entities or organizations that are maintaining their flow data private to comply with privacy regulations. Considering this, CS FL appears as a promising paradigm for ensuring that records are not shared outside of each organization/sensor. Moreover, as the number of clients is smaller, and the datasets sizes per client is greater, this configuration alleviates some of the non-iidness concerns. However, this paradigm shift also ushers in two pivotal challenges: First, as mentioned by Li et al. [97], flow conditions can be influenced by many random variables (e.g., daily schedules, weather, accidents), which makes it highly dynamic. As such, TFP demands real-time adaptation. Specifically, the FL process cannot afford to wait until a significant data pool accumulates. Implementing online

learning strategies becomes imperative to continuously update the model with incoming data streams. Second, the units responsible of collecting the data (*e.g.*, sensors) often grapple with limited computational and spatial resources [73]. As such, considering efficiency when designing solutions becomes crucial to ensure their practicality and adaptability to real world applications. In the following, we categorize TFP works into two groups, depending on the types of models employed : grid-based approaches relying on temporal models (*e.g.*, RNNs, LSTM, etc.) and graph-based approaches employing GNNs. We first introduce the metrics, then compare the literature on a high level before diving deeper into each work. Finally, We describe the datasets used for this task.

4.2.1 Metrics

The forecasting performance for traffic flow predictions is commonly measured as a mean absolute error (MAE), mean square error (MSE), root MSE (RMSE), and mean absolute percentage error (MAPE).

4.2.2 Grid based Approaches

In grid-based approaches, the input data into the model is a sequence of flow data per location over time (*e.g.*, average speed, number of vehicles). These approaches mainly focus on deep neural models such as RNN, GRU and LSTM to capture past historical traffic information as a predictor for future instances. Table 6 provides a comprehensive summary of the challenges addressed by the current body of literature in this category. Notably, our assessment reveals a notable gap in addressing the resource limitations inherent in sensor networks, encompassing constraints related to memory and computational capabilities. Moreover, with the exception of two notable works [119, 178], the overhead imposed by the proposed solutions remains largely unquantified.

Furthermore, the current array of works largely overlooks the critical aspect of robustness—neglecting to account for potential vulnerabilities wherein sensors may transmit compromised model updates due to malicious interference or system faults.

It is also imperative to highlight that, apart from the noteworthy contribution by Sepasgozar and Pierre [147], there remains a lack of comparative analyses with federated competitors beyond standalone FL—an aspect that deserves more attention in this field. A more comprehensive discussion of these shortcomings is expounded upon in Section 5.

Table 6: Summary of the challenges tackled by the reviewed Grid-Based works. The main takeaway is the absence of solutions tackling the resource constraints concerns.

	Privacy	Byzantine Resilience	Non-IIDness	Resource constraints	Overhead assessment	Federated competitors	Online Learning
FedGRU[106]	✗	✗	✓	✗	✗	✗	✗
Akallouch et al. [5]	✓	✗	✗	✗	✗	✗	✗
FedLSTM [153]	✓	✓	✗	✗	✗	✗	✗
BFRT[119]	✗	✓	✗	✗	✓	✗	✓
FedTSE[178]	✗	✗	✓	✗	✓	✗	✓
Fed-NTP[147]	✗	✗	✗	✗	✗	✓	✗
Zeng et al. [179]	✗	✗	✓	✗	✗	✗	✓

FedGRU [106] In their work, Liu et al. [106] introduced a federated learning algorithm for highway flow prediction, leveraging a GRU. Additionally, they extended the client sampling phase of federated learning through a joint-announcement protocol. This protocol allows willing clients to announce their participation, strategically aiming to reduce communication costs associated with the FedAvg aggregation step. To address the challenges posed by non-iid data, especially in the context of highly heterogeneous locations, the authors suggested clustering clients based on their location information into k clusters, utilizing a constrained K-means approach [21]. This results in a model per cluster, and ensemble learning is then applied to identify the optimal subset of global models, enhancing prediction accuracy.

Their methodology was rigorously tested using real-world data from the Caltrans Performance Measurement System (PeMS) [25], comprising 39,000 individual sensors monitoring the freeway system in real-time across major metropolitan areas of California. The results demonstrated comparable performance to the centralized baseline of the GRU model and various other centralized models.

Akallouch et al. [5] in this work, the authors address privacy concerns in the context of TFP. They tackle this issue by introducing LDP to provide theoretically provable privacy guarantees. Their approach involves training an LSTM model in a federated fashion, with clients transmitting noisy gradients where the noise is sampled from a Gaussian

distribution. Despite the common expectation of a performance drop with such privacy-preserving mechanisms, their experimental results defy this trend, revealing competitive performance comparable to centralized baselines.

FedLSTM [153] the authors strive to introduce both privacy and resilience guarantees into the context of TFP, operating under the distinctive assumption of computationally capable Roadside Units (RSUs) alongside traditional FL clients, represented by individual vehicles. In this scenario, the authors devise a protocol wherein clients engage with RSUs during each FL round, transmitting their model updates. These RSUs, equipped with uniform test sets, evaluate and validate the models through consensus mechanisms involving miners. The validated models are subsequently published on the blockchain, with local aggregation by clients occurring at a later stage.

To enhance privacy, the authors adopt a strategy akin to [5], leveraging LDP. However, they employ a distinct mechanism involving Laplacian noise, introduced to the count of vehicles in specific areas. The efficacy of this protocol is assessed using an LSTM and compared against various centralized baselines. Notably, the evaluation extends to scenarios with malicious users introducing poisoned models, showcasing the protocol’s resilience due to the presence of verifiers. This work contributes a unique perspective by considering both privacy and resilience aspects in TFP.

BFRT [119] In this contribution, the authors also employ blockchain technology to enhance accountability and verification aspects within FL. Specifically, they utilize a permissioned blockchain, implemented using Hyperledger Fabric [9], to establish a framework where clients can only submit their model updates with permission from a designated group of verifiers (referred to as peers). Acceptance is contingent upon approval by a group of orderers through a consensus algorithm. While sharing a conceptual similarity with [153], the distinguishing factor lies in the unique assumptions of this work. Here, the peers and orderers are presumed to be under the ownership of an institution, such as a service provider or government, with RSUs serving the role of clients in this setup.

Notably, the authors do not explicitly outline the criteria used by the verifiers (peers) to validate models. Evaluation of this protocol is conducted on a Delaware Department of Transportation (DelDOT) dataset, encompassing both GRU and LSTM architectures. A key distinction of this work, in contrast to previous ones [82, 106, 5, 153], is its focus on real-time TFP, even within the experimental settings, as opposed to offline learning scenarios. This real-time consideration adds a practical dimension to their exploration of FL in the context of TFP.

FedTSE [178] In FedTSE, authors proposed a framework for Travel State Estimation (TSE). They consider an LSTM model as the local training model for joint prediction of vehicular speed and traffic flow. A unique characteristic of FedTSE is that it relies on a deep reinforcement learning (DRL)-based algorithm to adjust model parameter uploading/downloading decisions such that it improves the estimation accuracy of local models and balances the trade-off between computation and communication cost. They evaluate their approach on the England Freeway Dataset which includes flow and speed for the entire year of 2014. They also consider three different aggregation strategies corresponding to a synchronous aggregation with the same number of epochs per client, namely, *FedTSE-Syn*; an asynchronous aggregation with different numbers of epochs, namely, *FedTSE-Asyn* and a weighted version of FedTSE-Asyn where clients that have trained more are given more weight.

Fed-NTP [147] proposed an LSTM model implemented in FL to predict network traffic based on the most influential features of network traffic flow in the Vehicular Ad-Hoc Network (VANET). Even though this work is *not* aiming in predicting vehicle traffic and focuses on network flow, it has some relevance with the existing models such as FedGRU [106]. Fed-NTP shows to outperform FedGRU on the same V2V dataset. Other similar works that have focused on network traffic prediction such as FedDA [181] exist in the literature. However, as they solely focus on network traffic problems and are not intended for mobility applications, we do not review them in the survey. We encourage the reader to refer to [78] for a full survey of network traffic prediction techniques in FL.

Zeng et al. [179] To our knowledge, this work stands as the sole work introducing a multi-task learning method for the TFP challenge within the FL framework. The motivation behind this approach is to craft personalized models capable of adjusting to the diverse forms of heterogeneity previously outlined (See Section 2). The authors’ methodology begins with a hierarchical clustering of clients’ (*i.e.*, sensors) local data based on various criteria such as weather conditions, time, and special events. This clustering ensures the presence of consistent clusters across clients. Subsequently, an LSTM model is trained in a federated manner for each cluster to capture temporal features and address time travel prediction. This process yields a dynamic graph delineating sub-segments connecting different stations, along with their respective travel times—this graph evolves over time. Interestingly, the authors propose a modification of the A^* algorithm to determine the optimal route utilizing this time-dependent graph. The experimental comparison with Google Maps’ route prediction demonstrates the advantages of this approach. Nonetheless, the paper lacks explicit details on the clustering criteria employed and the mechanisms for client consensus regarding these clusters.

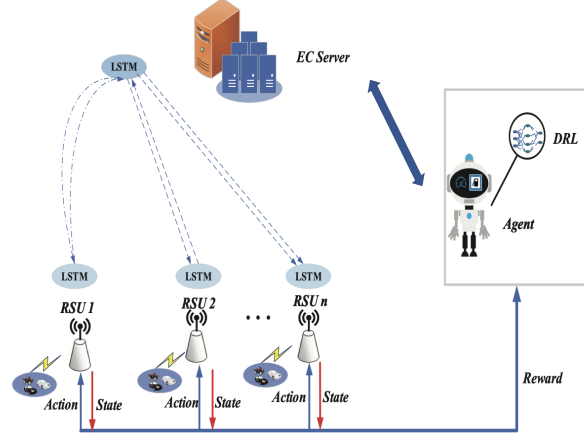


Figure 4: Figure from FedTSE [178] showing the interactions between Edge Computing (EC) Server and Road Side Units (RSU) acting as a cross-silo unit. The LSTM model weights are aggregated using FedAvg, and a Deep Reinforcement Learning (DRL) Agent to maximize reward.

Table 7: Number of parameters and performance of baseline centralized models on PEMS7M [173] dataset.

Model	Year	Number of Parameters	Performance (MAPE%)
STGCN [173]	2017	330K	5.02%
MTGNN [166]	2020	433K	5.02%
DCRNN [98]	2017	610K	5.33%

4.2.3 GNN Based Approaches

State-of-the-art multi-layer GNNs effectively address the spatial-temporal nature of traffic prediction in centralized settings. However, when applied in FL scenarios, GNN-based approaches encounter specific challenges. Notably, vertical FL poses a hurdle where each silo possesses a partial, overlapping view of the graph.

Motivated by intelligent transportation systems, such as RSUs, individual units exhibit unique traffic patterns across a city (See Figure 4). The primary challenge in using GNN models for traffic forecasting within FL arises from their abundance of parameters, contrasting with simpler models like ARIMA with significantly fewer parameters (on the order of hundreds). Table 7 shows the magnitude of parameters for baseline traffic prediction models over the PEMS7M dataset [173].

Moreover, the dynamic nature of these graphs over time demands adaptation through online learning, posing another significant challenge. Existing works aim to address these issues, as summarized in Table 8. While most studies tackle heterogeneity challenges, considerations regarding the cost of training and storing these complex models are rarely highlighted. Notably, [194] stands out by addressing these concerns, delegating the bulk of learning to a centralized server while clients handle and store manageable parameter subsets.

Furthermore, similar to grid-based approaches, the current literature lacks exploration into the models' robustness against faults and attacks that could compromise the learning process.

Table 8: Summary of the challenges tackled by the reviewed GNN-based works. Similarly to grid-based approaches, we note the absence of considerations w.r.t byzantine resilience.

	Privacy	Byzantine Resilience	Non-IIDness	Resource constraints	Overhead assessment	Federated competitors	Online Learning
Fed-STGRU[82]	✗	✗	✓	✗	✗	✗	✗
FedAGCN[136]	✗	✗	✓	✗	✗	✗	✗
FedSTN[177]	✗	✗	✗	✗	✓	✓	✗
CTFL[185]	✗	✗	✓	✗	✓	✗	✗
GOF-TTE[194]	✓	✗	✓	✓	✗	✓	✓

DST-GCN [160] The authors parallel FedGRU [106] by adopting a GRU model to predict temporal dependencies in TFP. However, they extend this approach by integrating a Graph Convolutional Network (GCN) to capture spatial features, recognizing the GCN’s aptitude for modeling inter-spatial characteristics. Augmenting this fusion, an attention mechanism is employed to identify and elevate critical spatio-temporal dependencies that dynamically evolve. Notably, the study promotes a client participation scheme based on both willingness and the evaluation of clients’ models prior to aggregation. Their experimentation with the PeMS dataset substantiates that this client selection strategy significantly enhances performance compared to traditional FL.

Fed-STGRU [82] Similar to DST-GCN [160], this work propose a combination of a GRU and a GCN to solve TFP. Moreover, the authors also focus on improving the clients selection process of FL through a clustering strategy, where they opt for the Fuzzy c-means algorithm. This alternative is highlighted as less computationally demanding compared to the constrained k-means proposed in [106].

In terms of performance, this work demonstrates superior results compared to a standalone FedAvg algorithm, showcasing the effectiveness of their hybrid GRU-GCN model and the efficiency of Fuzzy k-means for client clustering in the context of traffic flow prediction.

FedAGCN [136] In this work, the authors propose a comprehensive solution to address the efficiency of GCNs in handling spatial correlations, a critical aspect of traffic flow data. Previous methods, particularly those not based on graph models, rely on individual clients learning on their local view of the system, leading to a neglect of spatial associations of the traffic network’s topology. To overcome this limitation, the authors introduce GraphFed, a carefully devised algorithm aimed at enhancing communication efficiency while preserving the ability to capture spatial correlations. The key innovation lies in the division of the graph network into sub-graphs, with representative clients randomly selected for each sub-graph. These representatives aggregate data from all clients within their respective sub-graphs, facilitating the classification of the model’s parameters into global and local categories, the latter being tailored specifically to spatial relationships.

Importantly, the practicality of GraphFed is evident as it mitigates the need to store the entire traffic network at the client level, a cumbersome requirement for large traffic networks. By operating on disjoint regions of the traffic network, representative clients efficiently avoid the necessity of sharing local parameters with the FL server, ensuring that only global parameters are learned in a federated manner, thereby minimizing communication overhead. The authors also adopt the ADGCN model [137], enhancing its efficiency by refining the convolution operation. This meticulous approach guarantees the adeptness of the model in addressing spatial correlations and overcoming challenges associated with FL on a broad scale within traffic networks.

FedSTN [177] The FedSTN approach as formalized by Yuan et. al is a newly proposed solution to solving the TFP problem. To accomplish this task, authors proposed a Graph-based Representation Learning for CS FL using three main modules: a Recurrent Long-term capture Network (RLCN) module, an Attentive Mechanism Federated Network (AMFN) Module, and a Semantic Capture Network module (SCN). The RLCN module is responsible for learning long-term traffic behaviors as geometric time series data of fixed long-term interval p . The data as input into this module is comprised of inflow-outflow values for certain grid spaces, as initially computed for the data.

The AFMN module is responsible for learning short-term spatial-temporal features in a privacy-preserving manner. This module includes long-term contextual data such as meteorology and federated graph attention (FedGAT). Lastly, there is the SCN module, which takes into account Point of Interest Feature Components (POI) and non-euclidean connection relationships. Points of interest and their interactions have significant effects on TFP but are not incorporated into the raw time series prediction data. Furthermore, transportation networks also have methods of connection outside adjacent-grid connections, such as trains, highways, etc. These flows also have an important effect on TFP that is not incorporated elsewhere, so this module is to address this issue. The output of all these modules is then connected via an FC layer followed by a Tanh activation function for the final output.

CNFGNN [121] Meng et al. proposed a new architecture named Cross Node Federated Graph Neural Network (CNFGNN) to predict the flow. CNFGNN works by decomposing the problem into two stages: first, it uses an encoder-decoder network to extract temporal features locally (See right side of Figure 5), and a then a GNN to capture spatial relations across devices, as is illustrated in the left side of Figure 5. On each training step, the server processes a temporal encoding update, a partial gradient update, and an on-node graph embedding update. Each iteration updates the client side weights and then ships the model, hidden layer, and gradients back to the server.

One large issue of the proposed approach is the communication overhead in the training stage, with split learning requiring the global model to fetch all hidden states from each node, and ship gradients of node embeddings to each node. Then, it must receive gradients of node embedding and send gradients of hidden states in the back-propagation

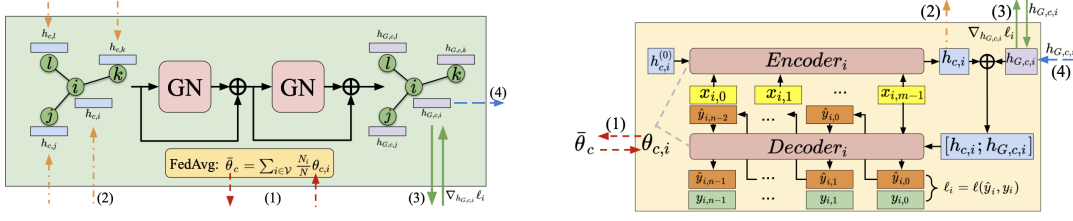


Figure 5: Figure from CNGNN [121] presents (a) the server-side graph neural network with a systematic overview of training steps: (1) Federated learning of on-node models. (2) Temporal encoding update. (3) Split Learning of GN. (b) Client i Auto-encoder architecture.

step. To mitigate this, they propose an alternative training approach in which the temporal encoder-decoder and the node embeddings are trained separately. First, the node embeddings are fixed and optimization is performed on the encoder-decoder. Then, after a fixed interval, the global model is updated fixing the node-level models. This drastically reduces communication overhead. The FL local models and the GNN model with only a local objective function. They perform alternating optimization to update clients' model weights with GNN model weights fixed and then update GNN model weights with the FL local model weights fixed, over multiple rounds.

CTFL [185] proposes a Clustering-based hierarchical and Two-step- optimized FL (CTFL), to overcome the large number of parameters that are needed for aggregation in the GNN-based models such as STGCN [173], DCRNN [98] and MTGNN [166]. CTFL employs a divide-and-conquer strategy, clustering clients based on the closeness of their local model parameters. It also accounts for optimization by applying a two-step strategy where the central server uploads only one representative local model update from each cluster, thus reducing the communication overhead associated with model update transmission in the FL.

GOF-TTE [194] This work addresses a different formulation of the TFP problem, namely, Time Travel Estimation (TTE), where the task is to predict the travel time needed to go from a point a to a point b considering spatial features (*e.g.*, road network map) and temporal features (*e.g.*, time of the day). [194] tackles this challenge within the context of taxi-driving scenarios, with potential applications to diverse driving types. In this work, generic spatio-temporal features are used to learn the global state of the network in a federated manner, allowing to obtain a global model. Subsequently, each client fine-tunes this model in locally (*i.e.*, Localized global model), before incorporating a personalized model trained on the clients' profile features.

The input architecture involves a dual graph representing road segments and intersections. These elements are translated into latent representations (*i.e.*, embeddings) before being past to a GCN layer, allowing for the capture of spatial dependencies. Spatial and temporal representations are then subjected to a cross-product operation. Temporal aspects, such as day of the week and time of the day, are encoded, with an attention mechanism considering both current and past temporal impacts. This cross product results in a global state of the road network.

In a subsequent phase, the global state is fine-tuned using clients' local data before being obfuscated with Laplacian noise through the DP mechanism. A personalized model is later introduced, comprising a fully connected layer of client profile feature embeddings (*e.g.*, frequently visited regions, average driving distance). This personalized model is integrated with the output of the localized global model to mitigate biases arising from the high heterogeneity of the data.

Notably, this work distinguishes itself from prior research by acknowledging the relationship between road segments and intersections, in contrast to road-segment-based solutions for instance. Furthermore, it introduces a novel personalization technique that accounts for the current global state of the network, enabling real-time personalized predictions—a combination of considerations not concurrently addressed in existing literature.

Finally, it is worth noting that other works similar to flow prediction exists which are focused on identifying travel modality inference (TMI). For instance [186] proposes MTSSFL which trains a deep neural network ensemble under a novel semi-supervised FL framework. It achieves a highly accurate score for a crowd-sourced TMI without depending on the availability of massive labeled data.

4.2.4 Datasets

For flow predictive approaches the following datasets are widely used to benchmark the comparison of various algorithms.

- **Taxi Datasets:** The following datasets contain taxi in-out flow data collected via GPS, which include pick-up time, drop-off time, and trip distance.
 - **TaxiBJ [189]** contains data collected in Beijing from four time intervals: 2013/7/1 to 2013/10/30, 2014/3/1 to 2014/6/30, 2015/3/1 to 2015/6/30, and 2015/11/1 to 2016/4/10. The data was collected at 30-minute intervals and include trajectories of over 34,000 taxis. It additionally contains meteorology data such as weather conditions, temperature, and wind speed.
 - **TaxiNYC [1]** contains data collected in New York City, which was collected between 2016/4/1 and 2016/6/30 and contains over 35 million records. A more extensive version of this dataset, collected from 2009 to 2022, is also available.
 - **TaxiPorto [120]** comprises a full year (from 2013/7/1 to 2014/6/30) of trajectories for all the 442 taxis running in the city of Porto, Portugal. It also contains information on the type of taxi call: central-based, stand-based, and demanded on a random street.
 - **T-Drive [197]** contains GPS traces collected in Beijing, China. The dataset was collected over a period of three years from 2008 to 2011, and it consists of over 10,000 taxi drivers' GPS trajectories. The dataset contains a total of 1.07 billion GPS points, covering approximately 150 million kilometers. That being said, the sample released by Microsoft is over a span of one week only, containing around 15 million GPS points and covering a total distance of trajectories that reaches 9 million kilometers.
- **METR-LA [75]:** is Los Angeles traffic collected using 207 sensors mounted around highways, and 1515 edges from 2012/3/1 to 2012/6/30.
- **PeMS** is a traffic flow dataset collected from California Transportation Agencies Performance Measurement System (PeMS).
 - **PeMS-BAY [98]:** It contains 325 nodes (traffic sensors) and 2369 edges in the Bay Area from 2017/1/1 to 2017/5/31.
 - **PeMSD7M:** It is a sub-sample of PeMS published as part of [173], also collected from PeMS. It covers 228 traffic sensors with a 5-minute sampling rate corresponding to 2012/5/1 to 2012/6/30.
- **NY-Bike [2]:** Spanning a period of ten years, from June 2013 to January 2023, this dataset includes comprehensive information about daily bike orders by people in New York City and is regularly updated. More specifically, it contains information about bike trips, such as duration, starting and ending point, and location, as well as details about bikers, including user type, gender, and year of birth.
- **Highways England network journey [35]:** This dataset consists of average journey time, speed and traffic flow information for 15-minute periods since April 2015 on all motorways and 'A' roads managed by Highways England.
- **DelDOT dataset [119]:** represents roadway traffic data collected by the Delaware Department of Transportation at a 5-minute time frequency.

4.3 Top-K Location-Based Recommendation Approaches

Location-based recommendation task, while resembling trajectory prediction, often entails lower demands. This task involve identifying a top-K set of relevant Points of Interest (POIs) for a user based on their historical POIs. In Table 9, we summarize the challenges addressed by the community for this use case. Notably, none of the reviewed works have considered an online learning setting. It can be argued that this task may not necessarily require online learning since the preferences of users might not significantly change over time, unlike the trajectory prediction task, which inherently encompasses a natural time dimension. However, we argue that exploring online learning could still be valuable. Moreover, recent research has demonstrated that users' preferences tend to evolve over time. While other challenges have been reasonably explored, it's noteworthy that several works have based their privacy solutions on not sharing all model parameters, but recent studies on similar models have shown this approach not to be entirely private [191, 176]. This is indicated in the table using the " \approx " symbol. Finally, similar to previous approaches/tasks, the exploration of robustness has been lacking. This is unfortunate, given that issues like shilling attacks [71] are significant concerns for recommendation systems and represent just one type of robustness concern not explored in this federated setting. Other types include malicious users fabricating fake profiles to manipulate recommendations, promote specific POIs, and introduce unfairness.

Given that the evaluation metrics and datasets align with those used in trajectory prediction approaches, the rest of this section focuses solely on the main works in this domain.

Table 9: Summary of the challenges tackled by the reviewed Top-K POI recommendation works. Byzantine resilience and online learning have not been tackled by current works.

	Privacy	Byzantine Resilience	Non-IIDness	Resource constraints	Overhead assessment	Federated competitors	Online Learning
FPL [10]	≈	✗	✗	✓	✗	✓	✗
DFL-PC [61]	✗	✗	✓	✗	✗	✗	✗
PREFER [60]	≈	✗	✗	✓	✓	✓	✗
MVFF [38]	≈	✗	✓	✗	✗	✗	✗
PriRec[26]	✓	✗	✓	✓	✗	✗	✗
PEPPER[14]	✗	✗	✓	✗	✓	✓	✗
FedPOIRec[134]	✓	✗	✓	✗	✓	✓	✗

FPL [10] This work extends the bayesian pair wise ranking (BPR) [143] algorithm to the federated setting. The authors train the sensitive user embeddings locally and provide users with the option to share these embeddings with a certain probability, while training the less sensitive parameters in a federated manner. This approach can be viewed as a flexible framework for factorization models, where clients can decide how many parameters to share while maintaining convergence of the model. They evaluate their approach on Foursquare dataset, considering different countries and levels of sparsity and compare themselves with a federated movies recommender [7].

DFL-PC [61] This work focuses on optimizing the aggregation process in federated learning. The authors pre-train deep models consisting of a GCN and a GRU connected to an MLP, which are used to estimate the system parameter space. These estimates are then optimized at the server level through a reinforcement learning algorithm. The evaluation results show that their solution surpasses traditional centralized deep models due to the added optimization step based on reinforcement learning.

In terms of system architecture, alternative solutions have been proposed for more efficient federated POI recommendation. The following works fall within this scope:

PREFER [60] Guo et al. proposed a two-set recommendation model training. First, for privacy purposes, they propose to train the user-dependent model parameters strictly locally. Subsequently, user share and aggregate less sensitive parameters (i.e., user-independent) in a multiple-edge server architecture, instead of remote cloud servers, with an aim to improve real-time response capability and reduce communication cost. They validate their approach both analytically and empirically on two standard POI recommendation models, namely, Distance2Pre [28] and PRME-G [44], and two check-in datasets, Foursquare [171] and Gowalla [27], and show the competitiveness of their approach with centralized and federated competitors.

MVFF [38] proposes a vertical federated learning framework for mobility data forecasting for CS FL applications where each organization holds a partial subset of data. Using a local learning model, each organization extracts the embedded spatio-temporal correlation between its locations. To account for global learning, a global model synchronizes with the local models to incorporate the correlation between all the organizations' locations.

PriRec [26] proposes a peer-to-peer approach to learn sensitive user embeddings, while less sensitive ones (*e.g.*, feature interaction model) are learned in a federated manner. This is achieved through the introduction of secret sharing [148] in the decentralized gradient descent (DGD) topology and considering geographical information when building this topology, that is, users closely geo-located learn (privately) collaboratively their respective user embeddings. As for the items' features, they are aggregated through a secure aggregation protocol (SA) [76]. Authors evaluate their approach on standard POI datasets' and compare it with a centralized factorization machine (FM), which it seems to compete with, while guaranteeing a better level of privacy.

PEPPER [14] emphasizes personalization of POI models through the aggregation step while using a gossip communication protocol to eliminate the central FL server. In this work, nodes gossip their models with their neighbors and aggregate them after evaluating their contribution. The authors also introduce a peer-sampling protocol that acts as a clustering over time, ensuring each node has similar users in its neighborhood. Results from their experiments show that fully decentralized FL can be competitive with centralized solutions, while offering scalability and personalization.

FedPOIRec [134] introduces a framework that leverages social relationships between users to make more personalized models. The authors first train a global model using FL and SMPC to protect the aggregation phase from a curious

server. A trusted third party is then tasked with finding similar users based on encrypted embeddings for personalized aggregation. The encryption is done using a leveled variant of the CKKS fully homomorphic encryption (FHE) scheme. This work also encompasses an adaptation of the setting of [10] and CESAR [154], a sequential recommender based on a convolutional neural network. The authors validate their approach through an evaluation on five versions (cities) of Foursquare as well as a formal privacy analysis.

Finally, another use case that has been explored in the context of location-based recommendation systems is the driver recommendation use case, as tackled by [159]. In this work, cab companies use federated learning to strengthen RSUs with the computational capability to develop an intelligent recommender system that recommends the appropriate driver for a subsequent trip. To this end, they consider both the driver’s stress and past behavioral patterns.

4.4 Other Approaches

4.4.1 Clustering Based Approaches

Most existing works in clustering in FL are focused on methods to identify and self-organize devices into communities are so to conduct model sharing within those communities. In [22], authors introduce and evaluate a hierarchical clustering for vision models, where the local model is a CNN that is trained under in a supervised fashion. The extensive evaluation presents the improvement that hierarchical clustering can bring to federated learning under a non-iid setting where each client holds partitioned data. In [85], the authors propose the dynamic Generative Adversarial Network (GAN)-based clustering in FL to improve the time series forecasting for cases such as cell tower handover prediction. Their proposed approach accounts for the adaptive clusters and non-iid data. IFCA proposed by [55], starts by randomly initializing k models, one per cluster. Each client assigns itself to a cluster at the start of each round of training by evaluating all k models on its local data and choosing the model with the lowest loss to train for m epochs. At the end of each round, the server performs federated averaging within each cluster of clients separately.

Although this theoretical line of work is receiving a great deal of attention from the FL research community we have seen almost no adaptation, except one, to the spatial-temporal models in FL. F-DEC [114] proposed a deep embedded clustering for urban community detection in FL. They expanded on the centralized model proposed by [47] and trained an autoencoder based on heatmap images of mobility trajectories transformed using the frequency of visits (where brighter pixels show more frequently visited areas). They then used a KL divergence loss for clustering similar heatmaps together. Furthermore, this work is the only early evidence found that measures the computational complexity of such algorithms when it is deployed on ordinary smartphones.

4.4.2 Privacy and Attacks in Spatial-Temporal FL Models

FL was initially designed to protect user privacy by sharing model parameters instead of data. However, research has shown that sharing these parameters can still reveal sensitive information, especially in models that use embeddings/latent features to capture user or point-of-interest semantics. To address this issue, researchers have proposed mainly three approaches: i) "a share less" policy, ii) injecting noise using techniques like DP, and iii) using cryptographic primitives (*e.g.*, FHE). Various studies have proposed solutions using each of these approaches, but each approach has its limitations. For instance, [60, 8] proposed sharing only user-independent embeddings to be learned in a federated manner while training user embeddings locally. [43] first proposed a practical attack, demonstrating that user check-ins can be easily inferred by a curious FL server based on POIs embeddings. Later on, they proposed to alleviate this attack by training these embeddings on noisy data generated using DP. They show that if the rest of the network (*i.e.*, non-sensitive layers) is frozen during this noisy training, and is pre-trained on real data, then the performance remains reasonable. Nevertheless, they did not quantify the impact of their attack nor the degree of protection provided by their solution. Another line of work, [134], opts for the third approach and uses SMPC to hide the individual contributions of the users from a curious FL server. Unfortunately, this approach opens the door to malicious users, whose goal would be to corrupt the learning, and who would be difficult to detect due to SMPC.

A more privacy-oriented solution was proposed by [26]. They aggregated less sensitive embeddings using SMPC and categorized sensitive embeddings into two parts: those related to POIs and those related to users. They used Local DP to add noise at the user level to the POI-related embeddings before sharing them with the server. They also proposed to share the user-dependent embeddings in a peer-to-peer fashion using secret sharing. To reduce overhead and the attack surface, they considered geographic information to build the peer neighborhood. Thus, a user shares its sensitive embeddings only with geographically close neighbors. This solution has competitive prediction performances with non-private solutions and undeniably provides more privacy guarantees. However, the impact of the assumption that sensitive embeddings can be safely shared with nearby users is still unclear.

Another type of attack that has received significant attention in recent mobility research is re-identification attacks [111, 50]. The fundamental concept behind such attacks is that a malicious service provider could exploit background knowledge to associate anonymized user traces with their respective owners, thereby compromising users' anonymity. To address this issue, Khalfoun et al. proposed a federated protocol for assessing the risk of re-identification on mobility data. This protocol involves training a re-identification model in a federated manner using users' traces and subsequently utilizing this model to select the optimal combination and hyperparameters of location privacy protection mechanisms (LPPMs) that can protect a user's privacy whenever they transmit their data to an untrusted Mobile Crowd Sensing Server (MCS). Notably, this solution appears to be the only privacy risk assessment mechanism for mobility data that does not necessitate the presence of a trusted curator, to the best of our knowledge. For a full survey of privacy and security techniques in FL see [125].

Based on our established taxonomy, an analysis of relevant studies, and a meta-analysis, we have observed that some key factors in certain tasks (e.g., privacy and non-iidness in recommendations, resource constraints and computational costs in TFP) are often not given sufficient consideration. Instead, other criteria are erroneously deemed more significant. For instance, in TFP, privacy is typically not a primary concern because the data collected from sensors is generally aggregated, however, current works have considered it more often than other criteria. Moreover, there are several still open-problems as well as a lack standardization, both in data and code. In light of these observations, we outline a proposed roadmap in the following section. This roadmap is designed to accelerate progress in FL-ST, with a focus on bridging the gap with real-world applications, which are currently underrepresented in the literature. We believe that this approach will not only streamline research in this field but also facilitate its practical implementation.

5 Discussion and open research challenges

Based on the review of the above papers, we see the following opportunities and roadmap for the research community to explore.

5.1 Semantic Location Embedding and Context-Awareness Modelling

One of the biggest opportunities that we see in continuing research on location and point trajectory predicting, is in regards to integrating more semantic and contextual information about types of places instead of focusing primarily on coordinates. For example, this contextual information can include information on whether a point in trajectory represents someone's workplace, their frequently visited locations, or a potential point of interest in a new town. While this is not limited to FL applications and previous works have often used open source maps to infer information about the type of places (e.g., popularity, socio-economic level), FL can bring a new level of anonymity and personalization to this integration. The semantic representation of the locations can then be learned over time on users' devices, maintaining users' privacy while allowing for better-personalized models.

5.2 Byzantine Resilient Spatio-Temporal Mobility Federated Learning

Within the realm of ML, Byzantine Resilience commonly denotes the ability to train an accurate statistical model amid the presence of arbitrary behaviors, commonly referred to as Byzantine users [59]. These behaviors manifest either due to faults or malicious users, encompassing scenarios like a compromised sensor in the TFP problem or a user aiming to promote specific items in a Top-K POI recommendation use case. In the best-case scenario, the trained model proves unusable, resulting in a cost loss incurred during its training. However, in the worst-case scenario, if the attack goes undetected (e.g., backdoors on GNN [193]), the model might exhibit concealed yet malicious behaviors. This predicament poses a significant challenge owing to the difficulty of interpreting model parameters.

Distinguishing between a Byzantine model update transmitted by a client or sensor and an honest but out-of-distribution client often presents a non-trivial task. This complexity amplifies in highly heterogeneous configurations typical of FL-ST mobility applications. Its pertinence highlights this challenge as a substantial research track for Distributed ML and FL. Surprisingly, this aspect remains relatively unexplored in the context of FL-ST mobility applications. Among the multitude of reviewed papers, only Meese et al. [119] and Tang et al. [153] have delved into such adversarial settings. However, neither of these works has considered the novel solutions proposed in this field, such as bucketing [66] and variance reduction techniques [41]. This gap is more concerning given the pronounced heterogeneity characteristic of these applications.

Consequently, there is an urgent imperative to delve into this challenge and contemplate environments accommodating Byzantine behaviors within FL clients. This pursuit should be seen as a fundamental step to bridge the gap between

research efforts and real world implementations. Specifically, it necessitates a comprehensive study and evaluation of existing Byzantine-resilient distributed learning algorithms, a deep understanding of their limitations, and subsequently adapting them to suit the inherent heterogeneity and dynamic evolution within the spatio-temporal context.

5.3 Communication Efficiency

The race to enhance the performance of machine learning (ML) models is driving exponential growth in their number of parameters across various fields. In the realm of ST mobility models, where the necessity to capture diverse dimensions (spatial, temporal, preferential, and characteristic) is paramount, this growth becomes even more pronounced. A striking illustration of this trend is the substantial increase of over 4000% observed between ARIMA, a parametric model, and recent GNNs, which have parameters in the order of hundreds of thousands (see Table 7).

Traditionally, in classical ML, concerns centered around the costs associated with training these models and the extensive data required for such endeavors. However, the emergence of FL has shifted the focus towards apprehensions about the expenses related to communicating and transmitting these increasingly expansive models. Surprisingly, a notable observation is that the majority of existing works do not adequately consider these challenges within their frameworks. That is, they do not consider both the cost of training as well as the cost of collecting these models by the central server.

Recognizing the significance of this issue, particularly in light of its comprehensive exploration in other fields, we advocate for heightened attention from the research community. We urge future work to delve into existing solutions like pruning [79] and quantization [91], emphasizing the need to quantify the cost-effectiveness of these approaches in the context of FL-ST mobility.

5.4 Trust, Fairness, and Accountability

In addition to trust and accountability, another challenge that we see spatial temporal mobility models will face under a federated setting is fairness. *That is to what extent the models that are trained on location traces are equitable?* Especially models that are designed for the purpose of mobility flow prediction and allocation of transport options. For instance, mobility demand prediction algorithms have been shown to offer higher service quality to neighborhoods with more white people [23]. Indeed, as recent evidence from the broader machine learning domain has shown, the systematic discrimination in making decisions against different groups has been shifted from people to autonomous algorithms [81, 68]. In many applications, discrimination may be defined by different protected attributes, such as race, gender, ethnicity, and religion, that directly prevent favorable outcomes for a minority group in societal resource allocation, education equality, employment opportunity, etc [146]. Measuring fairness of mobility models is a dimension that has been vastly overlooked in applications of spatial-temporal mobility models, with exception of a few works [169, 168, 53] and with little consensus on how fairness should be defined and measured for spatial-temporal applications. One way of controlling for fairness of mobility models under the FL setting is to create auditing systems that can infer information about the training without having access to location data of the devices or the global model at the FL server [115, 113]. We believe future work will focus on dynamic middle-wares that can leverage solutions such as clustered FL to offer interpretability of the underlying models [195, 55, 93] are crucial to transition exiting solutions from research to practice.

5.5 Standardisation and Reproducibility

As the landscape of spatio-temporal mobility research continuously evolves, it is increasingly crucial to establish a common ground for assessing its advancements. This involves not only pinpointing the most effective proposed solutions but also identifying the persistent challenges and shortcomings. Achieving this necessitates standardized datasets, uniform methods for data preprocessing and splitting, and ideally, a code base or metadata ensuring the reproducibility of each solution. Regrettably, similarly to other domains like recommendation systems [30, 31], this practice is seldom followed.

For instance, in studies related to the TFP task, there is a prevalent use of diverse datasets, sometimes from the same source but across disparate timeframes, different space and time discretizing methods as well as unclear distinct splitting methods. These splits often rely on various sampling techniques. Yet, as underlined in [122], the process of partitioning data into training, testing, and potentially validation sets significantly influences the measured performance. Another example, specifically in Top-K POI recommendation, involves the utilization of metrics based on parameters (*e.g.*, varying K values for accuracy at rank K). In FL, this challenge exacerbates due to the proliferation of parameters impacting the results (*e.g.*, the number of clients, sampling strategies, aggregation techniques).

Consequently, robustly comparing different works, even those evaluated on identical datasets, becomes impossible. This is likely why most reviewed works unfortunately lack substantial comparisons with their federated counterparts.

Moreover, considering that the distinct experimental procedures often stem from stochastic processes (client and data point sampling, stochastic learning), we stress the necessity for statistical tests to gauge the statistical significance of observed results.

We acknowledge that research contexts rarely align perfectly. Nevertheless, the need for standardized datasets and preprocessing methodologies in the field is undeniable. Existing surveys [196, 48] have proposed solutions in a centralized context. While the direct applicability of these solutions to FL may require adaptation, they provide valuable insights and methodologies that can serve as foundational building blocks for the FL paradigm. Moreover, there have been several software efforts [175, 62, 57, 132, 127, 19, 12] that furnish a set of libraries for standardized spatio-temporal data analysis, preprocessing, and visualization. Using these as building blocks might represent an essential step towards standardization and reproducibility in the field.

5.6 Realistic Cross-silo Spatial Temporal Datasets for Benchmarking

Existing approaches that we reviewed are mostly evaluated with ST data partitioned artificially. Nevertheless, the long-term development of this field still requires realistic and large-scale federated datasets to be made available to support experimental evaluations under settings close to practical applications. For instance, in the reviewed literature, there is a lack of research on how the geographic distribution of silos can lead to a geographically distinct flow of information. Establishing policy-based scenarios in order to guide how the data should be partitioned across silos to reflect real-world data ownership challenges is a direction that we believe the research community will be working closely with other stakeholders in the future.

5.7 Transition to Real-world Deployment Through Dedicated Frameworks

Finally, we believe that just as crowd-sensing research was successful a decade ago through frameworks such as AWARE [45], which reduced the burden of app development for data collection, frameworks specifically designed for federated mobility models will facilitate the transition from limited research to in-wild deployments. To achieve this transition, it is crucial to i) provide benchmarks for mobility applications and ii) develop mobility-centered federated learning frameworks, as was the case for graph applications [64] and IoT applications [187]. This will allow the research community to effectively evaluate and compare the performance of federated learning models on mobility data. We foresee that the transition between the current research efforts to real deployment will happen over stages where first multi-disciplinary research will focus on understanding users' attitudes towards using their location data for training models. After all, similar research on crowd-sensing applications has shown that location information is a top concern of users' involvement in these applications [3]. To the best of our knowledge, currently, there are no existing works in understanding users' privacy concerns when their data is not shared externally but is still used in creating predictive models. As a next step, we envision a slow transition between fully centralized models to decentralized models. Rather than training models focused on end-user prediction tasks, generative models that allow synthetic trace generation by learning from user mobility traces will be used to update and de-bias centralized datasets.

6 Conclusion

In this paper, we surveyed the federated learning models in the domain of mobility prediction as well as the widely used datasets for spatial-temporal models. We described the challenges that exist in applying common deep learning techniques in decentralized settings and discussed the opportunities for the research community to consider for future work. Our work indicates rapid growth and interest in this space, with promising future directions both in terms of theoretical frameworks and models, and practical applications and use cases. We hope both academics and practitioners find this survey useful for choosing the appropriate approach for their individual scenarios.

Acknowledgement

The publication of this work was made possible by generous support from the US National Science Foundation (NSF) award IIS-2304213.

References

- [1] 2009. The Official Website of the City of New York. <https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page>. Accessed: 2023-03-01.
- [2] 2013. City bike Official Web Page. <https://citibikenyc.com/system-data>. Accessed: 2023-03-01.
- [3] 2018. data usage privacy uk consumers. <https://ceo.digital/news/data-usage-privacy-uk-consumers>. Accessed: 2023-03-07.
- [4] 2019. Machine Learning Helps Payment Services Detect Fraud. <https://www.americanexpress.com/en-gb/business/trends-and-insights/articles/payment-services-fraud-detection-using-AI/>
- [5] Mohammed Akallouch, Oussama Akallouch, Khalid Fardousse, Afaf Bouhoute, and Ismail Berrada. 2022. Prediction and Privacy Scheme for Traffic Flow Estimation on the Highway Road Network. *Information* 13, 8 (2022), 381. <https://doi.org/10.3390/INF013080381>
- [6] Rasha Al-Huthaifi, Tianrui Li, Wei Huang, Jin Gu, and Chongshou Li. 2023. Federated learning in smart cities: Privacy and security survey. *Information Sciences* 632 (2023), 833–857. <https://doi.org/10.1016/J.INS.2023.03.033>
- [7] Muhammad Ammad-ud-din, Elena Ivannikova, Suleiman A. Khan, Were Oyomno, Qiang Fu, Kuan Eeik Tan, and Adrian Flanagan. 2019. Federated Collaborative Filtering for Privacy-Preserving Personalized Recommendation System. *CoRR* abs/1901.09888 (2019). arXiv:1901.09888 <http://arxiv.org/abs/1901.09888>
- [8] Hadis Anahideh and Abolfazl Asudeh. 2020. Fair Active Learning. *arXiv preprint arXiv:2001.01796* (2020).
- [9] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolic, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, Rui Oliveira, Pascal Felber, and Y. Charlie Hu (Eds.). ACM, 30:1–30:15. <https://doi.org/10.1145/3190508.3190538>
- [10] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Antonio Ferrara, and Fedelucio Narducci. 2021. How to put users in control of their data in federated top-N recommendation with learning to rank. In *SAC '21: The 36th ACM/SIGAPP Symposium on Applied Computing, Virtual Event, Republic of Korea, March 22-26, 2021*, Chih-Cheng Hung, Jiman Hong, Alessio Bechini, and Eunjee Song (Eds.). ACM, 1359–1362. <https://doi.org/10.1145/3412841.3442010>
- [11] Gowtham Atluri, Anuj Karpatne, and Vipin Kumar. 2018. Spatio-Temporal Data Mining: A Survey of Problems and Methods. *ACM Comput. Surv.* 51, 4 (2018), 83:1–83:41. <https://doi.org/10.1145/3161602>
- [12] Kay Werner Axhausen. 2007. Definition of movement and activity for transport modelling. In *Handbook of transport modelling*. Vol. 1. Emerald Group Publishing Limited, 329–343.
- [13] Nuria Rodríguez Barroso, Daniel Jiménez-López, María Victoria Luzón, Francisco Herrera, and Eugenio Martínez-Cámara. 2023. Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges. *Inf. Fusion* 90 (2023), 148–173. <https://doi.org/10.1016/J.INFFUS.2022.09.011>
- [14] Yacine Belal, Aurélien Bellet, Sonia Ben Mokhtar, and Vlad Nitu. 2022. PEPPER: Empowering User-Centric Recommender Systems over Gossip Learning. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 3 (2022), 101:1–101:27. <https://doi.org/10.1145/3550302>
- [15] Enrique Tomás Martínez Beltrán, Mario Quiles Pérez, Pedro Miguel Sánchez Sánchez, Sergio López Bernal, Gérôme Bovet, Manuel Gil Pérez, Gregorio Martínez Pérez, and Alberto Huertas Celdrán. 2023. Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges. *IEEE Commun. Surv. Tutorials* 25, 4 (2023), 2983–3013. <https://doi.org/10.1109/COMST.2023.3315746>
- [16] Sonia Ben Mokhtar, Antoine Boutet, Louafi Bouzouina, Patrick Bonnel, Olivier Brette, Lionel Brunie, Mathieu Cunche, Stephane D 'Alu, Vincent Primaault, Patrice Raveneau, Herve Rivano, and Razvan Stanica. 2017. PRIVA'MOV: Analysing Human Mobility Through Multi-Sensor Datasets. In *NetMob 2017*. HAL-Inria, Milan, Italy. <https://hal.inria.fr/hal-01578557>
- [17] Alastair R. Beresford and Frank Stajano. 2003. Location Privacy in Pervasive Computing. *IEEE Pervasive Comput.* 2, 1 (2003), 46–55. <https://doi.org/10.1109/MPRV.2003.1186725>

- [18] Daniel J. Beutel, Taner Topal, Akhil Mathur, Xinchu Qiu, Titouan Parcollet, and Nicholas D. Lane. 2020. Flower: A Friendly Federated Learning Research Framework. *CoRR* abs/2007.14390 (2020). arXiv:2007.14390 <https://arxiv.org/abs/2007.14390>
- [19] Geoff Boeing. 2017. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Comput. Environ. Urban Syst.* 65 (2017), 126–139. <https://doi.org/10.1016/j.compenvurbsys.2017.05.004>
- [20] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloé M Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. 2019. Towards Federated Learning at Scale: System Design. In *SysML 2019*.
- [21] Paul S Bradley, Kristin P Bennett, and Ayhan Demiriz. 2000. Constrained k-means clustering. *Microsoft Research, Redmond* 20, 0 (2000), 0.
- [22] Christopher Briggs, Zhong Fan, and Peter Andras. 2020. Federated learning with hierarchical clustering of local updates to improve training on non-IID data. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–9.
- [23] Anne Elizabeth Brown. 2018. *Ridehail revolution: Ridehail travel and equity in Los Angeles*. Ph. D. Dissertation. University of California, Los Angeles. <https://escholarship.org/uc/item/4r22m57k>
- [24] Sebastian Caldas, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. LEAF: A Benchmark for Federated Settings. *CoRR* abs/1812.01097 (2018). arXiv:1812.01097 <http://arxiv.org/abs/1812.01097>
- [25] Chao Chen, Karl Petty, Alexander Skabardonis, Pravin Varaiya, and Zhanfeng Jia. 2001. Freeway Performance Measurement System: Mining Loop Detector Data. *Transportation Research Record* 1748, 1 (2001), 96–102. <https://doi.org/10.3141/1748-12>
- [26] Chaochao Chen, Jun Zhou, Bingzhe Wu, Wenjing Fang, Li Wang, Yuan Qi, and Xiaolin Zheng. 2020. Practical Privacy Preserving POI Recommendation. *ACM Trans. Intell. Syst. Technol.* 11, 5 (2020), 52:1–52:20. <https://doi.org/10.1145/3394138>
- [27] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (San Diego, California, USA) (KDD '11)*. Association for Computing Machinery, New York, NY, USA, 1082–1090. <https://doi.org/10.1145/2020408.2020579>
- [28] Qiang Cui, Yuyuan Tang, Shu Wu, and Liang Wang. 2019. Distance2Pre: Personalized Spatial Preference for Next Point-of-Interest Prediction. In *Advances in Knowledge Discovery and Data Mining - 23rd Pacific-Asia Conference, PAKDD 2019, Macau, China, April 14-17, 2019, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 11441)*, Qiang Yang, Zhi-Hua Zhou, Zhiguo Gong, Min-Ling Zhang, and Sheng-Jun Huang (Eds.). Springer, 289–301. https://doi.org/10.1007/978-3-030-16142-2_23
- [29] Leylane Grazielle Ferreira da Silva, Djamel Fawzi Hadj Sadok, and Patricia Takako Endo. 2023. Resource optimizing federated learning for use with IoT: A systematic review. *J. Parallel Distributed Comput.* 175 (2023), 92–108. <https://doi.org/10.1016/j.jpdc.2023.01.006>
- [30] Maurizio Ferrari Dacrema, Simone Boglio, Paolo Cremonesi, and Dietmar Jannach. 2021. A Troubling Analysis of Reproducibility and Progress in Recommender Systems Research. *ACM Trans. Inf. Syst.* 39, 2 (2021), 20:1–20:49. <https://doi.org/10.1145/3434185>
- [31] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*, Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk (Eds.). ACM, 101–109. <https://doi.org/10.1145/3298689.3347058>
- [32] Manlio De Domenico, Antonio Lima, and Mirco Musolesi. 2013. Interdependence and predictability of human mobility and social interactions. *Pervasive and Mobile Computing* 9, 6 (2013), 798–807.
- [33] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3 (2013), 1376.
- [34] Zhaoyang Du, Celimuge Wu, Tsutomu Yoshinaga, Kok-Lim Alvin Yau, Yusheng Ji, and Jie Li. 2020. Federated Learning for Vehicular Internet of Things: Recent Advances and Open Issues. *IEEE Open J. Comput. Soc.* 1 (2020), 45–61. <https://doi.org/10.1109/OJCS.2020.2992630>
- [35] Highways England. 2014. Highways England network journey. <https://tris.highwaysengland.co.uk/>

- [36] Highways England. 2017. Traffic flow data. *Highways England*. [Online]. Available: <http://tris.highwaysengland.co.uk/detail/trafficflowdata> (2017).
- [37] Ecenaz Erdemir, Pier Luigi Dragotti, and Deniz Gündüz. 2021. Privacy-Aware Time-Series Data Sharing With Deep Reinforcement Learning. *IEEE Transactions on Information Forensics and Security* 16 (2021), 389–401. <https://doi.org/10.1109/TIFS.2020.3013200>
- [38] Fatima Zahra Errounda and Yan Liu. 2022. A Mobility Forecasting Framework with Vertical Federated Learning. In *46th IEEE Annual Computers, Software, and Applications Conferenc, COMPSAC 2022, Los Alamitos, CA, USA, June 27 - July 1, 2022*, Hong Va Leong, Sahra Sedigh Sarvestani, Yuuichi Teranishi, Alfredo Cuzzocrea, Hiroki Kashiwazaki, Dave Towey, Ji-Jiang Yang, and Hossain Shahriar (Eds.). IEEE, 301–310. <https://doi.org/10.1109/COMPSAC54236.2022.00050>
- [39] Castro Elizondo Jose Ezequiel, Martin Gjoreski, and Marc Langheinrich. 2022. Federated Learning for Privacy-Aware Human Mobility Modeling. *Frontiers Artif. Intell.* 5 (2022). <https://doi.org/10.3389/FRAI.2022.867046>
- [40] Zipei Fan, Xuan Song, Renhe Jiang, Qunjun Chen, and Ryosuke Shibasaki. 2019. Decentralized Attention-based Personalized Human Mobility Prediction. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 4 (2019), 133:1–133:26. <https://doi.org/10.1145/3369830>
- [41] Sadegh Farhadkhani, Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, and John Stephan. 2022. Byzantine Machine Learning Made Easy By Resilient Averaging of Momentums. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 6246–6283. <https://proceedings.mlr.press/v162/farhadkhani22a.html>
- [42] Jie Feng, Yong Li, Chao Zhang, Funing Sun, Fanchao Meng, Ang Guo, and Depeng Jin. 2018. DeepMove: Predicting Human Mobility with Attentional Recurrent Networks. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, Pierre-Antoine Champin, Fabien Gandon, Mounia Lalmas, and Panagiotis G. Ipeirotis (Eds.). ACM, 1459–1468. <https://doi.org/10.1145/3178876.3186058>
- [43] Jie Feng, Can Rong, Funing Sun, Diansheng Guo, and Yong Li. 2020. PMF: A Privacy-preserving Human Mobility Prediction Framework via Federated Learning. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 1 (2020), 10:1–10:21. <https://doi.org/10.1145/3381006>
- [44] Shanshan Feng, Xutao Li, Yifeng Zeng, Gao Cong, Yeow Meng Chee, and Quan Yuan. 2015. Personalized Ranking Metric Embedding for Next New POI Recommendation. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 2069–2075. <http://ijcai.org/Abstract/15/293>
- [45] Denzil Ferreira, Vassilis Kostakos, and Anind K. Dey. 2015. AWARE: mobile context instrumentation framework. *Frontiers ICT* 2 (2015), 6. <https://doi.org/10.3389/FICT.2015.00006>
- [46] Danielle L Ferreira, Bruno AA Nunes, Carlos Alberto V Campos, and Katia Obraczka. 2020. A Deep Learning Approach for Identifying User Communities Based on Geographical Preferences and Its Applications to Urban and Environmental Planning. 6, 3 (2020), 1–24.
- [47] Danielle L. Ferreira, Bruno A. A. Nunes, Carlos Alberto V. Campos, and Katia Obraczka. 2020. A Deep Learning Approach for Identifying User Communities Based on Geographical Preferences and Its Applications to Urban and Environmental Planning. *ACM Transactions on Spatial Algorithms and Systems* 6, 3 (May 2020), 1–24. <https://doi.org/10.1145/3380970>
- [48] Juliana Freire, Aline Bessa, Fernando Chirigati, Huy T. Vo, and Kai Zhao. 2016. Exploring What not to Clean in Urban Data: A Study Using New York City Taxi Trips. *IEEE Data Eng. Bull.* 39, 2 (2016), 63–77. <http://sites.computer.org/debull/A16june/p63.pdf>
- [49] Xingbo Fu, Binchi Zhang, Yushun Dong, Chen Chen, and Jundong Li. 2022. Federated Graph Machine Learning: A Survey of Concepts, Techniques, and Applications. *SIGKDD Explor.* 24, 2 (2022), 32–47. <https://doi.org/10.1145/3575637.3575644>
- [50] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. 2014. De-anonymization attack on geolocated data. *J. Comput. Syst. Sci.* 80, 8 (2014), 1597–1614. <https://doi.org/10.1016/J.JCSS.2014.04.024>
- [51] Dawei Gao, Daoyuan Chen, Zitao Li, Yuexiang Xie, Xuchen Pan, Yaliang Li, Bolin Ding, and Jingren Zhou. 2023. FS-Real: A Real-World Cross-Device Federated Learning Platform. *Proc. VLDB Endow.* 16, 12 (2023), 4046–4049. <https://doi.org/10.14778/3611540.3611617>

- [52] Qiang Gao, Fan Zhou, Goce Trajcevski, Kunpeng Zhang, Ting Zhong, and Fengli Zhang. 2019. Predicting Human Mobility via Variational Attention. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, Ling Liu, Ryan W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 2750–2756. <https://doi.org/10.1145/3308558.3313610>
- [53] Yanbo Ge, Christopher R Knittel, Don MacKenzie, and Stephen Zoepf. 2020. Racial discrimination in transportation network companies. *Journal of Public Economics* 190 (2020), 104205.
- [54] Mohammadhossein Ghahramani, MengChu Zhou, and Chi Tin Hon. 2018. Extracting significant mobile phone interaction patterns based on community structures. *IEEE Transactions on Intelligent Transportation Systems* 20, 3 (2018), 1031–1041.
- [55] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. 2022. An Efficient Framework for Clustered Federated Learning. *IEEE Trans. Inf. Theory* 68, 12 (2022), 8076–8091. <https://doi.org/10.1109/TIT.2022.3192506>
- [56] Bidisha Ghosh, Biswajit Basu, and Margaret O’Mahony. 2009. Multivariate Short-Term Traffic Flow Forecasting Using Time-Series Analysis. *IEEE Trans. Intell. Transp. Syst.* 10, 2 (2009), 246–254. <https://doi.org/10.1109/TITS.2009.2021448>
- [57] Anita Graser. 2019. Movingpandas: efficient structures for movement data in python. *GIForum* 1 (2019), 54–68.
- [58] Hao Guan, Pew-Thian Yap, Andrea Bozoki, and Mingxia Liu. 2024. Federated learning for medical image analysis: A survey. *Pattern Recognit.* 151 (2024), 110424. <https://doi.org/10.1016/J.PATCOG.2024.110424>
- [59] Rachid Guerraoui, Nirupam Gupta, and Rafael Pinot. 2023. Byzantine machine learning: A primer. *Comput. Surveys* (2023).
- [60] Yeting Guo, Fang Liu, Zhiping Cai, Hui Zeng, Li Chen, Tongqing Zhou, and Nong Xiao. 2021. PREFER: Point-of-interest REcommendation with efficiency and privacy-preservation via Federated Edge leaRning. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 1 (2021), 13:1–13:25. <https://doi.org/10.1145/3448099>
- [61] Zhiwei Guo, Keping Yu, Zhihan Lv, Kim-Kwang Raymond Choo, Peng Shi, and Joel J. P. C. Rodrigues. 2022. Deep Federated Learning Enhanced Secure POI Microservices for Cyber-Physical Systems. *IEEE Wirel. Commun.* 29, 2 (2022), 22–29. <https://doi.org/10.1109/MWC.002.2100272>
- [62] Salman Haidri, Yaksh J. Haranwala, Vania Bogorny, Chiara Renso, Vinicius Prado da Fonseca, and Amílcar Soares. 2022. PTRAIL - A python package for parallel trajectory data preprocessing. *SoftwareX* 19 (2022), 101176. <https://doi.org/10.1016/J.SOFTX.2022.101176>
- [63] Elizabeth Liz Harding, Jarno J Vanto, Reece Clark, L Hannah Ji, and Sara C Ainsworth. 2019. Understanding the scope and impact of the california consumer privacy act of 2018. *Journal of Data Protection & Privacy* 2, 3 (2019), 234–253.
- [64] Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Yu Rong, Peilin Zhao, Junzhou Huang, Murali Annavaram, and Salman Avestimehr. 2021. FedGraphNN: A Federated Learning System and Benchmark for Graph Neural Networks. *CoRR abs/2104.07145* (2021). arXiv:2104.07145 <https://arxiv.org/abs/2104.07145>
- [65] Chaoyang He, Songze Li, Jinhyun So, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Praneeth Vepakomma, Abhishek Singh, Hang Qiu, Li Shen, Peilin Zhao, Yan Kang, Yang Liu, Ramesh Raskar, Qiang Yang, Murali Annavaram, and Salman Avestimehr. 2020. FedML: A Research Library and Benchmark for Federated Machine Learning. *CoRR abs/2007.13518* (2020). arXiv:2007.13518 <https://arxiv.org/abs/2007.13518>
- [66] Lie He, Sai Praneeth Karimireddy, and Martin Jaggi. 2020. Byzantine-Robust Learning on Heterogeneous Datasets via Resampling. *CoRR abs/2006.09365* (2020). arXiv:2006.09365 <https://arxiv.org/abs/2006.09365>
- [67] Tianfu He, Jie Bao, Ruiyuan Li, Sijie Ruan, Yanhua Li, Li Song, Hui He, and Yu Zheng. 2020. What is the Human Mobility in a New City: Transfer Mobility Knowledge Across Cities. In *WWW ’20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, Yennun Huang, Irwin King, Tie-Yan Liu, and Maarten van Steen (Eds.). ACM / IW3C2, 1355–1365. <https://doi.org/10.1145/3366423.3380210>
- [68] Hoda Heidari, Michele Loi, Krishna P. Gummadi, and Andreas Krause. 2019. A Moral Framework for Understanding Fair ML through Economic Models of Equality of Opportunity. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* 2019, Atlanta, GA, USA, January 29-31, 2019*, danah boyd and Jamie H. Morgenstern (Eds.). ACM, 181–190. <https://doi.org/10.1145/3287560.3287584>
- [69] Neveen Mohammad Hijazi, Moayad Aloqaily, Mohsen Guizani, Bassem Ouni, and Fakhri Karray. 2024. Secure Federated Learning With Fully Homomorphic Encryption for IoT Communications. *IEEE Internet Things J.* 11, 3 (2024), 4289–4300. <https://doi.org/10.1109/JIOT.2023.3302065>

- [70] Chao Huang, Jianwei Huang, and Xin Liu. 2022. Cross-Silo Federated Learning: Challenges and Opportunities. *CoRR* abs/2206.12949 (2022). <https://doi.org/10.48550/ARXIV.2206.12949> arXiv:2206.12949
- [71] Chengzhi Huang and Hui Li. 2023. Single-User Injection for Invisible Shilling Attack against Recommender Systems. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21-25, 2023*, Ingo Frommholz, Frank Hopfgartner, Mark Lee, Michael Oakes, Mounia Lalmas, Min Zhang, and Rodrygo L. T. Santos (Eds.). ACM, 864–873. <https://doi.org/10.1145/3583780.3615062>
- [72] Haosheng Huang, Georg Gartner, Jukka Matthias Krisp, Martin Raubal, and Nico Van de Weghe. 2018. Location based services: ongoing evolution and research agenda. *J. Locat. Based Serv.* 12, 2 (2018), 63–93. <https://doi.org/10.1080/17489725.2018.1508763>
- [73] Juan Antonio Guerrero Ibáñez, Sherali Zeadally, and Juan Contreras-Castillo. 2018. Sensor Technologies for Intelligent Transportation Systems. *Sensors* 18, 4 (2018), 1212. <https://doi.org/10.3390/S18041212>
- [74] Wael Issa, Nour Moustafa, Benjamin P. Turnbull, Nasrin Sohrabi, and Zahir Tari. 2023. Blockchain-Based Federated Learning for Securing Internet of Things: A Comprehensive Survey. *ACM Comput. Surv.* 55, 9 (2023), 191:1–191:43. <https://doi.org/10.1145/3560816>
- [75] Hosagrahar V Jagadish, Johannes Gehrke, Alexandros Labrinidis, Yannis Papakonstantinou, Jignesh M Patel, Raghu Ramakrishnan, and Cyrus Shahabi. 2014. Big data and its technical challenges. *Commun. ACM* 57, 7 (2014), 86–94.
- [76] Tayyebeh Jahani-Nezhad, Mohammad Ali Maddah-Ali, Songze Li, and Giuseppe Caire. 2023. SwiftAgg+: Achieving Asymptotically Optimal Communication Loads in Secure Aggregation for Federated Learning. *IEEE J. Sel. Areas Commun.* 41, 4 (2023), 977–989. <https://doi.org/10.1109/JSAC.2023.3242702>
- [77] Renhe Jiang, Du Yin, Zhaonan Wang, Yizhuo Wang, Jiewen Deng, Hangchen Liu, Zekun Cai, Jinliang Deng, Xuan Song, and Ryosuke Shibasaki. 2021. DL-Traff: Survey and Benchmark of Deep Learning Models for Urban Traffic Prediction. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, Gianluca Demartini, Guido Zuccon, J. Shane Culpepper, Zi Huang, and Hanghang Tong (Eds.). ACM, 4515–4525. <https://doi.org/10.1145/3459637.3482000>
- [78] Weiwei Jiang. 2022. Cellular traffic prediction with machine learning: A survey. *Expert Syst. Appl.* 201 (2022), 117163. <https://doi.org/10.1016/J.ESWA.2022.117163>
- [79] Yuang Jiang, Shiqiang Wang, Víctor Valls, Bong Jun Ko, Wei-Han Lee, Kin K. Leung, and Leandros Tassiulas. 2023. Model Pruning Enables Efficient Federated Learning on Edge Devices. *IEEE Trans. Neural Networks Learn. Syst.* 34, 12 (2023), 10374–10386. <https://doi.org/10.1109/TNNLS.2022.3166101>
- [80] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista A. Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaïd Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Hang Qi, Daniel Ramage, Ramesh Raskar, Mariana Raykova, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. 2021. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* 14, 1-2 (2021), 1–210. <https://doi.org/10.1561/22000000083>
- [81] Maximilian Kasy and Rediet Abebe. 2021. Fairness, Equality, and Power in Algorithmic Decision-Making. In *FAccT '21: 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event / Toronto, Canada, March 3-10, 2021*, Madeleine Clare Elish, William Isaac, and Richard S. Zemel (Eds.). ACM, 576–586. <https://doi.org/10.1145/3442188.3445919>
- [82] Gaganbir Kaur, Surender K. Grewal, and Aarti Jain. 2023. Federated Learning based Spatio-Temporal framework for real-time traffic prediction. *CoRR* (2023). arXiv:rs.3.rs-2470634/v1 <https://doi.org/10.21203/rs.3.rs-2470634/v1>
- [83] Marc Keuschnigg, Selcan Mutgan, and Peter Hedström. 2019. Urban scaling and the regional divide. *Science advances* 5, 1 (2019), eaav0042.
- [84] Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak, and Vlad Nitu. 2021. EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 2 (2021), 68:1–68:25. <https://doi.org/10.1145/3463502>

- [85] Yeongwoo Kim, Ezeddin Al Hakim, Johan Haraldson, Henrik Eriksson, José Mairton B da Silva, and Carlo Fischione. 2021. Dynamic clustering in federated learning. In *ICC 2021-IEEE International Conference on Communications*. IEEE, 1–6.
- [86] Krzysztof W Kolodziej and Johan Hjelm. 2017. *Local positioning systems: LBS applications and services*. CRC press, Boca Raton.
- [87] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated Learning: Strategies for Improving Communication Efficiency. In *NIPS Workshop*.
- [88] Xiangjie Kong, Qiao Chen, Mingliang Hou, Hui Wang, and Feng Xia. 2023. Mobility trajectory generation: a survey. *Artif. Intell. Rev.* 56, Supplement 3 (2023), 3057–3098. <https://doi.org/10.1007/S10462-023-10598-X>
- [89] John Krumm. 2009. A survey of computational location privacy. *Pers. Ubiquitous Comput.* 13, 6 (2009), 391–399. <https://doi.org/10.1007/S00779-008-0212-5>
- [90] Fan Lai, Yinwei Dai, Sanjay Sri Vallabh Singapuram, Jiachen Liu, Xiangfeng Zhu, Harsha V. Madhyastha, and Mosharaf Chowdhury. 2022. FedScale: Benchmarking Model and System Performance of Federated Learning at Scale. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 11814–11827. <https://proceedings.mlr.press/v162/lai22a.html>
- [91] Natalie Lang, Elad Sofer, Tomer Shaked, and Nir Shlezinger. 2023. Joint Privacy Enhancement and Quantization in Federated Learning. *IEEE Trans. Signal Process.* 71 (2023), 295–310. <https://doi.org/10.1109/TSP.2023.3244092>
- [92] Jerry SH Lee, Kathleen M Darcy, Hai Hu, Yovanni Casablanca, Thomas P Conrads, Clifton L Dalgard, John B Freymann, Sean E Hanlon, Grant D Huang, Leonid Kvecher, et al. 2019. From discovery to practice and survivorship: building a national real-world data learning healthcare framework for military and veteran cancer patients. *Clinical Pharmacology & Therapeutics* 106, 1 (2019), 52–57.
- [93] Anran Li, Rui Liu, Ming Hu, Luu Anh Tuan, and Han Yu. 2023. Towards Interpretable Federated Learning. *CoRR* abs/2302.13473 (2023). <https://doi.org/10.48550/ARXIV.2302.13473> arXiv:2302.13473
- [94] Anliang Li, Shuang Wang, Wenzhu Li, Shengnan Liu, and Siyuan Zhang. 2020. Predicting Human Mobility with Federated Learning. In *SIGSPATIAL '20: 28th International Conference on Advances in Geographic Information Systems, Seattle, WA, USA, November 3-6, 2020*, Chang-Tien Lu, Fusheng Wang, Goce Trajcevski, Yan Huang, Shawn D. Newsam, and Li Xiong (Eds.). ACM, 441–444. <https://doi.org/10.1145/3397536.3422270>
- [95] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. 2023. A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. *IEEE Trans. Knowl. Data Eng.* 35, 4 (2023), 3347–3366. <https://doi.org/10.1109/TKDE.2021.3124599>
- [96] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine* 37, 3 (2020), 50–60.
- [97] Xiucheng Li, Gao Cong, Aixin Sun, and Yun Cheng. 2019. Learning Travel Time Distributions with Deep Generative Model. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 1017–1027. <https://doi.org/10.1145/3308558.3313418>
- [98] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=SJiHXGWAZ>
- [99] Dongliang Liao, Weiqing Liu, Yuan Zhong, Jing Li, and Guowei Wang. 2018. Predicting Activity and Location with Multi-task Context Aware Recurrent Neural Network. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 3435–3441. <https://doi.org/10.24963/IJCAI.2018/477>
- [100] Yuan Liao, Sonia Yeh, and Gustavo S. Jeuken. 2019. From individual to collective behaviours: exploring population heterogeneity of human mobility based on social media data. *EPJ Data Sci.* 8, 1 (2019), 34:1–34:22. <https://doi.org/10.1140/EPJDS/S13688-019-0212-X>
- [101] Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. Predicting the Next Location: A Recurrent Model with Spatial and Temporal Contexts. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*,

- February 12-17, 2016, Phoenix, Arizona, USA, Dale Schuurmans and Michael P. Wellman (Eds.). AAAI Press, 194–200. <https://doi.org/10.1609/AAAI.V30I1.9971>
- [102] Xiaoyuan Liu, Tianneng Shi, Chulin Xie, Qinbin Li, Kangping Hu, Haoyu Kim, Xiaojun Xu, Bo Li, and Dawn Song. 2022. UniFed: A Benchmark for Federated Learning Frameworks. *CoRR* abs/2207.10308 (2022). <https://doi.org/10.48550/ARXIV.2207.10308> arXiv:2207.10308
- [103] X. Liu and Y. Zhu. 2020. Privacy and Utility Preserving Trajectory Data Publishing for Intelligent Transportation Systems. *IEEE Access* 8 (2020), 176454–176466. <https://doi.org/10.1109/ACCESS.2020.3027299> Conference Name: IEEE Access.
- [104] Yang Liu, Tao Fan, Tianjian Chen, Qian Xu, and Qiang Yang. 2021. FATE: An Industrial Grade Platform for Collaborative Learning With Data Protection. *J. Mach. Learn. Res.* 22 (2021), 226:1–226:6. <http://jmlr.org/papers/v22/20-815.html>
- [105] Yiding Liu, Tuan-Anh Nguyen Pham, Gao Cong, and Quan Yuan. 2017. An Experimental Evaluation of Point-of-interest Recommendation in Location-based Social Networks. *Proc. VLDB Endow.* 10, 10 (2017), 1010–1021. <https://doi.org/10.14778/3115404.3115407>
- [106] Yi Liu, James J. Q. Yu, Jiawen Kang, Dusit Niyato, and Shuyu Zhang. 2020. Privacy-Preserving Traffic Flow Prediction: A Federated Learning Approach. *IEEE Internet Things J.* 7, 8 (2020), 7751–7763. <https://doi.org/10.1109/JIOT.2020.2991401>
- [107] Qingyue Long, Huandong Wang, Tong Li, Lisi Huang, Kun Wang, Qiong Wu, Guangyu Li, Yanping Liang, Li Yu, and Yong Li. 2023. Practical Synthetic Human Trajectories Generation Based on Variational Point Processes. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023*, Ambuj K. Singh, Yizhou Sun, Leman Akoglu, Dimitrios Gunopulos, Xifeng Yan, Ravi Kumar, Fatma Ozcan, and Jieping Ye (Eds.). ACM, 4561–4571. <https://doi.org/10.1145/3580305.3599888>
- [108] Xin Lu, Erik Wetter, Nita Bharti, Andrew J Tatem, and Linus Bengtsson. 2013. Approaching the limit of predictability in human mobility. *Scientific reports* 3, 1 (2013), 1–9.
- [109] María Victoria Luzón, Nuria Rodríguez Barroso, Alberto Argente-Garrido, Daniel Jiménez-López, Jose M. Moyano, Javier Del Ser, Weiping Ding, and Francisco Herrera. 2024. A Tutorial on Federated Learning from Theory to Practice: Foundations, Software Frameworks, Exemplary Use Cases, and Selected Trends. *IEEE CAA J. Autom. Sinica* 11, 4 (2024), 824–850. <https://doi.org/10.1109/JAS.2024.124215>
- [110] Mohamad Mansouri, Melek Önen, Wafa Ben Jaballah, and Mauro Conti. 2023. SoK: Secure Aggregation Based on Cryptographic Schemes for Federated Learning. *Proc. Priv. Enhancing Technol.* 2023, 1 (2023), 140–157. <https://doi.org/10.56553/POPETS-2023-0009>
- [111] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. 2017. AP-Attack: A Novel User Re-identification Attack On Mobility Datasets. In *Proceedings of the 14th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, Melbourne, Australia, November 7-10, 2017*, Tao Gu, Ramamohanarao Kotagiri, and Huai Liu (Eds.). ACM, 48–57. <https://doi.org/10.1145/3144457.3144494>
- [112] Pavel Masek, Jan Masek, Petr Frantik, Radek Fujdiak, Aleksandr Ometov, Jiri Hosek, Sergey Andreev, Petr Mlynek, and Jiri Misurec. 2016. A Harmonized Perspective on Transportation Management in Smart Cities: The Novel IoT-Driven Environment for Road Traffic Modeling. *Sensors* 16, 11 (2016), 1872. <https://doi.org/10.3390/S16111872>
- [113] Afra Mashhadi, Alex Kylo, and Reza M. Parizi. 2022. Fairness in Federated Learning for Spatial-Temporal Applications. *CoRR* abs/2201.06598 (2022). arXiv:2201.06598 <https://arxiv.org/abs/2201.06598>
- [114] Afra Mashhadi, Joshua Sterner, and Jeffrey Murray. 2021. Deep Embedded Clustering of Urban Communities Using Federated Learning. In *International Joint Conference on Neural Networks, IJCNN 2021, Shenzhen, China, July 18-22, 2021*. IEEE, 1–8. <https://doi.org/10.1109/IJCNN52387.2021.9534268>
- [115] Afra Mashhadi, Ali Tabarai, Yuting Zhan, and Reza M. Parizi. 2022. An Auditing Framework for Analyzing Fairness of Spatial-Temporal Federated Learning Applications. In *2022 IEEE World AI IoT Congress (AIIoT), Seattle, WA, USA, June 6-9, 2022*. IEEE, 699–707. <https://doi.org/10.1109/AIIOT54504.2022.9817283>
- [116] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *AISTATS*. 1273–1282.
- [117] H Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. 2016. Federated learning of deep networks using model averaging. *CoRR* abs/1602.05629 (2016). *arXiv preprint arXiv:1602.05629* (2016).
- [118] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *ICLR*.

- [119] Collin Meese, Hang Chen, Syed Ali Asif, Wanxin Li, Chien-Chung Shen, and Mark M. Nejad. 2022. BFRT: Blockchain Federated Learning for Real-time Traffic Flow Prediction. In *22nd IEEE International Symposium on Cluster, Cloud and Internet Computing, CCGRID 2022, Taormina, Italy, May 16-19, 2022*. IEEE, 317–326. <https://doi.org/10.1109/CCGRID54584.2022.00041>
- [120] Wendy Kan Meghan O’Connell, moreiraMatias. 2015. ECML/PKDD 15: Taxi Trajectory Prediction (I). <https://kaggle.com/competitions/pkdd-15-predict-taxi-service-trajectory-i>
- [121] Chuizheng Meng, Sirisha Rambhatla, and Yan Liu. 2021. Cross-Node Federated Graph Neural Network for Spatio-Temporal Data Modeling. In *KDD ’21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021*, Feida Zhu, Beng Chin Ooi, and Chunyan Miao (Eds.). ACM, 1202–1211. <https://doi.org/10.1145/3447548.3467371>
- [122] Zaiqiao Meng, Richard McCreadie, Craig Macdonald, and Iadh Ounis. 2020. Exploring Data Splitting Strategies for the Evaluation of Recommendation Models. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020*, Rodrygo L. T. Santos, Leandro Balby Marinho, Elizabeth M. Daly, Li Chen, Kim Falk, Noam Koenigstein, and Edleno Silva de Moura (Eds.). ACM, 681–686. <https://doi.org/10.1145/3383313.3418479>
- [123] Daniel Opoku Mensah, Godwin Badu-Marfo, Ranwa Al Mallah, and Bilal Farooq. 2022. eFedDNN: Ensemble based Federated Deep Neural Networks for Trajectory Mode Inference. In *IEEE International Smart Cities Conference, ISC2 2022, Pafos, Cyprus, September 26-29, 2022*. IEEE, 1–7. <https://doi.org/10.1109/ISC255366.2022.9922022>
- [124] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. 2021. PPFL: Privacy-Preserving Federated Learning with Trusted Execution Environments. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services (Virtual Event, Wisconsin) (MobiSys ’21)*. Association for Computing Machinery, New York, NY, USA, 94–108. <https://doi.org/10.1145/3458864.3466628>
- [125] Virraji Mothukuri, Reza M. Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. 2021. A survey on security and privacy of federated learning. *Future Gener. Comput. Syst.* 115 (2021), 619–640. <https://doi.org/10.1016/J.FUTURE.2020.10.007>
- [126] Yustus Eko Oktian and Sang-Gon Lee. 2023. Blockchain-Based Federated Learning System: A Survey on Design Choices. *Sensors* 23, 12 (2023), 5658. <https://doi.org/10.3390/S23125658>
- [127] Xavier Olive. 2019. traffic, a toolbox for processing and analysing air traffic data. *J. Open Source Softw.* 4, 39 (2019), 1518. <https://doi.org/10.21105/JOSS.01518>
- [128] Nuria Oliver, Bruno Lepri, Harald Sterly, Renaud Lambiotte, Sébastien Deletaille, Marco De Nadai, Emmanuel Letouzé, Albert Ali Salah, Richard Benjamins, Ciro Cattuto, Vittoria Colizza, Nicolas de Cordes, Samuel P. Fraiberger, Till Koebe, Sune Lehmann, Juan Murillo, Alex Pentland, Phuong N Pham, Frédéric Pivetta, Jari Saramäki, Samuel V. Scarpino, Michele Tizzoni, Stefaan Verhulst, and Patrick Vinck. 2020. Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle. *Science Advances* 6, 23 (2020), eabc0764. <https://doi.org/10.1126/sciadv.abc0764> arXiv:<https://www.science.org/doi/pdf/10.1126/sciadv.abc0764>
- [129] Nuria Oliver, Aleksandar Matic, and Enrique Frias-Martinez. 2015. Mobile network data for public health: opportunities and challenges. *Frontiers in public health* 3 (2015), 189.
- [130] Adil Oualid, Yassine Maleh, and Lahcen Moumoun. 2023. Federated learning techniques applied to credit risk management: A systematic literature review. *EDPACS* 68, 1 (2023), 42–56.
- [131] Sharnil Pandya, Gautam Srivastava, Rutvij Jhaveri, M Rajasekhara Babu, Sweta Bhattacharya, Praveen Kumar Reddy Maddikunta, Spyridon Mastorakis, Md Jalil Piran, and Thippa Reddy Gadekallu. 2023. Federated learning for smart cities: A comprehensive survey. *Sustainable Energy Technologies and Assessments* 55 (2023), 102987.
- [132] Luca Pappalardo, Filippo Simini, Gianni Barlacchi, and Roberto Pellungrini. 2022. scikit-mobility: A Python Library for the Analysis, Generation, and Risk Assessment of Mobility Data. *J. Stat. Softw.* 103, 4 (2022). <https://doi.org/10.18637/JSS.V103.I04>
- [133] Balázs Pejó. 2020. The Good, The Bad, and The Ugly: Quality Inference in Federated Learning. *CoRR* abs/2007.06236 (2020). arXiv:2007.06236 <https://arxiv.org/abs/2007.06236>
- [134] Vasileios Perifanis, George Drosatos, Giorgos Stamatelatos, and Pavlos S. Efraimidis. 2023. FedPOIRec: Privacy-preserving federated poi recommendation with social influence. *Inf. Sci.* 623 (2023), 767–790. <https://doi.org/10.1016/J.INS.2022.12.024>

- [135] Attia Qammar, Ahmad Karim, Huansheng Ning, and Jianguo Ding. 2023. Securing federated learning with blockchain: a systematic literature review. *Artif. Intell. Rev.* 56, 5 (2023), 3951–3985. <https://doi.org/10.1007/S10462-022-10271-9>
- [136] Tao Qi, Lingqiang Chen, Guanghui Li, Yijing Li, and Chenshu Wang. 2023. FedAGCN: A traffic flow prediction framework based on federated learning and Asynchronous Graph Convolutional Network. *Appl. Soft Comput.* 138 (2023), 110175. <https://doi.org/10.1016/J.ASOC.2023.110175>
- [137] Tao Qi, Guanghui Li, Lingqiang Chen, and Yanming Xue. 2022. ADGCN: An Asynchronous Dilation Graph Convolutional Network for Traffic Flow Prediction. *IEEE Internet Things J.* 9, 5 (2022), 4001–4014. <https://doi.org/10.1109/JIOT.2021.3102238>
- [138] Sita Rani, Aman Kataria, Sachin Kumar, and Prayag Tiwari. 2023. Federated learning for secure IoMT-applications in smart healthcare systems: A comprehensive review. *Knowl. Based Syst.* 274 (2023), 110658. <https://doi.org/10.1016/J.KNSYS.2023.110658>
- [139] Jinneng Rao, Song Gao, Yuhao Kang, and Qunying Huang. 2021. LSTM-TrajGAN: A Deep Learning Approach to Trajectory Privacy Protection. In *11th International Conference on Geographic Information Science, GIScience 2021, September 27-30, 2021, Poznań, Poland - Part I (LIPICs, Vol. 177)*, Krzysztof Janowicz and Judith Anne Versteegen (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 12:1–12:17. <https://doi.org/10.4230/LIPICs.GISCIENCE.2021.I.12>
- [140] Al-Huthaifi Rasha, Tianrui Li, Wei Huang, Jin Gu, and Chongshou Li. 2023. Federated learning in smart cities: Privacy and security survey. *Information Sciences* (2023).
- [141] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan. 2020. Adaptive Federated Optimization. *arXiv preprint arXiv:2003.00295* (2020).
- [142] K Dasaradharami Reddy and Thippa Reddy Gadekallu. 2023. A comprehensive survey on federated learning techniques for healthcare informatics. *Computational Intelligence and Neuroscience 2023* (2023).
- [143] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*, Jeff A. Bilmes and Andrew Y. Ng (Eds.). AUAI Press, 452–461. https://www.auai.org/uai2009/papers/UAI2009_0139_48141db02b9f0b02bc7158819ebfa2c7.pdf
- [144] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R. Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus H. Maier-Hein, Sébastien Ourselin, Micah J. Sheller, Ronald M. Summers, Andrew Trask, Daguang Xu, Maximilian Baust, and M. Jorge Cardoso. 2020. The future of digital health with federated learning. *npj Digit. Medicine* 3 (2020). <https://doi.org/10.1038/S41746-020-00323-1>
- [145] Amina Ben Salem, Bisma Khalfoun, Sonia Ben Mokhtar, and Afra J. Mashhadi. 2022. Quantifying fairness of federated learning LPPM models. In *MobiSys '22: The 20th Annual International Conference on Mobile Systems, Applications and Services, Portland, Oregon, 27 June 2022 - 1 July 2022*, Nirupama Bulusu, Ehsan Aryafar, Aruna Balasubramanian, and Junehwa Song (Eds.). ACM, 569–570. <https://doi.org/10.1145/3498361.3538788>
- [146] Prasanna Sattigeri, Samuel C. Hoffman, Vijil Chenthamarakshan, and Kush R. Varshney. 2019. Fairness GAN: Generating datasets with fairness properties using a generative adversarial network. *IBM J. Res. Dev.* 63, 4/5 (2019), 3:1–3:9. <https://doi.org/10.1147/JRD.2019.2945519>
- [147] Sanaz Shaker Sepasgozar and Samuel Pierre. 2022. Fed-NTP: A Federated Learning Algorithm for Network Traffic Prediction in VANET. *IEEE Access* 10 (2022), 119607–119616. <https://doi.org/10.1109/ACCESS.2022.3221970>
- [148] Adi Shamir. 1979. How to share a secret. *Commun. ACM* 22, 11 (1979), 612–613.
- [149] Sheng Shen, Tianqing Zhu, Di Wu, Wei Wang, and Wanlei Zhou. 2022. From distributed machine learning to federated learning: In the view of data privacy and security. *Concurr. Comput. Pract. Exp.* 34, 16 (2022). <https://doi.org/10.1002/CPE.6002>
- [150] Zhaosen Shi, Zeyu Yang, Alzubair Hassan, Fagen Li, and Xuyang Ding. 2023. A privacy preserving federated learning scheme using homomorphic encryption and secret sharing. *Telecommun. Syst.* 82, 3 (2023), 419–433. <https://doi.org/10.1007/S11235-022-00982-3>
- [151] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. 2011. Quantifying location privacy. In *2011 IEEE symposium on security and privacy*. IEEE, Oakland, CA, USA, 247–262.

- [152] Jinhyun So, Ramy E. Ali, Basak Güler, Jiantao Jiao, and Amir Salman Avestimehr. 2023. Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, 9864–9873. <https://doi.org/10.1609/AAAI.V37I8.26177>
- [153] Huimin Tang, Nianming Xue, and Gaoli Wang. 2022. Differentially Private Decentralized Traffic Flow Prediction Approach based on Federated Learning. In *Proceedings of the 2022 10th International Conference on Information Technology: IoT and Smart City, ICIT 2022, Shanghai, China, December 23-25, 2022*. ACM, 280–285. <https://doi.org/10.1145/3582197.3582244>
- [154] Jiayi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, February 5-9, 2018*, Yi Chang, Chengxiang Zhai, Yan Liu, and Yoelle Maarek (Eds.). ACM, 565–573. <https://doi.org/10.1145/3159652.3159656>
- [155] Eran Toch, Boaz Lerner, Eyal Ben-Zion, and Irad Ben-Gal. 2019. Analyzing large-scale human mobility data: a survey of machine learning methods and applications. *Knowledge and Information Systems* 58, 3 (2019), 501–523.
- [156] Lev Velykoivanenko, Kavous Salehzadeh Niksirat, Noé Zufferey, Mathias Humbert, Kevin Huguenin, and Mauro Cherubini. 2021. Are Those Steps Worth Your Privacy?: Fitness-Tracker Users’ Perceptions of Privacy and Utility. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4 (2021), 181:1–181:41. <https://doi.org/10.1145/3494960>
- [157] Madumitha Venkatasubramanian, Arash Habibi Lashkari, and Saqib Hakak. 2023. IoT Malware Analysis Using Federated Learning: A Comprehensive Survey. *IEEE Access* 11 (2023), 5004–5018. <https://doi.org/10.1109/ACCESS.2023.3235389>
- [158] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed.*, Cham: Springer International Publishing 10, 3152676 (2017), 10–5555.
- [159] Jayant Vyas, Debasis Das, and Sajal Kumar Das. 2020. Vehicular Edge Computing Based Driver Recommendation System Using Federated Learning. In *17th IEEE International Conference on Mobile Ad Hoc and Sensor Systems, MASS 2020, Delhi, India, December 10-13, 2020*. IEEE, 675–683. <https://doi.org/10.1109/MASS50613.2020.00087>
- [160] Hanqiu Wang, Rongqing Zhang, Xiang Cheng, and Liuqing Yang. 2022. Federated Spatio-Temporal Traffic Flow Prediction Based on Graph Convolutional Network. In *14th International Conference on Wireless Communications and Signal Processing, WCSP 2022, Nanjing, China, November 1-3, 2022*. IEEE, 221–225. <https://doi.org/10.1109/WCSP55476.2022.10039323>
- [161] Senzhang Wang, Jiannong Cao, and Philip Yu. 2020. Deep Learning for Spatio-Temporal Data Mining: A Survey. *IEEE Transactions on Knowledge and Data Engineering* 14, 8 (2020), 1–1. <https://doi.org/10.1109/TKDE.2020.3025580>
- [162] Shuang Wang, Bowei Wang, Shuai Yao, Jiangqin Qu, and Yuezheng Pan. 2022. Location prediction with personalized federated learning. *Soft Computing* (2022), 1–12.
- [163] Tongnian Wang, Yan Du, Yanmin Gong, Kim-Kwang Raymond Choo, and Yuanxiong Guo. 2023. Applications of federated learning in mobile health: scoping review. *Journal of Medical Internet Research* 25 (2023), e43006.
- [164] Jie Wen, Zhixia Zhang, Yang Lan, Zhihua Cui, Jianghui Cai, and Wensheng Zhang. 2023. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics* 14, 2 (2023), 513–535.
- [165] Marius Wernke, Pavel Skvortsov, Frank Dürr, and Kurt Rothermel. 2014. A classification of location privacy attacks and approaches. *Pers. Ubiquitous Comput.* 18, 1 (2014), 163–175. <https://doi.org/10.1007/S00779-012-0633-Z>
- [166] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020. Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash (Eds.). ACM, 753–763. <https://doi.org/10.1145/3394486.3403118>
- [167] Yuexiang Xie, Zhen Wang, Dawei Gao, Daoyuan Chen, Liuyi Yao, Weirui Kuang, Yaliang Li, Bolin Ding, and Jingren Zhou. 2023. FederatedScope: A Flexible Federated Learning Platform for Heterogeneity. *Proc. VLDB Endow.* 16, 5 (2023), 1059–1072. <https://doi.org/10.14778/3579075.3579081>

- [168] An Yan and Bill Howe. 2019. Fairness in Practice: A Survey on Equity in Urban Mobility. *IEEE Data Eng. Bull.* 42, 3 (2019), 49–63. <http://sites.computer.org/debull/A19sept/p49.pdf>
- [169] An Yan and Bill Howe. 2019. FairST: Equitable Spatial and Temporal Demand Prediction for New Mobility Systems. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2019, Chicago, IL, USA, November 5-8, 2019*, Farnoush Banaei Kashani, Goce Trajcevski, Ralf Hartmut Güting, Lars Kulik, and Shawn D. Newsam (Eds.). ACM, 552–555. <https://doi.org/10.1145/3347146.3359380>
- [170] Yunlu Yan, Chun-Mei Feng, Mang Ye, Wangmeng Zuo, Ping Li, Rick Siow Mong Goh, Lei Zhu, and C. L. Philip Chen. 2023. Rethinking Client Drift in Federated Learning: A Logit Perspective. *CoRR* abs/2308.10162 (2023). <https://doi.org/10.48550/ARXIV.2308.10162> arXiv:2308.10162
- [171] Dingqi Yang, Daqing Zhang, Vincent W. Zheng, and Zhiyong Yu. 2015. Modeling User Activity Preference by Leveraging User Spatial Temporal Characteristics in LBSNs. *IEEE Trans. Syst. Man Cybern. Syst.* 45, 1 (2015), 129–142. <https://doi.org/10.1109/TSMC.2014.2327053>
- [172] Mang Ye, Xiuwen Fang, Bo Du, Pong C. Yuen, and Dacheng Tao. 2024. Heterogeneous Federated Learning: State-of-the-art and Research Challenges. *ACM Comput. Surv.* 56, 3 (2024), 79:1–79:44. <https://doi.org/10.1145/3625558>
- [173] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 3634–3640. <https://doi.org/10.24963/IJCAI.2018/505>
- [174] Fuxun Yu, Zirui Xu, Zhuwei Qin, and Xiang Chen. 2022. Privacy-preserving federated learning for transportation mode prediction based on personal mobility data. *High-Confidence Computing* 2, 4 (2022), 100082.
- [175] Qing Yu and Jian Yuan. 2022. TransBigData: A Python package for transportation spatio-temporal big data processing, analysis and visualization. *J. Open Source Softw.* 7, 71 (2022), 4021. <https://doi.org/10.21105/JOSS.04021>
- [176] Wei Yuan, Chaoqun Yang, Quoc Viet Hung Nguyen, Lizhen Cui, Tieke He, and Hongzhi Yin. 2023. Interaction-level Membership Inference Attack Against Federated Recommender Systems. In *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, Ying Ding, Jie Tang, Juan F. Sequeda, Lora Aroyo, Carlos Castillo, and Geert-Jan Houben (Eds.). ACM, 1053–1062. <https://doi.org/10.1145/3543507.3583359>
- [177] Xiaoming Yuan, Jiahui Chen, Jiayu Yang, Ning Zhang, Tingting Yang, Tao Han, and Amir Taherkordi. 2023. FedSTN: Graph Representation Driven Federated Learning for Edge Computing Enabled Urban Traffic Flow Prediction. *IEEE Trans. Intell. Transp. Syst.* 24, 8 (2023), 8738–8748. <https://doi.org/10.1109/TITS.2022.3157056>
- [178] Xiaoming Yuan, Jiahui Chen, Ning Zhang, Chunsheng Zhu, Qiang Ye, and Xuemin Sherman Shen. 2022. FedTSE: Low-Cost Federated Learning for Privacy-Preserved Traffic State Estimation in IoV. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications Workshops, INFOCOM 2022 - Workshops, New York, NY, USA, May 2-5, 2022*. IEEE, 1–6. <https://doi.org/10.1109/INFOCOMWKSHP54753.2022.9798018>
- [179] Tengchan Zeng, Jianlin Guo, Kyeong Jin Kim, Kieran Parsons, Philip V. Orlik, Stefano Di Cairano, and Walid Saad. 2021. Multi-Task Federated Learning for Traffic Prediction and Its Application to Route Planning. In *IEEE Intelligent Vehicles Symposium, IV 2021, Nagoya, Japan, July 11-17, 2021*. IEEE, 451–457. <https://doi.org/10.1109/IV48863.2021.9575211>
- [180] Yuting Zhan, Alex Kylo, Afra Mashhadi, and Hamed Haddadi. 2022. Privacy-Aware Human Mobility Prediction via Adversarial Networks. *CoRR* abs/2201.07519 (2022). arXiv:2201.07519 <https://arxiv.org/abs/2201.07519>
- [181] Chuanting Zhang, Shuping Dang, Basem Shihada, and Mohamed-Slim Alouini. 2021. Dual Attention-Based Federated Learning for Wireless Traffic Prediction. In *40th IEEE Conference on Computer Communications, INFOCOM 2021, Vancouver, BC, Canada, May 10-13, 2021*. IEEE, 1–10. <https://doi.org/10.1109/INFOCOM42981.2021.9488883>
- [182] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. 2020. BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning. In *2020 USENIX Annual Technical Conference, USENIX ATC 2020, July 15-17, 2020*, Ada Gavrilovska and Erez Zadok (Eds.). USENIX Association, 493–506. <https://www.usenix.org/conference/atc20/presentation/zhang-chengliang>

- [183] Chao Zhang, Keyang Zhang, Quan Yuan, Haoruo Peng, Yu Zheng, Tim Hanratty, Shaowen Wang, and Jiawei Han. 2017. Regions, Periods, Activities: Uncovering Urban Dynamics via Cross-Modal Representation Learning. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, Rick Barrett, Rick Cummings, Eugene Agichtein, and Evgeniy Gabrilovich (Eds.). ACM, 361–370. <https://doi.org/10.1145/3038912.3052601>
- [184] Chao Zhang, Keyang Zhang, Quan Yuan, Luming Zhang, Tim Hanratty, and Jiawei Han. 2016. GMove: Group-Level Mobility Modeling Using Geo-Tagged Social Media. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 1305–1314. <https://doi.org/10.1145/2939672.2939793>
- [185] Chenhan Zhang, Shiyao Zhang, Shui Yu, and James J. Q. Yu. 2022. Graph-Based Traffic Forecasting via Communication-Efficient Federated Learning. In *IEEE Wireless Communications and Networking Conference, WCNC 2022, Austin, TX, USA, April 10-13, 2022*. IEEE, 2041–2046. <https://doi.org/10.1109/WCNC51071.2022.9771883>
- [186] Chenhan Zhang, Yuanshao Zhu, Christos Markos, Shui Yu, and James J. Q. Yu. 2022. Toward Crowdsourced Transportation Mode Identification: A Semisupervised Federated Learning Approach. *IEEE Internet Things J.* 9, 14 (2022), 11868–11882. <https://doi.org/10.1109/JIOT.2021.3132056>
- [187] Fengda Zhang, Kun Kuang, Zhaoyang You, Tao Shen, Jun Xiao, Yin Zhang, Chao Wu, Yueting Zhuang, and Xiaolin Li. 2020. Federated Unsupervised Representation Learning. arXiv:2010.08982 [cs.LG]
- [188] Junwei Zhang, Min Gao, Junliang Yu, Lei Guo, Jundong Li, and Hongzhi Yin. 2021. Double-Scale Self-Supervised Hypergraph Learning for Group Recommendation. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, Gianluca Demartini, Guido Zuccon, J. Shane Culpepper, Zi Huang, and Hanghang Tong (Eds.). ACM, 2557–2567. <https://doi.org/10.1145/3459637.3482426>
- [189] Junbo Zhang, Yu Zheng, and Dekang Qi. 2017. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, Satinder Singh and Shaul Markovitch (Eds.). AAAI Press, 1655–1661. <https://doi.org/10.1609/AAAI.V31I1.10735>
- [190] Qi Zhang, Tiancheng Wu, Peichen Zhou, Shan Zhou, Yuan Yang, and Xiulang Jin. 2022. Felicitas: Federated Learning in Distributed Cross Device Collaborative Frameworks. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, Aidong Zhang and Huzefa Rangwala (Eds.). ACM, 4502–4509. <https://doi.org/10.1145/3534678.3539039>
- [191] Shijie Zhang, Wei Yuan, and Hongzhi Yin. 2023. Comprehensive privacy analysis on federated recommender system against attribute inference attacks. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [192] Ying Zhang, Zhiwen Yu, Minling Dang, En Xu, Bin Guo, Yuxuan Liang, Yifang Yin, and Roger Zimmermann. 2023. Predictability in Human Mobility: From Individual to Collective (Vision Paper). *ACM Transactions on Spatial Algorithms and Systems* (2023).
- [193] Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Zhenqiang Gong. 2021. Backdoor Attacks to Graph Neural Networks. In *SACMAT '21: The 26th ACM Symposium on Access Control Models and Technologies, Virtual Event, Spain, June 16-18, 2021*, Jorge Lobo, Roberto Di Pietro, Omar Chowdhury, and Hongxin Hu (Eds.). ACM, 15–26. <https://doi.org/10.1145/3450569.3463560>
- [194] Zhiwen Zhang, Hongjun Wang, Zipei Fan, Jiyuan Chen, Xuan Song, and Ryosuke Shibasaki. 2022. GOF-TTE: Generative Online Federated Learning Framework for Travel Time Estimation. *IEEE Internet Things J.* 9, 23 (2022), 24107–24121. <https://doi.org/10.1109/JIOT.2022.3190864>
- [195] Fengpan Zhao, Yan Huang, Akshita Maradapu Vera Venkata Sai, and Yubao Wu. 2020. A Cluster-based Solution to Achieve Fairness in Federated Learning. In *2020 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCloud/SocialCom/SustainCom)*. 875–882. <https://doi.org/10.1109/ISPA-BDCloud-SocialCom-SustainCom51426.2020.00135>
- [196] Kai Zhao, Sasu Tarkoma, Siyuan Liu, and Huy T. Vo. 2016. Urban human mobility data mining: An overview. In *2016 IEEE International Conference on Big Data (IEEE BigData 2016), Washington DC, USA, December 5-8, 2016*, James Joshi, George Karypis, Ling Liu, Xiaohua Hu, Ronay Ak, Yinglong Xia, Weijia Xu, Aki-Hiro Sato, Sudarsan Rachuri, Lyle H. Ungar, Philip S. Yu, Rama Govindaraju, and Toyotaro Suzumura (Eds.). IEEE Computer Society, 1911–1920. <https://doi.org/10.1109/BIGDATA.2016.7840811>

-
- [197] Yu Zheng. 2011. T-Drive trajectory data sample. <https://www.microsoft.com/en-us/research/publication/t-drive-trajectory-data-sample/> T-Drive sample dataset.
- [198] Yu Zheng, Xing Xie, Wei-Ying Ma, et al. 2010. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.* 33, 2 (2010), 32–39.
- [199] Zhaohua Zheng, Yize Zhou, Yilong Sun, Zhang Wang, Boyi Liu, and Keqiu Li. 2022. Applications of federated learning in smart cities: recent advances, taxonomy, and open challenges. *Connect. Sci.* 34, 1 (2022), 1–28. <https://doi.org/10.1080/09540091.2021.1936455>
- [200] Juncen Zhu, Jiannong Cao, Divya Saxena, Shan Jiang, and Houda Ferradi. 2023. Blockchain-empowered Federated Learning: Challenges, Solutions, and Future Directions. *ACM Comput. Surv.* 55, 11 (2023), 240:1–240:31. <https://doi.org/10.1145/3570953>

7 Appendix

7.1 Survey Review Methodology

To create the summary of the surveys and their topics, we crawled the citation (bib file) of all the published articles with the keywords ‘Federated Learning’ and ‘Survey’ or ‘Review’ in the title for each year. We used the publish or perish tool for this crawling and used Google Scholar as the platform. We downloaded the abstracts of all these surveys from those papers and fed them to a Large Language Model (GPT 3.5) for thematic categorization of the topics. We followed a similar process for counting the number of papers per year that were published on FL-ST topics (non-surveys).