

Federated Diversified Rule Discovery for Privacy-Preserving Cross-Enterprise Data Analytics

Enzo Makinen

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
enzomakinen82@missouri.edu

Niklas Bay

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
niklas1978@binghamton.edu

Zhudong Xie

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence,
KS, USA.
zhudongxie137@ku.edu

Ishaan D. Martin

Department of Computer Science, University of Houston, Houston, TX, USA.
ishaanmartin660@uh.edu

Abstract

Cross-enterprise data collaboration has become indispensable for uncovering systemic insights in sectors such as healthcare, finance, and supply chain management. Rule discovery, long valued for its interpretability and operational relevance, is challenged when data silos must remain physically and legally separate. Federated learning offers a privacy-conscious paradigm, yet federated rule mining introduces distinctive tensions between pattern coverage, rule diversity, privacy guarantees, and cross-organizational governance. This paper presents a system-oriented framework for federated diversified rule discovery that intertwines local pattern extraction, secure aggregation, and multi-objective diversification mechanisms derived from embedding-guided top-k discovery techniques. The discussion centers on architectural trade-offs, including the cost of communication against the richness of global rule sets, the balance between local autonomy and global consistency, and the integration of differential privacy with aggregation protocols. Diversification is examined as a structural device for reducing redundancy, elevating actionability, and distributing inferential power across participants. Governance, fairness, and regulatory alignment are woven into the analysis, emphasizing data sovereignty, incentive design, and the right to contest algorithmic outputs. Further investigation covers robustness under adversarial and heterogeneous conditions, sustainable deployment in resource-constrained environments, and the systemic policy implications of large-scale federated rule analytics. The work does not propose a specific algorithm but rather articulates a reference architecture and critical evaluation lens for constructing federated rule discovery pipelines that are simultaneously diverse, privacy-preserving, and institutionally sustainable.

Keywords

Federated learning, rule discovery, diversification, privacy preservation, cross-enterprise analytics, governance, fairness, system architecture.

1. Introduction

Contemporary data ecosystems are characterized by valuable information distributed across organizational boundaries. Healthcare consortia seek to identify clinical pathways from electronic health records held by separate hospitals; financial alliances aim to detect transnational fraud patterns; and manufacturing supply chains require joint analysis of maintenance logs and quality metrics from independent suppliers. The collaborative mining of such multiparty data can yield interpretable rules that drive high-stakes decisions, yet the predominant architectural assumption of centralized data curation founders against privacy regulations, trade-secret protections, and the sheer reluctance of enterprises to relinquish control over sensitive assets. Federated learning has surfaced as a foundational response, enabling the co-construction of statistical models without raw data leaving local premises [1]. Transferring the federated ethos to rule discovery, however, introduces complexities beyond model averaging, because rules are discrete, symbolic artifacts whose semantic coherence can dissolve during naive aggregation.

The classical knowledge discovery literature has established association rule mining as a pillar of descriptive analytics, providing human-readable implications such as “if condition X holds, then outcome Y is likely” [2]. The interpretability of rules makes them particularly appealing in regulated domains where decisions must be justified. Nevertheless, traditional rule mining algorithms operate under the assumption that a single data custodian has unrestricted access to a complete table, often leading to a deluge of redundant or trivial patterns when applied without diversification controls. Furthermore, the release of exact rule statistics from individual enterprises can leak membership information or business-sensitive distributions, calling for rigorous privacy-preserving instruments. Differential privacy has been formalized as a rigorous mathematical guarantee that bounds the influence of any single data record on released statistics [3], and secure aggregation protocols have enabled servers to compute sums over client updates without inspecting individual contributions [4]. Early adaptations to privacy-sensitive association rule mining demonstrated feasibility for vertically partitioned data by using cryptographic protocols [5], underscoring both the potential and the computational cost of distributed rule inference.

The productivity of cross-enterprise rule analytics hinges on overcoming two interrelated deficits: the lack of diversity in discovered rule sets and the tension between transparency and privacy. Rule diversity ensures that the collective output spans disparate regions of the data space, avoids near-duplicate articulations of the same underlying phenomenon, and provides a variegated portfolio of actionable hypotheses. Meanwhile, embedding-based diversification strategies recently introduced in centralized settings have shown that user-guided vector representations can dramatically accelerate and focus the search for diverse top-k rules [6]. Extending this philosophy to a federated topology requires a systemic rethinking of how local embedding signals are shared, how global diversity is negotiated when participants have heterogeneous data distributions, and how the overall infrastructure can satisfy cross-jurisdictional data governance frameworks. This paper engages precisely these system-level dimensions, offering a comprehensive analysis of federated diversified rule discovery as an architectural problem rather than an algorithmic one. The discussion is organized around structural trade-offs, governance imperatives, privacy integration, robustness, fairness, and

long-term sustainability, aiming to equip researchers and practitioners with a holistic blueprint for designing and evaluating such systems.

2. Background and Related Work

Federated rule discovery sits at the intersection of several mature research streams. Federated learning originally targeted communication-efficient training of deep neural networks by exchanging model updates instead of raw data [1], and its landscape has since expanded to encompass a broad family of decentralized computation schemes involving heterogeneous clients, non-identically distributed data, and adversarial environments. The security of such systems relies heavily on secure aggregation, which computes sums of client-provided vectors while masking individual inputs, often through additive secret sharing and pairwise blinding [4]. Differential privacy is layered on top to prevent inference attacks on aggregated outputs; the sensitivity of a query determines the magnitude of noise calibrated according to a privacy budget, and composition theorems track cumulative leakage across multiple releases [3], [8]. These primitives are equally applicable to statistical queries about rule support or confidence, yet their direct application to complex discrete structures demands careful sensitivity analysis, because a single transaction can alter the support of many candidate rules simultaneously.

Pattern discovery has historically been framed as a search for frequent itemsets and their transformation into association rules [2]. Subsequent research expanded the taxonomy to include decision rules, subgroup discovery, and sequence rules, all sharing the goal of producing human-comprehensible conditional statements. The sheer volume of patterns unearthed by exhaustive mining quickly motivated research on condensed representations, such as closed itemsets and minimal generators, but these still leave room for semantic redundancy. Diversification arose as a post-processing or search-integrated mechanism to select subsets of rules that jointly maximize some measure of novelty or disjoint coverage. Recent work has proposed embedding-based guidance, in which items, transactions, and user preferences are projected into a continuous vector space, enabling fast retrieval of diverse top- k rules without exhaustive enumeration [6]. This development demonstrates that diversity can be treated as a first-class optimization objective rather than an afterthought, but it assumes centralized access to the full transactional corpus, a condition that rarely holds in cross-enterprise consortia.

Privacy-preserving data mining has a long lineage of cryptographic protocols for rule induction across multiple parties holding horizontally or vertically partitioned tables. For horizontally partitioned data, where enterprises possess the same attribute schema but disjoint rows, secure set intersection and secure sum computations have been employed to compute global supports without revealing individual transaction details [11]. These approaches, while theoretically robust, often incur substantial communication overhead and do not naturally incorporate diversification constraints. More modern federated analytics frameworks incorporate differential privacy directly into the computation of frequent patterns, using local or distributed noise addition to protect individual participation. The federated learning surveys summarize the broader system challenges, including stragglers, bandwidth constraints, and client selection, which equally afflict federated rule mining [7], [9], [10]. The gap, therefore, lies in the absence of an integrated system architecture that simultaneously addresses rule diversity, privacy, interpretability, and cross-organizational governance within a single design philosophy.

3. System Architecture for Federated Diversified Rule Discovery

An effective federated diversified rule discovery framework must orchestrate a loose federation of autonomous data controllers, each operating on its own proprietary transactional store, with a logically centralized but physically distributed coordination service. The architecture distinguishes three functional layers: local rule mining and embedding, secure aggregation and diversification orchestration, and global rule curation and delivery. At each enterprise, a local analytics engine derives candidate rules alongside auxiliary statistics such as support, confidence, lift, and an embedding representation of the rules in a shared semantic space. This embedding space is crucial for the subsequent phase of diversification, because it encodes the proximity and overlap among patterns, allowing the system to quantify diversity as the dispersion of rule vectors. The embedding model itself may be pre-trained on external corpora or co-learned across participants through a federated embedding alignment protocol, which itself must respect privacy budgets.

The second layer conducts privacy-protected aggregation of rule metadata. Instead of transmitting raw rules, which could inadvertently disclose proprietary combinations or rare event patterns, each client submits differentially private summaries, such as noisy support counts or aggregated embedding centroids, through a secure aggregation channel. The aggregator then merges these contributions to form a global candidate pool and applies diversification algorithms that select a representative set of rules optimizing a composite score that balances statistical strength, novelty, and inter-rule dissimilarity. The diversification step can be informed by user-guided embeddings that allow domain experts to steer the search toward specific facets, such as clinical outcomes or financial risk indicators, thereby aligning the global output with consortium-wide analytic goals. Communication topology can vary from flat star architectures, where a central server orchestrates rounds, to hierarchical federations that aggregate within sub-networks before a final merge, a design that reduces latency and accommodates jurisdictional data flow restrictions.

Several structural trade-offs emerge at this layer. Increasing the frequency of federation rounds improves the freshness of the global rule set but strains bandwidth and amplifies the cumulative privacy loss, demanding careful privacy accounting via the composition properties of differential privacy. Clients with strongly skewed local distributions may contribute rules that are globally insignificant yet locally valuable; the system must decide whether such rules should be preserved as localized annotations or filtered out to satisfy a uniform diversity criterion. Likewise, the richness of the embedding space directly influences the fidelity of diversification: low-dimensional spaces may collapse semantically distinct rules, while high-dimensional spaces heighten communication costs and complicate private embedding alignment. These tensions are not resolvable through a one-size-fits-all parameterization; they require the architecture to expose configurable policies that organizations can negotiate during consortium formation. Interpretable rule models built on Bayesian or decision-list formalisms can further serve as quality filters, ensuring that only statistically robust and human-readable rules propagate to the global stage [12].

4. Diversification Mechanisms and Quality Assurance

Diversification in federated rule discovery is a multi-objective endeavor that must reconcile statistical merit with semantic coverage. Unchecked, a federated aggregation of locally frequent rules tends to converge on a small set of high-support patterns that characterize dominant cultural or demographic majorities in the data, thereby suppressing rare yet critical rules about minority subgroups or emerging anomalies. Embedding-guided diversification counters this by encouraging the selection of rules whose vector representations are

maximally apart while still satisfying minimal quality thresholds. The approach builds on the insight that rule antecedents, consequents, and their relational bindings can be encoded as dense vectors where spatial distance approximates conceptual dissimilarity [6]. In the federated extension, the global coordinator can construct a joint embedding space either by aggregating locally learned projections or by sharing a common projection matrix that enterprises adapt through localized fine-tuning, a process that itself warrants privacy budgeting.

Quality assurance extends beyond diversification scores. Rules must be locally faithful, globally interpretable, and immune to spurious correlations that arise from distributional shifts across sites. One practical mechanism is to require each candidate rule to pass a stability test: its support and confidence should exhibit bounded variation across a subset of test nodes that were not used in the mining phase. The system can inject controlled noise into these statistics during aggregation and use hypothesis testing frameworks to decide whether a pattern meets a consortium-defined significance level. Fairness-aware rule construction adds another dimension; rules that inadvertently encode discriminatory biases, such as associating protected attributes with unfavorable outcomes, must be flagged or suppressed. Fairness constraints can be integrated directly into the diversification objective by penalizing rule sets that exhibit differential validity or disparate impact across demographic slices, extending principles from fair classification to descriptive pattern mining [13]. The architecture should support configurable fairness metrics that align with the legal and ethical standards of the participating jurisdictions, acknowledging that fairness is inherently contextual and contested.

A further nuance involves the temporal stability of rules. In dynamic enterprise environments such as fraud detection or predictive maintenance, a rule that is diverse and fair today may become stale or even misleading as data distributions drift. Federated diversified rule discovery pipelines therefore benefit from continuous monitoring modules that track concept drift across clients and trigger selective re-issuance of mining rounds only when the embedding topology shifts significantly. This event-driven refresh strategy conserves communication and privacy budgets while preserving the relevance of the discovered rule portfolio.

5. Privacy and Security Foundations for Cross-Enterprise Deployment

Privacy preservation in federated rule discovery must address both the inference threats during the aggregation protocol and the risks incurred by publishing the final set of rules. The threat surface includes honest-but-curious servers that attempt to reconstruct individual transactions from aggregated rule statistics, malicious clients that submit manipulated data to steer the global rule set, and external adversaries that correlate the released rules with auxiliary information. Differential privacy, applied at the client level, provides a principled defense: each enterprise adds carefully calibrated noise to its local support counts or rule embeddings, rendering any single transaction's presence indistinguishable up to a privacy parameter epsilon. The overall privacy expenditure is tracked across rounds, and the system can enforce a global privacy budget by halting further participation when the cumulative loss exceeds a consortium-agreed threshold. Deep learning with differential privacy [8] and foundational expositions of the privacy framework [10] offer concrete noise calibration strategies that can be adapted to discrete rule statistics by analyzing the sensitivity of counting queries over itemsets.

Secure aggregation protocols prevent the central server from ever observing individual noisy updates; they compute the sum of client contributions over a masked cryptographic protocol

that reveals only the aggregate. In federated rule discovery, this sum might represent the total support of a rule across all enterprises, which, together with the sum of transaction counts, allows the global coordinator to estimate confidence without accessing individual enterprise profiles. The security of such schemes can be bolstered by verifiable computation techniques that allow clients to audit that the aggregation was performed correctly without exposing their masked inputs. However, secure aggregation alone does not defend against data poisoning attacks, where a malicious participant submits artificially distorted rule statistics to promote or suppress certain patterns. Robust aggregation methods, such as coordinate-wise median or trimmed-mean operators, can mitigate isolated poisoning but must be carefully integrated with diversification to avoid inadvertently discarding minority-interest rules that merely appear anomalous. Adversarial federated learning studies have demonstrated that backdoor insertion through rule manipulation is a realistic threat, because a few crafted transactions can bias global hypothesis spaces [14]. The system architecture should therefore incorporate anomaly detection on client contribution vectors and support transparent audit logs that enable forensic analysis without violating privacy.

Heterogeneity across enterprises further complicates the privacy-utility landscape. Clients with vastly different dataset sizes or attribute granularities may require individualized privacy budgets, and the noise addition strategy must account for the fact that a rule that is rare globally may be fully supported by a single enterprise, creating a high-sensitivity disclosure risk. Techniques such as personalized differential privacy or per-client sampling allow the system to tailor protection levels to data sensitivity and organizational risk appetite. Cross-silo federated learning, where the number of participants is small to moderate, permits more sophisticated trust models, including the use of trusted execution environments at the aggregator to perform plaintext diversification on encrypted inputs, though this introduces hardware dependency and supply-chain trust assumptions. The choice of trust model—cryptographic, statistical, or hardware-based—is a systemic decision that shapes the entire architecture and must be negotiated early in the consortium’s lifecycle [15].

6. Governance, Fairness, and Policy Implications

Governance arrangements in federated diversified rule discovery are as critical as the underlying technology, because they define how decision rights, accountability, and value are distributed among participants. A cross-enterprise consortium must agree upon protocols for data contribution, rule ownership, intellectual property derived from discovered patterns, and mechanisms for dispute resolution. The economic perspective emphasizes that privacy is an intermediate good whose protection and disclosure involve complex trade-offs, and institutions are more likely to participate when they perceive a favorable ratio of insight gain to information leakage [16]. Incentive-compatible designs can align contributions with rewards: enterprises that contribute high-quality rules or rare data patterns that prove commercially valuable might receive monetary compensation, reduced membership fees, or privileged access to consortium-curated analytics. Contract-theoretic incentive mechanisms developed for federated learning can be adapted to the rule discovery context by mapping rule utility to participant contributions [19].

Fairness extends beyond technical metrics into procedural and distributive dimensions. Procedural fairness requires that the rule discovery process be transparent, contestable, and free from coercion, meaning that any enterprise or data subject whose behavioral patterns are encoded in the resulting rules must have avenues for redress. The right to explanation enshrined in data protection regulations such as the General Data Protection Regulation

establishes that individuals subjected to automated decisions based on rules should be able to understand and challenge the logic involved, though the scope of this right remains contested [17]. Federated rule systems can address this by exposing rule provenance trails that trace a global rule back to the aggregation of noisy local statistics without revealing individual-level data, thus providing a form of algorithmic due process. Fairness audits conducted by independent oversight bodies can evaluate whether the discovered rule portfolio perpetuates systemic biases against protected groups, and the system can be mandated to re-run diversification with adjusted objective weights if disparities exceed predefined thresholds.

Environmental sustainability is an increasingly prominent governance concern. Federated analytics involve iterative rounds of communication, model retraining, and heavy cryptographic operations that collectively impose a significant carbon footprint. Energy-aware scheduling, sparse embedding compression, and the use of green data centers can mitigate this impact, yet the policy dimension requires that consortia establish sustainability reporting standards that capture the full life-cycle energy cost of federated rule discovery projects [18]. Data sovereignty considerations demand that rule discovery pipelines respect geographical restrictions, such as prohibitions on transferring raw data across borders, a requirement that can be met by colocating aggregation servers within each region and performing hierarchical diversification that first pools local rules within jurisdictions before inter-regional merge. These governance layers convert the technical architecture into a socio-technical infrastructure that is operationally viable and legally defensible.

7. Robustness, Sustainability, and Deployment Considerations

Deploying federated diversified rule discovery in production environments surfaces robustness challenges that are often underestimated in laboratory prototypes. System components must withstand client churn, where enterprises join or leave the federation dynamically, and must handle straggling clients whose computational delays can stall global rounds. Asynchronous aggregation protocols, where the coordinator updates the global rule set whenever a quorum of contributions arrives, offer resilience but complicate the diversification pass, because the candidate pool evolves in real time and embedding alignment may drift. The system can maintain a bounded staleness guarantee through versioned rule snapshots that capture the state after each refresh cycle, allowing clients that later contribute to retroactively validate how their input influenced the global outcome.

Large-scale system design lessons from federated learning indicate that operational robustness requires device selection strategies that prioritize clients with sufficient computational capacity and network quality, a practice that can inadvertently exclude smaller enterprises with limited infrastructure [20]. To counter this, the architecture can support tiered participation: lightweight clients submit coarse-grained rule summaries without embedding computations, while heavyweight clients perform full local diversification and embedding. This tiered approach broadens inclusion and ensures that the federation does not degenerate into a platform dominated by a few resource-rich actors. Model and embedding compression, using quantization or sparsification, can dramatically lower communication payloads while retaining sufficient fidelity for downstream diversification.

Sustainability also involves the long-term maintainability of the embedding models and rule ontologies. As business terminologies evolve, the semantic space in which rules are embedded must adapt, and the system should schedule periodic re-indexing of the embedding space using federated fine-tuning. Version control of rule sets, akin to software artifact management, emerges as a necessity in regulated industries where decisions based on

obsolete rules could have legal repercussions. Auditability is enhanced by immutably logging the diversification decisions, privacy budgets consumed, and consensus outcomes on a distributed ledger, though the energy footprint of such logging must be weighed against trust benefits. Forward-looking configurations can envision self-tuning federations that use meta-learning to automatically adjust diversification thresholds, privacy budgets, and refresh frequencies in response to observed rule utility and drift metrics, gradually reducing human governance overhead while preserving alignment with consortium values.

8. Conclusion

Federated diversified rule discovery represents a confluence of systems thinking, privacy engineering, and organizational governance that extends well beyond simple algorithmic adaptation. This paper has articulated the architectural spine of such systems, emphasizing how local rule mining, embedding-guided diversification, and secure aggregation can be woven into a resilient fabric that respects data sovereignty while generating high-value, interpretable cross-enterprise knowledge. Structural trade-offs between communication cost, privacy leakage, and diversity resolution pervade every design decision, and the analysis underscored that diversification is not merely a post-processing filter but a central coordinating principle that shapes the entire pipeline, from local embedding alignment to global fairness constraints. Privacy preservation relies on a layered strategy integrating differential privacy, secure aggregation, and robustness to poisoning, all tuned to the distinctive sensitivity profile of discrete rule statistics. Governance mechanisms, including incentive alignment, procedural fairness, sustainability accounting, and the right to explanation, are not peripheral add-ons but constitutive elements that determine whether a federated rule discovery infrastructure will gain and sustain the trust of participating enterprises and the publics they serve. The deployment outlook calls for modular, tiered architectures that accommodate heterogeneous resources, dynamic membership, and evolving regulatory landscapes, while embedding versioning and auditability as first-class features. Future systems will need to navigate the tension between automated self-optimization and human oversight, potentially through federated meta-learning loops that continuously refine diversification policies. As cross-enterprise data collaboration becomes increasingly indispensable for tackling complex societal challenges, the principles outlined here offer a foundation for building federated analytics ecosystems that are simultaneously intelligent, responsible, and institutionally durable.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics* (pp. 1273–1282). PMLR.
2. Agrawal, R., & Srikant, R. (1994). Fast algorithms for mining association rules. In *Proceedings of the 20th International Conference on Very Large Data Bases* (pp. 487–499). Morgan Kaufmann Publishers Inc.
3. Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference* (pp. 265–284). Springer, Berlin, Heidelberg.
4. Bonawitz, K., Ivanov, V., Kreuter, Vaidya, J., & Clifton, C. (2002). Privacy preserving association rule mining in vertically partitioned data. In *Proceedings of the eighth ACM*

SIGKDD international conference on Knowledge discovery and data mining (pp. 639–644). ACM.

5. Han, Z., Chen, W., Han, Y., Mao, R., & Qin, J. (2026). Fast Diversified Top-k Rule Discovery via User-Guided Embeddings. *IEEE Transactions on Knowledge and Data Engineering*.
6. 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.
7. 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). ACM.
8. 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.
9. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
10. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
11. Letham, B., Rudin, C., McCormick, T. H., & Madigan, D. (2015). Interpretable classifiers using rules and Bayesian analysis: Building a better stroke prediction model. *Annals of Applied Statistics*, 9(3), 1350–1371.
12. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
13. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics* (pp. 2938–2948). PMLR.
14. Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM workshop on artificial intelligence and security* (pp. 1–11). ACM.
15. Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492.
16. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2), 76–99.
17. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). ACL.
18. Kang, J., Xiong, Z., Niyato, D., Xie, S., & Zhang, J. (2019). Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory. *IEEE Internet of Things Journal*, 6(6), 10700–10714.

19. Nishio, T., & Yonetani, R. (2019). Client selection for federated learning with heterogeneous resources in mobile edge. In ICC 2019-2019 IEEE International Conference on Communications (pp. 1–7). IEEE.