Business Value of Privacy-Preserving Mobility Analytics in Platform Operations: A Federated Risk-Utility Framework

Jiahui Ding[1], Fafa (Monica) Yan[2, *]

[1] School of Business, Beijing Technology and Business University, Beijing, China 100048

[2] School of Economics and Management, Guangzhou Institute of Science and Technology, Guangzhou, China 510540

* Email: yanff@mail3.sysu.edu.cn (Corresponding Author)

Received 10 March 2026 · Revised 24 April 2026 · Accepted 8 June 2026 · Published 30 September 2026

Abstract

Mobility platforms use location and movement information to forecast demand, allocate capacity, and improve service, but the same information can expose sensitive patterns about individuals. Federated analytics can reduce the need to pool raw data, although model updates, participation records, and released outputs may still leak information. This article develops a risk–utility framework for evaluating the business value of privacy-preserving mobility analytics in platform operations. The framework links a defined operational decision to expected service gains, deployment costs, residual incident losses, and separately enforced privacy, security, and service constraints. It distinguishes centralized data pooling, federated learning, and a federated design with secure aggregation and a properly specified differential privacy mechanism. A hypothetical monthly decision example compares the alternatives using stated assumptions: their respective expected incremental values are 104, 112, and 106 thousand US dollars, making ordinary federation the highest-valued option in the baseline scenario. The ranking changes when the severity of modeled incidents increases, and privacy requirements may independently exclude an option. A separate one-release Gaussian mechanism illustrates how a formal user-level privacy statement depends on the unit of protection, clipping, the fixed denominator, and the trust boundary; it does not establish a privacy guarantee for multi-round federated training. The calculations are analytical illustrations rather than results from an implemented system or observed platform data. The article concludes with an evaluation protocol addressing causal business impact, heterogeneous participants, fairness, privacy accounting, and deployment governance.

Keywords: mobility analytics; federated learning; differential privacy; platform operations; privacy–utility trade-off; business value; location privacy; secure aggregation

Business Value of Privacy-Preserving Mobility Analytics in Platform Operations: A Federated Risk-Utility Framework

1. Introduction

Digital mobility platforms coordinate riders, drivers, vehicles, couriers, and other resources across space and time. To plan an operation, a platform must anticipate where requests will arise, where capacity will be available, and how long matching or movement will take. These choices affect completed trips, waiting times, empty travel, operating expense, and the quality of service experienced in different neighborhoods. Forecasting demand is therefore valuable only to the extent that it changes decisions or outcomes; a more accurate model does not automatically produce a financially or socially preferable allocation (Chen, Thakuriah, & Ampountolas, 2021; Ma et al., 2022; Rahman & Rifaat, 2021; Wang, Zheng, et al., 2024).

The information needed for these tasks can be sensitive. A sequence of locations may reveal repeated visits, commuting routines, workplace patterns, or other personal circumstances. Research on human mobility shows both the regularity of movement and the identifying power of sparse spatiotemporal observations (de Montjoye et al., 2013; González et al., 2008; Song et al., 2010). Removing names or replacing them with pseudonyms is not necessarily sufficient when records can be linked with outside information. Platform operators must therefore consider not only the value of prediction but also what information is collected, who can observe it, how long it is retained, and what inferences are possible from the outputs.

Federated learning offers a different arrangement for analysis: participating devices or organizations perform computations locally and exchange selected updates instead of routinely transferring their raw records to one central training repository. Mobility applications have examined federated next-location prediction, travel-mode classification, traffic forecasting, and ride-hailing demand estimation (Ezequiel et al., 2022; Feng et al., 2020; Liao et al., 2024; Yu et al., 2022; Zhang, Zhang, et al., 2021). The architecture can reduce some forms of raw-data movement, but it does not by itself ensure confidentiality or guarantee that a user's participation or behavior cannot be inferred. The data may also remain exposed at a client, and an update or final model can reveal information under certain adversarial conditions (Hatamizadeh et al., 2023; Shokri et al., 2017; Yin et al., 2022).

This distinction is important for business evaluation. A federated arrangement may incur additional communication, coordination, device, engineering, and auditing costs. Stronger privacy controls can alter predictive quality or operational coverage, while potentially reducing certain exposures or enabling collaboration that would otherwise be unavailable. The relationship is context dependent: a model that improves a high-volume area may worsen performance in a smaller community, and a privacy mechanism that protects an external recipient may not protect information from a trusted processing server. Comparisons that report only aggregate accuracy or only whether raw records moved do not capture these decision consequences.

This paper asks how a platform can evaluate privacy-preserving mobility analytics as an operational investment while respecting non-negotiable protections. It develops a federated risk–utility framework with four linked elements: an explicit mobility decision and baseline; a map of participants, data flows, protected units, and adversaries; a business value calculation with transparent cost and incident assumptions; and a feasibility screen for privacy, security, service, and governance requirements. A compact numerical example shows how changes in incident severity can reverse a ranking. A separate mathematical example explains a limited differential privacy construction without claiming that a full federated training system has been implemented.

This methodological paper reports no proprietary data, field experiment, or measured return. It develops reproducible illustrations and an evaluation agenda for later testing by a platform and its partners.

2. Literature Review

2.1. Mobility analytics and platform operations

Human mobility research has documented temporal and spatial regularities in movement, providing a foundation for forecasting and aggregate network analysis (González et al., 2008; Song et al., 2010). Yet predictability depends on the data source, temporal and spatial resolution, and the population represented. A forecast of demand in one area cannot automatically be transferred to another without considering differences in geography, available transport, events, and participation. Origin–destination forecasting and research on ride-hailing demand demonstrate the usefulness of spatiotemporal information while also illustrating the importance of the chosen prediction target (Chen, Thakuriah, & Ampountolas, 2021; Wang, Zheng, et al., 2024).

Operational value requires a decision that uses the forecast. A platform can change how it positions vehicles, schedules drivers, allocates an incentive budget, plans charging, or communicates expected waiting times. Work on pricing and demand–supply gaps shows why the distribution of demand and capacity matters alongside the level of aggregate demand (Ma et al., 2022; Rahman & Rifaat, 2021). Studies of ride-sourcing demand and shared-ride patterns also emphasize that observed requests and realized trips reflect supply constraints and spatial variation, making them imperfect proxies for latent demand (Dean & Kockelman, 2021; Krueger et al., 2023).

Mobility data have applications beyond ride-hailing, including public transit, logistics, and multimodal planning. Different tasks require different information: a county-level flow matrix, a short-term zone demand forecast, and an individual's next-location prediction have different privacy and utility implications. A privacy-preserving release should be assessed against the intended operational use, not solely against a generic similarity score. Research comparing anonymized mobility representations and their suitability for specific applications supports this task-oriented view (Gibbs et al., 2026; Mattos et al., 2024).

2.2. Federated mobility analytics

Federated learning coordinates training across participants whose data remain under local control. General treatments describe its communication, optimization, privacy, and systems challenges (Kairouz et al., 2021; Li et al., 2020; Yang et al., 2019). The term covers cross-device settings with many users and cross-silo settings with a smaller number of organizations. These settings differ in participation, trust, legal authority, computational capacity, and the degree to which a user's data may appear at more than one participant. Keeping raw records local can be useful, but it is an architectural property rather than a complete privacy definition.

Mobility research has explored privacy-aware next-location modeling and transportation mode prediction, as well as cross-silo approaches for human movement (Feng et al., 2020; Liu et al., 2021; Yu et al., 2022). More recent work applies federated approaches to traffic graphs, vehicle-flow prediction, passenger flow, and ride-hailing demand (Feng et al., 2024; Hu et al., 2024; Liao et al., 2024; Rocco di Torrepadula et al., 2025; Zhang, Zhang, et al., 2021). A study of shared mobility also demonstrates practical challenges in coordinating matching and analytics across organizations (Wei et al., 2024). These contributions show the feasibility of studying mobility tasks without a single raw-data repository, but their predictive results and threat models are not interchangeable across platforms.

Statistical heterogeneity is a central limitation. Cities and service providers may have different road networks, traveler populations, fleet sizes, sampling rates, and behavior; training data are often not independent and identically distributed. Research on non-identically distributed data, adaptive federated optimization, and personalized cross-silo learning addresses some of these challenges (Huang et al., 2021; Zhang, Hong, et al., 2021; Zhu et al., 2021). In an operation, heterogeneity can cause a global model to work well on high-volume

participants while offering little value or poorer service to a smaller participant. Accuracy should therefore be evaluated by participant and relevant service group as well as in aggregate.

2.3. Privacy risks and protection mechanisms

Location privacy research demonstrates that apparently anonymous mobility traces can be distinctive and potentially linkable to individuals (de Montjoye et al., 2013; Shokri et al., 2011). Threats in federated learning extend beyond raw data transfer: an adversary observing model updates, model outputs, or changes across rounds may attempt to infer training membership or reconstruct attributes. The practical strength of such attacks depends on the model, information available to the adversary, batch construction, and other assumptions; a successful attack in one setting is not proof that every federated system is equally vulnerable (Hatamizadeh et al., 2023; Li et al., 2023; Shokri et al., 2017).

Secure aggregation allows a server to learn an aggregate of participant contributions while limiting access to individual contributions under the protocol's security assumptions (Bonawitz et al., 2017). It does not automatically hide the aggregate, eliminate leakage across repeated rounds, or prevent an adversarial participant from influencing a model. Research has demonstrated possible inference from aggregated gradients under particular assumptions (Wang, Chang, et al., 2024). Threat modeling must therefore address the server, participants, recipients of the released model, collusion, dropout, participation metadata, and the possibility of malicious updates.

Differential privacy (DP) gives a mathematical bound on how much the distribution of a specified released output can change when one protected unit's data are changed under a defined adjacency relation (Dwork & Roth, 2014; Dwork et al., 2006). It is not the probability of a breach, and the parameter ε is not a measure of how many individuals have been identified. Practical training methods use clipping, randomization, and privacy accounting to bound the effect of protected contributions (Abadi et al., 2016). Location-oriented methods such as geo-indistinguishability use related ideas for location releases under their own definitions and distance metrics (Andrés et al., 2013). A privacy statement is meaningful only when its protected unit, mechanism, outputs, parameters, and assumptions are stated.

Privacy mechanisms may affect utility unevenly. A recent study of local differential privacy for federated mobility data finds that the trade-off depends on the data and the selected analytics, with low-frequency flows especially susceptible to noise in the scenarios studied (Gibbs et al., 2026). Research on differentially private mobility forecasting examines where perturbation is applied and how that choice affects predictive output (Arcolezi et al., 2022). Neither work implies a universal privacy–accuracy curve for every platform. The platform's operational decisions, the distribution of demand, and the treatment of small groups determine whether observed differences matter in practice.

2.4. Business value and the research gap

Federated mobility studies often emphasize prediction, communication, or technical privacy mechanisms (Du et al., 2025; Feng et al., 2024; Hu et al., 2024; Weng et al., 2025; Yin et al., 2022). A platform investment decision also requires a common baseline, operational gains, implementation costs, and residual risks.

This article contributes a synthesis rather than claiming to invent federated learning, differential privacy, or mobility forecasting. It separates hard privacy and service requirements from optional monetary risk valuation; traces the path from data and model outputs to operational decisions; and sets out a common comparison of alternative architectures. The numerical examples are designed to expose assumptions and thresholds rather than to estimate a real platform's performance. This design also makes it possible to identify when the data needed for a defensible choice do not yet exist.

3. Platform Setting and Analytical Architecture

3.1. Operational setting and participants

Consider a platform that allocates service capacity across geographic zones and time intervals. In a ride-hailing application, the capacity could be available driver time; in logistics, it could be courier availability or vehicle loading slots. A forecast estimates requests, travel times, or a related demand–supply gap, and an allocation rule chooses a feasible response. The reference case is an existing operational policy without the proposed additional mobility analytics. This baseline can still use ordinary platform data and retain existing controls; the model evaluates incremental changes rather than assuming that the platform previously had no privacy or security risk.

The participating entities may include user devices, transport operators, platform business units, infrastructure providers, a coordinating server, and authorized recipients of an aggregate or trained model. In a cross-silo design, each operator retains its own records and computes a local update. In a cross-device design, devices may contribute directly subject to availability and permission. Figure 1 summarizes the conceptual information and decision flows. It is not a description of an existing deployment, and no specific organization is assumed to participate.

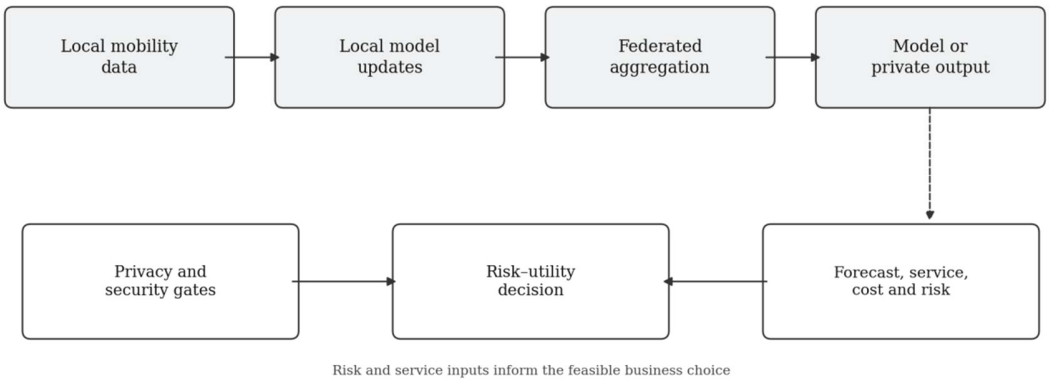


Figure 1. Conceptual Data Flows and Decision Inputs in the Federated Risk–Utility Framework

The purpose of the system should be specific: for example, a 30-minute forecast of zone-level requests to support fleet positioning. Data sources might include trip origins and destinations, timestamps, fleet availability, road conditions, and public contextual information. Before any use, the platform must establish the authority and permissions needed for the purpose and limit collection and retention to what the approved design requires. The existence of a technical privacy mechanism does not establish a lawful basis for data collection or justify a use that has not been authorized.

The architecture should distinguish at least four boundaries. The first is the boundary between a person and the local device or organization holding that person's information. The second is the boundary between a participant and the coordinator receiving an update or aggregate. The third is the boundary between the coordinator and downstream model users. The fourth is the boundary between the prediction and the operational action. A mechanism may protect one boundary while leaving another exposed; specifying these boundaries prevents a general "privacy-preserving" label from concealing an important disclosure.

3.2. Decision alternatives and scope

The framework compares four stylized alternatives. Alternative O continues the baseline policy without the proposed new analytics. Alternative C pools specified raw records at a central location for training, subject to the same legal and organizational feasibility screen as other alternatives. Alternative F trains

collaboratively while retaining raw records at participants, without assuming a formal DP guarantee. Alternative P combines a federated arrangement with secure aggregation and a specified DP mechanism for the outputs covered by that mechanism. P is a family of designs, not a guarantee that every feature or output is private.

Each alternative should be evaluated for the same service, planning horizon, geographic coverage, data access permissions, and baseline. C can be included as a counterfactual for research, but it is not eligible for deployment when pooling is prohibited or unauthorized. If F or P cannot meet a minimum participation threshold or satisfy an approved operating condition, that option is also infeasible. Alternative O remains a legitimate comparator when the proposed analytics produce too little net value or acceptable protections cannot be established.

Local model training can be represented by an objective of the form $F(\theta) = \Sigma_k w_k F_k(\theta)$, where k identifies a participant, F_k is its local loss, and nonnegative weights w_k sum to one. This objective is a modeling convention: it does not guarantee convergence, privacy, or equitable performance. A participant with more observations can dominate a size-weighted model, while equal weighting can reduce aggregate performance or reward participants with few or low-quality records. The choice should reflect the target population and be accompanied by participant-level performance reporting (Huang et al., 2022; Kairouz et al., 2021; Zhu et al., 2021).

In practice, local objectives can change with season, geography, data availability, and service policy. Devices can be unavailable or drop out, creating stale updates and biased participation. Asynchronous federated methods address some timing issues, but they do not remove the need to check which populations are represented in an update or evaluation sample (Chen, Mao, & Ma, 2021). A site-level model might be preferable to a single global model in some locations; personalization should be assessed against operational impact, communication costs, and the possibility of additional privacy leakage (Huang et al., 2021; Zhang, Hong, et al., 2021).

3.3. Threat model and control layers

For each use case, the threat model should describe what an adversary can observe and control. A curious coordinator may see an aggregate and metadata; a malicious participant may submit a crafted update; an outside recipient may query a released model; and a compromised endpoint may expose local records. The model should specify whether parties can collude, whether multiple rounds are visible, what side information is available, and which infrastructure is trusted. No single attack simulation exhausts this space, and low measured attack success under one test is not proof that a system is secure.

Controls should be layered according to the stated threat. Data minimization and access controls restrict collection and use. Secure aggregation can limit access to individual updates where protocol assumptions hold. A correctly implemented DP mechanism can limit the effect of one protected unit on a defined output. Integrity checks, model validation, and monitoring address different problems, such as corrupted contributions and operational failures. Secure aggregation can make individual inspection harder, so a design must explicitly reconcile confidentiality with abuse detection and robust aggregation instead of assuming that each control independently solves the other's problem (Bonawitz et al., 2017; Chen et al., 2025; Wang, Chang, et al., 2024).

Table 1 compares architectural choices and their protection boundaries. Federation, secure aggregation, and differential privacy should be reported separately; none establishes the guarantees of the others.

Table 1. Architecture Choices and Protection Boundaries

Approach	Basic arrangement	Interpretation
Central pooling (C)	Designated raw records can move to a central repository.	Privacy depends on access, processing controls and any additional mechanism.
Federation (F)	Participants retain raw records and exchange selected updates.	Updates, metadata and final models may still leak; no DP bound is implied.
Secure aggregation	A protocol limits access to individual inputs under stated assumptions.	The aggregate and repeated releases may still expose information.
Differential privacy	A specified release randomizes a bounded contribution.	The protected unit, outputs, parameters, accounting and observer must be defined.
Baseline (O)	Current operational policy without the proposed new analytics.	Existing security and privacy exposures remain.

3.4. Protected units and observable outputs

A protected unit could be one location record, one trip, one device, one account, one organization, or an individual across an entire observation window. These choices provide different guarantees. Record-level clipping does not automatically protect a person who contributes many trips; device-level protection may not cover the same person on several devices; and cross-silo protection may not cover a person whose records occur at multiple providers. The chosen unit must match the risk being discussed and the actual data architecture.

The observable outputs must be enumerated as carefully as the input. A final model, an intermediate update, a validation statistic, a hyperparameter selection result, a participation list, a zone-level dashboard, and a sequence of revised predictions can all reveal information. If each release accesses the same protected data, its privacy impact must be included in the accounting when a DP guarantee is claimed. Publicly stated bounds, side information, data-dependent selection, and the policy for low-participation groups can also affect the validity or interpretation of a guarantee.

4. Federated Risk–Utility Framework

4.1. From forecast to operational utility

Let a denote an analytics architecture, z a service zone, and t a decision period. The model produced under a yields a forecast that informs a feasible action $x_{zt}(a)$, such as assigning driver capacity or setting a dispatch schedule. Let D_{zt} be realized demand and let $\text{completed}_{zt}(a)$ be the trips or orders actually served under that action. In a simplified capacity-limited model, $\text{completed}_{zt}(a)$ cannot exceed either realized demand or effective available capacity. More detailed dispatch, travel time, and network constraints can replace this approximation.

Let v_{zt} be the contribution margin per additional completed service, excluding costs already counted separately, and let K_a represent changes in repositioning, overtime, incentives, or related operational spending. The expected incremental operational gain G_a over the baseline can be represented as:

$$G_a = E[\Sigma_{z,t}\, v_{zt} \times (\text{completed}_{zt}(a) - \text{completed}_{zt}(O)) - K_a]. \quad (1)$$

The equation does not imply that all completed trips have equal value, that every additional forecast improvement creates an additional trip, or that observed trips equal unconstrained demand. The evaluation must specify the demand and capacity model, accounting treatment, time horizon, and any operational

constraints. Revenue and variable costs should not be counted twice. The quality of a forecast is a possible input to G_a; it is not itself a monetary benefit.

Waiting time, cancellations, coverage, driver idle time, and missed service should be reported alongside the financial metric. A platform can monetize a penalty only when it has a defensible accounting basis and when the same consequence is not already embedded in contribution margin or another cost. Privacy impacts that cannot be quantified or should not be traded against profit remain separate decision criteria. This multi-outcome view is particularly important where a small area or a low-volume participant receives worse service despite a better overall average (Dean & Kockelman, 2021; Hong et al., 2021; Huang et al., 2022).

4.2. Cost and residual risk

Let C_a include incremental engineering, local computation, network transmission, coordination, privacy accounting, secure aggregation, monitoring, independent review, and other operating costs attributable to alternative a. Capital spending can be converted to a consistent per-period cost using a stated amortization assumption; this should not be mixed with a one-time undiscounted benefit. If participants bear device, staff, or electricity costs that the platform does not pay directly, the accounting boundary should identify both the platform's cost and the broader partner cost when relevant.

For a simplified incident class e, let q_{ae} be the probability of the event within the evaluation period and L_{ae} its conditional incremental loss if it occurs. The expected incremental incident cost is $\Sigma_e\, q_{ae}\, L_{ae}$, provided the event classes and loss definitions are specified so that overlap is handled correctly. A full risk model may use a joint distribution, including dependence between events and costs. The probability and loss are scenario inputs to be estimated from evidence or expert judgment; they cannot be read directly from ε or from a model's forecast error.

The resulting risk-adjusted incremental value V_a is:

$$V_a = G_a - C_a - \Sigma_e\, q_{ae}\, L_{ae}. \quad (2)$$

Equation (2) is one decision aid, not a complete valuation of privacy. A breach event, a misuse of location data, discrimination, or loss of autonomy can have consequences that are hard to estimate or monetize. If an organization adds a separate privacy-related monetary term, it must define what it covers and avoid including the same consequence again in L_{ae}. Where a credible monetary estimate is unavailable, the analysis should report a scenario range and retain unpriced harms in the feasibility and review process.

Expected value is sensitive to rare high-impact events. Decision makers can supplement V_a with a loss distribution, stress scenarios, and a tail-risk measure, provided the assumptions and limitations of each are clear. A single expected-loss number does not show whether the incident probability comes from a validated model or from a speculative judgment. In a data-scarce setting, sensitivity analysis can identify the combinations of probability and severity under which the preferred option changes; it cannot establish the true likelihood of an attack (Wang, Chang, et al., 2024; Weng et al., 2025).

4.3. Feasibility, privacy, and service constraints

The preferred design is the feasible alternative with the highest value under the selected evaluation criteria. Let $A_{feasible}$ be the set of designs that satisfy the applicable purpose, data access permissions, security requirements, privacy assurances, operational restrictions, and independently established service obligations. A stylized decision is:

$$a^* = \arg\max_{a \in A_{feasible}} V_a. \quad (3)$$

$A_{feasible}$ is not defined by the model alone. It must be established through the organization's authorization, legal, security, and governance processes. For a specific release, a policy may require a user-level privacy bound, an approved access arrangement, or a minimum number of independent contributors. If C is not

authorized, it is excluded even when its estimated value is largest. If no new design meets the applicable requirements, O can be selected or the proposal can be revised.

Service requirements may specify a maximum wait-time disparity, minimum coverage, or a limit on service degradation for a defined area. Such thresholds require careful selection and evidence; a single fairness metric cannot capture every distributional harm. Measuring a protected group may itself require sensitive information, so the organization needs a defensible, protected measurement process and should avoid collecting more sensitive data simply to make a metric convenient. Fair federated learning research illustrates the tension between overall performance and performance across participants (Hong et al., 2021; Huang et al., 2022).

The framework also supports participation and consent constraints. A mobility provider may agree to contribute for a specified planning task but not for individualized advertising or law-enforcement use. Such boundaries should be reflected in contracts, access controls, data retention, and technical enforcement as appropriate, rather than treated as an optional cost. Organizations may decide that certain location-derived inferences should not be produced even if their statistical error is low and their estimated business value is high.

4.4. Uncertainty and decision robustness

Let Θ represent uncertain values for participation, forecast-to-decision conversion, operating costs, incident probabilities, and losses. A scenario analysis computes $V_a(\Theta)$ over a documented set of plausible parameter combinations. A conservative organization may compare the worst plausible value of each feasible action, while another may examine the distribution of net values or the frequency with which an action ranks first. These outputs answer different questions; none substitutes for evidence about whether the chosen range is credible.

Scenario construction should preserve relationships between inputs. Better privacy controls might reduce one incident exposure but increase operational cost; participant dropout can affect both accuracy and secure-aggregation thresholds; and a common network failure can affect several partners simultaneously. Changing each variable independently without a plausible joint story can generate unrealistic combinations. A decision can be labeled robust only relative to the stated uncertainty set, not universally robust to all attacks, market changes, or implementation defects.

The comparison should also distinguish correlation from causal effect. Platforms with better data infrastructure may have better outcomes for reasons unrelated to federation. Users who opt into data collection may differ from nonparticipants, and a forecast-based intervention can change future observed demand through pricing or availability. Business gains require a credible counterfactual, such as a randomized operational trial where appropriate or a carefully designed quasi-experiment. A simulation can be useful for planning, but its outputs remain conditional on the simulated environment.

4.5. Information needed for an operational decision

A practical decision record should report the service being optimized, the baseline, participants and data rights, model version, privacy unit, threat model, feasible actions, expected operational effect, cost assumptions, risk scenarios, and unresolved uncertainties. The record should identify who can inspect a raw contribution, an unnoised aggregate, a noisy release, and the final model. It should also list any outputs outside the scope of the DP accounting and any dependencies on a trusted coordinator.

5. Illustrative Risk–Utility Analysis

5.1. Hypothetical platform and assumptions

Consider a fictional platform evaluating one month of zone-level demand forecasting for service allocation. The numbers in this section are invented for explanation and are not derived from a dataset, a field trial, or measured effects of differential privacy. Alternative O is the reference case. Alternatives C, F, and P are assumed to be technically and institutionally feasible for the numerical comparison; if a required privacy or authorization condition fails, the corresponding alternative must be removed before the ranking is used.

The expected incremental operational gain G_a is 240 thousand US dollars for C, 218 for F, and 195 for P. These values are assumed net of changes in routine service-delivery and repositioning expenses, but before the separate architecture costs C_a. The architecture costs are 40, 50, and 62 thousand dollars, respectively. The lower assumed gain for P is a scenario assumption about its combined operational performance, not a theorem that privacy protection always reduces value or that the stated differential is representative of real platforms.

For illustration, suppose the only modeled incremental incident class has probabilities of 0.040, 0.028, and 0.015 for C, F, and P, respectively, with conditional incremental losses of 2,400, 2,000, and 1,800 thousand dollars. These probabilities and losses are hypothetical business risk inputs; they are not inferred from ε, secure aggregation, historical breach data, or an evaluated system. The baseline's existing risks are outside this incremental illustration. Table 2 provides the complete assumptions, with all monetary amounts expressed in thousands of US dollars per month.

Table 2. Hypothetical Monthly Inputs and Values (US$ Thousands)

Alternative	Gain G	Cost C	Event q	Loss L	Value V
O baseline	0	0	—	—	0
C centralized	240	40	0.040	2,400	104
F federated	218	50	0.028	2,000	112
P enhanced	195	62	0.015	1,800	106

5.2. Expected value comparison

Applying equation (2), the expected incident losses for C, F, and P are 0.040 × 2,400 = 96; 0.028 × 2,000 = 56; and 0.015 × 1,800 = 27. Their corresponding values are 240 − 40 − 96 = 104; 218 − 50 − 56 = 112; and 195 − 62 − 27 = 106. The baseline O has incremental value zero by definition. Under these assumptions and with all alternatives in the feasible set, F ranks first, followed by P and C. The ranking is not evidence that ordinary federation is generally preferable.

The comparison makes the trade-off visible. Relative to F, P is assumed to reduce operational gain by 23 and add 12 in architecture cost, while reducing expected incident loss by 29. Its net value is therefore 6 lower than F at the baseline assumptions. This calculation omits unpriced privacy and social consequences. If an applicable condition requires a privacy guarantee that only P satisfies among the modeled options, P is selected even though F has the higher estimated monetary value. If P also fails a required condition, the choice is between O and any other eligible design.

A potential increase in collaboration can instead make the privacy-enhanced option more valuable. Suppose, solely as an alternative scenario, that P enables authorized participation by an additional partner and increases its operational gain from 195 to 210, with its other values unchanged. Its value becomes 210 − 62 − 27 = 121, exceeding F's 112. This scenario is not evidence that a partner would participate or that its data would

produce a 15-thousand-dollar gain; it shows where a measurable collaboration effect would enter the accounting.

5.3. Sensitivity to incident severity

Let κ multiply the conditional incident losses for all three new architectures while their probabilities and other assumptions remain unchanged. It is a sensitivity parameter, not a change in the probability of a breach and not a measure of privacy protection. The value functions are $V_C = 200 - 96\kappa$, $V_F = 168 - 56\kappa$, and $V_P = 133 - 27\kappa$, with $V_O = 0$. Table 3 reports several scenarios. At $\kappa = 1.5$, the values of C, F, and P are 56, 84, and 92.5, making P the highest-valued feasible option among those modeled.

Table 3. Sensitivity to Incident Severity (US$ Thousands)

κ	C	F	P	O	Highest value
0.5	152.0	140.0	119.5	0.0	Centralized
1.0	104.0	112.0	106.0	0.0	Federated
1.5	56.0	84.0	92.5	0.0	Enhanced
5.0	-280.0	-112.0	-2.0	0.0	Baseline

C and F have equal value at $\kappa = 0.80$; F and P are equal at $\kappa = 35/29$, approximately 1.207. With the stated assumptions and O available, C has the largest value for $0 \leq \kappa < 0.80$, F for $0.80 < \kappa < 35/29$, and P for $35/29 < \kappa < 133/27$. Above $\kappa = 133/27$, approximately 4.926, O is highest because all three new designs have negative expected incremental values. At the boundary values the corresponding alternatives tie. The regimes depend entirely on the assumed linear severity scaling and do not describe a universal relationship.

When credible parameter ranges cross these thresholds, the ranking is unstable and further evidence or a pilot may be appropriate. Varying one input at a time is useful for explanation, but actual costs, participation, and risks may be correlated.

5.4. Limits of the numerical comparison

The three architecture labels summarize many possible implementations. In practice, a centralized system may use strong access and DP controls, and a federated system may use weak controls or retain sensitive intermediate outputs. The assumed probabilities and costs therefore describe particular hypothetical scenarios, not inherent properties of the labels. Security posture, data access, participant behavior, incident response, and the value and sensitivity of the protected information can all change the ranking.

6. A Separate Differential Privacy Illustration

6.1. A one-release mechanism

To make a formal privacy statement concrete, consider a separate hypothetical release of one aggregate vector. A fixed public roster contains m = 1,000 slots, each corresponding to at most one protected user who affects no other slot. For each slot i, a local procedure depending only on that slot's data and public inputs produces a vector u_i clipped to have Euclidean norm at most C = 1; an absent contribution is represented by a zero vector. The released statistic is the fixed-denominator mean plus independent Gaussian noise in each coordinate:

$$M(D) = (1/m) \times \Sigma_i u_i + Z, \text{ where } Z \sim N(0, \sigma^2 I). \quad (4)$$

The protected adjacency relation changes one slot's complete contribution between a bounded vector and zero. Under these assumptions, the L2 sensitivity of the mean is at most $\Delta = C/m = 0.001$. The fixed roster

and denominator are public and do not themselves hide the fact that a named roster member exists. The construction protects the specified contribution to the released statistic; it does not protect public enrollment, unprotected metadata, or other outputs.

For $\varepsilon = 0.8$ and $\delta = 10^{-5}$, the classical sufficient Gaussian bound gives $\sigma \geq \Delta \times \sqrt{[2 \ln(1.25/\delta)]}/\varepsilon$, or approximately 0.00606 per coordinate for the stated Δ. The bound applies under its stated conditions, including $0 < \varepsilon < 1$, and can be conservative relative to other valid analyses. The numerical parameters specify one mechanism and one release; they are not a recommended privacy budget for any real population. This example follows standard sensitivity-based DP analysis (Dwork & Roth, 2014; Dwork et al., 2006).

If adjacency instead allows replacement of any bounded vector by another bounded vector, the sensitivity can be as large as $2C/m$, so the corresponding sufficient σ is approximately 0.01211 with the same ε and δ. The relevant adjacency must be fixed before reporting the guarantee. If the denominator is data dependent, if a person affects several slots, or if a participant contributes an unclipped vector, the simple sensitivity argument no longer applies as stated. A practical deployment would require a valid enforcement and identity design, including an analysis of users appearing across devices or providers.

6.2. Trust boundary and repeated releases

A secure-aggregation protocol may let the coordinator obtain the sum without observing each individual vector, subject to its threat model, threshold, collusion assumptions, and correct implementation (Bonawitz et al., 2017). If the coordinator then adds the Gaussian noise, the formal DP claim above applies to the released noisy statistic as seen by its recipients. The coordinator that observes the exact pre-noise aggregate is inside the trusted boundary for that central-DP statement. Protecting the aggregate from that coordinator requires a different design, such as appropriately verified distributed noise generation, and a separate security and privacy analysis.

Clipping also requires attention. If a client can submit an arbitrary vector, the stated sensitivity bound is not guaranteed merely because the protocol specification asks for clipping. An implementation must enforce or justify the bound using an appropriate trust or verification mechanism; secure aggregation alone can hide a violating contribution from the server. Participant selection, dropout, threshold failures, and participation metadata require their own analysis. Security and privacy properties should be validated for the actual protocol rather than inferred from an architectural diagram.

A multi-round federated training process is different from this one-release example. Repeated access to protected data and publication of intermediate or final outputs require composition and an appropriate accountant. Sampling-based amplification, if used, relies on the actual sampling process and its assumptions. Hyperparameter selection and validation releases must be included where they consume privacy budget. DP-SGD and other mechanisms provide relevant methods, but their guarantees depend on how clipping, noise, sampling, and accounting are implemented (Abadi et al., 2016; Weng et al., 2025).

The example does not prove that alternative P in Section 5 achieves $\varepsilon = 0.8$, nor does it calculate the accuracy of a trained model or convert ε into an incident probability. It illustrates the information a reviewer needs before accepting a DP claim. Likewise, failure of a tested reconstruction or membership attack under one scenario does not establish a formal DP guarantee. A responsible evaluation reports both the formal scope of the guarantee and implementation-specific attack and security tests (Hatamizadeh et al., 2023; Li et al., 2023; Shokri et al., 2017).

6.3. Utility and equity under privacy protection

The noise standard deviation of a released mean is one component of utility, not a complete description of operational quality. Its relative impact can be large for low-volume zones or rare origin–destination pairs, especially when the resulting statistic is used to allocate a scarce vehicle. Clipping may also introduce bias when high-contribution users are more common in certain zones. Evaluation should therefore report error

and decision outcomes for a range of volumes, times, participant types, and geographic groups, subject to protections for the evaluation data itself (Arcolezi et al., 2022; Gibbs et al., 2026; Mattos et al., 2024).

A platform can consider alternatives such as a coarser time interval, a larger geographic cell, improved sampling, or a different approved task when the privacy–utility trade-off is unfavorable. These choices can change the purpose and distribution of benefits: a coarser map might conceal useful service gaps even as it improves aggregate accuracy. They should not be treated as automatic remedies. A reduction in a single error metric or a smaller privacy parameter is useful evidence only in relation to the operational question, protected unit, and affected population.

7. Managerial Implications and Governance

7.1. Business case and contracting

A credible business case begins with the decision rather than the available data. The platform should define the operational objective, the policy currently used, the action that can change, and the smallest set of data required to evaluate that action. A useful question is not simply whether a federated forecast has lower mean absolute error, but whether it improves completed service, waiting time, or capacity utilization after the costs of acting on the forecast are included. Demand and supply constraints can make an improvement operationally irrelevant in one zone and valuable in another (Krueger et al., 2023; Ma et al., 2022; Rahman & Rifaat, 2021).

In a multi-operator federation, a technical protocol must be supported by a clear allocation of responsibilities. Participants should document permitted purposes, local processing, model and update ownership, retention, access to aggregate outputs, incident reporting, withdrawal procedures, and how system changes are approved. They should also decide who pays for computation, connectivity, privacy auditing, and remediation. A partner may rationally refuse participation if the private cost or perceived exposure exceeds its benefit even when the aggregate platform value is positive.

Cost allocation and benefit distribution can change the incentives to contribute. A partner serving a low-volume area may bear the cost of participation while receiving little improvement from a global model; a large partner may gain disproportionately from others' data. Transparent reporting by participant, accompanied by appropriate privacy protection for the reports, can support negotiation. A shared model should be reviewed for both the aggregate operating result and whether particular partners or service areas consistently experience a loss (Huang et al., 2021; Huang et al., 2022; Zhu et al., 2021).

7.2. Privacy assurance and security operations

An assurance case should document what is protected, from whom, for how long, and by which mechanisms. The organization should retain evidence about clipping enforcement, the noise-generation process, the privacy accountant, aggregate thresholds, versioned code, and the treatment of validation outputs when a DP claim is made. It should also identify what a coordinator, client, service provider, or model recipient can observe. Explanations to users should be accurate about these boundaries and avoid implying that location information becomes harmless merely because it stays on a device.

Operational security remains necessary throughout the system. Local storage and endpoints need access controls; model updates require authentication and integrity protection; servers and key material require secure administration; and monitoring must cover protocol failures and suspicious changes. Secure aggregation, differential privacy, and adversarial robustness address different failure modes. In particular, a mechanism that hides individual updates from the coordinator can complicate inspection for malicious updates, so an approved design must explicitly account for this trade-off (Bonawitz et al., 2017; Chen et al., 2025; Wang, Chang, et al., 2024).

Data retention and withdrawal require attention to both local data and derived artifacts. Deleting raw observations from one device does not automatically remove their influence from a model trained previously, and removal from a released aggregate may not be technically possible. Organizations should communicate these practical limits and implement the deletion, retraining, access, and retention procedures required for their setting. A formal DP guarantee is not a replacement for a comprehensive lifecycle plan.

7.3. Service quality, participation, and equity

Federated participation can be uneven because of connectivity, battery constraints, working hours, institutional capacity, or willingness to share. The participants represented in an update may therefore differ from the customers or drivers affected by the operational decision. A model that performs well for frequent users or well-connected partners can produce a misleading picture of overall service. Evaluation should distinguish data coverage, model accuracy, operational outcomes, and participation rates instead of assuming that one measure establishes representativeness.

The platform should define service-level guardrails before comparing designs. Examples include maximum acceptable changes in waiting time or cancellations across operational zones, limits on concentrating empty travel in particular areas, and minimum coverage for participating partners. These are examples of governance choices, not universal thresholds. Protected assessment methods should be used for small groups, and uncertainty intervals should accompany subgroup results when data are limited. Research on fairness in federated systems highlights the importance of evaluating disparate participant outcomes (Hong et al., 2021; Huang et al., 2022).

Where the data indicate a persistent disadvantage, a platform can test alternative aggregation weights, personalized models, different capacity policies, or additional local data collection only if authorized. Each remedy can create new costs and privacy consequences. Weighting smaller groups more heavily may improve one service metric while affecting another; personalized models may increase the number of released artifacts that require privacy accounting. The decision record should show these trade-offs and the roles authorized to accept them.

7.4. Deployment stages and monitoring

A staged deployment can reduce the risk of confusing an analytical prototype with a reliable operational system. First, teams can map the use case and obtain relevant approvals. Second, they can evaluate models and privacy mechanisms in an isolated environment. Third, they can run a limited pilot with predefined stopping rules and protected logging. Finally, they can monitor an approved release for drift in participation, prediction performance, incident indicators, and service outcomes. A pilot is not a substitute for a security review or an applicable formal privacy analysis.

Monitoring should be proportionate to the purpose and should not recreate a centralized repository of the sensitive data the design sought to avoid. Aggregated or differentially private monitoring outputs can help, but they require their own privacy accounting and sufficient sample sizes. A sharp decline in local participation, an increase in protocol failures, or a shift in local forecast error can signal that an earlier business case no longer applies. The platform should specify rollback criteria, the responsible owner, and a process for independent review.

8. Evaluation Agenda and Limitations

8.1. Proposed empirical evaluation

The analysis presented here does not implement a federated model or use real platform observations. A future evaluation should predefine the mobility task, time horizon, participant population, permitted data, and the baseline operating policy. It should compare centralized, local-only, federated, and privacy-enhanced alternatives only when each can be evaluated under appropriate permissions. Temporal splits and, where

feasible, geographic or participant holdouts can test whether a model generalizes beyond the data on which it was developed.

Predictive metrics should include task-relevant errors and calibration, with breakdowns by geography, time, demand volume, and participating organization. Operational metrics should include completed service, cancellations, waiting time, travel or repositioning effort, and costs. These measures should be linked to the same unit and horizon as the financial model. A forecast that improves average error but changes no feasible dispatch decision should not be credited with a speculative revenue gain.

To estimate business impact, a platform could randomize an approved dispatch policy across comparable operational units or time windows where this is safe, ethical, and operationally feasible. Spillovers between adjacent zones, shared drivers, seasonality, and changes in demand can invalidate simple comparisons, so the design should anticipate interference and use appropriate uncertainty estimates. When randomization is infeasible, a documented quasi-experimental or simulation approach can be used with explicit assumptions; simulated outcomes should not be described as observed commercial benefits.

Privacy evaluation should combine a formal review of the claimed mechanism with targeted implementation and attack tests. The review should cover the protected unit, adjacency, contribution limits, randomness, composition, data-dependent sampling, outputs, trust boundaries, and parameter selection. Attack studies can examine membership, reconstruction, or inference under specified adversary knowledge, but their measured success rates are conditional on that test. Both positive and negative findings should be reported with limitations and without asserting that the system is invulnerable (Hatamizadeh et al., 2023; Li et al., 2023; Shokri et al., 2017; Wang, Chang, et al., 2024).

8.2. Sensitivity and failure testing

A useful study would vary privacy budget, clipping bound, number of participants, model architecture, communication rate, and the amount of heterogeneity. It should also test participant dropout, stale updates, small cohorts, infrastructure interruptions, and malicious or malformed contributions. Results should include computational and communication costs and the operational effects of degraded performance, not only model accuracy. Asynchronous and personalized federated methods can inform such tests but do not remove the need for site-specific evaluation (Chen, Mao, & Ma, 2021; Huang et al., 2021; Zhang, Hong, et al., 2021).

Ablation studies can compare ordinary federation, secure aggregation, the specified DP mechanism, and their combination. The comparison must account for mechanisms that cannot be isolated cleanly: adding secure aggregation can change monitoring capability, and a DP release can change the utility of later decisions. A central-DP mechanism should be evaluated against its stated trusted server assumption; a local-DP mechanism has a different observer model and may have different accuracy. Conflating the two would produce an invalid comparison (Gibbs et al., 2026; Yin et al., 2022).

The business model should be stress-tested against different values of incremental gain, implementation cost, incident severity, and participation. Scenarios should include a no-deployment option and constraints that exclude an otherwise profitable architecture. The analysis should document the range over which a recommendation remains stable and distinguish an uncertainty range based on evidence from an arbitrary illustrative range. In a real deployment, an independent team should review the assumptions most likely to change the decision.

8.3. Limitations

The proposed framework is conceptual and does not establish the causal effect of federation on revenue, incident probability, customer trust, or equity. The illustrative monetary values and the single-release privacy example are constructed for clarity. The former does not estimate real risk; the latter is not a complete federated learning implementation or an end-to-end assurance statement. A multi-round system, a data-

dependent denominator, or a user appearing at multiple participants requires a new sensitivity and accounting analysis.

9. Conclusion

This article develops a framework for assessing the business value of privacy-preserving mobility analytics in platform operations. It connects forecasts to actual allocation decisions, compares incremental operational gains with implementation costs and modeled incident losses, and places privacy, security, authorization, and service requirements within a separate feasibility screen. The framework distinguishes raw-data localization, secure aggregation, and differential privacy, each of which protects a different boundary under different assumptions.

In the hypothetical monthly example, ordinary federation has the highest risk-adjusted value under the baseline assumptions, while a privacy-enhanced design becomes preferable when the modeled severity of incidents is sufficiently higher. The result changes again when all new alternatives have negative expected value or when a mandatory requirement removes an option. A separate Gaussian example shows how a one-release privacy bound depends on the contribution unit, adjacency, clipping, fixed denominator, and trusted observer. Neither example constitutes empirical validation.

For future application, platforms need defensible operational counterfactuals, documented privacy accounting, participant-level performance analysis, security testing, and governance that protects authorized use throughout the data lifecycle. The framework is intended to make those requirements and their economic implications explicit so that an organization can make a more informed decision about whether, where, and how to use federated mobility analytics.

Acknowledgement

We acknowledge the researchers whose published work informed this study. Any remaining errors are our responsibility.

Author Contributions

Conceptualization, J.H.D. and F.F.Y.; Methodology, J.H.D. and F.F.Y.; Formal Analysis, J.H.D. and F.F.Y.; Investigation, J.H.D.; Data Curation, J.H.D.; Software, J.H.D.; Validation, J.H.D. and F.F.Y.; Resources, J.H.D. and F.F.Y.; Visualization, J.H.D.; Writing – Original Draft, J.H.D.; Writing – Review & Editing, J.H.D. and F.F.Y.; Supervision, F.F.Y.; Project Administration, J.H.D. and F.F.Y. We have read and approved the final manuscript and take responsibility for the work.

Funding

This research received no external funding.

Conflicts of Interest

We declare no conflicts of interest.

Ethical Approval and Consent

This analytical study reports no human-participant, personal-data, or animal research. Any future study using personal mobility data should obtain the approvals and permissions applicable to that study.

Data Availability Statement

The numerical illustrations use only the hypothetical values and formulas reported in Sections 5 and 6 and Tables 2–3. No confidential platform dataset was used.

Declaration of Generative AI Use

Generative AI assisted us with drafting and editorial preparation. We take full responsibility for the accuracy of the content and references and for the final manuscript.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (pp. 308–318). https://doi.org/10.1145/2976749.2978318

Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., & Palamidessi, C. (2013). Geo-indistinguishability: Differential privacy for location-based systems. In Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security - CCS '13 (pp. 901–914). https://doi.org/10.1145/2508859.2516735

Arcolezi, H. H., Couchot, J. F., Renaud, D., Al Bouna, B., & Xiao, X. (2022). Differentially private multivariate time series forecasting of aggregated human mobility with deep learning: Input or gradient perturbation? Neural Computing and Applications, 34(16), 13355–13369. https://doi.org/10.1007/s00521-022-07393-0

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 1175–1191). https://doi.org/10.1145/3133956.3133982

Chen, G., Tong, M., Yin, J., Wang, M., Cao, J., & Wang, H. (2025). SecureGraphFL: A privacy-preserving and attack-resilient federated learning framework for traffic prediction. IEEE Internet of Things Journal, 12(21), 44988–44999. https://doi.org/10.1109/jiot.2025.3599568

Chen, L., Thakuriah, P., & Ampountolas, K. (2021). Short-term prediction of demand for ride-hailing services: A deep learning approach. Journal of Big Data Analytics in Transportation, 3(2), 175–195. https://doi.org/10.1007/s42421-021-00041-4

Chen, M., Mao, B., & Ma, T. (2021). FedSA: A staleness-aware asynchronous federated learning algorithm with non-IID data. Future Generation Computer Systems, 120, 1–12. https://doi.org/10.1016/j.future.2021.02.012

de Montjoye, Y. A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013). Unique in the crowd: The privacy bounds of human mobility. Scientific Reports, 3(1), 1376. https://doi.org/10.1038/srep01376

Dean, M. D., & Kockelman, K. M. (2021). Spatial variation in shared ride-hail trip demand and factors contributing to sharing: Lessons from Chicago. Journal of Transport Geography, 91, 102944. https://doi.org/10.1016/j.jtrangeo.2020.102944

Du, R., He, D., Song, Y., Ding, Z., Chan, S., & Li, X. (2025). Privacy-preserving short-term travel demand forecasting based on federated learning. IEEE Transactions on Vehicular Technology, 74(7), 11436–11449. https://doi.org/10.1109/tvt.2025.3547823

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In Theory of cryptography: Third Theory of Cryptography Conference (TCC 2006) (pp. 265–284). Springer. https://doi.org/10.1007/11681878_14

Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3-4), 211–487. https://doi.org/10.1561/0400000042

Ezequiel, C. E. J., Gjoreski, M., & Langheinrich, M. (2022). Federated learning for privacy-aware human mobility modeling. Frontiers in Artificial Intelligence, 5, 867046. https://doi.org/10.3389/frai.2022.867046

Feng, J., Rong, C., Sun, F., Guo, D., & Li, Y. (2020). PMF: A privacy-preserving human mobility prediction framework via federated learning. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 4(1), Article 10. https://doi.org/10.1145/3381006

Feng, X., Sun, H., Liu, S., Guo, J., & Zheng, H. (2024). Federated meta-learning on graph for traffic flow prediction. IEEE Transactions on Vehicular Technology, 73(12), 19526–19538. https://doi.org/10.1109/tvt.2024.3441759

Gibbs, H., Musolesi, M., Cheshire, J., & Eggo, R. M. (2026). Impact of federated data with local differential privacy for human mobility modeling. EPJ Data Science, 15(1), 17. https://doi.org/10.1140/epjds/s13688-025-00611-4

González, M. C., Hidalgo, C. A., & Barabási, A. L. (2008). Understanding individual human mobility patterns. Nature, 453(7196), 779–782. https://doi.org/10.1038/nature06958

Hatamizadeh, A., Yin, H., Molchanov, P., Myronenko, A., Li, W., Dogra, P., Feng, A., Flores, M. G., Kautz, J., Xu, D., & Roth, H. R. (2023). Do gradient inversion attacks make federated learning unsafe? IEEE Transactions on Medical Imaging, 42(7), 2044–2056. https://doi.org/10.1109/tmi.2023.3239391

Hong, J., Zhu, Z., Yu, S., Wang, Z., Dodge, H. H., & Zhou, J. (2021). Federated adversarial debiasing for fair and transferable representations. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (pp. 617–627). https://doi.org/10.1145/3447548.3467281

Hu, N., Liang, W., Zhang, D., Xie, K., Li, K., & Zomaya, A. Y. (2024). FedGCN: A federated graph convolutional network for privacy-preserving traffic prediction. IEEE Transactions on Sustainable Computing, 9(6), 925–935. https://doi.org/10.1109/tsusc.2024.3395350

Huang, W., Li, T., Wang, D., Du, S., Zhang, J., & Huang, T. (2022). Fairness and accuracy in horizontal federated learning. Information Sciences, 589, 170–185. https://doi.org/10.1016/j.ins.2021.12.102

Huang, Y., Chu, L., Zhou, Z., Wang, L., Liu, J., Pei, J., & Zhang, Y. (2021). Personalized cross-silo federated learning on non-IID data. Proceedings of the AAAI Conference on Artificial Intelligence, 35(9), 7865–7873. https://doi.org/10.1609/aaai.v35i9.16960

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., … Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1-2), 1–210. https://doi.org/10.1561/2200000083

Krueger, R., Bierlaire, M., & Bansal, P. (2023). A data fusion approach for ride-sourcing demand estimation: A discrete choice model with sampling and endogeneity corrections. Transportation Research Part C: Emerging Technologies, 152, 104180. https://doi.org/10.1016/j.trc.2023.104180

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/msp.2020.2975749

Li, Z., Wang, L., Chen, G., Zhang, Z., Shafiq, M., & Gu, Z. (2023). E2EGI: End-to-end gradient inversion in federated learning. IEEE Journal of Biomedical and Health Informatics, 27(2), 756–767. https://doi.org/10.1109/jbhi.2022.3204455

Liao, Z., Li, S., Zhao, Y., Liu, Y., Liang, W., & Wan, S. (2024). Predicting ride-hailing passenger demand: A POI-based adaptive clustering federated learning approach. Future Generation Computer Systems, 156, 168–178. https://doi.org/10.1016/j.future.2024.02.026

Liu, Y., Fan, Z., Song, X., & Shibasaki, R. (2021). FedVoting: A cross-silo boosting tree construction method for privacy-preserving long-term human mobility prediction. Sensors, 21(24), 8282. https://doi.org/10.3390/s21248282

Ma, H., Fang, F., & Parkes, D. C. (2022). Spatio-temporal pricing for ridesharing platforms. Operations Research, 70(2), 1025–1041. https://doi.org/10.1287/opre.2021.2178

Mattos, E. P. D., Domingues, A. C. S. A., Silva, F. A., Ramos, H. S., & Loureiro, A. A. F. (2024). Protect your data and I'll rank its utility: A framework for utility analysis of anonymized mobility data for smart city applications. Ad Hoc Networks, 163, 103567. https://doi.org/10.1016/j.adhoc.2024.103567

Rahman, M. H., & Rifaat, S. M. (2021). Using spatio-temporal deep learning for forecasting demand and supply-demand gap in ride-hailing system with anonymised spatial adjacency information. IET Intelligent Transport Systems, 15(7), 941–957. https://doi.org/10.1049/itr2.12073

Rocco di Torrepadula, F., Fisichella, M., Di Martino, S., & Mazzocca, N. (2025). FedFlow: A personalized federated learning framework for passenger flow prediction. Machine Learning, 114(7), 166. https://doi.org/10.1007/s10994-025-06795-0

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In 2017 IEEE Symposium on Security and Privacy (SP) (pp. 3–18). https://doi.org/10.1109/sp.2017.41

Shokri, R., Theodorakopoulos, G., Le Boudec, J. Y., & Hubaux, J. P. (2011). Quantifying location privacy. In 2011 IEEE Symposium on Security and Privacy (pp. 247–262). https://doi.org/10.1109/sp.2011.18

Song, C., Qu, Z., Blumm, N., & Barabási, A. L. (2010). Limits of predictability in human mobility. Science, 327(5968), 1018–1021. https://doi.org/10.1126/science.1177170

Wang, N., Zheng, L., Shen, H., & Li, S. (2024). Ride-hailing origin-destination demand prediction with spatiotemporal information fusion. Transportation Safety and Environment, 6(2), tdad026. https://doi.org/10.1093/tse/tdad026

Wang, Z., Chang, Z., Hu, J., Pang, X., Du, J., Chen, Y., & Ren, K. (2024). Breaking secure aggregation: Label leakage from aggregated gradients in federated learning. In IEEE INFOCOM 2024 - IEEE Conference on Computer Communications (pp. 151–160). https://doi.org/10.1109/infocom52122.2024.10621090

Wei, S., Zhang, Y., Zhou, Z., Zhang, T., & Xu, K. (2024). FedSM: A practical federated shared mobility system. Proceedings of the VLDB Endowment, 17(12), 4445–4448. https://doi.org/10.14778/3685800.3685896

Weng, S., Gou, Y., Zhang, L., & Imran, M. A. (2025). Evaluating privacy loss in differential privacy based federated learning. Future Generation Computer Systems, 172, 107848. https://doi.org/10.1016/j.future.2025.107848

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), Article 12. https://doi.org/10.1145/3298981

Yin, X., Zhu, Y., & Hu, J. (2022). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. ACM Computing Surveys, 54(6), 1–36. https://doi.org/10.1145/3460427

Yu, F., Xu, Z., Qin, Z., & Chen, X. (2022). Privacy-preserving federated learning for transportation mode prediction based on personal mobility data. High-Confidence Computing, 2(4), 100082. https://doi.org/10.1016/j.hcc.2022.100082

Zhang, C., Zhang, S., Yu, J. J. Q., & Yu, S. (2021). FASTGNN: A topological information protected federated learning approach for traffic speed forecasting. IEEE Transactions on Industrial Informatics, 17(12), 8464–8474. https://doi.org/10.1109/tii.2021.3055283

Zhang, X., Hong, M., Dhople, S., Yin, W., & Liu, Y. (2021). FedPD: A federated learning framework with adaptivity to non-IID data. IEEE Transactions on Signal Processing, 69, 6055–6070. https://doi.org/10.1109/tsp.2021.3115952

Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. Neurocomputing, 465, 371–390. https://doi.org/10.1016/j.neucom.2021.07.098